Optimal Intermittent Particle Filter

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

Abstract—The problem of the optimal allocation (in the expected mean square error sense) of a measurement budget for particle filtering is addressed. We propose three different optimal intermittent filters, whose optimality criteria depend on the information available at the time of decision making. For the first, the stochastic program filter, the measurement times are given by a policy that determines whether a measurement should be taken based on the measurements already acquired. The second, called the offline filter, determines all measurement times at once by solving a combinatorial optimization program before any measurement acquisition. For the third one, which we call online filter, each time a new measurement is received, the next measurement time is recomputed to take all the information that is then available into account. We prove that in terms of expected mean square error, the stochastic program filter outperforms the online filter, which itself outperforms the offline filter. However, these filters are generally intractable. For this reason, the filter estimate is approximated by a particle filter. Moreover, the mean square error is approximated using a Monte-Carlo approach, and different optimization algorithms are compared to approximately solve the combinatorial programs (a random trial algorithm, greedy forward and backward algorithms, a simulated annealing algorithm, and a genetic algorithm). Finally, the performance of the proposed methods is illustrated on two examples: a tumor motion model and a common benchmark for particle filtering.

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

I. Introduction

STOCHASTIC nonlinear dynamical systems have shown their ability to model a number of real-world problems [1], [2], [3]. Particle filtering is a popular approach to estimate the state of such systems from a set of noisy measurements [4]. This tool has been largely used, among others, in computer vision [5], [6], [7], positioning and navigation [8], [9], chemistry [10], [11], mechanics [12], robotics [13], and medicine [14]. In practice, performing measurements may be difficult due to energy consumption, economic 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 radiation [15].

Under such constraints, the problem is to select the best times to measure the system. 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 square error* (MSE) over the complete time horizon.

Intuitively, there are various reasons why regular measurement times may be suboptimal. For example, if the system is (almost) deterministic but there is a large uncertainty in the initial state, it may be better to take measurements as early as possible. However, if the signal amplitude varies greatly over time, it may be worthwhile to take measurements when the signal-to-noise ratio is greatest.

First, the problem of optimal measurement budget allocation is formalized as a nonlinear multistage *stochastic program*. The measurement times are selected one after the other and each measurement time is selected taking the measurements already observed into account.

Next, an easier related problem that we call the *offline program* is presented. It corresponds to selecting all the measurement times at once and offline, i.e., before any measurement acquisition. In the case of real-time applications where the time between each measurement is short, being able to calculate the measurement times offline can be decisive. However, filtering a signal using the measurement times given by this offline program gives, on average, poorer filtering performance because in that case, the information present in the measurements already acquired is ignored.

Then, we develop an online adaptation of the Offline problem which we call the *online program*. This gives better filtering performance, at the cost of a higher computational cost, as computations must be carried out online.

A. Related work

In recent decades, the problem of estimating the state of a system subject to missing measurements has been extensively studied in both the linear case (see e.g., [16], [17], [18] and references there in) and the nonlinear case (see e.g., [19], [20], [21] and references there in). The latter use the particle filtering framework. In these previous studies, missing measurements occur randomly. In this paper, on the contrary, we aim to choose the measurement times optimally. The measurement times (or equivalently, the missing measurement times), are, for our case, parameters to be optimized.

The *sensor scheduling* (or sensor management) *problem* is closely related to ours and has been widely studied. It consists in choosing at each time step one (or a few) active sensors among a set in order to minimize the variance of the estimation error [22], [23], [24], [25], [26], [27], [28], [29], [30], [31].

In some of the works on the sensor scheduling problem [23], [25], [30], [31], a cost is (or can be) associated to the use of each sensor and the objective is to minimize a combination of the error variance and the cost of the sensors. In this case, we speak of the sensor scheduling problem *with a sensing cost*.

The problem studied in this article is a constrained version of the sensor scheduling problem without sensing cost. Indeed,

consider the sensor scheduling problem where an active sensor must be selected among two available sensors, one corresponding to a measurement acquisition and the other corresponding to no measurements, i.e., it returns a value independent of the system state. By adding a constraint on the maximum number of uses of the measuring sensor, we find the problem of optimal allocation of a measurement budget.

In addition, the problem we address can be reduced to the sensor scheduling problem with a sensing cost in the following way. Consider the sensor scheduling problem with a sensing cost with two available sensors: the first one is equivalent to not measuring (this sensor is costless but returns a value independent of the system state); the second sensor has a cost $c > 0$. Then by choosing a sensing cost c such that the obtained solution respects the measurement budget (e.g., by using a binary search on c), one can reduce our constrained problem to a sensor scheduling problem with a sensing cost.

In the linear and Gaussian case, [22] proposes two algorithms to solve the sensor scheduling problem. The first one provides an optimal solution but can be computationally expensive. The second one provides a suboptimal solution and depends on a tuning parameter to make a trade-off between the quality of the solution and the simplicity of the problem. In that setting, the estimate is given by the Kalman filter. Using the monotonicity and concavity of the Riccati equation, a condition for the non-optimality of the initialization of a schedule is derived. The exact algorithm uses this condition to prune the search tree of all possible sensor schedules. On the other hand, the suboptimal algorithm uses a relaxed condition, and the optimality gap is proven to be upper bounded by a linear function of the relaxation parameter. Given the central role played by the Riccati equation in the proposed method, this approach seems difficult to generalize to the nonlinear and non-Gaussian case. For a literature review on the sensor scheduling problem in the linear and Gaussian case, see references in [22].

The sensor scheduling problem with a sensing cost has received attention when the state space is a finite set, i.e., the dynamics is a finite hidden Markov model (HMM) [23]. In that article, the estimate is computed using a HMM state filter and the optimal sensor scheduling policy is obtained by stochastic dynamic programming. They propose two efficient methods to obtain a suboptimal solution: the first one is a one-step look-ahead approximation, the other one is based on Lovejoy's algorithm for partially observable Markov decision processes. The proposed methods rely on the fact that the state space and the measurement space are finite sets, which is not the case in the problem we study.

In the case of a continuous state space, [24] proposes a Gaussian process global optimization method used as a black-box optimizer to dynamically select sensors in order to minimize the error at the end of a time horizon. The approach is model free (the system dynamics is unknown) and only the last measurement is used to estimate the state of the system, the previous measurements being used to choose the last active sensor. On the contrary, in our work, the (known) system dynamics is used to choose when to measure. Another difference is that we want to minimize the average MSE.

For nonlinear but Gaussian dynamics, a sparsity promoting approach is proposed in [25] to choose a set of active sensors at each time step. A trade-off is made between the number of active sensors and the current posterior Cramér-Rao lower bound (PCRLB). This is justified by the fact that the PCRLB is a lower bound on the mean square estimation error. The main differences with our approach are that we tackle the general case of a non-Gaussian dynamics, we do not consider a one-step look-ahead approximation, and we directly minimize the MSE instead of the PCRLB.

The sensor scheduling problem has also been studied when the set of sensors is continuous (e.g., to optimize the spatial position of a mobile sensor). For this problem, [26], [27] propose an algorithm using a simulation-based gradient approximation. Since the core of the method is based on the gradient of the objective function with respect to the selected sensor, this approach is not applicable when the set of available sensors is finite.

For the sensor scheduling problem in the nonlinear and non-Gaussian case, an optimal one-step look-ahead policy is proposed in [28], [29] to maximize the Kullbach Leibler divergence between filtering and prediction densities. Intuitively, it maximizes the information obtained at the current measurement time. Again, the main difference with this work is that we do not consider the one-step look-ahead approximation, and we minimize the average MSE.

In [30], [31], the sensor scheduling problem with a sensing cost for nonlinear and non-Gaussian dynamics is formalized as a continuous partially observable Markov decision process. The Q-function is approximated by the policy rollout method. This requires a base policy, i.e., a heuristic, to choose future actions in order to estimate the cumulated costs that will follow the next action. For example, in [30], the available sensors are spatially distributed to track a moving target and the base policy selects the sensor closest to the estimated position of the target. As the authors acknowledge "The choice of a base policy may have a significant impact on the performance of the rollout policy." On the contrary, our method does not require a priori knowledge of a heuristic to select the measurement times.

Table I summarizes the state of the art concerning the sensor scheduling problem.

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 [32] and continuous-time [33], [34] settings.

B. Contribution

This paper addresses the problem of optimal allocation of a measurement budget in the discrete-time nonlinear case with disturbance and measurement noise processes following arbitrary distributions.

In [35], we have proposed an initial intermittent filter (which is reported as Problem 2 here). Our contributions in the present paper are that (i) we show that the problem of optimal intermittent filtering can be modeled as a stochastic program

TABLE I: Summary of the state of the art of the sensor scheduling problem. When the sensors are chosen by a one-step look-ahead policy, the mention "(one step)" is indicated in the column "Objective function".

