



HAL
open science

Independent resampling sequential Monte Carlo algorithms

Roland Lamberti, Yohan Petetin, François Desbouvries, François Septier

► **To cite this version:**

Roland Lamberti, Yohan Petetin, François Desbouvries, François Septier. Independent resampling sequential Monte Carlo algorithms. *IEEE Transactions on Signal Processing*, 2017, 65 (20), pp.5318 - 5333. 10.1109/TSP.2017.2726971 . hal-01593422

HAL Id: hal-01593422

<https://hal.science/hal-01593422>

Submitted on 15 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Independent Resampling Sequential Monte Carlo Algorithms

Roland Lamberti, Yohan Petetin, François Desbouvries, and François Septier

Abstract

Sequential Monte Carlo algorithms, or Particle Filters, are Bayesian filtering algorithms which propagate in time a discrete and random approximation of the a posteriori distribution of interest. Such algorithms are based on Importance Sampling with a bootstrap resampling step which aims at struggling against weights degeneracy. However, in some situations (informative measurements, high dimensional model), the resampling step can prove inefficient. In this paper, we revisit the fundamental resampling mechanism which leads us back to Rubin's static resampling mechanism. We propose an alternative rejuvenation scheme in which the resampled particles share the same marginal distribution as in the classical setup, but are now independent. This set of independent particles provides a new alternative to compute a moment of the target distribution and the resulting estimate is analyzed through a CLT. We next adapt our results to the dynamic case and propose a particle filtering algorithm based on independent resampling. This algorithm can be seen as a particular auxiliary particle filter algorithm with a relevant choice of the first-stage weights and instrumental distributions. Finally we validate our results via simulations which carefully take into account the computational budget.

Index Terms

Sequential Monte Carlo algorithms; Particle Filters; Importance Sampling; Auxiliary Particle Filter; Resampling.

I. INTRODUCTION

Let $\{X_k \in \mathbb{R}^m\}_{k \geq 0}$ (resp. $\{Y_k \in \mathbb{R}^n\}_{k \geq 0}$) be a hidden (resp. observed) process. Let $X_{0:k}$, say, denote $\{X_i, 0 \leq i \leq k\}$, $x_{0:k} = \{x_i, 0 \leq i \leq k\}$, and let $p(x)$ (resp. $p(x|y)$), say, denote the probability density function (pdf) of random variable (r.v.) X (resp. of X given $Y = y$); capital letters are used for r.v. and lower case ones for their realizations. We assume that $\{(X_k, Y_k)\}_{k \geq 0}$ is a Hidden Markov chain, i.e. that

$$p(x_{0:k}, y_{0:k}) = p(x_0) \prod_{i=1}^k f_i(x_i | x_{i-1}) \prod_{i=0}^k g_i(y_i | x_i). \quad (1)$$

Roughly speaking, pdf $f_k(x_k | x_{k-1})$ describes the dynamical evolution of the Markovian hidden process $\{X_k\}_{k \geq 0}$ between time $k-1$ and time k while the likelihood $g_k(y_k | x_k)$ describes the relation at time k between an observation

R. Lamberti, Y. Petetin and F. Desbouvries are with Samovar, Telecom Sudparis, CNRS, Université Paris-Saclay, 9 rue Charles Fourier, 91011 Evry, France

F. Septier is with Telecom Lille & UMR CNRS CRISTAL, Rue Guglielmo Marconi, 59653 Villeneuve d'Ascq, France

y_k and the associated hidden state x_k . We address the problem of computing a moment of some function $f(\cdot)$ w.r.t. the filtering pdf $p(x_k|y_{0:k})$, i.e. the pdf of the hidden state given the past observations:

$$\Theta_k = \int f(x_k)p(x_k|y_{0:k})dx_k. \quad (2)$$

As is well known, Θ_k can be exactly computed only in very specific models, and one needs to resort to approximations in the general case. In this paper, we focus on a popular class of approximations called sequential Monte Carlo (SMC) algorithms or Particle Filters (PF), see e.g. [1]–[3] PF propagate over time a set of N Monte Carlo (MC) weighted samples $\{w_k^i, x_k^i\}_{i=1}^N$ which defines a discrete approximation $\sum_{i=1}^N w_k^i \delta_{x_k^i}$ of $p(x_k|y_{0:k})$ and enables to compute an estimate $\hat{\Theta}_k$ of Θ_k :

$$\hat{\Theta}_k = \sum_{i=1}^N w_k^i f(x_k^i). \quad (3)$$

More precisely, the computation of the set $\{w_k^i, x_k^i\}_{i=1}^N$ is based on the sequential application of the Importance Sampling (IS) mechanism [4]. This mechanism consists in sampling particles according to an importance distribution and next weighting these samples in order to correct the discrepancy between the target and the importance distribution. However the direct sequential application of the IS mechanism in model (1) fails in practice since after a few time steps most weights get close to zero, while only a few particles have non negligible weights. Consequently IS alone becomes more and more inefficient since a lot of computational effort is devoted to sampling particles which will hardly contribute to the estimate $\hat{\Theta}_k$ in (3).

As is well known, a traditional rescue against weights degeneracy consists in *resampling* the particles (- either at each time step or depending on some criterion such as the number of efficient particles [5] [6] [7] [8]), i.e. of re-drawing each particle with a probability equal to its weight. This yields the class of Sampling Importance Resampling (SIR) algorithms [9] [1] [10] [11]. This resampling (i.e., bootstrap) mechanism has proved to be beneficial in the long run, but its instantaneous effects are mitigated; though the resampling step indeed discards particles with low weights (such particles are likely never to be resampled), particles with significant weights are resampled several times, which results in dependency among the resampled points and support shrinkage. Consequently, particle filters based on the resampling mechanism can give poor results in some Markovian models (1), such as informative models where the likelihood $g_k(y_k|x_k)$ is sharp. Our aim in this paper is thus to revisit this key rejuvenation scheme in order to design new PF algorithms which would keep the benefits of the resampling mechanism, while avoiding the local impoverishment of the resulting MC approximation of the filtering distribution.

To that end we begin with revisiting the SIR mechanism at one single time step $k \rightarrow k+1$. This leads us back to an analysis of Rubin’s static SIR mechanism [12, §2] [13] [9] [14, §9.2], in which, roughly speaking, one obtains samples x^j approximately drawn from a target distribution p by drawing intermediate samples $\{\tilde{x}^i\}_{i=1}^N$ from an instrumental distribution q , and next selecting x^j among $\{\tilde{x}^i\}_{i=1}^N$ with a probability proportional to $\frac{p(\tilde{x}^i)}{q(\tilde{x}^i)}$. We first observe that the samples $\{x^j\}$ produced by this SIR mechanism are dependent and marginally distributed from some compound pdf $\tilde{q}_N = \phi(p, q, N)$ which takes into account the effects of both pdfs p and q . Here the dependency is detrimental, because samples that would be i.i.d from \tilde{q}_N would produce, whichever the number of sampled and resampled particles, a moment estimate with reduced variance; this result is further illustrated by a central limit

theorem (CLT) which is compared to the existing CLTs for the static IS estimate (based on the pre-resampling samples $\{\tilde{x}^i\}_{i=1}^N$), on the one hand, and for the SIR estimate (based on the post-resampling ones $\{x^j\}_{j=1}^{M_N}$), on the other hand.

We next propose a procedure to obtain i.i.d. samples from \tilde{q}_N , which leads to the computation of two point estimates of $\Theta = \int f(x)p(x)dx$. The first one is based on unweighted i.i.d. samples and is an improved version of the classical (i.e., dependent) SIR estimate; the second one is based on post-resampling-weighted i.i.d. samples and can be seen as new IS estimate, based on the compound pdf \tilde{q}_N . Finally we adapt these results to the sequential computation of Θ_k in model (1). We thus propose two new PF algorithms. One of them has an interesting interpretation in terms of Auxiliary Particle Filter (APF); more precisely, that algorithm naturally produces a relevant importance mixture distribution from which it is easy to sample. We finally illustrate our results via simulations, and carefully compare our algorithms with existing ones in terms of Root Mean Square Error (RMSE) and computational cost. The rest of this paper is organized as follows. Section II is devoted to the static case. In section III we address the sequential case, and derive new PF based on the results of section II. In section IV we perform simulations and discuss implementation issues, and we end the paper with a conclusion.

II. IS WITH RESAMPLING VIEWED AS A COMPOUND IS SCHEME

As recalled in the introduction, resampling from time to time is a standard rescue when applying IS in the sequential case. In this section we thus focus on one such time step $k \rightarrow k + 1$. This amounts to revisiting Rubin's static SIR mechanism (see section (II-A)), which consists in resampling points $\{x_i\}_{i=1}^{M_N}$ from the weighted distribution $\sum_{i=1}^N w_i \delta_{\tilde{x}_i}$ where $\tilde{x}_i \stackrel{i.i.d.}{\sim} q$ and the pre-resampling weights $w_i \propto \frac{p(\tilde{x}_i)}{q(\tilde{x}_i)}$ with $\sum_{i=1}^N w_i = 1$. As is well known, when $N \rightarrow \infty$ the resampled points $\{x_i\}_{i=1}^{M_N}$ become asymptotically i.i.d. from the target distribution p . For finite N however, these samples are dependent and drawn from some pdf \tilde{q}_N which differs from p and can indeed be seen as a compound IS density $\tilde{q}_N = \phi(p, q, N)$ produced by the succession of the sampling (S), weighting (W) and resampling (R) steps. We discuss on the benefits of drawing independent samples from \tilde{q}_N (see section II-B), and next on reweighting these independent samples with post-resampling weights $w'_i \propto \frac{p(x_i)}{\tilde{q}_N(x_i)}$ (see section II-C). In all this section we assume the scalar case for simplicity. We end the section with a summary (see section II-D).

A. The dependent SIR mechanism

Let us begin with a brief review of Rubin's classical SIR sampling mechanism and of the properties of the sampled and resampled particles.

1) *Properties of the sampled particles $\{\tilde{x}^i\}_{i=1}^N$* : In the context of this paper we first recall the principle of IS. Let $p(x)$ be a probability density function and assume that we want to compute

$$\Theta = \int f(x)p(x)dx = E_p(f(X)). \quad (4)$$

In the Bayesian framework $p(x)$ is generally only known up to a constant, i.e. $p(x) \propto p_u(x)$ (subscript u is for unnormalized) and it is not possible to obtain samples directly drawn from $p(x)$. A solution is to introduce

an importance distribution $q(x)$ which satisfies $q(x) > 0$ when $p(x) > 0$ and to rewrite Θ as the ratio of two expectations w.r.t. q ,

$$\Theta = \frac{\int f(x) \frac{p_u(x)}{q(x)} q(x) dx}{\int \frac{p_u(x)}{q(x)} q(x) dx} = \frac{E_q(f(X) \frac{p(X)}{q(X)})}{E_q(\frac{p(X)}{q(X)})}. \quad (5)$$

Next, each expectation is approximated by a Monte Carlo method based on N i.i.d. samples $(\tilde{x}^1, \dots, \tilde{x}^N)$ drawn from $q(\cdot)$; the IS estimate of Θ is given by

$$\hat{\Theta}_N^{\text{IS}} = \sum_{i=1}^N w^i f(\tilde{x}^i) = E_{\hat{p}}(f(X)) \quad (6)$$

where

$$\hat{p}(x) = \sum_{i=1}^N w^i \delta_{\tilde{x}^i}(x) \quad (7)$$

and where w^i (the i -th normalized importance weight) reads

$$w^i = \frac{\frac{p_u(\tilde{x}^i)}{q(\tilde{x}^i)}}{\sum_{j=1}^N \frac{p_u(\tilde{x}^j)}{q(\tilde{x}^j)}} = \frac{\frac{p(\tilde{x}^i)}{q(\tilde{x}^i)}}{\sum_{j=1}^N \frac{p(\tilde{x}^j)}{q(\tilde{x}^j)}}. \quad (8)$$

As is well known [4], under mild assumptions

$$\hat{\Theta}_N^{\text{IS}} \xrightarrow{\text{a.s.}} \Theta, \quad (9)$$

and a CLT is available too ($\mathcal{D} \rightarrow$ denotes the convergence in distribution):

$$\sqrt{N}(\hat{\Theta}_N^{\text{IS}} - \Theta) \xrightarrow{\mathcal{D}} \mathcal{N}\left(0, E_q\left(\frac{p^2(X)}{q^2(X)}(f(X) - \Theta)^2\right)\right). \quad (10)$$

2) *Properties of the resampled particles* $\{x^i\}_{i=1}^{M_N}$: From (9) and (10), \hat{p} can be seen as a discrete approximation of the target density p , and one expects that for large N , (re)sampling from \hat{p} would produce samples approximately drawn from p . This is the rationale of Rubin's SIR mechanism [12, §2], [13], [9], [14, §9.2]. More precisely, let us as above draw N i.i.d. samples \tilde{x}^i from q , and next M_N i.i.d samples x^i from \hat{p} in (7). It is indeed well-known (see [9] [12]) that when $N \rightarrow \infty$, each r.v. x^i produced by this mechanism converges in distribution to $p(\cdot)$, so Rubin's technique can be seen as a two-step sampling mechanism which transforms samples drawn from q into samples (approximately) drawn from p .

