WEIGHTING A RESAMPLED PARTICLE IN SEQUENTIAL MONTE CARLO

L. Martino[◇], V. Elvira[†], F. Louzada[◇]

[†] Dep. of Signal Theory and Communic., Universidad Carlos III de Madrid, Leganés (Spain).
[◇] Institute of Mathematical Sciences and Computing, Universidade de São Paulo, São Carlos (Brazil).

ABSTRACT

The Sequential Importance Resampling (SIR) method is the core of the Sequential Monte Carlo (SMC) algorithms (a.k.a., particle filters). In this work, we point out a suitable choice for weighting properly a resampled particle. This observation entails several theoretical and practical consequences, allowing also the design of novel sampling schemes. Specifically, we describe one theoretical result about the sequential estimation of the marginal likelihood. Moreover, we suggest a novel resampling procedure for SMC algorithms called partial resampling, involving only a subset of the current cloud of particles. Clearly, this scheme attenuates the additional variance in the Monte Carlo estimators generated by the use of the resampling.

Index Terms— Importance Sampling; Sequential Importance Resampling; Sequential Monte Carlo; Particle Filtering.

1. INTRODUCTION

Sequential Monte Carlo (SMC) methods have become essential tools for Bayesian analysis in statistical signal processing [2, 6, 7, 10, 15]. SMC algorithms are based on the importance sampling technique [5, 4, 9, 12, 16] and its sequential version known as Sequential Importance Sampling (SIS) [8, 11]. Another essential piece of SMC is the application of resampling procedures [3, 8]. The combination of SIS and resampling is often referred as Sequential Importance Resampling (SIR).

Since the unnormalized importance weight of a resampled particle cannot be computed analytically using the standard IS weight definition, in the classical SIR formulation, the users consider only the estimators involving normalized weights. The concept of the unnormalized weight of a resampled particle is usually not considered, i.e., its computation is avoided and omitted [6, 7, 8].

In this work, we propose a proper unnormalized importance weight for a resampled particle which provides unbiased IS estimators, following the Liu's definition of a proper

weighted sample [11, Section 2.5.4]. The introduction of this unnormalized proper weight for a resampled particle yields several interesting consequences. Here, we describe two of them.

First, we point out that all the estimators derived in the SIS approach can also be employed in SIR using the introduced weight definition of a resampled particle. We show it considering the estimation of the marginal likelihood (a.k.a., Bayesian evidence or partition function) [8, 15, 16]. In SIS, there are two possible estimators of the marginal likelihood which are completely equivalent [15]. Using the proper unnormalized weight for a resampled particle, we show that also in SIR we can employ two equivalent estimators of the marginal likelihood. They coincide with the estimators in SIS as special case, when no resampling is applied. Furthermore, we describe an alternative resampling procedure for SMC algorithms, called *partial resampling*, involving only a subset of the current population of particles. This scheme attenuates the additional variance in the Monte Carlo estimators generated by the use of resampling steps. Moreover, it can reduce the loss of diversity in the population of particle due to the application of the resampling.

2. IMPORTANCE SAMPLING

Let us denote the target probability density function (pdf) as $\bar{\pi}(\mathbf{x}) = \frac{1}{Z}\pi(\mathbf{x})$ (known up to a normalizing constant) with

$$\mathbf{x} = x_{1:D} = [x_1, x_2, \ldots, x_D] \in \mathcal{X} \subseteq \mathbb{R}^{D \times \eta},$$

where $x_d \in \mathbb{R}^\eta$ for all $d = 1, \ldots, D$. We consider the Monte Carlo approximation of complicated integrals involving the target $\bar{\pi}(\mathbf{x})$ and a square-integrable function $h(\mathbf{x})$, e.g.,

$$I = E_\pi[h(\mathbf{X})] = \int_{\mathcal{X}} h(\mathbf{x})\bar{\pi}(\mathbf{x})d\mathbf{x}, \tag{1}$$

where $\mathbf{X} \sim \bar{\pi}(\mathbf{x})$. In general, generating samples directly from the target $\bar{\pi}(\mathbf{x})$ is impossible. Thus, one usually considers a (simpler) proposal pdf, $q(\mathbf{x})$. Indeed, the equality

$$\begin{aligned} E_\pi[h(\mathbf{X})] &= E_q[h(\mathbf{X})w(\mathbf{X})], \\ &= \frac{1}{Z}\int_{\mathcal{X}} h(\mathbf{x})\frac{\pi(\mathbf{x})}{q(\mathbf{x})}q(\mathbf{x})d\mathbf{x}, \end{aligned} \tag{2}$$

where $w(\mathbf{x}) = \frac{\pi(\mathbf{x})}{q(\mathbf{x})} : \mathcal{X} \rightarrow \mathbb{R}$, suggests an alternative procedure. Indeed, we can draw N samples (also called *particles*) $\mathbf{x}_1, \ldots, \mathbf{x}_N$ from $q(\mathbf{x})$,[1] and then assign to each sample the following weights

$$w(\mathbf{x}_n) = \frac{\pi(\mathbf{x}_n)}{q(\mathbf{x}_n)}, \quad n = 1, \ldots, N. \tag{3}$$