Dynamics	Sensing cost	Objective function	References
linear and Gaussian	no	average MSE	[22]
hidden Markov model	yes	average MSE	[23]
model-free	no	terminal root MSE	[24]
nonlinear and Gaussian	yes	posterior Cramér-Rao lower bound (one step)	[25]
nonlinear and non-Gaussian	no	average MSE	[26] [27]
		Kullback-Leibler divergence between prediction and filtering densities (one-step)	[28] [29]
	yes	average MSE	[30] [31]

(see Problem 1); (ii) based on this, we propose a recursive version of optimal intermittent filter (see Problem 3); (iii) we prove the Theorem 1 that states that the expected mean square filtering error is smaller for the stochastic program filter than for the online filter, which itself is smaller than for the offline filter; (iv) new optimization algorithms are tested (greedy forward, greedy backward, and simulated annealing); and (v) the results are presented (among others) on a new model inspired by real-world tumor tracking applications.

Overall, the two main contributions of this paper are to propose an efficient algorithm to solve the problem of optimal measurement times selection; and to show the interest of intermittent measurements in particle filtering.

C. Paper outline

The rest of this paper is organized as follows: Section II presents how to define filtering with intermittent measurements (Subsection II-A); different versions of optimal intermittent filters (Subsection II-B) and how to compute them numerically with different optimization algorithms (subsections II-C and II-D). Results and discussion are presented in Section III, where a model of tumor motion is proposed for benchmarking (Subsection III-A1), in addition to a common benchmark for particle filters (Subsection III-A2), the optimization algorithms are compared (Subsection III-B), and the filtering performance of the offline and online particle filters are compared with a regular particle filter (Subsection III-C). Finally, Section IV concludes and discusses possible improvements and perspectives.

A MATLAB (MathWorks, Natick, Massachusetts, USA) implementation of the presented algorithms and the code that generate the figures is available on GITHUB at GITHUB.COM/AMAURYGOUVERNEUR/OPTIMAL_MEASUREMENT_TIMES_FOR_PARTICLE_FILTERING.

II. MATERIALS AND METHODS

A. Intermittent filter

A discrete time stochastic nonlinear dynamic system describes the evolution of a state $x(t)$ over the finite time horizon $t = 0, \ldots, T$. 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 precisely, only N measurements are available, at times $t_1, \ldots, t_N \in \{0, \ldots, T\}$. This is modelled as

$$x(t+1) = f_t(x(t), w(t)) \quad \text{for } t = 0, \ldots, T-1, \quad (1)$$
$$y(t_j) = g_{t_j}(x(t_j), v(t_j)) \quad \text{for } j = 1, \ldots, N, \quad (2)$$
$$z(t) = h_t(x(t)) \quad \text{for } t = 0, \ldots, T, \quad (3)$$
$$x(0) \sim \mathcal{F}, \quad (4)$$

with $t_i \neq t_j$ if $i \neq j$, and $x(t) \in \mathbb{R}^n$, $y(t) \in \mathbb{R}^m$ 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 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 3-dimensional space.

Let us introduce some notations and definitions. We use the so-called Matlab notation, for $j \leq k$, $t_{j:k} := \{t_j, t_{j+1}, \ldots, t_k\}$ and $y(t_{j:k}) := \{y(t_j), y(t_{j+1}), \ldots, y(t_k)\}$.

We define the *intermittent filter* estimate,

$$\hat{z}(t|y(t_{1:j})) := \mathbb{E}_{x(0), w(0), \ldots, w(t-1)}[z(t)|y(t_k), \forall t_k \leq t], \quad (5)$$

where $\mathbb{E}_X[\cdot]$ holds for the expectation operator according to the random variable X. The name "intermittent" emphasizes that measurements are not accessible at each time step. Subsection II-C1 will explain how an intermittent particle filter can be used to calculate an approximation of the intermittent filter estimate.

Let's define the expected filtering error variance,

$$\mathcal{E}[t|y(t_{1:j})] :=$$
$$\mathbb{E}_{x(0), w(0), \ldots, w(t-1)}[\|z(t) - \hat{z}(t|y(t_{1:j}))\|^2 | y(t_k), \forall t_k \leq t], \quad (6)$$

where $\| \cdot \|$ is the Euclidean norm. A numerical approach to estimate this quantity is described in Paragraph II-C2.

Remark 1. *Thanks to the condition $\forall t_k \leq t$ in definitions (5) and (6), it follows that if $t_k \leq t < t_{k+1} \leq t_j$, it holds that $\hat{z}(t|y(t_{1:j})) = \hat{z}(t|y(t_{1:k}), y(t_{k+1:j})) = \hat{z}(t|y(t_{1:k}))$, and $\mathcal{E}[t|y(t_{1:j})] = \mathcal{E}[t|y(t_{1:k}), y(t_{k+1:j})] = \mathcal{E}[t|y(t_{1:k})]$. This ensures the causality of the filter by avoiding future measurements of influencing the present estimation. Otherwise, these two quantities are random variables because they depend on $y(t_1), \ldots, y(t_j)$ and therefore on $v(t_1), \ldots, v(t_j)$, which are random.*

Remark 2. *Instead of the estimator of $z(t)$ defined in* (5), *one could use any estimator that can be defined from the distribution of $x(t)|(y(t_k), \forall t_k \leq t)$, for example, the* maximum a posteriori estimator *could replace* (5) *using the method from [36]. This flexibility is allowed because we will use a particle filter which estimates the whole distribution of $x(t)|(y(t_k), \forall t_k \leq t)$.*

Moreover, since we are going to use black-box optimization algorithms, another quality criterion could be used instead of the error variance (6), *for example, the squared Euclidean norm could be replaced by the $1-$norm or the $\infty-$norm.*

B. Optimal intermittent filters

An *optimal intermittent filter* is an intermittent filter whose measurement times have been chosen according to a certain optimality criterion. In this section, we present three different optimality criteria for selecting the measurement times. Each optimality criterion induces a different optimal intermittent filter.

1) Stochastic program filter: In the stochastic program filter, each measurement time t_{j+1} is chosen taking the already acquired measurements $y(t_{1:j})$ into account. Then, each measurement time t_{j+1} is the solution of an optimization program.

More formally, the *stochastic program filter* is the intermittent filter whose measurement times are the solution of the following stochastic program.

Problem 1 (Stochastic Program).

$$
V_0(\emptyset; \emptyset) = \min_{t_1} \left\{ \sum_{t=0}^{t_1-1} \mathcal{E}[t|\emptyset] \right.
$$
$$
+ \mathbb{E}_{y(t_1)} \left[\min_{t_2} \left\{ \sum_{t=t_1}^{t_2-1} \mathcal{E}[t|y(t_1)] + \cdots \right. \right.
$$
$$
+ \mathbb{E}_{y(t_{N-1})} \left[\min_{t_N} \left\{ \sum_{t=t_{N-1}}^{t_N-1} \mathcal{E}[t|y(t_{1:N-1})] \right. \right.
$$
$$
\left. \left. \left. + \mathbb{E}_{y(t_N)} \left[\sum_{t=t_N}^{T} \mathcal{E}[t|y(t_{1:N})] \right] \right\} \right] \right) \cdots \left. \right\} \right] \left. \right\},
$$

subject to $0 \leq t_1 < t_2 < \cdots < t_N \leq T$.

Introducing the notation $t_0 := 0$, the problem can be stated recursively as,

$$
V_j(t_{1:j}; y(t_{1:j})) = \min_{t_{j+1}} \left\{ \sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})] \right.
$$
$$
\left. + \mathbb{E}_{y(t_{j+1})} \left[V_{j+1}(t_{1:j+1}; y(t_{1:j+1})) \right] \text{ with } t_j < t_{j+1} \leq T \right\},
$$
(7)

with the terminal condition,

$$
V_N(t_{1:N}; y(t_{1:N})) = \sum_{t=t_N}^{T} \mathcal{E}[t|y(t_{1:N})].
$$
(8)

Solving this multistage stochastic program requires finding an optimal policy for each t_{j+1}. Such optimal policies

are functions that associate the next optimal measurement time with the previously acquired measurements, i.e., it is $(y(t_1), \ldots, y(t_j)) \mapsto t_{j+1}$.

Under certain restrictive assumptions, including linearity and finite stochastic outcomes, stochastic multistage programs can be solved using scenario trees and duality [37]. For more complicated problems such as ours (nonlinear, continuous random variables), one approach is to use reinforcement learning. In this paper, we propose another approach, which is to optimize the variables all in once and offline, i.e., before any measurement acquisition. We use this approximation and we call it the offline approach.

2) Offline filter: The offline filter is optimal (in the sense of the expected MSE) if all the measurement times must be chosen before receiving any measurement. Each measurement time t_j is independent of the measurements that preceded it, $y(t_{1:j-1})$.

Formally, the *offline filter* is the intermittent filter whose measurement times are the solution of the following offline program.

Problem 2 (Offline Program).

$$
J_0(\emptyset; \emptyset) := \min_{0 \leq t_1 < \cdots < t_N \leq T} \mathbb{E}_{y(t_{1:N})} \left[\sum_{t=0}^{T} \mathcal{E}[t|y(t_{1:N})] \right].
$$

Because each optimal measurement time t_j is independent of the previously acquired measurements $y(t_{1:j-1})$, all these optimal measurement times $t_{1:N}$ can be computed at once, before any measurement acquisition. This can be a decisive advantage for certain real-world applications if the time between two discrete time steps is too short to solve an optimization program. Unlike Problem 1, here the optimization variables are no longer functions but natural numbers.