This convergence result can be completed by a CLT which involves the estimate of Θ based on the unweighted set $\{(\frac{1}{M_N}, x^i)\}_{i=1}^{M_N}$:

$$\hat{\Theta}_{M_N}^{\text{SIR}} = \frac{1}{M_N} \sum_{i=1}^{M_N} f(x^i). \quad (11)$$

Let $N \rightarrow \infty$, let M_N be a non decreasing sequence with $M_N \rightarrow \infty$, and let $\lim_{M_N} \frac{N}{M_N} = \alpha > 0$ (possibly ∞); then under mild conditions (see e.g. [14, §9])

$$\sqrt{M_N}(\hat{\Theta}_{M_N}^{\text{SIR}} - \Theta) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \text{var}_p(f(X)) + \alpha^{-1} E_q\left(\left(\frac{p^2(X)}{q^2(X)}(f(X) - \Theta)^2\right)\right)). \quad (12)$$

If $\alpha \rightarrow \infty$ then the asymptotic variance tends to $\text{var}_p(f(X))$, which shows that the SIR estimate asymptotically has the same behavior as a crude Monte Carlo estimate directly deduced from M_N samples according to the target distribution $p(\cdot)$, provided the number N of intermediate samples is large compared to M_N .

However, for computational reasons, the number of samples N and M_N should not be too large in practice. Consequently we now focus on the samples produced by the SIR procedure from a non asymptotical point of view and we have the following result (the proof is given in the Appendix).

Proposition 1: Let us consider the samples $\{x^i\}_{i=1}^{M_N}$ produced by the SIR mechanism described above. Then these samples are identically distributed according to a pdf \tilde{q}_N , with

$$\tilde{q}_N(x) = Nh_N(x)q(x), \quad (13)$$

$$h_N(x) = \int \int \frac{\frac{p(x)}{q(x)}}{\frac{p(x)}{q(x)} + \sum_{l=1}^{N-1} \frac{p(x^l)}{q(x^l)}} \prod_{l=1}^{N-1} q(x^l) dx^l. \quad (14)$$

So for fixed sample size N , the SIR mechanism produces dependent samples $\{x^i\}_{i=1}^{M_N}$ distributed from \tilde{q}_N (these samples are independent given the intermediate set $\{\tilde{x}^i\}_{i=1}^N$, but become dependent when this conditioning is removed). In practice, this dependency results in support shrinkage since, by construction, an intermediate sample \tilde{x}^i can be resampled several times, and $\{x^i\}_{i=1}^{M_N}$ is a subset of $\{\tilde{x}^i\}_{i=1}^N$. For instance let $M_N = N$. If we assume that $w^j = 1$ for some j and $w^i = 0$ for $i \neq j$, then $x^i = \tilde{x}^j$ for all i . By contrast, if $w^i = 1/N$ for all i , then the average number of different samples $\{x^i\}_{i=1}^N$ is approximately $N/3$ [15]. Nevertheless the resampling step remains useful in a dynamic setup (see section III): even though locally it leads to an impoverishment of the diversity, this step is critical for recreating diversity at the next time step.

B. The independent SIR mechanism

Observe that the two factors in (13) reflect the effects of the sampling and resampling step: pdf q is used in the S step, while $h_N(x)$, which can be interpreted as the conditional expectation of a normalized importance weight when its associated particle is x , results from the (W,R) steps. So particles drawn from \tilde{q}_N are likely to be in regions where 1) q is large (since these particles have first been sampled); and 2) which have also been resampled because their associated weight was large enough. Now our objective is to propose an alternative mechanism which, in the sequential case, will produce the same positive effect as the classical SIR mechanism (i.e. fighting against weight degeneracy by eliminating the samples with weak importance weights), while ensuring the diversity of the final support. Such a support diversity is ensured if we draw samples *independently* from the continuous pdf $\tilde{q}_N(\cdot)$. We first study the potential benefits of this sampling mechanism (see section II-B1) and next discuss its implementation (see section II-B2).

1) Statistical properties: Let us now assume that we have at our disposal a set of M_N i.i.d. samples $\{\bar{x}^i\}_{i=1}^{M_N}$ drawn from $\tilde{q}_N(\cdot)$ defined in (13) (14). Before addressing the practical computation of such a set (see section II-B2), let us study its properties by considering the crude estimate of Θ based on these M_N i.i.d samples:

$$\hat{\Theta}_{M_N}^{\text{I-SIR}} = \frac{1}{M_N} \sum_{i=1}^{M_N} f(\bar{X}^i). \quad (15)$$

(I in notation I-SIR stands for independent). Our aim is to compare $\hat{\Theta}_{M_N}^{\text{I-SIR}}$ to $\hat{\Theta}_{M_N}^{\text{SIR}}$, and more generally $\hat{\Theta}_{M_N}^{\text{IS}}$, $\hat{\Theta}_{M_N}^{\text{SIR}}$ and $\hat{\Theta}_{M_N}^{\text{I-SIR}}$. We first have the following result (the proof is given in the Appendix).

Proposition 2: Let us consider the three estimates $\widehat{\Theta}_N^{\text{IS}}$, $\widehat{\Theta}_{M_N}^{\text{SIR}}$ and $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ defined in (6), (11) and (15) respectively. Then

$$\mathbb{E}(\widehat{\Theta}_N^{\text{IS}}) = \mathbb{E}(\widehat{\Theta}_{M_N}^{\text{SIR}}) = \mathbb{E}(\widehat{\Theta}_{M_N}^{\text{I-SIR}}), \quad (16)$$

$$\text{var}(\widehat{\Theta}_{M_N}^{\text{SIR}}) = \text{var}(\widehat{\Theta}_{M_N}^{\text{I-SIR}}) + \frac{M_N - 1}{M_N} \text{var}(\widehat{\Theta}_N^{\text{IS}}). \quad (17)$$

Equation (17) ensures that an estimate based on independent samples obtained from \tilde{q}_N outperforms the classical SIR estimate; the gain of $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ w.r.t. $\widehat{\Theta}_{M_N}^{\text{SIR}}$ depends on the variance of $\text{var}(\widehat{\Theta}_N^{\text{IS}})$. On the other hand it is well known (see e.g. [14, p. 213]) that $\text{var}(\widehat{\Theta}_{M_N}^{\text{SIR}}) = \text{var}(\widehat{\Theta}_N^{\text{IS}}) + \mathbb{E}(\text{var}(\widehat{\Theta}_{M_N}^{\text{SIR}} | \{\tilde{x}^i\}_{i=1}^N))$; so both $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ and $\widehat{\Theta}_N^{\text{IS}}$ are preferable to $\widehat{\Theta}_{M_N}^{\text{SIR}}$.

On the other hand, comparing the variance of $\widehat{\Theta}_N^{\text{IS}}$ to that of $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ is more difficult, because we have to compare $\frac{1}{M_N} \text{var}_{\tilde{q}_N}(f(\overline{X}))$ to $\text{var}(\sum_{i=1}^N w^i(\tilde{X}^1, \dots, \tilde{X}^N) f(\tilde{X}_i))$ where $\tilde{X}^i \stackrel{\text{i.i.d.}}{\sim} q(\cdot)$. However, we have the following CLT (the proof is given in the Appendix).

Theorem 1: Let us consider the independent SIR estimate defined in (15). Let assume that $N \rightarrow \infty$, M_N is a non decreasing sequence with $M_N \rightarrow \infty$ and $\lim_{N \rightarrow \infty} \frac{N}{M_N} = \alpha > 0$. Then $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ satisfies

$$\sqrt{M_N}(\widehat{\Theta}_{M_N}^{\text{I-SIR}} - \Theta) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \text{var}_p(f(X))). \quad (18)$$

Let us comment this result. First Theorem 1 enables again to compare $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ to $\widehat{\Theta}_{M_N}^{\text{SIR}}$. Comparing (12) and (18) confirms (17), since the asymptotic variance of $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ is always lower than that of $\widehat{\Theta}_{M_N}^{\text{SIR}}$. Also note that in the independent case the asymptotic variance of $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ no longer depends on $\alpha > 0$.

Next Theorem 1 also gives some elements for comparing $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ to $\widehat{\Theta}_N^{\text{IS}}$. Let for simplicity $M_N = N$. Then the comparison of both estimates relies on that of the asymptotic variances in (10) and (18):

$$\sigma_\infty^{2,\text{IS}}(q) = \mathbb{E}_q \left(\frac{p^2(X)}{q^2(X)} (f(X) - \Theta)^2 \right), \quad (19)$$

$$\sigma_\infty^{2,\text{I-SIR}} = \text{var}_p(f(X)). \quad (20)$$

For a given target pdf $p(\cdot)$ and function $f(\cdot)$, $\sigma_\infty^{2,\text{IS}}(q)$ depends on the importance pdf $q(\cdot)$ and is well known [16, §2.9] [4, Theorem 3] to be minimum for $q^*(x) \propto p(x)|f(x) - \Theta|$; for that q^* , $\sigma_\infty^{2,\text{IS}}(q^*) = (\int |f(x) - \Theta| p(x) dx)^2 \leq \text{var}_p(f(X))$, so $\widehat{\Theta}_N^{\text{IS}}(q^*)$ outperforms $\widehat{\Theta}_N^{\text{I-SIR}}$ for large values of N . On the other hand for other importance distributions $\sigma_\infty^{2,\text{IS}}(q)$ may become larger than $\sigma_\infty^{2,\text{I-SIR}}$. Also note that the variances in (19) and (20) depend on function $f(\cdot)$; on the other hand, for large N , $\widehat{\Theta}_N^{\text{I-SIR}}$ has the same behavior as a crude estimate built from samples drawn from $p(\cdot)$ and so is adapted for a large class of functions $f(\cdot)$.

2) *Sampling procedure:* It remains to describe a procedure to obtain i.i.d. samples from \tilde{q}_N . Algorithm 1 ensures that the final samples $\{x^i, \dots, x^{M_N}\}$ are drawn independently from \tilde{q}_N .

Compared to the classical SIR procedure, the independent SIR algorithm described in Algorithm 1 relies on a sampling step of $N \times M_N$ intermediate samples \tilde{x} and M_N independent resampling steps. Consequently, for a given budget of sampling and resampling steps, the independent procedure should be compared with a classical SIR one in which we sample $N \times M_N$ points and resample M_N of them. In this last case, we obtain M_N dependent samples drawn from $\tilde{q}_{N \times M_N}$. First, using (12) with $\alpha = \lim_{N \rightarrow \infty} \frac{N \times M_N}{M_N} = \infty$, we see that both estimates $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ and

Input: an importance distribution q , N and M_N

Result: $\{x^i\}_{i=1}^{M_N} \stackrel{\text{i.i.d.}}{\sim} \tilde{q}_N$

for $1 \leq i \leq M_N$ **do**

for $1 \leq j \leq N$ **do**

S. $\tilde{x}^{i,j} \sim q(\cdot)$;

W. $w^{i,j} \propto p_u(\tilde{x}^{i,j})/q(\tilde{x}^{i,j})$, $\sum_{j=1}^N w^{i,j} = 1$;

end

R. $x^i \sim \sum_{j=1}^N w^{i,j} \delta_{\tilde{x}^{i,j}}$

end

Algorithm 1: The independent SIR algorithm

$\hat{\Theta}_{M_N}^{\text{SIR}}$ with $N \times M_N$ intermediate samples have the same asymptotic behavior. However the independent procedure can be easily parallelized because the resampling steps are by nature independent contrary to the SIR procedure where the $N \times M_N$ intermediate samples are directly resampled.