If the target function $\pi(\mathbf{x})$ is normalized, i.e., $Z = 1$, $\bar{\pi}(\mathbf{x}) = \pi(\mathbf{x})$, a natural (unbiased) IS estimator [11, 16] is defined as

$$\widehat{I} = \frac{1}{N} \sum_{n=1}^{N} w(\mathbf{x}_n) h(\mathbf{x}_n) \xrightarrow[N \rightarrow \infty]{} I, \tag{4}$$

where $\mathbf{x}_n \sim q(\mathbf{x})$, $n = 1, \ldots, N$. If the normalizing constant Z is unknown, defining the normalized weights,

$$\bar{w}(\mathbf{x}_n) = \frac{w(\mathbf{x}_n)}{\sum_{i=1}^{N} w(\mathbf{x}_i)}, \quad n = 1, \ldots, N, \tag{5}$$

an alternative self-normalized (biased) IS estimator [11, 16] is

$$\bar{I} = \sum_{n=1}^{N} \bar{w}(\mathbf{x}_n) h(\mathbf{x}_n) \xrightarrow[N \rightarrow \infty]{} I. \tag{6}$$

Moreover, an unbiased estimator of marginal likelihood, $Z = \int_{\mathcal{X}} \pi(\mathbf{x}) d\mathbf{x}$, is given by

$$\widehat{Z} = \frac{1}{N} \sum_{i=1}^{N} w(\mathbf{x}_i) \xrightarrow[N \rightarrow \infty]{} Z. \tag{7}$$

2.1. Concept of proper weighted sample

Although the weights of Eq. (3) are broadly used in the literature, the concept of a *properly weighted sample*, suggested in [16, Section 14.2] and in [11, Section 2.5.4], can be used to construct more general weights. More specifically, following the definition in [11, Section 2.5.4], a set of weighted samples is considered *proper* with respect to the target π if, for any square integrable function h,

$$E_q[w(\mathbf{x}_n) h(\mathbf{x}_n)] = c E_\pi[h(\mathbf{x}_n)], \quad \forall n = \{1, \ldots, N\}, \tag{8}$$

where c is a constant value, also independent from the index n, and the expectation of the left hand side is performed, in general, w.r.t. to the joint pdf of $w(\mathbf{x})$ and $\mathbf{x}$, i.e., $q(w, \mathbf{x})$. Namely, the weight $w(\mathbf{x})$, (for a given value of $\mathbf{x}$), could even be considered a random variable.

3. IMPORTANCE WEIGHT OF A RESAMPLED PARTICLE

Let us consider the following multinomial resampling procedure [3, 6, 7]:

---

[1] We assume that $q(\mathbf{x}) > 0$ for all $\mathbf{x}$ where $\bar{\pi}(\mathbf{x}) \neq 0$, and $q(\mathbf{x})$ has heavier tails than $\bar{\pi}(\mathbf{x})$.

1. Draw N particles $\mathbf{x}_n \sim q(\mathbf{x})$ and weight them with $w(\mathbf{x}_n) = \frac{\pi(\mathbf{x}_n)}{q(\mathbf{x}_n)}$, with $n = 1, \ldots, N$.

2. Draw one particle $\widetilde{\mathbf{x}}' \in \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$ from the discrete probability mass

$$\widehat{\pi}(\mathbf{x}|\mathbf{x}_{1:N}) = \sum_{n=1}^{N} \bar{w}(\mathbf{x}_n) \delta(\mathbf{x} - \mathbf{x}_n), \tag{9}$$

where $\bar{w}(\mathbf{x}_n) = \frac{w(\mathbf{x}_n)}{\sum_{i=1}^{N} w(\mathbf{x}_i)}$.

Question 1. What is the distribution of the resampled particle $\widetilde{\mathbf{x}}'$ (not conditioned to $\mathbf{x}_{1:N}$)? We can easily write its corresponding density as

$$\widetilde{q}(\mathbf{x}) = \int_{\mathcal{X}^N} \left[\prod_{i=1}^{N} q(\mathbf{x}_i) \right] \widehat{\pi}(\mathbf{x}|\mathbf{x}_{1:N}) d\mathbf{x}_{1:N}. \tag{10}$$

where $\widehat{\pi}$ is given in Eq. (9). However, the integral above cannot be compute analytically.