3) Online filter: To enhance filtering performance, once the measurements $y(t_{1:j})$ have been acquired, we may want to use them to recompute the next measurement time t_{j+1} including the information previously acquired, i.e., the previous measurements. This requires solving N different optimization programs, one for each measurement time.

Formally, the *online filter* is the intermittent filter for which each measurement time t_{j+1} (for $j = 0, \ldots, N-1$) is the solution of the following $(j+1)^{\text{th}}$ online program:

Problem 3 ($(j+1)^{\text{th}}$ Online Program). *When measurements $y(t_{1:j})$ have been acquired, the next measurement time t_{j+1} is computed by solving*

$$
J_j(t_{1:j}; y(t_{1:j})) := \min_{t_{j+1} < \cdots < t_N \leq T} \mathbb{E}_{y(t_{j+1:N})} \left[\right.
$$
$$
\left. \sum_{t=t_j}^{T} \mathcal{E}[t|y(t_{1:N})] \right], \text{ such that } t_{j+1} > t_j,
$$
(9)

where $t_0 := 0$ but the constraint is replaced by $t_1 \geq 0$ when $j = 0$.

In this $(j+1)^{\text{th}}$ program, only the measurement time t_{j+1} is used by the online filter. To avoid ambiguities, it will sometimes be noted t_{j+1}^o. The other variables of the $(j+1)^{\text{th}}$ program, i.e., the $t_{j+2:N}$, are not the measurement times of

the online filter because they will be replaced by the solutions of the next programs. Such online programs can be seen as a recursive version of the offline program.

Remark 3. *Setting $j = 0$ in Problem 3 gives exactly Problem 2. Consequently, the optimal t_1s are the same for these two problems.*

Remark 4. *If the dynamical system (1)-(4) is linear and Gaussian, i.e., functions $f_t(\cdot, \cdot), g_t(\cdot, \cdot)$ and $h_t(\cdot)$ are linear and distributions of $w(t), v(t)$ and $x(0)$ are Gaussian, then for given measurement times $t_1, \ldots, t_N$, the optimal filtered estimate $\hat{z}(t|y(t_{1:j}))$ is given explicitly by the Kalman filter [38, Theorem 2].*

In addition, the variance of the filtering error $\mathcal{E}[t|y(t_{1:j})]$ is independent of the measurements $y(t_{1:j})$, which implies that the Kalman gains can be computed offline (see, e.g., [32, p.3] for details). Consequently, the expectation operators in Problems 1, 2 and 3 are equivalent to identity operators, which gives the same problem three times. In conclusion, in the linear Gaussian case, the optimal solutions of Problems 1, 2 and 3 are the same.

If $t_{1:j}$ are the j first measurement times and $y(t_{1:j})$ the corresponding measurements, we denote $F_j(t_{1:j}; y(t_{1:j}))$ the remaining cost-to-go from time t_j using the online filter. For $j = 0, \ldots, N - 1$, it is

$$
\begin{aligned}
F_j(t_{1:j}; y(t_{1:j})) = &\sum_{t=t_j}^{t_{j+1}^o - 1} \mathcal{E}[t|y(t_{1:j})] \\
&+ \mathbb{E}_{y(t_{j+1}^o)}\Bigg[\sum_{t=t_{j+1}^o}^{t_{j+2}^o - 1} \mathcal{E}[t|y(t_{1:j}), y(t_{j+1}^o)] + \cdots \\
&+ \mathbb{E}_{y(t_{N-1}^o)}\Bigg[\sum_{t=t_{N-1}^o}^{t_N^o - 1} \mathcal{E}[t|y(t_{1:j}), y(t_{j+1:N-1}^o)] \\
&+ \mathbb{E}_{y(t_N^o)}\Bigg[\sum_{t=t_N^o}^{T} \mathcal{E}[t|y(t_{1:j}), y(t_{j+1:N}^o)] \Bigg] \Bigg] \cdots \Bigg],
\end{aligned}
$$

and for $j = N$, it is

$$
F_N(t_{1:N}; y(t_{1:N})) = \sum_{t=t_N}^{T} \mathcal{E}[t|y(t_{1:N})]. \tag{10}
$$

With this definition, the following recursive relation holds,

$$
\begin{aligned}
F_j(t_{1:j}; y(t_{1:j})) = &\sum_{t=t_j}^{t_{j+1}^o - 1} \mathcal{E}[t|y(t_{1:j})] \\
&+ \mathbb{E}_{y(t_{j+1}^o)}\Big[F_{j+1}(t_{1:j}, t_{j+1}^o; y(t_{1:j}), y(t_{j+1}^o)) \Big]. \tag{11}
\end{aligned}
$$

In the following theorem, we establish that the expected mean square filtering error is smaller or equal for the stochastic program filter than for the online filter, which itself is smaller or equal than for the offline filter.

Theorem 1. $V_0(\emptyset; \emptyset) \le F_0(\emptyset; \emptyset) \le J_0(\emptyset; \emptyset).$

Proof. The proof is available in the appendix. □
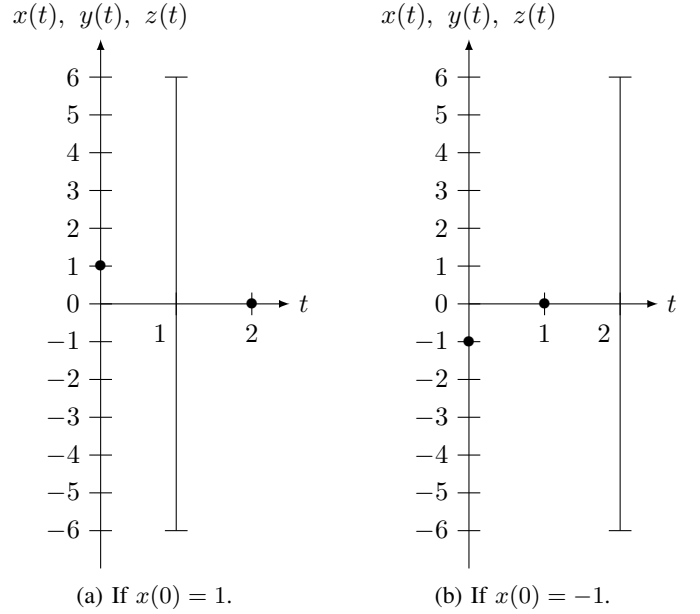

(a) If $x(0) = 1$.   (b) If $x(0) = -1$.

Fig. 1: Illustration of the system (12) through (14). Vertical lines represent uniform distributions $\mathcal{U}([-6, 6])$. Each of scenarios (a) and (b) has a 50/50 chance of occurring, depending on whether $x(0) = 1$ or -1.

The difference between Problems 1 and 3 may not seem obvious. Indeed, they both require computing the measurement times online by solving N optimization problems. To clarify the difference, we propose an example where these two problems give different solutions.

Example 1. *Consider three time steps, i.e., $T = 2$, with $N = 2$ measurements and consider the system illustrated in Figure 1 and summarized as*

$$
y(0) = z(0) = x(0) = \begin{cases} 1, & \text{with probability } 1/2 \\ -1, & \text{with probability } 1/2, \end{cases} \tag{12}
$$

$$
y(1) = z(1) = x(1) = \begin{cases} w(0), & \text{if } x(0) = 1 \\ 0, & \text{if } x(0) = -1, \end{cases} \tag{13}
$$

$$
y(2) = z(2) = x(2) = \begin{cases} 0, & \text{if } x(1) \ne 0 \\ w(1), & \text{if } x(1) = 0, \end{cases} \tag{14}
$$

where $w(0)$ and $w(1)$ independently follow a uniform distribution between -6 and 6, i.e., $w(0), w(1) \sim \mathcal{U}([-6, 6])$.

First observe that it is possible to make no filtering error. To do this, set the first measurement time to $t_1 = 0$ and then, if $y(0) = 1$, set the second to $t_2 = 1$ and on the contrary, if $y(0) = -1$, set the second to $t_2 = 2$. With this schedule, a measurement is acquired whenever there is an uncertainty, which guarantees that there is no filtering error. It is the solution of Problem 1.

However, we will show that the solutions of both Problems 2 and 3 are to choose $t_1 = 1$ and $t_2 = 2$ which results in a filtering error at time $t = 0$.

Doing an exhaustive search for Problem 2 shows that choosing $t_1 = 1$ and $t_2 = 2$ is optimal. This is due to the fact that if there is no measurement at $t = 1$ or at $t = 2$, a

large error will be made half the time because of the uniformly distributed $x(t)$. Then, because the optimal t_1s for Problems 3 and 2 are the same (see Remark 3), $t_1 = 1$ is also optimal for the former. But then, the only remaining possibility for t_2 is $t_2 = 2$. This shows that the optimal solutions of both Problems 2 and 3 are $t_1 = 1$ and $t_2 = 2$ which gives a filtering error at time $t = 0$.