C. Reweighting the independent samples?

We finally discuss the final weights which are attributed to the resampled particles. In the SIR procedure, each final sample is weighted by $1/M_N$. From an IS point of view, this weighting traduces the fact that the final samples become drawn from the target distribution $p(\cdot)$ and independent when $N \rightarrow \infty$ [12]. Moreover the convergence results of $\hat{\Theta}_{M_N}^{\text{I-SIR}}$ to Θ (see e.g. [17] [14]) confirm that these weights are valid from an asymptotical point of view. In the independent SIR procedure, the only difference is that the final samples are independent, even from a non-asymptotical point of view.

Now, if N is finite, one can wonder if weights $1/M_N$ are optimal. In Algorithm 1, samples $\{X^i\}_{i=1}^{M_N}$ are independent and sampled from \tilde{q}_N . Consequently, for a given N , \tilde{q}_N can be seen as a post-resampling compound importance distribution $\tilde{q}_N = \phi(p, q, N)$, and a final sample x^i should be weighted by a post-resampling weight proportional to $p_u(x^i)/\tilde{q}_N(x^i)$. This yields a new estimate $\hat{\Theta}_{M_N}^{\text{I-SIR-w}}$ of (4) (superscript w stands for weighted)

$$\hat{\Theta}_{M_N}^{\text{I-SIR-w}} = \sum_{i=1}^{M_N} \frac{\frac{p_u(x^i)}{\tilde{q}_N(x^i)}}{\sum_{j=1}^{M_N} \frac{p_u(x^j)}{\tilde{q}_N(x^j)}} f(x^i), \quad (21)$$

which coincides with the IS estimate (6) with importance distribution $\tilde{q}_N(\cdot)$. It is difficult to compare $\hat{\Theta}_{M_N}^{\text{I-SIR}}$ and $\hat{\Theta}_{M_N}^{\text{I-SIR-w}}$ because the expression of the weights in this last case depends on N . However, it is interesting to note that contrary to $\hat{\Theta}_{M_N}^{\text{I-SIR}}$, M_N impacts the bias of the estimate $\hat{\Theta}_{M_N}^{\text{I-SIR-w}}$. For example, if we set $N = 1$ (so $q_N = q$) and M_1 is arbitrary then $\hat{\Theta}_{M_1}^{\text{I-SIR-w}}$ coincides with the IS estimate with M_1 i.i.d. samples drawn from q while the unweighted estimate $\hat{\Theta}_{M_1}^{\text{I-SIR}}$ is a crude estimate of $\int f(x)q(x)dx$ and is not adapted for the estimation of Θ . More generally, using the delta method to approximate $E(\hat{\Theta}_{M_N}^{\text{I-SIR}})$ and $E(\hat{\Theta}_{M_N}^{\text{I-SIR-w}})$ [18] we observe that

$$E(\hat{\Theta}_{M_N}^{\text{I-SIR}}) = E(\hat{\Theta}_N^{\text{IS}}) \approx \Theta - \frac{1}{N} E_q \left(\frac{p^2(X)}{q^2(X)} (f(X) - \Theta) \right), \quad (22)$$

$$E(\hat{\Theta}_{M_N}^{\text{I-SIR-w}}) \approx \Theta - \frac{1}{M_N} E_{\tilde{q}_N} \left(\frac{p^2(X)}{\tilde{q}_N^2(X)} (f(X) - \Theta) \right). \quad (23)$$

So for a fixed number of sampled points N , we see that in the unweighted case the bias of $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ is independent of M_N . By contrast, whichever N the bias of $E(\widehat{\Theta}_{M_N}^{\text{I-SIR-w}})$ tends to 0 as $M_N \rightarrow \infty$.

Finally, it remains to compute $p_u(x^i)/\tilde{q}_N(x^i)$ in practice. In general, \tilde{q}_N in (13) is not available in close form because it relies on the integral $h_N(x)$ in (14). However, the $N \times M_N$ intermediate samples which have been used in Algorithm 1 can be recycled to approximate the conditional expectation $h_N(x)$. For a given x and using the intermediate samples $\tilde{x}^{i,j}$ of Algorithm 1, a crude Monte Carlo estimate of $h_N(x)$ reads

$$\hat{h}_N(x) = \sum_{i=1}^{M_N} \frac{\frac{p_u(x)}{q(x)}}{\frac{p_u(x)}{q(x)} + \sum_{j=1}^{N-1} \frac{p_u(\tilde{x}^{i,j})}{q(\tilde{x}^{i,j})}}. \quad (24)$$

Importance weights $\frac{p_u(x)}{\tilde{q}_N(x)}$ in (21) can be approximated by $\frac{p_u(x)}{N\hat{h}_N(x)q(x)}$. Note that the computation of these approximated weights do not require extra computational cost since $p_u(\tilde{x}^{i,j})/q(\tilde{x}^{i,j})$ has already been computed in Algorithm 1 to obtain i.i.d. samples.

D. Summary

In summary, we now have at our disposal four estimates to compute Θ in (4) from an importance distribution $q(\cdot)$. $\widehat{\Theta}_N^{\text{IS}}$ and $\widehat{\Theta}_{M_N}^{\text{SIR}}$ are deduced from the IS and Rubin's SIR mechanisms, respectively. $\widehat{\Theta}_{M_N}^{\text{SIR}}$ relies on unweighted dependent samples from \tilde{q}_N . Using unweighted independent samples from \tilde{q}_N produces the estimate $\widehat{\Theta}_{M_N}^{\text{I-SIR}}$ which outperforms $\widehat{\Theta}_{M_N}^{\text{SIR}}$ and possibly $\widehat{\Theta}_N^{\text{IS}}$; it also becomes asymptotically independent of the choice of the initial importance distribution $q(\cdot)$ according to theorem 1. This estimate does not suffer from the support impoverishment caused by the resampling step. On the other hand it requires a larger computational cost which, however, can be exploited in order to associate to the i.i.d. samples post-resampling importance weights based on the $\tilde{q}_N(x)$. We thus obtain a weighted estimate $\widehat{\Theta}_{M_N}^{\text{I-SIR-w}}$ which can be seen as the estimate deduced from the IS mechanism based on the compound IS distribution $\tilde{q}_N(x)$. We will compare these estimates via simulations and will take into account their computational cost in Section IV-A.

III. INDEPENDENT RESAMPLING BASED PF

We now adapt the results of Section II to the Bayesian filtering problem. In section III-A we briefly recall the principle of classical SIR algorithms which are based on dependent resampling. Our SIR algorithm with independent resampling and unweighted samples is proposed in section III-B. However, computing the post-resampling weights is more challenging here than in the static case because the pdf \tilde{q}_N of the static case becomes a sum of N terms which should be computed for each final sample. So in section III-C we revisit the algorithm of section III-B in terms of APF. We first observe that the independent SIR algorithm can be seen as the first step of an APF algorithm since it implicitly draws samples from a mixture pdf. Making full use of the APF methodology enables us to weight our final samples.

A. Classical SIR algorithms (based on dependent resampling)

We now assume that we are given some hidden Markov model (1) and we briefly recall how Θ_k in (2) can be computed recursively via PF. PF relies on the sequential application of the normalized IS mechanism described in

Section II-A for the target distribution $p(x_{0:k}|y_{0:k})$ which is known up to a constant according to (1). Let $q(x_{0:k})$ be an importance distribution ($q(x_{0:k})$ can depend on $y_{0:k}$ but this dependency is not written here to avoid notational burden). Starting from N weighted trajectories $x_{0:k-1}^i$ sampled from $q(x_{0:k-1})$, we first extend each trajectory $x_{0:k-1}^i$ by a particle \tilde{x}_k^i sampled from $q(x_k|x_{0:k-1}^i)$ and next update the old weights w_{k-1}^i via

$$w_k^i \propto w_{k-1}^i \frac{f_k(\tilde{x}_k^i|x_{k-1}^i)g_k(y_k|\tilde{x}_k^i)}{q(\tilde{x}_k^i|x_{0:k-1}^i)}, \sum_{i=1}^N w_k^i = 1. \quad (25)$$

Unfortunately, it is well-known that this direct sequential application of IS leads to weight degeneracy: after a few iterations only few weights w_k^i have a non null value [19]. A traditional rescue consists in resampling, either systematically or according to some criterion such as the Effective Sample Size [5] [6] which is approximated by $1/\sum_{i=1}^N (w_k^i)^2$. The corresponding algorithm is given in Algorithm 2 and we shall assume that the size N of the MC approximation remains constant throughout the iterations. Finally Algorithm 2 enables to compute two estimates of Θ_k :

$$\hat{\Theta}_{N,k}^{\text{SIS}} = \sum_{i=1}^N w_k^i f(\tilde{X}_k^i), \quad (26)$$

$$\hat{\Theta}_{N,k}^{\text{SIR}} = \frac{1}{N} \sum_{i=1}^N f(X_k^i). \quad (27)$$

As is well known, the pre-resampling estimator $\hat{\Theta}_{N,k}^{\text{SIS}}$ is preferable to the post-resampling one $\hat{\Theta}_{N,k}^{\text{SIR}}$ and should be used in practice; but $\hat{\Theta}_{N,k}^{\text{SIR}}$ is recalled here because it will be compared below to the independent resampling estimator (32).

Input: $q(x_k|x_{0:k-1})$, y_k , $\{w_{k-1}^i, x_{0:k-1}^i\}_{i=1}^N$

Result: $\{w_k^i, x_{0:k}^i\}_{i=1}^N$

for $1 \leq i \leq N$ **do**

S. $\tilde{x}_k^i \sim q(x_k|x_{0:k-1}^i)$;
W. $w_k^i \propto w_{k-1}^i \frac{f_k(\tilde{x}_k^i|x_{k-1}^i)g_k(y_k|\tilde{x}_k^i)}{q(\tilde{x}_k^i|x_{0:k-1}^i)}$, $\sum_{i=1}^N w_k^i = 1$;
end

end

if *Resampling* **then**

for $1 \leq i \leq N$ **do**

R. $l^i \sim \text{Pr}(L = l | \{x_{0:k-1}^j, \tilde{x}_k^j\}_{j=1}^N) = w_k^l$
Set $x_{0:k}^i = (x_{0:k-1}^{l^i}, \tilde{x}_k^{l^i})$, $w_k^i = \frac{1}{N}$
end

else

Set $\{x_k^i\}_{i=1}^N = \{\tilde{x}_k^i\}_{i=1}^N$

end

Algorithm 2: The classical SIR algorithm (based on dependent resampling)

In practice, it remains to choose the conditional importance distribution $q(x_k|x_{0:k-1})$. A popular solution consists in choosing $q(x_k|x_{0:k-1}) = f_k(x_k|x_{k-1})$, since this pdf is part of model (1) and is generally easy to sample from;

another one is the so-called optimal conditional importance distribution $q(x_k|x_{0:k-1}) = p(x_k|x_{k-1}, y_k)$ which takes into account the new observation y_k and for which weights w_k^i no longer depend on the sampled particles $\{\tilde{x}_k^i\}_{i=1}^N$. The optimal conditional importance distribution is generally not available in closed form but some approximation techniques have been proposed, see e.g. [19] [20] [21]. The choice of the importance distribution will be not discussed in this paper and does not impact the proposed methodology. Finally, let us mention that convergence results are also available for the PF presented in Algorithm 2, see e.g. [22] [23] [17] [14]. Some of them are based on the recursive application of the CLTs recalled in Section II.