Question 2. Can we obtain a proper importance weight associated to the resampled particle $\widetilde{\mathbf{x}}$? As a consequence of the previous observations, we are not able to evaluate the corresponding standard importance weight, $w(\widetilde{\mathbf{x}}) = \frac{\pi(\widetilde{\mathbf{x}})}{\widetilde{q}(\widetilde{\mathbf{x}})}$.

For solving this issue, let us consider N resampled particles $\widetilde{\mathbf{x}}_1, \ldots, \widetilde{\mathbf{x}}_N$ independently obtained by the resampling procedure above. In SMC and adaptive IS applications, [4, 5, 6, 7], the *unnormalized* importance weights of $\widetilde{\mathbf{x}}_1, \ldots, \widetilde{\mathbf{x}}_N$ are not usually needed, but only the *normalized* ones. Thus, a well known proper strategy [6, 7, 8] in this case is to set

$$w(\widetilde{\mathbf{x}}_1) = w(\widetilde{\mathbf{x}}_2) = \ldots = w(\widetilde{\mathbf{x}}_N), \tag{11}$$

and, as a consequence, the normalized weights are

$$\bar{w}(\widetilde{\mathbf{x}}_1) = \bar{w}(\widetilde{\mathbf{x}}_2) = \ldots = \bar{w}(\widetilde{\mathbf{x}}_N) = \frac{1}{N}. \tag{12}$$

The reason why this approach is suitable lies on the Liu's definition of proper importance weights in Section 2.1. Indeed, considering the random variable $\widetilde{\mathbf{X}} \sim \widehat{\pi}(\mathbf{x}|\mathbf{x}_{1:N})$, we have

$$
\begin{aligned}
E_{\widehat{\pi}}[h(\widetilde{\mathbf{X}})|\mathbf{x}_{1:N}] &= \int h(\mathbf{x}) \widehat{\pi}(\mathbf{x}|\mathbf{x}_{1:N}) d\mathbf{x}, \\
&= \sum_{n=1}^{N} \bar{w}(\mathbf{x}_n) h(\mathbf{x}_n) = \bar{I}. \tag{13}
\end{aligned}
$$

where $\mathbf{x}_n \sim q(\mathbf{x})$ for $n = 1, \ldots, N$, are considered fixed in the expectation $E_{\widehat{\pi}}[h(\widetilde{\mathbf{X}})|\mathbf{x}_{1:N}]$. Thus, The self-normalized IS estimator using the resampled particles is

$$\widetilde{I} = \frac{1}{N} \sum_{n=1}^{N} h(\widetilde{\mathbf{x}}_n) \xrightarrow[N \rightarrow \infty]{} E_{\widehat{\pi}}[h(\widetilde{\mathbf{X}})|\mathbf{x}_{1:N}] = \bar{I}. \tag{14}$$

Hence, we can write for

$$\widetilde{I} \xrightarrow[N\to\infty]{} \overline{I} \xrightarrow[N\to\infty]{} I, \qquad (15)$$

due to Eqs. (13)-(14). This is the proof that the choice $\bar{w}(\widetilde{\mathbf{x}}_n) = \frac{1}{N}$, for all n, is proper by Liu's definition.

However, for several theoretical and practical reasons (some of them discussed below), it is interesting to define also a proper *unnormalized* importance weight of a resampled particle. Let us consider the following definition.

Definition 1. *A proper choice for an unnormalized importance weight (following Section 2.1) of a resampled particle $\widetilde{\mathbf{x}} \in \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$ is*

$$w(\widetilde{\mathbf{x}}) = \widehat{Z} = \frac{1}{N}\sum_{i=1}^{N} w(\mathbf{x}_i). \qquad (16)$$

Considering N independent resampled particles, Note that with this definition we again have

$$w(\widetilde{\mathbf{x}}_1) = w(\widetilde{\mathbf{x}}_2) = \ldots = w(\widetilde{\mathbf{x}}_N),$$

so that also $\bar{w}(\widetilde{\mathbf{x}}_n) = \frac{1}{N}$, for all $n = 1, \ldots, N$. The unnormalized weight in Eq. (16) suggests that a resampled particle $\widetilde{\mathbf{x}}$ "condenses" the statistical information provided by $\mathbf{x}_1, \ldots, \mathbf{x}_N$.

Remark 1. *The previous definition allows is to estimate Z using the resampled particles as well. Indeed,*

$$\widetilde{Z} = \frac{1}{N}\sum_{i=1}^{N} w(\widetilde{\mathbf{x}}_i) = \frac{1}{N}\left(N\widehat{Z}\right) = \widehat{Z}, \qquad (17)$$

is an unbiased estimator of Z (equivalent to $\widehat{Z}$).