This example illustrates that the optimal time t_1 of Problem 3 is computed by assuming that the time t_2 cannot depend on $y(t_1)$.

C. Numerical approximations

In this section, we present the numerical approximations which are used to make Problems 2 and 3 tractable. Firstly, we briefly explain how to approximate numerically the estimate $\hat{z}(t|y(t_{1:j}))$ using a particle filter and then how to approximate the objective functions of these two problems. We reserve the term "intermittent filter" for the ideal filter and use "intermittent particle filter" for its approximation. Let us emphasize that Theorem 1 applies only in the ideal case.

1) Approximation of the estimate: The estimate $\hat{z}(t|y(t_{1:j}))$ is not straightforward to compute but can be approximated using a particle filter.

Essentially, a particle filter algorithm alternates between (i) a *prediction step* (also called mutation), used to compute the estimate at the next time step from the estimate at the current step; and (ii) a *correction step* (also called selection) that updates the current 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 $t \neq t_k$ for all k. For reasons of numerical stability, the particle weights are computed and stored in the logarithmic domain (see, e.g., [39, Section 4.3.1] for details).

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

In general, the approximation provided by the particle filter is biased when the number of particles is finite (the bias being $O(1/\#\text{ particles})$) [40, Section 3.2]. However, we neglect this aspect in this article and leave the study of the impact of this bias for future work.

In what follows, we will make a slight abuse of notation by using the same notation $\hat{z}(t|y(t_{1:j}))$ for the quantity (5) and for its approximation calculated by the particle filter.

2) Approximation of the objective functions: The objective functions of Problems 2 and 3 contain an expectation operator, which makes these functions challenging to evaluate. To tackle this problem, we use a Monte Carlo approximation of the expectation.

Because the objective function of Problem 2 is the same as the objective function of Problem 3 for $j = 0$ (see Remark 3), we focus on the objective function of Problem 3 for any j, which includes Problem 2.

The Monte Carlo approximation of the objective function of Problem 3 is represented in Figure 2. Assuming that the first j measurements $y(t_{1:j})$ have already been acquired and that future measurement times $t_{j+1:N}$

$$t_{j+1:N}$$

$$\boxed{\text{Model}}$$

$$\{y^k(t_{j+1:N})\}^{k=1,\ldots,K}$$

$$y(t_{1:j}) \longrightarrow \boxed{\text{Particle Filter}} \qquad \{z^k(t)\}_{t=t_j,\ldots,T}^{k=1,\ldots,K}$$

$$\{\hat{z}^k(t)\}_{t=t_j,\ldots,T}^{k=1,\ldots,K}$$

$$\boxed{\text{Mean Square Error}}$$

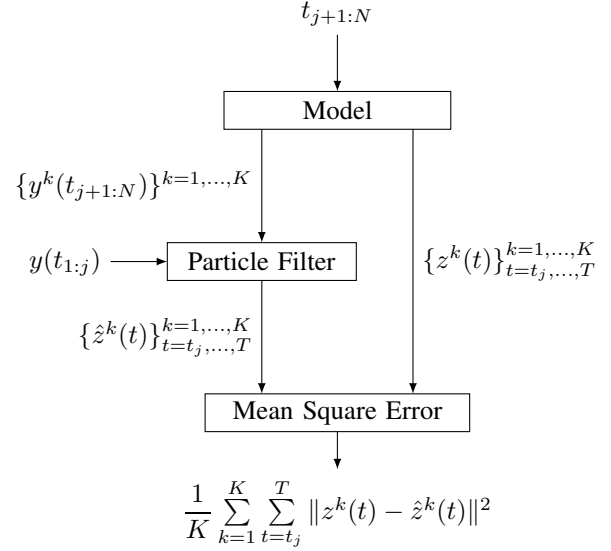$$\frac{1}{K}\sum_{k=1}^{K}\sum_{t=t_j}^{T}\|z^k(t) - \hat{z}^k(t)\|^2$$

Fig. 2: Representation of the Monte Carlo algorithm that estimates $\mathbb{E}_{y(t_{j+1:N})}[\sum_{t=t_j}^{T} \mathcal{E}[t|y(t_{1:N})]]$. The outputs generated by the Model block are drawn according to Equations (1) through (4) and the distributions of $v(t)$ and $w(t)$. The particle filter computes an estimate $\hat{z}^k(t) = \hat{z}^k(t|y(t_{1:j}), y^k(t_{j+1:N}))$ of $z^k(t)$ from previous intermittent measurements. The Mean Square Error block computes the mean squared of the difference between its inputs; it is the right-hand side of Equation (15).

are fixed, one can simulate K realizations of the state $\{x^k(t)\}_{t=0,\ldots,T}^{k=1,\ldots,K}$ drawn according to Equations (1) and (4). The corresponding quantities that remain to be estimated $\{z^k(t)\}_{t=t_j,\ldots,T}^{k=1,\ldots,K}$ can be computed using Equation (3). Similarly, corresponding measurements $\{y^k(t_l)\}_{l=j+1,\ldots,N}^{k=1,\ldots,K}$ can be simulated thanks to Equation (2). For each k, the simulated measurement sequence can be concatenated to the known measurements to give a complete measurement sequence $y(t_1), \ldots, y(t_j), y^k(t_{j+1}), \ldots, y^k(t_N)$, from which the particle filter computes the estimates $\{\hat{z}^k(t|y(t_1), \ldots, y^k(t_N))\}_{t=t_j,\ldots,T}$ of $\{z^k(t)\}_{t=t_j,\ldots,T}$. Then, thanks to Definition (6), the Monte Carlo estimator is computed as

$$\mathbb{E}_{y(t_{j+1:N})}\left[\sum_{t=t_j}^{T}\mathcal{E}[t|y(t_{1:N})]\right] = \mathbb{E}_{y(t_{j+1:N})}\Bigg[$$

$$\sum_{t=t_j}^{T}\mathbb{E}_{x(0),w(0),\ldots,w(T-1)}\Bigg[$$

$$\|z(t) - \hat{z}(t|y(t_{1:N}))\|^2\Big|y(t_l), \forall t_l \leq t\Bigg]\Bigg]$$

$$\approx \frac{1}{K}\sum_{k=1}^{K}\sum_{t=t_j}^{T}\|z^k(t) - \hat{z}^k(t|y(t_{1:j}), y^k(t_{j+1:N}))\|^2. \quad (15)$$

Computing this Monte Carlo approximation requires running the particle filter K times, which can be computationally expensive.

The fact that we have access only to an approximation of the objective function requires particular attention during the optimization. Note that since the particle filter provides a biased approximation, the approximation of the cost function may be biased as well. However, accounting for this bias is outside the scope of this paper and is left to future work.

D. Optimization algorithms

Solving the combinatorial optimization Problem 2 (or equivalently, in light of Remark 3, Problem 3 with $j = 0$) with an exhaustive search would require evaluating the objective function $\frac{(T+1)!}{(T+1-N)!N!}$ times, i.e., the number of subsets of size N in a set of size $T + 1$. This is computationally intractable for large N and T values.

In this section, we present five different heuristic optimization algorithms to find an approximate solution for Problems 2 and 3 efficiently.

1) Random trial algorithm (RT): The random trial algorithm samples measurement time sets uniformly at random and evaluates their corresponding costs. The measurement time set with minimum cost is returned.

2) Greedy forward algorithm (GF): The greedy forward algorithm begins with an empty measurement set. Sequentially, it adds the measurement times one at a time such that, at each iteration, the added measurement time minimizes the cost. It stops when the set contains N measurement times.

3) Greedy backward algorithm (GB): The greedy backward algorithm works in the opposite direction from the greedy forward algorithm. It begins with a complete measurement set. Sequentially, it takes out the measurement times one at a time such that, at each iteration, the dropped measurement time minimizes the cost. It stops when the set contains N measurement times.

4) Simulated annealing algorithm (SA): The implementation of the SA algorithm is done according to [41, Section 1.2.3]. The initial temperature parameter is $T^\circ = 10$ and the geometric temperature decay rate is 0.9, i.e., after each generation, $T^\circ := 0.9 \cdot T^\circ$. If a mutation reduces the cost, it is kept. If it increases the cost, it is kept with probability $\exp(-(\text{cost increase})/T^\circ)$. The mutations are executed by prohibiting the duplication of measurement times and the mutation probability is 0.1 by measurement time.

5) Genetic algorithm (GA): The implementation of the genetic algorithm is based on [42]. The selection of the individuals kept for the next generation is done according to stochastic universal sampling [42, Section 5.4]. To ensure that the number of measurement times remains constant over generations, a count preserving crossover operator is implemented [43, Section 3.6]. It prevents the production of individuals $(t_{j+1}, \ldots, t_N)$ for which $t_k = t_l$ with $k \neq l$. For the same reason, the mutations are executed by prohibiting the duplication of measurement times. The algorithm uses sigma scaling with a unitary sigma coefficient [42, Section 5.4]. Crossover probability is 1 and mutation probability per measurement time is 0.003.