B. An alternative SIR algorithm (based on independent resampling)

Let us first adapt Proposition 1 to the sequential context. So we address the conditional distribution given $\{x_{0:k-1}^i\}_{i=1}^N$ of the resampled particles X_k^i and we have the following result (the proof is omitted).

Proposition 3: Let us consider the samples $\{X_k^i\}_{i=1}^N$ produced by the SIR mechanism of Algorithm 2. Let

$$p_{i,k}(x) = w_{k-1}^i f_k(x|x_{k-1}^i) g_k(y_k|x), \quad (28)$$

$$q_{i,k}(x) = q(x|x_{0:k-1}^i). \quad (29)$$

Then given the initial trajectories $\{x_{0:k-1}^i\}_{i=1}^N$, the new samples $\{X_k^i\}_{i=1}^N$ are identically distributed according to a pdf $\tilde{q}_{N,k}$ which reads

$$\tilde{q}_{N,k}(x) = \sum_{i=1}^N h_{i,k}(x) q_{i,k}(x), \quad (30)$$

where $h_{i,k}(x)$ coincides with the conditional expectation (given $(X_k^i = x)$) of the i -th importance weight at time k ,

$$h_{i,k}(x) = \int \int \frac{\frac{p_{i,k}(x)}{q_{i,k}(x)}}{\frac{p_{i,k}(x)}{q_{i,k}(x)} + \sum_{l \neq i} \frac{p_{l,k}(x^l)}{q_{l,k}(x^l)}} \prod_{l \neq i} q_{l,k}(x^l) dx^l. \quad (31)$$

Note that in this proposition we focus on the distribution of X_k^i given $\{x_{0:k-1}^i\}_{i=1}^N$. Given $\{x_{0:k-1}^i, \tilde{x}_k^i\}_{i=1}^N$, the new samples $\{X_k^i\}_{i=1}^N$ are independent; when we remove the dependency in $\{\tilde{x}_k^i\}_{i=1}^N$, $\{X_k^i\}_{i=1}^N$ become identically distributed according to $\tilde{q}_{N,k}$ but are dependent (a same particle can be resampled several times).

Since $\tilde{q}_{N,k}$ is a pdf, a procedure which would produce samples conditionally i.i.d. from $\tilde{q}_{N,k}$ would enable us to keep the advantage of the resampling step, i.e. to recreate diversity for the next time iteration while avoiding local impoverishment of the support. Except in a particular case which will be described later, sampling directly from $\tilde{q}_{N,k}(x)$ is difficult for an arbitrary conditional importance distribution $q(x_k|x_{0:k-1})$. We thus propose a procedure similar to Algorithm 1 but adapted to the dynamical context. The SIR algorithm with independent resampling is given by Algorithm 3. Note that a difference with Algorithm 2 is that the distribution of the discrete index L^i now depends on i .

We now propose a new estimate $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ of Θ which is based on the set $\{X_k^i\}_{i=1}^N$ produced by Algorithm 3:

$$\hat{\Theta}_{N,k}^{\text{I-SIR}} = \frac{1}{N} f(X_k^i). \quad (32)$$

Input: $q(x_k|x_{0:k-1}), y_k, \{w_{k-1}^i, x_{0:k-1}^i\}_{i=1}^N$
Result: $\{w_k^i, x_{0:k}^i\}_{i=1}^N$
for $1 \leq i \leq N$ **do**
 for $1 \leq j \leq N$ **do**
 S. $\tilde{x}_k^{i,j} \sim q(x_k|x_{0:k-1}^i);$
 W. $w_k^{i,j} \propto \frac{f_k(\tilde{x}_k^{i,j}|x_{k-1}^i)g_k(y_k|\tilde{x}_k^{i,j})}{q(\tilde{x}_k^{i,j}|x_{0:k-1}^i)}, \sum_{j=1}^N w_k^{i,j} = 1;$
 end
 R. $l^i \sim \Pr(L = l | \{x_{0:n-1}^j, \tilde{x}_n^{i,j}\}_{j=1}^N) = w_k^{i,l}$
 Set $x_{0:k}^i = (x_{0:k-1}^{l^i}, \tilde{x}_k^{i,l^i}), w_k^i = \frac{1}{N}$
end

Algorithm 3: A SIR algorithm based on independent resampling

Comparing (32) with (27), remember that the samples $\{X_k^i\}_{i=1}^N$ share the same pdf $\tilde{q}_{N,k}$, but that in (32) they are now independent given $\{x_{0:k-1}^i\}_{i=1}^N$. Starting from a dataset $\{x_{0:k-1}^i\}_{i=1}^N$, it is ensured that $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ outperforms $\hat{\Theta}_k^{\text{SIR}}$ since

$$\mathbb{E}(\hat{\Theta}_{N,k}^{\text{I-SIR}} | \{x_{0:k-1}^i\}_{i=1}^N) = \mathbb{E}(\hat{\Theta}_n^{\text{SIR}} | \{x_{0:k-1}^i\}_{i=1}^N), \quad (33)$$

$$\text{var}(\hat{\Theta}_{N,k}^{\text{SIR}} | \{x_{0:k-1}^i\}_{i=1}^N) = \text{var}(\hat{\Theta}_n^{\text{I-SIR}} | \{x_{0:k-1}^i\}_{i=1}^N) + \frac{N-1}{N} \text{var}(\hat{\Theta}_{N,k}^{\text{SIS}} | \{x_{0:k-1}^i\}_{i=1}^N). \quad (34)$$

Of course, computing $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ via the samples produced by Algorithm 3 requires an extra computational cost. This point will be discussed in detail in our Simulations section, but for the moment let us make two comments: first, this algorithm can be seen as an alternative resampling scheme which ensures the diversity of the resampled support without changing the conditional distribution of the final samples; if resampling needs to be performed rarely, then the independent resampling procedure may be used only when necessary. On the other hand, we will see that $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ can also provide an interesting alternative to $\hat{\Theta}_{N,k}^{\text{SIS}}$ but requires an extra computational cost; so if we want to perform the independent resampling procedure at each time step we will decrease the number N of particles associated with $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ in order to reach the same computational cost associated with $\hat{\Theta}_{N,k}^{\text{SIS}}$.

Remark 1: Note that the idea of using extra MC samples has already been proposed in the context of Island PFs [24]. The idea behind this class of techniques is to exploit parallel architectures, and the rationale is as follows. Instead of considering a unique set of N particles, the method consists in dividing the population of N samples into N_1 sets of N_2 samples such as $N_1 N_2 = N$. It is well known that such a configuration does not improve the classical PF with N samples, but it has the advantage to split the associated computational cost when parallel architectures are available. In other words, the objective of the PFs is not to struggle against the support impoverishment.

C. Interpretation of the independent sampling scheme in terms of APF

At this point, we have seen that it was possible to obtain an estimate of Θ_k based on i.i.d. samples from the conditional pdf $\tilde{q}_{N,k}$. As in the static case, we now wonder whether the final weights $1/N$ used to compute $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ (see eq. (32)) are optimal when N is finite. To this end we would like to make use of the expression of $\tilde{q}_{N,k}$ to

propose an alternative weighting mechanism. At first glance, the computation of a weight which would rely on (30)-(31) seems compromised because $\tilde{q}_{N,k}$ involves a sum of N terms which should be computed for each N final sample x_k^i . As we will see, the interpretation of the independent SIR algorithm as a particular first step of an APF algorithm will help circumvent this limitation. Let us first begin with a brief presentation of APF filters.

1) *A brief presentation of APF:* In model (1), the filtering density at time k can be written in terms of that at time $k - 1$,

$$p(x_k|y_{0:k}) \propto g_k(y_k|x_k) \int f_k(x_k|x_{k-1})p(x_{k-1}|y_{0:k-1})dx_{k-1}. \quad (35)$$

Plugging an MC approximation $\{w_{k-1}^i, x_{k-1}^i\}_{i=1}^N$ of $p(x_{k-1}|y_{0:k-1})$ into (35) yields

$$\begin{aligned} \hat{p}(x_k|y_{0:k}) &\propto g_k(y_k|x_k) \sum_{i=1}^N w_{k-1}^i f_k(x_k|x_{k-1}^i), \\ &\propto \sum_{i=1}^N w_{k-1}^i p(y_k|x_{k-1}^i) p(x_k|x_{k-1}^i, y_k), \end{aligned} \quad (36)$$

where $p(y_k|x_{k-1}) = \int f_k(x_k|x_{k-1})g_k(y_k|x_k)dx_{k-1}$ and $p(x_k|x_{k-1}, y_k) \propto f_k(x_k|x_{k-1})g_k(y_k|x_k)$. Sampling from $\hat{p}(x_k|y_{0:k})$ in (36) leads to a particular SMC algorithm referred to as the FA-APF [25]. However sampling directly from $\hat{p}(x_k|y_{0:k})$ is not necessarily possible because $p(y_k|x_{k-1}^i)$ or $p(x_k|x_{k-1}^i, y_k)$ are often unavailable. To that end it has been proposed [25] to obtain samples from an instrumental mixture pdf

$$\bar{q}(x_k) = \sum_{i=1}^N \mu(x_{0:k-1}^i) \tau(x_k|x_{0:k-1}^i) \quad (37)$$

and to use IS in augmented dimension; finally APF aims at targeting the mixture pdf $\hat{p}(x_k|y_{0:k})$ in (36) which, itself, targets the filtering distribution $p(x_k|y_{0:k})$. The resulting algorithm is displayed below.

Input: $\mu(x_{0:k-1}), \tau(x_k|x_{0:k-1}), y_k, \{w_{k-1}^i, x_{0:k-1}^i\}_{i=1}^N$

Result: $\{w_k^i, x_{0:k}^i\}_{i=1}^N$

for $1 \leq i \leq N$ **do**

R. $l^i \sim \text{Pr}(L = l | \{x_{0:k-1}^i\}_{i=1}^N) = \mu(x_{0:k-1}^l)$
S. $x_k^i \sim \tau(x_k|x_{0:k-1}^l)$;
W. $w_k^i \propto \frac{w_{k-1}^i f_k(x_k^i|x_{k-1}^i) g_k(y_k|x_k^i)}{\mu(x_{0:k-1}^i) \tau(x_k^i|x_{0:k-1}^i)}, \sum_{i=1}^N w_k^i = 1$;
 Set $x_{0:k}^i = (x_{0:k-1}^i, x_k^i)$

end

Algorithm 4: The APF algorithm

Let us comment the choice of the instrumental distribution $\bar{q}(x_k)$ in (37). Compared to the SIS algorithm of paragraph III-A we see that there is an additional degree of freedom, $\mu(x_{0:k-1})$, which is called the first stage weight; $\tau(x_k|x_{0:k-1}^i)$ refers to a given conditional importance distribution. Generally, the objective of the first stage weights is to avoid the computational waste induced by the resampling step of the SIR algorithm by pre-selecting trajectories at time $k - 1$ which are in accordance with the new observation y_k . Designing this pdf $\bar{q}(x_k)$ is critical and classical approximations of the predictive likelihood such as the likelihood taken at the mode of the transition

pdf (i.e. $\mu(x_{0:k-1}^i) \propto w_{k-1}^i g_k(y_k | \phi(x_{k-1}^i))$) where $\phi(x_{k-1}^i)$ is the mode of $f_{k|k-1}(x_k | x_{k-1}^i)$) can actually damage the performance of the estimate. This is why it is often suggested in practice to build a first-stage weight as close as possible to $w_{k-1} p(y_k | x_{k-1})$, although this problem is generally difficult [26] [27] due to the computation of the predictive likelihood $p(y_k | x_{k-1})$. It remains to choose the importance distribution $\tau(x_k | x_{0:k-1})$; as in the SIR algorithm, one generally tries to approximate the optimal importance distribution $p(x_k | x_{k-1}, y_k)$. Finally note that similarly to classical IS, the FA-APF setting is not necessarily optimal from an asymptotic point of view even if it performs very well in practice [28].