When the RT, SA or GA algorithm is used to solve the N online Problems 3, the solution found for the j^{th} problem can be added to the initial population when solving the $(j + 1)^{\text{th}}$ problem. This can speed up the convergence of the algorithm but also reduce the exploration capacity of these methods. In this work, we do not use this acceleration trick.

As mentioned in Subsection II-C, the objective functions of Problems 2 and 3 can be evaluated approximately only. This makes minimization difficult if it is too sensitive to bad cost function estimations. To face this issue, the SA algorithm and the GA return the best individual of the last generation instead of the best individual in all generations.

In Section III, the performances of these five algorithms are compared. It will show that the GA outperforms the other algorithms on the studied example.

III. RESULTS AND DISCUSSION

A. Results parameters

In this section, first we present two dynamical systems used for the results, then we describe the parameters of the simulations and the performance indicators.

1) Tumor motion model: To illustrate the performances of our method, we study a model describing one-dimensional tumor motion. As mentioned in introduction, the problem of tracking tumors with X-rays requires doing the best of each X-ray acquisition in order to spare healthy tissues.

The duration between each discrete time step is δ. The tumor position to estimate $z(t)$ is a shifted sinusoidal signal with a time-varying amplitude $a(t)$, a time-varying shift $b(t)$, and a constant frequency ω. Both $a(t)$ and $b(t)$ are bounded random walks. Bounds ensure that values remain realistic over time. In addition, the constant oscillation frequency ω is picked uniformly at random at the beginning of the process. Each measurement $y(t)$ is a noisy version of the position $z(t)$.

Let us define the state $x(t) = \begin{bmatrix} a(t) & b(t) & \omega(t) \end{bmatrix}^\top \in \mathbb{R}^3$ and the process noise $w(t) = \begin{bmatrix} w_a(t) & w_b(t) \end{bmatrix}^\top \in \mathbb{R}^2$. The dynamic of the system is defined as follows:

$$
\begin{aligned}
x(t+1) &= \begin{bmatrix} a(t+1) \\ b(t+1) \\ \omega(t+1) \end{bmatrix} \\
&= f(x(t), w(t)) \\
&:= \begin{bmatrix} \text{clip}(a(t) + w_a(t), \underline{a}, \bar{a}) \\ \text{clip}(b(t) + w_b(t), \underline{b}, \bar{b}) \\ \omega(t) \end{bmatrix},
\end{aligned} \tag{16}
$$

$$
y(t) = g_t(x(t), v(t)) = a(t)\sin(\omega(t)t\delta) + b(t) + v(t), \tag{17}
$$

$$
z(t) = h_t(x(t)) = a(t)\sin(\omega(t)t\delta) + b(t), \tag{18}
$$

$$
x(0) = \begin{bmatrix} a(0) \\ b(0) \\ \omega(0) \end{bmatrix} \sim \mathcal{U}\left([\underline{a}, \bar{a}] \times [\underline{b}, \bar{b}] \times [\underline{\omega}, \bar{\omega}]\right). \tag{19}
$$

Equations (16), (17) and (18) hold for $t = 0, \ldots, T - 1$, $t \in \{t_1, \ldots, t_N\}$, and $t = 0, \ldots, T$, respectively. Noises $w_a(t)$, $w_b(t)$ and $v(t)$ are independent zero mean Gaussian noises and have a unitary standard deviation, i.e., $\sigma_a = \sigma_b = \sigma_v = 1 \, [mm]$. The clipping function is defined as $\text{clip}(x, \underline{x}, \bar{x}) := \min(\max(x, \underline{x}), \bar{x})$. Equation (19) indicates

that the initial state is uniformly distributed at random on the indicated domain.

We consider $N = 11$ measurements in the range of 0 to $T = 30$. To make this model more realistic, the values of $\underline{a}$, $\bar{a}$, $\underline{b}$, and $\bar{b}$ are inspired from [44, Table 1] and the values of $\underline{\omega}$ and $\bar{\omega}$ are inspired from [45, Table 1]. The parameters of the system are

$$\underline{a} = 8.8 \ [mm], \ \bar{a} = 24 \ [mm],$$
$$\underline{b} = -5.8 \ [mm], \ \bar{b} = 5.8 \ [mm],$$
$$\underline{\omega} = 1.3 \ [rad/s], \ \bar{\omega} = 2.1 \ [rad/s],$$
$$\delta = 0.25 \ [s], \ \sigma_a = \sigma_b = \sigma_v = 1 \ [mm],$$
$$T = 30, \ N = 11.$$

The importance density used by the particle filter is the prior density, except for $\omega(t)$ where we use $\omega(t+1) = \text{clip}(\omega(t) + w_\omega(t), \underline{\omega}, \bar{\omega})$ with $w_\omega(t)$, an independent zero mean Gaussian noise with standard deviation $\sigma_\omega = 0.005$.

2) A common benchmark for particle filters: We also test our method on the following widely studied model [4], [35], [46], [47], [48]:

$$x(t+1) = \frac{x(t)}{2} + \frac{25x(t)}{1 + x(t)^2} + 8\cos(1.2t) + w(t), \quad (20)$$

$$y(t) = \frac{x(t)^2}{20} + v(t), \quad (21)$$

$$z(t) = x(t), \quad (22)$$

$$x(0) \sim \mathcal{N}(0, 5^2), \quad (23)$$

where $x(t), y(t), z(t), w(t), v(t) \in \mathbb{R}$. Equations (20), (21) and (22) hold for $t = 0, \ldots, T-1$, $t \in \{t_1, \ldots, t_N\}$, and $t = 0, \ldots, T$, respectively. Quantities $w(t)$, $v(t)$ and $x(0)$ are randomly distributed according to independent zero mean Gaussian distributions with standard deviations $\sigma_w = 1$, $\sigma_v(t) = \sin(0.25t) + 2$ and $\sigma_{x_0} = 5$, respectively. As for the tumor model, $N = 11$ measurements are allowed and the horizon is set at $T = 30$. For this model, the importance density used by the particle filter is the prior density.

3) Performance indicators: To test the proposed methods, we simulate many realizations of $\{y(t)\}_{t=0,\ldots,T}$ and $\{z(t)\}_{t=0,\ldots,T}$ according to system dynamic on which our filtering method is applied. Then, the filtering mean square error, $\text{MSE} = \frac{1}{T+1}\sum_{t=0}^{T}\|z(t) - \hat{z}(t)\|^2$, is computed and is compared with the filtering mean square error MSE_{REG} obtained with regular measurement times,

$$t_{1:N} = \left\{ \text{Round}\left[\frac{kT}{N-1}\right] \middle| k = 0, \ldots, N-1 \right\}, \quad (24)$$

where $\text{Round}[\cdot]$ is the rounding operator. More precisely, we look at the gain

$$G := \log_{10}\left(\frac{\text{MSE}_{\text{REG}}}{\text{MSE}}\right). \quad (25)$$

It illustrates the merits of using optimal measurement times instead of regular ones. A positive gain indicates that the considered intermittent particle filter outperforms the regular one.

As performance indicators, we look at the mean and median gain and the proportion of positive gain over all the simulations. Notice that the mean gains can be seen as an estimation of the expected gain and the proportion of positive gains as an estimation of the probability of outperforming the regular measurement filter.

4) Optimization and simulations parameters: As the measurement times that are the solution of the online Problem 3 are a function of the previous measurements, the optimization algorithm has to be run on each simulation. This increases the computation time significantly. Therefore, the solutions of Problems 2 and 3 are studied separately and the number of simulations and the optimization parameters have been set accordingly. The values of the optimization parameters and the simulation parameters are given in Table II. Note that the number of particles used for the optimization is different from the number of particles for the filtering.

	Parameters	Offline	Online
Optimization	# draws K	1000	200
	# particles	200	100
	pop. size	50	30
	# generations	25	15
Simulations	# simulations	100000	500
	# particles	1000	1000

TABLE II: Values of the model parameters, test parameters, and optimization parameters for the offline Problem 2 and the online Problem 3.

B. Comparison of optimization algorithms

The different optimization algorithms proposed in Subsection II-D are tested on the offline Problem 2 for the tumor motion model described in Section III-A1.

For all the optimization algorithms, the evolution of the cost of Problem 2 with respect to the number of cost function evaluations is illustrated in Figure 3. Because the number of cost function evaluations of the GF and GB algorithm is fixed, they are represented by points in the figure.

One can observe that all five optimization algorithms return a measurement subset with associated expected MSE significantly lower than the regularly spaced measurements. The final expected MSE of the SA, RT, GF and GB optimization algorithms all lie close to 6. Note that the GF and GB algorithms give good results for a relatively small number of evaluations of the cost function. Consequently, they can be a good option if computing resources are limited. The GA and SA algorithms perform notably better with a final minimum cost around 5.5. One can note that the difference between GA average and minimum cost decreases over generations before reaching quasi-convergence, meaning that most of the individuals are identical, which is the desired convergence behavior for a GA.

As a result of these observations, the GA is used as the optimization algorithm in the rest of the paper.

C. Filtering performance

In this subsection, we analyze the filtering performance obtained using the offline (Problem 2) and online (Problem 3)
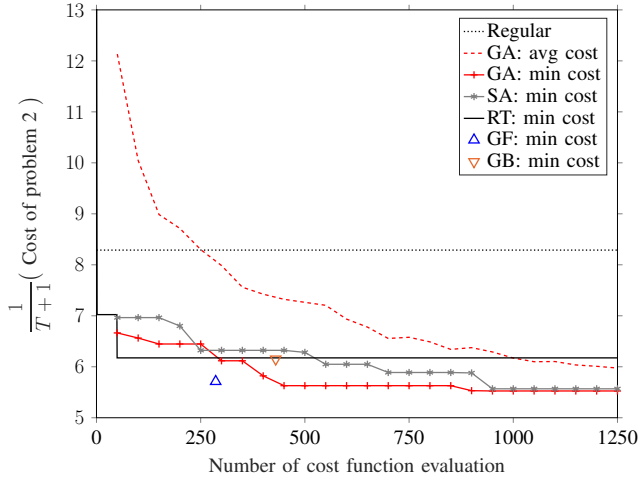
Fig. 3: Evolution of the minimum cost of the offline Problem 2 with respect to the number of cost function evaluations for the different optimization algorithms: genetic algorithm (GA), simulated annealing (SA), random trial (RT), greedy forward (GF), and greedy backward (GB). It illustrates the quality of minimization for given computational resources. The evolution of the average cost of the population of the GA algorithm at each generation and the cost of regularly spaced measurement times are shown as well.

measurement times. First, we report the results on the tumor motion model (Subsection III-A1), then we present the results for the common benchmark (Subsection III-A2).

1) Results for the tumor motion model (Subsection III-A1): The histogram of the gains G for the offline Problem 2 and for the online Problem 3 can be found in Figure 4. The offline gain is computed over 100000 simulations and the online gain is computed over 500 simulations. The vertical line corresponds to the null gain. These histograms can be interpreted as approximations of the probability density of the gains G. The proportion of positive gain for the GA algorithm can be read as the area of the histogram above 0. The corresponding performance indicators are reported in Table III for both the offline particle filter (related to Problem 2) and online particle filter (related to Problems 3).

The mean and median gains are positive for our offline and online methods, which indicates that they generally outperform the regular particle filter. In addition, the proportion of positive gain indicates that our offline and online methods outperform the regular particle filter in 69.3% and 76.3% of the cases, respectively.

	Offline	Online
Mean gain ($\pm$ std)	0.142 ($\pm$ 0.269)	0.238 ($\pm$ 0.296)
Median gain	0.134	0.248
Proportion positive gain	69.3 %	76.3 %

TABLE III: Performance indicators for the offline and online solutions. The considered system, the performance indicators, and the optimization and simulation parameters are described in Subsections III-A1, III-A3, and III-A4, respectively.

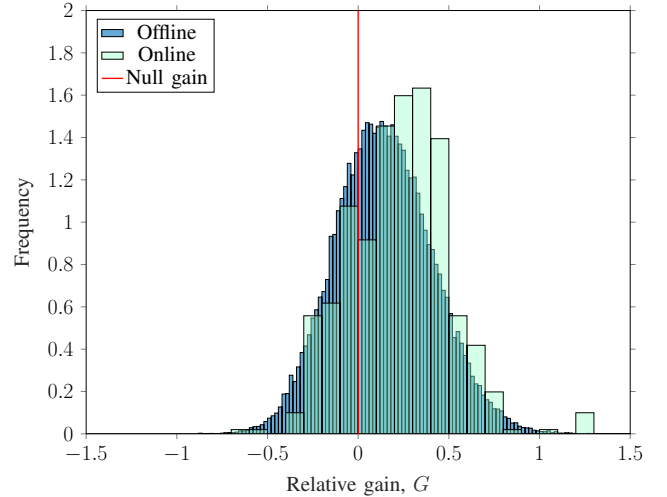Figure 5 represents a particular trajectory $z(t)$ and the

Fig. 4: Histograms of the gain G (see Equation (25)) for the tumor motion model (Subsection III-A1). The histograms have been obtained over 100000 and 500 simulations for the offline and online particle filters, respectively.

trajectory filtered by a particle filter with regular measurement times (see Equation (24)), as well as with the offline and online particle filters. In all three cases, the measurement times are indicated. As expected, we obtain better filtering performance using the online particle filter (the gain is $G_{\text{online}} = 0.20$) than the offline particle filter (the gain is $G_{\text{offline}} = 0.15$), the largest mean square filtering error being obtained with the regular measurements.
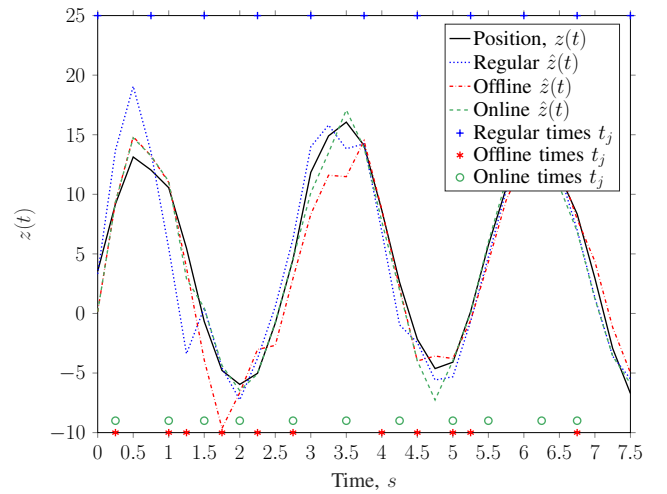

Fig. 5: Comparison of a particular realization $z(t)$ with the filtered values $\hat{z}(t)$ obtained with a particle filter with regular measurement times and both the offline and the online particle filters. The gains on the complete sequence are $G_{\text{offline}} = 0.223$ and $G_{\text{online}} = 0.616$. Results are simulated from the tumor motion model (Subsection III-A1).

2) Results for the common benchmark (Subsection III-A2): Figure 6 represents the histogram of the offline and online gains G for the common benchmark. The corresponding performance indicators are reported in Table IV. It is observed

that both the offline and online methods outperform the particle filter with regular measurement times. In addition, the performance difference between the offline and online methods is important. The fact that the performance gap between the offline and online particle filters is larger for this common benchmark than for the tumor motion model can be related to Remark 4. Indeed, if the model were linear and Gaussian, there would be no performance gap.
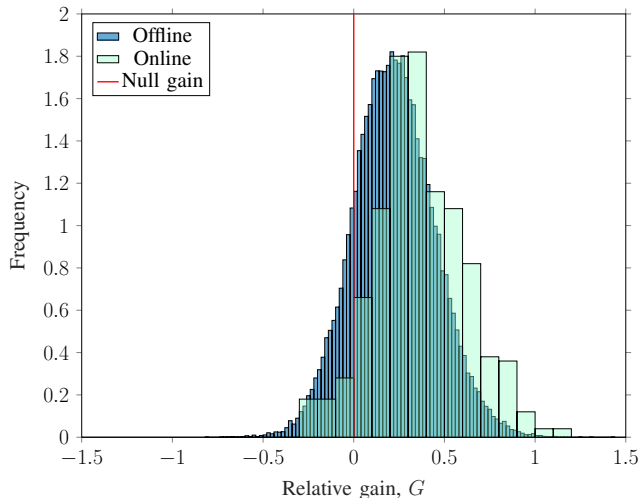

Fig. 6: Histograms of the gain G (see Equation (25)) for the common benchmark (Subsection III-A2). The histograms have been obtained over 100000 and 500 simulations for the offline and online particle filters, respectively.

	Offline	Online
Mean gain ($\pm$ std)	0.221 ($\pm$ 0.225)	0.365 ($\pm$ 0.253)
Median gain	0.216	0.340
Proportion positive gain	84.3 %	93.8 %

TABLE IV: Performance indicators for the offline and online solutions. The considered system, the performance indicators, and the optimization and simulation parameters are described in Subsections III-A2, III-A3, and III-A4, respectively.

Finally, Figure 7 shows a trajectory of the common benchmark as well as the tracking obtained by the regular, offline and online particle filters. The measurement times are also displayed.

IV. CONCLUSION

The problem of the optimal intermittent particle filter has been addressed. This consists in selecting the measurement times of the particle filter according to a certain optimality criterion.

We have proposed three variants of expected mean square error minimization giving rise to three different intermittent filters: the stochastic program filter, the offline filter, and the online filter.

A theorem has been proven that compares the performance of these three intermittent filters.