2) *Independent resampling as the first step of a canonical APF algorithm:* Let us now turn to the interpretation of our independent resampling procedure in terms of APF. Let us observe that $\tilde{q}_{N,k}$ in (30) can be rewritten as

$$\tilde{q}_{N,k}(x) = \sum_{i=1}^N \int h_{i,k}(x) q_{i,k}(x) dx \times \frac{h_{i,k}(x) q_{i,k}(x)}{\int h_{i,k}(x) q_{i,k}(x) dx} \quad (38)$$

and so can be seen as one particular mixture pdf $\bar{q}(x_k)$ in (37), in which the weights $\mu^{\text{ind}}(x_{0:k-1}^i)$ are given by $\int h_{i,k}(x) q_{i,k}(x) dx$ and the components $\tau^{\text{ind}}(x_k | x_{0:k-1}^i)$ by $\frac{h_{i,k}(x) q_{i,k}(x)}{\int h_{i,k}(x) q_{i,k}(x) dx}$. We now verify that the couple of samples (l^i, x_k^i) produced by the independent resampling algorithm (Algorithm 3) can indeed be seen as an augmented sample according to $\tilde{q}_{N,k}(x)$ in (38):

- given $\{x_{0:k-1}^j\}_{j=1}^N$ and $\{\tilde{x}_k^{i,j}\}_{j=1}^N$, $\Pr(L^i = l) = w_k^{i,l}$. Since $\tilde{x}_k^{i,j} \sim q_i(x)$, the distribution of l^i given $\{x_{0:k-1}^j\}_{j=1}^N$ becomes $\Pr(L^i = l) = \mathbb{E}(w_k^l | \{x_{0:k-1}^j\}_{j=1}^N) = \int h_{l,k}(x) q_{l,k}(x) dx$;
- given $\{x_{0:k-1}^j\}_{j=1}^N$, $\{\tilde{x}_k^{i,j}\}_{j=1}^N$ and l^i , $x_k^i = \tilde{x}_k^{i,l^i}$. Removing the dependency in $\{\tilde{x}_k^{i,j}\}_{j=1}^N$, the distribution of x_k^i given $\{x_{0:k-1}^j\}_{j=1}^N$ and l^i becomes $\frac{h_{l^i,k}(x) q_{l^i,k}(x)}{\int h_{l^i,k}(x) q_{l^i,k}(x) dx}$.

In summary, our independent resampling procedure is nothing but the first step of one particular APF algorithm, because the pdf $\tilde{q}_{N,k}(x)$ from which we draw i.i.d. samples (given $\{w_{k-1}^i, x_{0:k-1}^i\}_{i=1}^N$) coincides with the mixture pdf (38), which itself constitutes a class of instrumental distributions $\bar{q}(x_k)$ in (37) parametrized by $q(x_k | x_{0:k})$.

In order to appreciate the relevance of that particular solution let us comment on the choice of the first-stage weights $\mu^{\text{ind}}(x_{0:k-1}^i)$ and distributions $\tau^{\text{ind}}(x_k | x_{0:k-1}^i)$:

- at time $k - 1$, trajectories $\{x_{0:k-1}^i\}_{i=1}^N$ are first resampled according to a first stage weight which coincides with the expectation of the importance weights w_k^i of the SIR algorithm defined in (25). In other words, these trajectories are preselected in such a way that the new importance weight w_k^i which would be affected in the weighting step of the SIR algorithm will tend to be large;
- once a trajectory $x_{0:k-1}^i$ has been selected, it is not ensured that its associated weight w_k^i will indeed be large. By sampling according to a pdf proportional to $h_{i,k}(x) q_{i,k}(x)$, the objective is to produce a sample in the region where $h_{i,k}(x)$ (the conditional expectation of the importance weight w_k^i , given that $(X_k^i = x)$) and the distribution $q_{i,k}(x)$ are large.

Consequently, the mixture pdf $\tilde{q}_{N,k}(x)$ appears as a natural instrumental candidate for the APF when the objective is to pre-select the trajectories and to extend them in accordance with the given conditional importance distributions $q_{i,k}(x) = q(x | x_{0:k-1}^i)$ used in the SIR algorithm. If the SIR algorithm IS densities $q_{i,k}(x)$ coincide with the optimal importance distribution $p(x | x_{k-1}^i, y_k)$, then one can see easily that our canonical APF instrumental pdf (38)

reduces to the target mixture (36) (since $h_{k,i}$ in (31) is reduced to a term proportional to $w_{k-1}^i p(y_k | x_{k-1}^i)$) and the independent SIR procedure to the FA-APF algorithm. In that case one can sample from $\tilde{q}_{N,k}$ very efficiently (since (38) is a known mixture) and the resulting estimate outperforms the SIR estimate $\hat{\Theta}_{N,k}^{\text{SIR}}$ with optimal conditional importance distribution [14] [15]. In the case where the FA-APF algorithm is not available, it remains possible to sample from the mixture pdf $\tilde{q}_{N,k}(x)$ in (38) as soon as we can sample from the root pdf $q_{i,k}(x)$, even when $\mu^{\text{ind}}(x_{0:k-1}^i)$ cannot be computed, or one cannot sample from $\tau^{\text{ind}}(x_k | x_{0:k-1}^i)$.

3) *Reweighting the independent samples?*: We can finally use this APF interpretation in order to reweight our conditional independent samples $\{x_k^i\}_{i=1}^N$. Since $\tilde{q}_{N,k}$ can be seen as a mixture (37) with parameters $\mu^{\text{ind}}(x_{0:k-1}^i)$ and $\tau^{\text{ind}}(x_k | x_{0:k-1}^i)$, $\mu^{\text{ind}}(x_{0:k-1}^i) \times \tau^{\text{ind}}(x_k | x_{0:k-1}^i)$ reduces to $h_{i,k}(x) q_{i,k}(x)$. Finally when we target mixture (36), the second-stage weights associated with the independent samples x_k^i produced by Algorithm 3 read

$$w_k^i \propto \frac{w_{k-1}^i f_{k|k-1}(x_k^i | x_{k-1}^i) g_k(y_k | x_k^i)}{h_{l^i,k}(x_k^i) q_{l^i,k}(x_k^i)}, \sum_{i=1}^N w_k^i = 1. \quad (39)$$

We thus obtain a new estimate of Θ_k ,

$$\hat{\Theta}_{N,k}^{\text{I-SIR-w}} = \sum_{i=1}^N w_k^i f(x_k^i) \quad (40)$$

where w_k^i are defined in (39). The practical computation of these final weights relies on that of $h_{i,k}(x)$ in (31), which can be approximated by recycling the extra samples $\tilde{x}^{i,j}$ generated in Algorithm 3,

$$\hat{h}_l(x) = \sum_{i=1}^N \frac{\frac{p_{l,k}(x)}{q_{l,k}(x)}}{\frac{p_{l,k}(x)}{q_{l,k}(x)} + \sum_{j \neq l} \frac{p_{j,k}(\tilde{x}^{i,j})}{q_{j,k}(\tilde{x}^{i,j})}}. \quad (41)$$

D. Summary

Let us summarize the discussions of section III. When the objective is to compute Θ_k in (2) we have several options:

- 1) using the classical SIR algorithm (see Algorithm 2) in which we compute $\hat{\Theta}_{N,k}^{\text{SIS}}$ defined in (26). The resampling step which follows the computation of this estimate produces a conditionally dependent unweighted set of particles sampled from $\tilde{q}_{N,k}$;
- 2) an alternative to avoid the local impoverishment induced by the traditional resampling step is to perform Algorithm 3 and to compute estimate $\hat{\Theta}_{N,k}^{\text{I-SIR}}$. This estimate is still based on an unweighted set of particles marginally sampled from $\tilde{q}_{N,k}$ but these samples have become conditionally independent;
- 3) finally, the samples produced by Algorithm 3 can also be seen as the result of a sampling procedure according to a partial APF instrumental mixture pdf (38). Using further the APF methodology with mixture $\tilde{q}_{N,k}$ it is possible to target mixture (36) which itself is an approximation of $p(x_k | y_{0:k})$. This leads to estimate $\hat{\Theta}_{N,k}^{\text{I-SIR-w}}$ in (40), in which the weights (39) are estimated by recycling the extra samples produced by Algorithm 3.

These three estimates are now going to be compared (in terms of performances and computational cost) in the next section.

IV. SIMULATIONS

We now validate our discussions through computer-generated experiments. In section IV-A we first illustrate the results of Section II and we compare the classical resampling mechanism to the independent one with both unweighted and weighted samples. We also discuss the computational cost associated with our independent resampling mechanism.

In section IV-B we next perform simulations in the ARCH model. On the one hand, the FA-APF algorithm can be computed in this model [25]. On the other hand, remember that our weighted estimate (40) can be interpreted as the estimate deduced from a particular APF which uses the instrumental mixture pdf $\tilde{q}_{N,k}$ in (38), from which it is always possible to sample from (with an extra computational cost). Thus the estimate deduced from the FA-APF algorithm is used as a benchmark and enables us to analyze the relevance of the instrumental pdf $\tilde{q}_{N,k}$ in the APF algorithm.

Next in section IV-C we compute our independent estimates for a target tracking problem with range-bearing measurements. Our estimates are compared to those obtained from the classical SIR algorithm, for a given computational budget measured via the number of sampling operations; this means that we compare $\hat{\Theta}_{M,k}^{\text{I-SIR}}$ and $\hat{\Theta}_{M,k}^{\text{I-SIR-w}}$ (M is the number of particles after the independent resampling step) to $\hat{\Theta}_{N,k}^{\text{SIS}}$ in which $N = \frac{M^2+M}{2}$. Thus all estimates are based on $M^2 + M$ sampling operations (we do not distinguish if we sample according to a continuous or a discrete distribution). The relative performances of the estimates are analyzed in function of the parameters of the state-space model.

Finally in section IV-D we compute our estimates in models where the dimension m of the hidden state is large and we analyze their performances w.r.t. classical PF estimates in function of the dimension m and with a fixed number of sampling operations. Finally throughout this section our simulations are averaged over $P = 1000$ MC runs, we set $f(x) = x$ in (2) and we use an averaged Root Mean Square Error (RMSE) criterion, defined as

$$\text{RMSE}(\hat{\Theta}) = \frac{1}{T} \sum_{k=1}^T \left(\frac{1}{P} \sum_{p=1}^P \|\hat{\Theta}_{k,p} - x_{k,p}\|^2 \right)^{1/2} \quad (42)$$

where $x_{k,p}$ is the true state at time k for the p -th realization, $\hat{\Theta}_{k,p}$ is an estimate of $x_{k,p}$ and T is the time length of the scenario.

A. Comparison of static sampling procedures

Let us first consider the (static) Bayesian estimation problem in which we look for computing

$$\Theta = \text{E}(X|y) = \int xp(x|y)dx \quad (43)$$

via the techniques described in Section II. We assume that $p(x|y)$ is known up to a constant, $p(x|y) \propto p(x)p(y|x)$ where $p(x) = \mathcal{N}(x; 0; \sigma_x^2)$ and $p(y|x) = \mathcal{N}(y; x, \sigma_y^2)$ with $\sigma_x^2 = 10$ and $\sigma_y^2 = 3$. We chose the IS distribution $q(x) = p(x)$. For a given number of final samples N , we compute six estimates: the estimate $\hat{\Theta}_N^{\text{SIS}}$ deduced from the IS mechanism with importance distribution $q(\cdot)$; the estimate $\hat{\Theta}_N^{\text{SIR}}$ deduced from the SIR mechanism with N intermediate samples and $M_N = N$ final samples; our estimate $\hat{\Theta}_N^{\text{I-SIR}}$ based on N unweighted independent

samples drawn from \tilde{q}_N (see (15)); our estimate $\hat{\Theta}_N^{\text{I-SIR-w}}$ based on N weighted independent samples from \tilde{q}_N (see (21)). Remember that the computation of the independent resampling mechanism is based on the sampling of N^2 intermediate particles and N resampling steps and thus requires an extra computational cost w.r.t. the dependent one. Consequently, we also compute $\hat{\Theta}_N^{\text{SIR-2}}$ based on the classical SIR procedure with N^2 intermediate samples and N (dependent) resampling steps; in other words this estimate relies on N dependent samples obtained from \tilde{q}_{N^2} . Finally, we would like to observe the effects of weighting the final samples in the dependent resampling case; so we compute $\hat{\Theta}_N^{\text{SIR-w}}$ which relies on extra samples to approximate the weight proportional to $p(x, y)/\tilde{q}_N(x)$.