Then, different algorithms to compute the measurement times of the offline and online particle filters have been
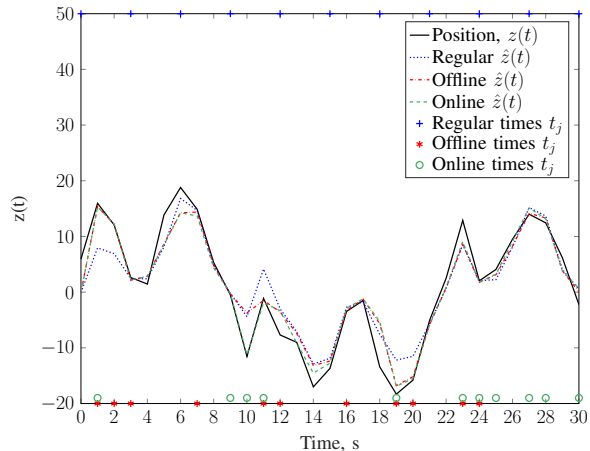

Fig. 7: Comparison of a particular realization $z(t)$ with the filtered values $\hat{z}(t)$ obtained with a particle filter with regular measurement times and both the offline and the online particle filters. The gains on the complete sequence are $G_{\text{offline}} = 0.173$ and $G_{\text{online}} = 0.283$. Results are simulated from the common benchmark (Subsection III-A2).

proposed: the random trial, greedy forward, greedy backward, simulated annealing, and genetic algorithms. These different optimization algorithms were compared on a tumor movement model inspired by real data. The genetic and simulated annealing algorithms are the two optimizers showing the best performance.

Finally, on this same tumor model, our offline and online particle filters were compared with a particle filter with regular measurement times. The offline and online particle filters outperform the regular particle filter in 69.3% and 84.3% of the cases, respectively. For a second example, the regular particle filter is outperformed in 76.3% and 93.8% of the cases by the offline and online particle filters, respectively.

Further work will focus on robustness analyses to determine how model uncertainties affect the performance of the proposed methods for situations in which the system dynamics are known only approximately. In addition, we will use a reinforcement learning approach to solve the stochastic program and test the stochastic program particle filter. Other continuations of this work will consist of experimental validation on real-world data, for example, in the case of mobile tumor tracking. It would also be interesting to investigate a continuous time version of the problem.

Overall, our results demonstrate the added value of selecting measurement times according to system dynamics for the optimal estimation of a state from limited measurements.

APPENDIX

PROOF OF THEOREM 1

The structure of the proof consists in (i) observing that once all the measurements are known, cost-to-go is equal for the three filters; then (ii) we show that if the cost-to-go when choosing the measurement time t_{j+1} is smaller with the stochastic program filter than with the online filter, which itself

is smaller than with the offline filter, then this is also the case when choosing the measurement time t_j; and finally (iii) a backward recursive argument shows that this is also true for the initial cost-to-go.

This structure is applied to each inequality separately.

Proof of the first inequality: From Relations (8) and (10) it follows that for all $t_{1:N}$ and $y(t_{1:N})$,

$$V_N(t_{1:N}; y(t_{1:N})) = F_N(t_{1:N}; y(t_{1:N})). \tag{26}$$

For a fixed $j \in \{0, \ldots, N-1\}$, assume that for all $t_{1:j+1}$ and $y(t_{1:j+1})$ it holds that

$$V_{j+1}(t_{1:j+1}; y(t_{1:j+1})) \leq F_{j+1}(t_{1:j+1}; y(t_{1:j+1})). \tag{27}$$

Then, taking the expectation with respect to $y(t_{j+1})$ and adding a same term on both sides, we have,

$$\sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})] + \mathbb{E}_{y(t_{j+1})}\left[V_{j+1}(t_{1:j+1}; y(t_{1:j+1}))\right] \leq$$
$$\sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})] + \mathbb{E}_{y(t_{j+1})}\left[F_{j+1}(t_{1:j+1}; y(t_{1:j+1}))\right],$$

for all $t_{1:j+1}$ and $y(t_{1:j})$. Taking the minimum over all t_{j+1} on the left-hand side and fixing t_{j+1} to t^o_{j+1} on the right-hand side, we have,

$$\min_{t_{j+1}}\left\{ \sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})]\right.$$
$$\left. + \mathbb{E}_{y(t_{j+1})}\left[V_{j+1}(t_{1:j+1}; y(t_{1:j+1}))\right]\right\} \leq$$
$$\sum_{t=t_j}^{t^o_{j+1}-1} \mathcal{E}[t|y(t_{1:j}))]$$
$$+ \mathbb{E}_{y(t^o_{j+1})}\left[F_{j+1}(t_{1:j}, t^o_{j+1}; y(t_{1:j}), y(t^o_{j+1}))\right],$$

for all $t_{1:j}$ and $y(t_{1:j})$. From Relation (7) on the left-hand side and Relation (11) on the right-hand side, we have

$$V_j(t_{1:j}; y(t_{1:j})) \leq F_j(t_{1:j}; y(t_{1:j})), \tag{28}$$

for all $t_{1:j}$ and $y(t_{1:j})$.

But then, because Relation (27) implies (28) and because of (26), we have $V_0(\emptyset; \emptyset) \leq F_0(\emptyset; \emptyset)$.

Proof of the second inequality: It follows from a similar argument. We define

$$J_N(t_{1:N}; y(t_{1:N})) := F_N(t_{1:N}; y(t_{1:N})). \tag{29}$$

Then, for a fixed $j \in \{0, \ldots, N-1\}$, assume that for all $t_{1:j+1}$ and $y(t_{1:j+1})$ it holds that

$$F_{j+1}(t_{1:j+1}; y(t_{1:j+1})) \leq J_{j+1}(t_{1:j+1}; y(t_{1:j+1})). \tag{30}$$

Rewriting the right-hand side according to (9) gives

$$F_{j+1}(t_{1:j+1}; y(t_{1:j+1})) \leq$$
$$\min_{t_{j+2}<\cdots<t_N\leq T} \mathbb{E}_{y(t_{j+2:N})}\left[\sum_{t=t_{j+1}}^{T} \mathcal{E}[t|y(t_{1:N})]\right].$$

The right-hand side is upper bounded by

$$\mathbb{E}_{y(t_{j+2:N})}\left[\sum_{t=t_{j+1}}^{T} \mathcal{E}[t|y(t_{1:N})]\right],$$

for all $t_{j+2:N}$. Then one can write

$$F_{j+1}(t_{1:j+1}; y(t_{1:j+1})) \leq$$
$$\mathbb{E}_{y(t_{j+2:N})}\left[\sum_{t=t_{j+1}}^{T} \mathcal{E}[t|y(t_{1:N})]\right],$$

for all $t_{1:N}$ and $y(t_{1:j+1})$. Taking the expectation with respect to $y(t_{j+1})$ and adding a same term on both sides gives

$$\sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})] + \mathbb{E}_{y(t_{j+1})}\left[F_{j+1}(t_{1:j+1}; y(t_{1:j+1}))\right] \leq$$
$$\sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})] + \mathbb{E}_{y(t_{j+1:N})}\left[\sum_{t=t_{j+1}}^{T} \mathcal{E}[t|y(t_{1:N})]\right], \tag{31}$$

for all $t_{1:N}$ and $y(t_{1:j})$. Using the fact that the first term of the right-hand side is independent of $y(t_{j+1:N})$, and then thanks to Remark 1, the right-hand side can be rewritten successively as

$$\mathbb{E}_{y(t_{j+1:N})}\left[\sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})] + \sum_{t=t_{j+1}}^{T} \mathcal{E}[t|y(t_{1:N})]\right] =$$
$$\mathbb{E}_{y(t_{j+1:N})}\left[\sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:N})] + \sum_{t=t_{j+1}}^{T} \mathcal{E}[t|y(t_{1:N})]\right].$$

By combining the two sums on this right-hand side, the inequality (31) becomes

$$\sum_{t=t_j}^{t_{j+1}-1} \mathcal{E}[t|y(t_{1:j})] + \mathbb{E}_{y(t_{j+1})}\left[F_{j+1}(t_{1:j+1}; y(t_{1:j+1}))\right] \leq$$
$$\mathbb{E}_{y(t_{j+1:N})}\left[\sum_{t=t_j}^{T} \mathcal{E}[t|y(t_{1:N})]\right],$$

for all $t_{1:N}$ and $y(t_{1:j})$. We set $t_{j+1:N}$ to the value that minimizes this right-hand side. In so doing, $t_{j+1} = t^o_{j+1}$ and the inequality becomes

$$\sum_{t=t_j}^{t^o_{j+1}-1} \mathcal{E}[t|y(t_{1:j})]$$
$$+ \mathbb{E}_{y(t^o_{j+1})}\left[F_{j+1}(t_{1:j}, t^o_{j+1}; y(t_{1:j}), y(t^o_{j+1}))\right] \leq$$
$$\min_{t_{j+1}<\cdots<t_N\leq T} \mathbb{E}_{y(t_{j+1:N})}\left[\sum_{t=t_j}^{T} \mathcal{E}[t|y(t_{1:N})]\right],$$

for all $t_{1:j}$ and $y(t_{1:j})$. Now, using Relation (11) on the left-hand side and Relation (9) on the right-hand side, we obtain

$$F_j(t_{1:j}; y(t_{1:j})) \leq J_j(t_{1:j}; y(t_{1:j})), \tag{32}$$

for all $t_{1:j}$ and $y(t_{1:j})$.

Finally, because Relation (30) implies Relation (32) and because of (29), we have $F_0(\emptyset; \emptyset) \leq J_0(\emptyset; \emptyset)$, which concludes the proof.

References

[1] L. Arnold and R. Lefever, *Stochastic Nonlinear Systems in Physics, Chemistry, and Biology: Proceedings of the Workshop Bielefeld, Fed. Rep. of Germany, October 5–11, 1980.* Springer Science & Business Media, 2012, vol. 8.

[2] G. Adomian, *Nonlinear stochastic systems theory and applications to physics.* Springer Science & Business Media, 1988, vol. 46.

[3] J. A. Freund and T. Pöschel, *Stochastic processes in physics, chemistry, and biology.* Springer Science & Business Media, 2000, vol. 557.

[4] M. S. 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.

[5] C. Hue, J.-P. Le Cadre, and P. Pérez, "A particle filter to track multiple objects," in *Proceedings 2001 IEEE Workshop on Multi-Object Tracking.* IEEE, 2001, pp. 61–68.

[6] J. U. Cho, S. H. Jin, X. Dai Pham, J. W. Jeon, J. E. Byun, and H. 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.

[7] H. Andreasson, A. Treptow, and T. 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.

[8] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. 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.

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

[10] A. V. Shenoy, J. Prakash, K. McAuley, V. Prasad, and S. Shah, "Practical issues in the application of the particle filter for estimation of chemical processes," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 2773–2778, 2011.

[11] M. S. Jha, M. Bressel, B. Ould-Bouamama, and G. Dauphin-Tanguy, "Particle filter based hybrid prognostics of proton exchange membrane fuel cell in bond graph framework," *Computers & Chemical Engineering*, vol. 95, pp. 216–230, 2016.

[12] F. Cadini, E. Zio, and D. Avram, "Monte carlo-based filtering for fatigue crack growth estimation," *Probabilistic Engineering Mechanics*, vol. 24, no. 3, pp. 367–373, 2009.

[13] S. Thrun, "Particle filters in robotics," in *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence.* Morgan Kaufmann Publishers Inc., 2002, pp. 511–518.

[14] A. A. Rahni, E. Lewis, M. 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.

[15] G. C. Sharp, S. B. Jiang, S. Shimizu, and H. Shirato, "Prediction of respiratory tumour motion for real-time image-guided radiotherapy," *Physics in Medicine & Biology*, vol. 49, no. 3, p. 425, 2004.

[16] S. Sun, L. Xie, W. Xiao, and Y. C. Soh, "Optimal linear estimation for systems with multiple packet dropouts," *Automatica*, vol. 44, no. 5, pp. 1333–1342, 2008.

[17] K. Rutledge, S. Z. Yong, and N. Ozay, "Finite horizon constrained control and bounded-error estimation in the presence of missing data," *Nonlinear Analysis: Hybrid Systems*, vol. 36, p. 100854, 2020.

[18] Y. Liang, T. Chen, and Q. Pan, "Optimal linear state estimator with multiple packet dropouts," *IEEE Transactions on Automatic Control*, vol. 55, no. 6, pp. 1428–1433, 2010.

[19] W. Li, Z. Wang, Y. Yuan, and L. Guo, "Particle filtering with applications in networked systems: a survey," *Complex & Intelligent Systems*, vol. 2, no. 4, pp. 293–315, 2016.

[20] J. Chen, J. Li, and L. Ma, "Nonlinear filtering with multiple packet dropouts," in *2010 International Conference on Image Analysis and Signal Processing.* IEEE, 2010, pp. 83–88.

[21] C. Yang, H. Fang, and B. Shi, "Particle filter with markovian packet dropout and time delay," *Journal of the Franklin Institute*, vol. 356, no. 1, pp. 675–696, 2019.

[22] M. P. Vitus, W. Zhang, A. Abate, J. Hu, and C. J. Tomlin, "On efficient sensor scheduling for linear dynamical systems," *Automatica*, vol. 48, no. 10, pp. 2482–2493, 2012.

[23] V. Krishnamurthy, "Algorithms for optimal scheduling and management of hidden markov model sensors," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1382–1397, 2002.

[24] R. Garnett, M. A. Osborne, and S. J. Roberts, "Bayesian optimization for sensor set selection," in *Proceedings of the 9th ACM/IEEE international conference on information processing in sensor networks*, 2010, pp. 209–219.

[25] X. Yang and R. Niu, "Adaptive sensor selection for nonlinear tracking via sparsity-promoting approaches," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1966–1982, 2018.

[26] S. Singh, B.-N. Vo, A. Doucet, and R. Evans, "Stochastic approximation for optimal observer trajectory planning," in *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, vol. 6. IEEE, 2003, pp. 6313–6318.

[27] S. S. Singh, N. Kantas, B.-N. Vo, A. Doucet, and R. J. Evans, "Simulation-based optimal sensor scheduling with application to observer trajectory planning," *Automatica*, vol. 43, no. 5, pp. 817–830, 2007.

[28] A. Doucet, B.-N. Vo, C. Andrieu, and M. Davy, "Particle filtering for multi-target tracking and sensor management," in *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002.(IEEE Cat. No. 02EX5997)*, vol. 1. IEEE, 2002, pp. 474–481.

[29] S. Singh, B.-N. Vo, A. Doucet, and R. Evans, "Variance reduction for monte carlo implementation of adaptive sensor management," in *Proceedings of the international conference on information fusion*, 2004, pp. 901–908.

[30] Y. He and E. K. Chong, "Sensor scheduling for target tracking: A monte carlo sampling approach," *Digital Signal Processing*, vol. 16, no. 5, pp. 533–545, 2006.

[31] Y. Li, L. W. Krakow, E. K. Chong, and K. N. Groom, "Approximate stochastic dynamic programming for sensor scheduling to track multiple targets," *Digital Signal Processing*, vol. 19, no. 6, pp. 978–989, 2009.

[32] A. Aspeel, D. Dasnoy, R. M. Jungers, and B. Macq, "Optimal intermittent measurements for tumor tracking in x-ray guided radiotherapy," *Medical Imaging 2019: Image-Guided Procedures, Robotic Interventions, and Modeling*, vol. 10951, p. 109510C, 2019.

[33] A. Sano and M. Terao, "Measurement optimization in optimal process control," *Automatica*, vol. 6, no. 5, pp. 705–714, 1970.

[34] A. Aksenov, P.-O. Amblard, O. Michel, and C. Jutten, "Optimal measurement times for a small number of measures of a brownian motion over a finite period," *arXiv preprint arXiv:1902.06126*, 2019.

[35] A. Aspeel, A. Gouverneur, R. M. Jungers, and B. Macq, "Optimal measurement budget allocation for particle filtering," *27th IEEE International Conference on Image Processing (accepted for publication), arXiv preprint arXiv:2005.08557*, 2020.

[36] S. Saha, Y. Boers, H. Driessen, P. K. Mandal, and A. Bagchi, "Particle based map state estimation: A comparison," in *2009 12th International Conference on Information Fusion.* IEEE, 2009, pp. 278–283.

[37] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory.* SIAM, 2014.

[38] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[39] J. Lien, U. J. Ferner, W. Srichavengsup, H. Wymeersch, and M. Z. Win, "A comparison of parametric and sample-based message representation in cooperative localization," *International Journal of Navigation and Observation*, vol. 2012, 2012.

[40] A. Doucet and A. M. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," *Handbook of nonlinear filtering*, vol. 12, no. 656-704, p. 3, 2009.

[41] D. Delahaye, S. Chaimatanan, and M. Mongeau, "Simulated annealing: From basics to applications," in *Handbook of Metaheuristics.* Springer, 2019, pp. 1–35.

[42] M. Mitchell, *An introduction to genetic algorithms.* MIT press, 1998.

[43] A. Umbarkar and P. Sheth, "Crossover operators in genetic algorithms: A review." *ICTACT journal on soft computing*, vol. 6, no. 1, 2015.

[44] J. Dhont, J. Vandemeulebroucke, M. Burghelea, K. Poels, T. Depuydt, R. Van Den Begin, C. Jaudet, C. Collen, B. Engels, T. Reynders *et al.*, "The long-and short-term variability of breathing induced tumor motion in lung and liver over the course of a radiotherapy treatment," *Radiotherapy and Oncology*, vol. 126, no. 2, pp. 339–346, 2018.

[45] R. Wuyts, E. Vlemincx, K. Bogaerts, I. Van Diest, and O. Van den Bergh, "Sigh rate and respiratory variability during normal breathing and the role of negative affectivity," *International Journal of Psychophysiology*, vol. 82, no. 2, pp. 175–179, 2011.

[46] B. P. Carlin, N. G. Polson, and D. S. Stoffer, "A monte carlo approach to nonnormal and nonlinear state-space modeling," *Journal of the American Statistical Association*, vol. 87, no. 418, pp. 493–500, 1992.

[47] G. Kitagawa, "Monte carlo filter and smoother for non-gaussian non-linear state space models," *Journal of computational and graphical statistics*, vol. 5, no. 1, pp. 1–25, 1996.

[48] V. Kadirkamanathan, P. Li, M. H. Jaward, and S. 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.