In Fig. 1 we display the distance of each estimate w.r.t. the true expectation $E(X|Y = y)$ in function of the number of samples N . As expected, the estimate $\hat{\Theta}_N^{\text{I-SIR}}$ based on N independent samples drawn from \tilde{q}_N outperforms the estimate $\hat{\Theta}_N^{\text{SIR}}$ which is computed from N dependent samples drawn from \tilde{q}_N . However, an interesting result is that $\hat{\Theta}_N^{\text{I-SIR}}$ also outperforms $\hat{\Theta}_N^{\text{SIS}}$. It means that the distribution \tilde{q}_N produced by the SIR mechanism is more adapted than the prior $q(x) = p(x)$, which is not surprising since \tilde{q}_N uses implicitly the observation y through the resampling mechanism of intermediate samples. Of course, the computation of $\hat{\Theta}_N^{\text{I-SIR}}$ requires an extra computational cost but it is interesting to note that the size of the final support is the same in the three cases. We finally compare the estimates based on the same computational cost. When N increases, these estimates have the same asymptotical behavior. It can be seen that the estimate $\hat{\Theta}_N^{\text{SIR-2}}$ based on N samples drawn from \tilde{q}_{N^2} outperforms $\hat{\Theta}_N^{\text{I-SIR}}$. However, when our i.i.d. samples are weighted by a term proportional to $p(x, y)/\tilde{q}_N(x)$ in an IS perspective, our estimate $\hat{\Theta}_N^{\text{I-SIR-w}}$ has the best performance whatever N . We finally note that contrary to the independent procedure, weighting the samples when they are dependent does not improve the performance when compared to the estimate based on dependent and unweighted samples; indeed, $\hat{\Theta}_N^{\text{SIR-w}}$ is not any better than $\hat{\Theta}_N^{\text{SIR}}$. The performances of these algorithms are also presented in terms of RMSE (w.r.t. to the true value of X) in Table I.

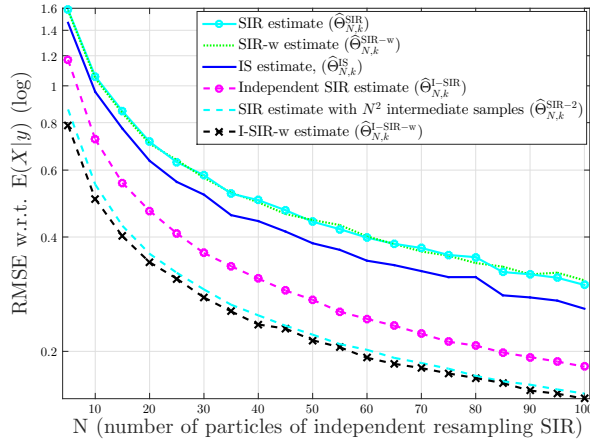


Fig. 1. Static linear and Gaussian model - $\sigma_x^2 = 10$, $\sigma_y^2 = 3$ - Bayesian estimates of $E(X|y)$ based on the independent resampling mechanism outperform the estimates based on the traditional IS and SIR mechanisms although they require an extra computational cost. When the computational cost is fixed, the estimate based on weighted i.i.d. samples from \tilde{q}_N outperforms the estimates based on identically distributed samples from \tilde{q}_{N^2} .

N	$\hat{\Theta}_N^{\text{SIR}}$	$\hat{\Theta}_N^{\text{SIR-w}}$	$\hat{\Theta}_N^{\text{SIS}}$	$\hat{\Theta}_N^{\text{I-SIR}}$	$\hat{\Theta}_N^{\text{SIR-2}}$	$\hat{\Theta}_N^{\text{I-SIR-w}}$
20	1.6844	1.6819	1.6542	1.5951	1.5618	1.5610
40	1.5925	1.5981	1.5763	1.5606	1.5446	1.5410
60	1.5752	1.5777	1.5637	1.5442	1.5395	1.5335
80	1.5623	1.5639	1.5530	1.5345	1.5309	1.5293
100	1.5519	1.5504	1.5410	1.5320	1.5290	1.5290

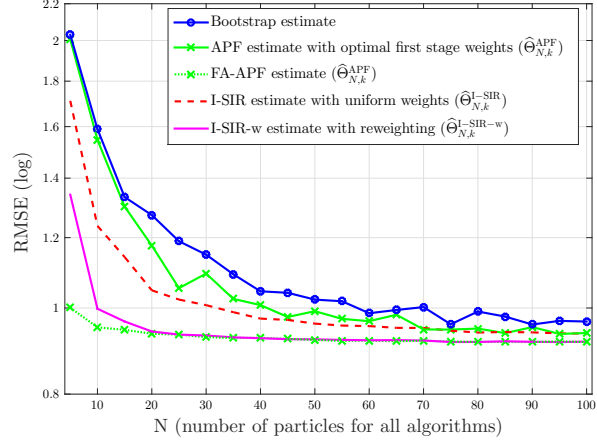
TABLE I

STATIC LINEAR AND GAUSSIAN MODEL - RMSE VALUES OF EACH ESTIMATE.


B. Comparison with APF algorithms

We now focus on the interpretation of our independent resampling algorithm in terms of APF. We study the ARCH model which is a particular hidden Markov model (1) in which $f_k(x_k|x_{k-1}) = \mathcal{N}(x_k; 0; \beta_0 + \beta_1 x_{k-1}^2)$ and $g_k(y_k|x_k) = \mathcal{N}(y_k; x_k; R)$. We set $R = 1$, $\beta_0 = 3$ and $\beta_1 = 0.75$. In this model one can compute $p(y_k|x_{k-1}) = \mathcal{N}(y_k; 0; R + \beta_0 + \beta_1 x_{k-1}^2)$ and $p(x_k|x_{k-1}, y_k) = \mathcal{N}(x_k; \frac{\beta_0 + \beta_1 x_{k-1}^2}{R + \beta_0 + \beta_1 x_{k-1}^2} y_k; \frac{R(\beta_0 + \beta_1 x_{k-1}^2)}{R + \beta_0 + \beta_1 x_{k-1}^2})$; consequently, it is possible to obtain i.i.d. samples from the target mixture (36) and thus to compute the estimate $\hat{\Theta}_{N,k}^{\text{FA}}$ based on the FA-APF algorithm. Remember that the FA-APF can also be seen as a particular case of our independent resampling Algorithm 3 in which the importance distribution $q(x_k|x_{0:k-1})$ coincides with $p(x_k|x_{k-1}, y_k)$ (see section III-C2). However this setting can be implemented in specific models only, while Algorithm 3 can be used with any importance distribution $q(x_k|x_{0:k-1})$, while keeping the same interpretation as the FA-APF (see our discussion in section III-C2). So we also compute our estimates $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ and $\hat{\Theta}_{N,k}^{\text{I-SIR-w}}$ which can be seen as an estimate deduced from the APF in which the importance mixture (37) coincides with $\tilde{q}_{N,k}$. We finally compute the estimate $\hat{\Theta}_{N,k}^{\text{APF}}$ which is deduced from the APF with $\mu(x_{0:k-1}) \propto w_{k-1} p(y_k|x_{k-1})$ and $\tau(x_k|x_{0:k-1}) = f_k(x_k|x_{k-1})$; with this configuration, the particles are pre-selected with the so-called optimal first stage weight and sampled from the transition pdf.

The RMSE of each estimate is displayed in Fig. 2(a) as a function of the number of samples N . Interestingly enough, our weighted independent resampling algorithm which produces $\hat{\Theta}_{N,k}^{\text{I-SIR-w}}$ has the same performances as the FA-APF algorithm when $N \geq 15$, without using the predictive likelihood $p(y_k|x_{k-1})$ nor the optimal importance distribution $p(x_k|x_{k-1}, y_k)$. It means that the mixture pdf \tilde{q}_N which has been interpreted in section III-C2 is indeed as relevant as the target mixture (36); so in general models where the FA-APF is no longer computable, one can expect that our estimate $\hat{\Theta}_{N,k}^{\text{I-SIR-w}}$ would give a performance close to that deduced from FA-APF. Indeed, one advantage of the mixture pdf \tilde{q}_N deduced from the resampling mechanism is that its interpretation does not depend on the importance distribution $q_{i,k}$ which has been chosen and that it is possible to sample from it in general hidden Markov models (1). We also observe that re-weighting the final samples is beneficial w.r.t. attributing uniform weights. In order to analyze the behavior of the weights associated to our estimate $\hat{\Theta}_{N,k}^{\text{I-SIR-w}}$, we compute the normalized effective sample size defined as $N_{\text{norm,eff}} = \frac{1}{N \sum_{i=1}^N (w_k^i)^2}$. In Fig. 2(b), we display the time-averaged normalized effective sample size. It can be observed that $N_{\text{norm,eff}}$ tends to 1 as N increases, meaning that these weights tend to become uniform, so estimates $\hat{\Theta}_{N,k}^{\text{I-SIR}}$ and $\hat{\Theta}_{N,k}^{\text{I-SIR-w}}$ become close when N is large.



(a) RMSE



(b) Averaged normalized effective sample size

Fig. 2. ARCH model - $R = 1$, $\beta_0 = 3$ and $\beta_1 = 0.75$ - (a) The estimate based on the independent resampling mechanism with a final reweighting has the same performances as the estimate deduced from the FA-APF. The final reweighting mechanism is beneficial when compared to the use of uniform weights - (b) When N is large, final weights associated to the estimate $\hat{\Theta}_{N,k}^{\text{I-SIR-w}}$ tend to be uniform.

C. Tracking from range-bearing measurements

We now study the performance of our algorithms in a tracking scenario with range-bearing measurements. We look for estimating the state vector $X_k = [p_{x,k}, \dot{p}_{x,k}, p_{y,k}, \dot{p}_{y,k}]^T$ (position and velocity in Cartesian coordinates) of a target from noisy range-bearing measurements y_k . The pdfs in model (1) associated with this tracking problem are

$$f_k(x_k|x_{k-1}) = \mathcal{N}(x_k; \mathbf{F}x_{k-1}; \mathbf{Q}) \text{ and } g_k(y_k|x_k) = \mathcal{N}(y_k; \left(\begin{array}{c} \sqrt{p_{x,k}^2 + p_{y,k}^2} \\ \arctan \frac{p_{y,k}}{p_{x,k}} \end{array} \right); \mathbf{R}) \text{ where } \tau = 1, \mathbf{R} = \begin{pmatrix} \sigma_\rho^2 & 0 \\ 0 & \sigma_\theta^2 \end{pmatrix},$$

$$\mathbf{F} = \begin{pmatrix} 1 & \tau & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \tau \\ 0 & 0 & 0 & 1 \end{pmatrix}, \mathbf{Q} = \sigma_Q^2 \begin{pmatrix} \frac{\tau^3}{3} & \frac{\tau^2}{2} & 0 & 0 \\ \frac{\tau^2}{2} & \tau & 0 & 0 \\ 0 & 0 & \frac{\tau^3}{3} & \frac{\tau^2}{2} \\ 0 & 0 & \frac{\tau^2}{2} & \tau \end{pmatrix}.$$

The conditional importance distribution used to sample particles is the transition pdf $q(x_k|x_{0:k-1}) = f_k(x_k|x_{k-1})$; so the importance weights w_k^i at time k are proportional to $w_{k-1}^i g(y_k|x_k^i)$. We compute $\hat{\Theta}_{N,k}^{\text{SIS}}$ (see (26)), $\hat{\Theta}_{M,k}^{\text{I-SIR}}$ (see (32)), $\hat{\Theta}_{M,k}^{\text{I-SIR-w}}$ (see (40)) with $N = \frac{M^2+M}{2}$ to set the number of sampling operations. We also compare these estimates with $\hat{\Theta}_{N,k}^{\text{IPF}}$ deduced from the Island PF with 5 islands and $N/5$ particles per island.

The results are displayed for two set of parameters. Fig. 3(a) corresponds to the case where $\sigma_Q = \sqrt{10}$, $\sigma_\rho = 0.25$ and $\sigma_\theta = \frac{\pi}{720}$ while Fig 3(b) corresponds to a very informative case where $\sigma_Q = \sqrt{10}$, $\sigma_\rho = 0.05$ and $\sigma_\theta = \frac{\pi}{3600}$. For the first configuration, we observe that $\hat{\Theta}_{M,k}^{\text{I-SIR-w}}$ outperforms the other estimates and improves $\hat{\Theta}_{M,k}^{\text{I-SIR}}$ which does not rely on weighted samples. Compared to the classical SIS estimate, $\hat{\Theta}_{M,k}^{\text{I-SIR}}$ gives better performance as long as the number of samples M is weak ($M < 30$, so $N < 465$) but is next outperformed when the number of samples is large. As shown in Fig. 3(b), when the observations become informative, $\hat{\Theta}_{M,k}^{\text{I-SIR}}$ gives the best performances. Contrary to $\hat{\Theta}_{N,k}^{\text{SIS}}$ and $\hat{\Theta}_{N,k}^{\text{IPF}}$ our estimate does not suffer from the degeneration of the importance weights. Indeed when the measurements are informative (and so the likelihood is sharp), few importance weights have a non null value. However, the independent resampling procedure ensures the diversity of the final samples when we use uniform weights. Concerning $\hat{\Theta}_{M,k}^{\text{I-SIR-w}}$, remember that it relies on the MC approximation (41). A close analysis of (41) when the likelihood is sharp shows that the final weights tend to be null except that of the particle with the larger likelihood; consequently, in this case the estimate $\hat{\Theta}_{M,k}^{\text{I-SIR-w}}$ is affected by the lack of diversity.

D. High dimensional problems

We finally study the impact of the dimension of the hidden state X_k . We consider a state vector of dimension $m = 4 \times l$, $x_k = [p_{x,k}^1, \dot{p}_{x,k}^1, p_{y,k}^1, \dot{p}_{y,k}^1, \dots, p_{x,k}^l, \dot{p}_{x,k}^l, p_{y,k}^l, \dot{p}_{y,k}^l]^T$. Each component $x_k^l = [p_{x,k}^l, \dot{p}_{x,k}^l, p_{y,k}^l, \dot{p}_{y,k}^l]^T$ evolves independently from all the other components, according to $f_{k|k-1}(x_k|x_{k-1}) = \mathcal{N}(x_k; \mathbf{F}x_{k-1}; \mathbf{Q})$ where $\tau = 1$,

$$\mathbf{F} = \begin{pmatrix} 1 & \tau & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \tau \\ 0 & 0 & 0 & 1 \end{pmatrix}, \mathbf{Q} = \sigma_Q^2 \begin{pmatrix} \frac{\tau^3}{3} & \frac{\tau^2}{2} & 0 & 0 \\ \frac{\tau^2}{2} & \tau & 0 & 0 \\ 0 & 0 & \frac{\tau^3}{3} & \frac{\tau^2}{2} \\ 0 & 0 & \frac{\tau^2}{2} & \tau \end{pmatrix}.$$

Each component is observed independently via $g_k(y_k|x_k) = \mathcal{N}(y_k; \mathbf{H}x_k; \mathbf{R})$ where

$$\mathbf{H} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \mathbf{R} = \begin{pmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{pmatrix}.$$

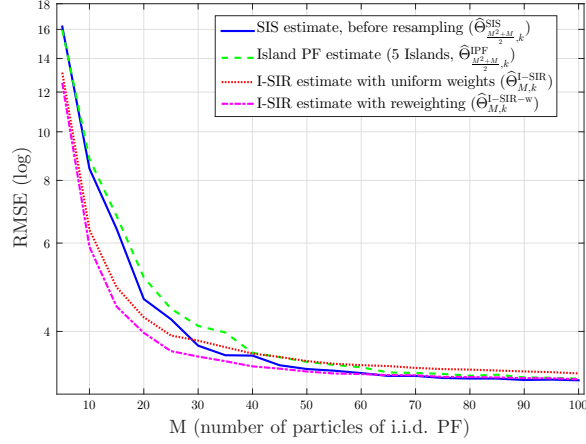

(a) $\sigma_Q = \sqrt{10}$, $\sigma_\rho = 0.25$ and $\sigma_\theta = \frac{\pi}{720}$ (b) $\sigma_Q = \sqrt{10}$, $\sigma_\rho = 0.05$ and $\sigma_\theta = \frac{\pi}{3600}$

Fig. 3. Target tracking model from range-bearing measurements - (a) the independent resampling procedure with final weighting outperforms the other estimates and is particularly interesting when the number of final samples is weak - (b) in the informative case, all estimates suffer from the degeneration of the importance weights except that based on the unweighted independent resampling algorithm. To achieve the same performances as $\hat{\Theta}_{M,k}^{I-SIR-w}$ with $M = 20$, the classical PF uses $N = (50^2 + 50)/2 = 1275$ samples

Again, we compute the estimate based on classical PF $\hat{\Theta}_{N,k}^{SIS}$ (see (26)). It is well known that the PF tends to degenerate when the dimension of the hidden state increases. We also compute $\hat{\Theta}_{M,k}^{I-SIR}$ (see (32)), $\hat{\Theta}_{M,k}^{I-SIR-w}$ (see (40)) with $N = \frac{M^2+M}{2}$ for $M = 100$ and $M = 1000$ as a function of the dimension m to see how the dimension impacts our estimate and the classical PF estimate.

The results are displayed in Fig. 4. It can be seen that the estimates $\hat{\Theta}_{M,k}^{I-SIR}$ and $\hat{\Theta}_{M,k}^{I-SIR-w}$ outperform $\hat{\Theta}_{M,k}^{I-SIS}$ more and more significantly as the dimension increases, due to the local impoverishment phenomenon. First, $\hat{\Theta}_{M,k}^{I-SIR-w}$ outperforms $\hat{\Theta}_{M,k}^{I-SIR}$ as long as the dimension of the hidden state is low ($m = 4$ and $m = 8$); when m increases, the estimate based on weighted samples from \tilde{q}_N limits the degeneration phenomenon w.r.t. that based

on weighted samples from q but using unweighted samples when the dimension is large ensures the diversity and gives better performances. Note that the dependent and independent SIR algorithms give approximately the same performance when m is low but the gap between the dependent and the independent SIR estimates increases with the dimension.




Fig. 4. Multi-dimensional linear Gaussian model - $\sigma_Q^2 = 25$, $\sigma_x^2 = 4$ and $\sigma_y^2 = 4$. The estimates of interest are compared as a function of the dimension m of the hidden state x_k for a fixed number of sampling operations. The independent resampling mechanism limits the impact of the large dimension m and estimate $\hat{\Theta}_{1000,k}^{I-SIR}$ when $m = 46$ has the same performance as $\hat{\Theta}_{500500,k}^{I-SIR}$ when $m = 32$.

V. CONCLUSION

SMC algorithms in Hidden Markov models are based on the sequential application of the IS principle. However the direct sequential application of the IS principle leads to the degeneration of the weights, against which multinomial resampling has been proposed. This rejuvenation scheme, which is now routinely used in SIR algorithms, enables to discard particles (or trajectories) with low weights, but particles with large weights will be resampled several times, which leads to dependency and support degeneracy. In this paper we thus revisited the resampling step used in the classical SIR algorithms. We first addressed the static case, showed that the particles sampled by Rubin's SIR mechanism are dependent samples drawn from some pdf \tilde{q}_N , and proposed an alternative sampling mechanism which produces independent particles drawn from that same marginal pdf \tilde{q}_N . This set of independent samples enables us to build a moment estimator which outperforms the classical SIR-based one, both from a non-asymptotical and an asymptotical points of view. Finally the succession of the sampling, weighting and resampling steps indeed transforms an elementary instrumental pdf q into a compound importance distribution $\tilde{q}_N = \phi(p, q, N)$, which leads us to reweight the (originally unweighted) resampled particles x^i by post-resampling weights proportional to $\frac{p(x^i)}{q(x^i)}$. Such post-resampling weights cannot be computed exactly, but can easily be estimated by recycling the extra MC samples which were needed for producing the independently resampled particles.

We next adapted this methodology to the dynamic case, in order to estimate a moment of interest in an hidden Markov model. The computation of the post-resampling weights is more challenging than in the static case, but reinterpreting our independent resampling scheme as the first step of a particular APF algorithm enables us to make full use of the APF methodology and so to reweight the final samples via the second-stage APF weights. Finally we validated our discussions by computer-generated experiments and carefully took into account the computational budget. Simulations in model where the FA-APF algorithm is computable show that the independent resampling gives a performance close to the FA-APF algorithm. Consequently, it confirms the relevance of the instrumental mixture pdf used implicitly by the independent resampling PF which can be used in any hidden Markov model since it not require to compute the predictive likelihood nor the optimal importance distribution. Finally independent PF gives very satisfying results when applied in highly informative models which are challenging for classical PF and limits the degeneration phenomenon in high dimensional models.

APPENDIX

PROOF OF PROPOSITION 1

Let A be any Borel set. Let $\mathbb{1}_A(x) = 1$ if $x \in A$ and 0 otherwise. Then for any l , $1 \leq l \leq M_N$,

$$\begin{aligned}
& \Pr(X^l \in A) \\
&= \int_{\mathbb{R}^N} \left[\sum_{i=1}^N w^i(\tilde{x}^1, \dots, \tilde{x}^N) \mathbb{1}_A(\tilde{x}^i) \right] \prod_{j=1}^N q(\tilde{x}^j) d\tilde{\mathbf{x}}^{1:N} \\
&= \sum_{i=1}^N \int_{\mathbb{R}^N} w^i(\tilde{x}^1, \dots, \tilde{x}^N) \mathbb{1}_A(\tilde{x}^i) \prod_{j=1}^N q(\tilde{x}^j) d\tilde{\mathbf{x}}^{1:N} \\
&= \sum_{i=1}^N \int_A \left[\int_{\mathbb{R}^{N-1}} w^i(\tilde{x}^1, \dots, \tilde{x}^N) \prod_{\substack{j=1 \\ j \neq i}}^N q(\tilde{x}^j) d\tilde{\mathbf{x}}^{1:i-1, i+1:N} \right] q(\tilde{x}^i) d\tilde{x}^i \\
&= \sum_{i=1}^N \int_A h_N(\tilde{x}^i) q(\tilde{x}^i) d\tilde{x}^i \\
&= \int_A N h_N(\tilde{x}) q(\tilde{x}) d\tilde{x},
\end{aligned}$$

so X^l has pdf \tilde{q}_N w.r.t. Lebesgue measure.

PROOF OF PROPOSITION 2

Let X^i (for any i , $1 \leq i \leq M_N$) be produced by the classical SIR mechanism. Then

$$\mathbb{E}(f(X^i) | \tilde{\mathbf{x}}^{1:N}) = \hat{\Theta}_N^{\text{IS}}. \quad (44)$$

So $\mathbb{E}(\hat{\Theta}_{M_N}^{\text{SIR}} | \tilde{\mathbf{x}}^{1:N}) = \hat{\Theta}_N^{\text{IS}}$, and $\mathbb{E}(\hat{\Theta}_{M_N}^{\text{SIR}}) = \mathbb{E}(\hat{\Theta}_N^{\text{IS}})$. On the other hand $\mathbb{E}(\hat{\Theta}_{M_N}^{\text{I-SIR}}) = \mathbb{E}(\hat{\Theta}_{M_N}^{\text{SIR}})$, whence (16). Next

$$\text{var}(\hat{\Theta}_{M_N}^{\text{SIR}}) = \frac{1}{M_N^2} \sum_{i=1}^{M_N} \text{var}(f(X^i)) + \frac{1}{M_N^2} \sum_{\substack{k,l=1 \\ k \neq l}}^{M_N} \text{Cov}(f(X^k), f(X^l)). \quad (45)$$

in which $X^i \sim \tilde{q}_N$ for all i . The first term is equal to $\text{var}(\hat{\Theta}_{M_N}^{I-SIR})$. Let us compute the second term. For all $k, l, 1 \leq k, l, \leq M_N$ with $k \neq l$, $E(f(X^k)f(X^l)|\tilde{\mathbf{X}}^{1:N}) = (\hat{\Theta}_N^{IS})^2$, so $E(f(X^k)f(X^l)) = E(E(f(X^k)f(X^l)|\tilde{\mathbf{X}}^{1:N})) = E((\hat{\Theta}_N^{IS})^2)$. Using (44) again, we conclude that $\text{Cov}(f(X^k), f(X^l)) = \text{var}(\hat{\Theta}_N^{IS})$, whence (17).

PROOF OF THEOREM 1

We first introduce the following notations:

$$\Theta(f) = \int f(x)p(x)dx, \quad (46)$$

$$\hat{\Theta}_N^{IS}(f) = \sum_{i=1}^N \frac{\frac{p(X^i)}{q(X^i)}}{\sum_{j=1}^N \frac{p(X^j)}{q(X^j)}} f(X^i), \quad X^i \stackrel{\text{i.i.d.}}{\sim} q(\cdot), \quad (47)$$

$$\hat{\Theta}_{M_N}^{I-SIR}(f) = \frac{1}{M_N} \sum_{i=1}^{M_N} f(\bar{X}^i), \quad \bar{X}^i \stackrel{\text{i.i.d.}}{\sim} \tilde{q}_N(\cdot), \quad (48)$$

and we will assume that $E(\hat{\Theta}_N^{IS}(f^2))$ is finite.

Using $E(\hat{\Theta}_{M_N}^{I-SIR}(f)) = E(\hat{\Theta}_N^{IS}(f))$, we have

$$\sqrt{M_N} (\hat{\Theta}_{M_N}^{I-SIR}(f) - \Theta(f)) = A_N + B_N, \quad (49)$$

$$A_N = \sqrt{M_N} (\hat{\Theta}_{M_N}^{I-SIR}(f) - E(\hat{\Theta}_{M_N}^{I-SIR}(f))), \quad (50)$$

$$B_N = \frac{\sqrt{M_N}}{\sqrt{N}} E(\sqrt{N}(\hat{\Theta}_N^{IS}(f) - \Theta(f))). \quad (51)$$

Our objective is to show that A_N converges to a centered Gaussian distribution with variance $\text{var}_p(f(X))$ and that B_N converges to 0.

Convergence of B_N

We have recalled (see (9)) that under mild assumptions [4]

$$\sqrt{N}(\hat{\Theta}_N^{IS}(f) - \Theta(f)) \xrightarrow{\mathcal{D}} \mathcal{N}\left(0, E_q\left(\frac{p^2(X)}{q^2(X)}(f(X) - \Theta)^2\right)\right).$$

According to Theorem 9.1.10 in [14], $E(|\sqrt{N}(\hat{\Theta}_N^{IS}(f) - \Theta(f))|^2)$ is bounded and so its upper bound is finite.

According to the corollary of Theorem 25.12 in [29], it is ensured that $\sqrt{N}E((\hat{\Theta}_N^{IS}(f) - \Theta(f))) \rightarrow 0$; consequently

$$\frac{\sqrt{M_N}}{\sqrt{N}} E(\sqrt{N}(\hat{\Theta}_N^{IS}(f) - \Theta(f))) \rightarrow 0. \quad (52)$$

Convergence of A_N

A_N reads

$$\sqrt{M_N} \left(\frac{1}{M_N} \sum_{i=1}^{M_N} f(\bar{X}^i) - E(f(\bar{X}^i)) \right). \quad (53)$$

To prove the convergence when $N \rightarrow \infty$, we need a CLT for triangular arrays and we use the version presented in Theorem 9.5.13 of [14]. The required assumptions are:

- 1) $\{\bar{X}^i\}_{i=1}^{M_N}$ are independent;
- 2) $\frac{1}{M_N} \sum_{i=1}^{M_N} E(f^2(\bar{X}^i)) - (E(f(\bar{X}^i)))^2 \rightarrow \text{var}_p(f(X))$;

3) for any positive C , $\frac{1}{M_N} \sum_{i=1}^{M_N} \mathbb{E}(f^2(\bar{X}^i) \mathbb{1}_{|f(\bar{X}^i)| \geq C}) \rightarrow \Theta(f^2 \mathbb{1}_{|f| \geq C})$.

Assumption 1) is satisfied since $\{\bar{X}^i\}_{i=1}^{M_N}$ are i.i.d. from \tilde{q}_N . Next, $\mathbb{E}(f(\bar{X}^i)) = \mathbb{E}_{\tilde{q}_N}(f(\bar{X}))$ which coincides with $\mathbb{E}(\hat{\Theta}_N^{\text{IS}}(f))$. Using again Theorem 9.1.10 of [14] and Theorem 25.12 of [29], $\mathbb{E}(\hat{\Theta}_N^{\text{IS}}(f)) \rightarrow \Theta(f)$ when $N \rightarrow \infty$. With the same argument, $\mathbb{E}(f^2(\bar{X}^i)) \rightarrow \Theta(f^2)$. Consequently, assumption 2) is satisfied since

$$\frac{1}{M_N} \sum_{i=1}^{M_N} \mathbb{E}(f^2(\bar{X}^i)) - (\mathbb{E}(f(\bar{X}^i)))^2 = \mathbb{E}_{\tilde{q}_N}(f^2(\bar{X})) - (\mathbb{E}_{\tilde{q}_N}(f(\bar{X})))^2 \rightarrow \Theta(f^2) - (\Theta(f))^2 = \text{var}_p(f(X)).$$

Finally, $\mathbb{E}(f^2(\bar{X}^i) \mathbb{1}_{|f(\bar{X}^i)| \geq C}) = \mathbb{E}(\hat{\Theta}_N^{\text{IS}}(f^2 \mathbb{1}_{|f| \geq C}))$ which converges to $\Theta(f^2 \mathbb{1}_{|f| \geq C})$ and assumption 3) is satisfied. Consequently,

$$\sqrt{M_N} \left(\frac{1}{M_N} \sum_{i=1}^{M_N} f(\bar{X}^i) - \mathbb{E}(f(\bar{X}^i)) \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \text{var}_p(f(X))). \quad (54)$$

Combining (52), (54) and (49) we obtain (18).

REFERENCES

- [1] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/ non-Gaussian Bayesian state estimation," *IEE Proceedings-F*, vol. 140, no. 2, pp. 107–113, April 1993.
- [2] A. Doucet, N. de Freitas, and N. Gordon, Eds., *Sequential Monte Carlo Methods in Practice*, ser. Statistics for Engineering and Information Science. New York: Springer Verlag, 2001.
- [3] 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, February 2002.
- [4] J. Geweke, "Bayesian inference in econometric models using Monte Carlo integration," *Econometrica*, vol. 57, no. 6, pp. 1317–1339, November 1989.
- [5] A. Kong, J. S. Liu, and W. H. Wong, "Sequential imputations and bayesian missing data problems," *Journal of the American Statistical Association*, vol. 89, no. 425, pp. 278–88, March 1994.
- [6] J. S. Liu and R. Chen, "Blind deconvolution via sequential imputation," *Journal of the American Statistical Association*, vol. 90, no. 430, pp. 567–76, June 1995.
- [7] J. S. Liu, "Metropolized independent sampling with comparisons to rejection sampling and importance sampling," *Statistics and Computing*, vol. 6, pp. 113–119, 1996.
- [8] J. Cornebise, É. Moulines, and J. Olsson, "Adaptive methods for sequential importance sampling with application to state-space models," *Statistics and Computing*, vol. 18, no. 4, pp. 461–480, 2008.
- [9] A. F. M. Smith and A. E. Gelfand, "Bayesian statistics without tears : a sampling-resampling perspective." *The American Statistician*, vol. 46, no. 2, pp. 84–87, 1992.
- [10] R. Douc, O. Cappé, and É. Moulines, "Comparison of resampling schemes for particle filtering," in *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis (ISPA)*, Zagreb, Croatia, September 2005.
- [11] T. Li, M. Bolić, and P. M. Djuric, "Resampling Methods for Particle Filtering: Classification, implementation, and strategies," *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 70–86, May 2015.
- [12] D. B. Rubin, "Using the SIR algorithm to simulate posterior distributions," in *Bayesian Statistics III*, M. H. Bernardo, K. M. Degroot, D. V. Lindley, and A. F. M. Smith, Eds. Oxford: Oxford University Press, 1988.
- [13] A. E. Gelfand and A. F. M. Smith, "Sampling based approaches to calculating marginal densities," *Journal of the American Statistical Association*, vol. 85, no. 410, pp. 398–409, 1990.
- [14] O. Cappé, É. Moulines, and T. Rydén, *Inference in Hidden Markov Models*. Springer-Verlag, 2005.
- [15] Y. Petetin and F. Desbouvries, "Optimal SIR algorithm vs. fully adapted auxiliary particle filter: a non asymptotical analysis," *Statistics and computing*, vol. 23, no. 6, pp. 759–775, 2013.
- [16] T. Hesterberg, "Advances in importance sampling," Ph.D. dissertation, Stanford University, 1988.
- [17] P. del Moral, *Feynman-Kac formulae. Genealogical and interacting particle systems with applications*, ser. Probability and its applications. New-York: Springer, 2004.

- [18] J. S. Liu, *Monte Carlo strategies in scientific computing*. Springer, 2001.
- [19] A. Doucet, S. J. Godsill, and C. Andrieu, “On sequential Monte Carlo sampling methods for Bayesian filtering,” *Statistics and Computing*, vol. 10, pp. 197–208, 2000.
- [20] R. van der Merwe, A. Doucet, N. de Freitas, and E. Wan, “The Unscented Particle Filter.” *Advances in Neural Information Processing Systems*, 2000.
- [21] S. Saha, P. K. Manda, Y. Boers, H. Driessen, and A. Bagchi, “Gaussian proposal density using moment matching in SMC methods.” *Statistics and Computing*, vol. 19, no. 2, pp. 203–208, 2009.
- [22] D. Crisan and A. Doucet, “A survey of convergence results on particle filtering methods for practitioners,” *IEEE Transactions on Signal Processing*, vol. 50, no. 3, pp. 736–46, March 2002.
- [23] N. Chopin, “Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference,” *The Annals of Statistics*, vol. 32, no. 6, pp. 2385–2411, 2004.
- [24] C. Vergé, C. Dubarry, P. Del Moral, and É. Moulines, “On parallel implementation of sequential monte carlo methods: the island particle model,” *Statistics and Computing*, vol. 25, no. 2, pp. 243–260, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s11222-013-9429-x>
- [25] M. K. Pitt and N. Shephard, “Filtering via simulation : Auxiliary particle filter,” *Journal of the American Statistical Association*, vol. 94, no. 446, pp. 590–99, June 1999.
- [26] A. M. Johansen and A. Doucet, “A note on the auxiliary particle filter,” *Statistics and Probability Letters*, vol. 78, no. 12, pp. 1498–1504, September 2008.
- [27] N. Whiteley and A. M. Johansen, “Recent developments in auxiliary particle filtering,” in *Inference and Learning in Dynamic Models*, Barber, Cemgil, and Chiappa, Eds. Cambridge University Press, 2010.
- [28] R. Douc, É. Moulines, and J. Olsson, “Optimality of the auxiliary particle filter,” *Probability and Mathematical Statistics*, vol. 29, no. 1, pp. 1–28, 2009.
- [29] P. Billingsley, *Probability and Measure*, ser. Wiley Series in Probability and Statistics. Wiley, 1995. [Online]. Available: <https://books.google.de/books?id=z39jQgAACAAJ>