4. APPLICATION IN SIR

Let recall $\mathbf{x} = x_{1:D} = [x_1, x_2, \ldots, x_D] \in \mathcal{X} \subseteq \mathbb{R}^{D\times\eta}$ where $x_d \in \mathbb{R}^\eta$ for all $d = 1, \ldots, D$ and let us consider a target pdf $\bar{\pi}(\mathbf{x})$ factorized as

$$\bar{\pi}(\mathbf{x}) \propto \pi(\mathbf{x}) = \gamma_1(x_1)\prod_{d=2}^{D}\gamma_d(x_d|x_{1:d-1}), \quad (18)$$

where $\gamma_1(x_1)$ is a marginal pdf and $\gamma_d(x_d|x_{1:d-1})$ are conditional pdfs. We also denote the joint probability of $[x_1, \ldots, x_d]$,

$$\bar{\pi}_d(x_{1:d}) = \frac{1}{Z_d}\pi_d(x_{1:d}), \qquad (19)$$

where

$$\pi_d(x_{1:d}) = \gamma_1(x_1)\prod_{j=2}^{d}\gamma_j(x_j|x_{1:j-1}).$$

Clearly, $\bar{\pi}(\mathbf{x}) \equiv \bar{\pi}_D(x_{1:D})$. We can also consider a proposal pdf decomposed in the same fashion,

$$q(\mathbf{x}) = q_1(x_1)q_2(x_2|x_1)\cdots q_D(x_D|x_{1:D-1}).$$

In a batch IS scheme, given the n-th sample $\mathbf{x}_n = x_{1:D}^{(n)} \sim q(\mathbf{x})$, we assign the importance weight

$$
\begin{aligned}
w(\mathbf{x}_n) &= \frac{\pi(\mathbf{x}_n)}{q(\mathbf{x}_n)} \\
&= \frac{\gamma_1(x_1^{(n)})\gamma_2(x_2^{(n)}|x_1^{(n)})\cdots\gamma_D(x_D^{(n)}|x_{1:D-1}^{(n)})}{q_1(x_1^{(n)})q_2(x_2^{(n)}|x_1^{(n)})\cdots q_D(x_D^{(n)}|x_{1:D-1}^{(n)})}.
\end{aligned}
$$

The previous expression suggests a recursive procedure for computing the importance weights. Indeed, in a sequential Importance sampling (SIS) approach [6, 7], we can write

$$w_d^{(n)} = w_{d-1}^{(n)}\beta_d^{(n)} = \prod_{j=1}^{d}\beta_j^{(n)}, \qquad n = 1, \ldots, N, \quad (20)$$

where we have set

$$w_1^{(n)} = \beta_1^{(n)} = \frac{\pi(x_1^{(n)})}{q(x_1^{(n)})} \quad \text{and} \quad \beta_d^{(n)} = \frac{\gamma_d(x_d^{(n)}|x_{1:d-1}^{(n)})}{q_d(x_d^{(n)}|x_{1:d-1}^{(n)})}, \qquad (21)$$

for $d = 2, \ldots, D$. Clearly, $w(\mathbf{x}_n) \equiv w_D^{(n)}$. The estimator of the normalizing constant $Z_d = \int_{\mathbb{R}^{d\times n}} \pi_d(x_{1:d})dx_{1:d}$ at the d-th iteration is

$$
\begin{aligned}
\widehat{Z}_d &= \frac{1}{N}\sum_{n=1}^{N} w_d^{(n)} = \frac{1}{N}\sum_{n=1}^{N} w_{d-1}^{(n)}\beta_d^{(n)}, \qquad (22) \\
&= \frac{1}{N}\sum_{n=1}^{N}\left[\prod_{j=1}^{d}\beta_j^{(n)}\right]. \qquad (23)
\end{aligned}
$$

Again, $Z \equiv Z_D$ and $\widehat{Z} \equiv \widehat{Z}_D$. However, an alternative formulation is often used [7, 8]

$$
\begin{aligned}
\overline{Z}_d &= \prod_{j=1}^{d}\left[\sum_{n=1}^{N}\bar{w}_{j-1}^{(n)}\beta_j^{(n)}\right] = \prod_{j=1}^{d}\left[\frac{\sum_{n=1}^{N} w_j^{(n)}}{\sum_{n=1}^{N} w_{j-1}^{(n)}}\right], \\
&= \widehat{Z}_1\prod_{j=2}^{d}\left[\frac{\widehat{Z}_j}{\widehat{Z}_{j-1}}\right] = \widehat{Z}_1\frac{\widehat{Z}_2}{\widehat{Z}_1}\cdots\frac{\widehat{Z}_d}{\widehat{Z}_{d-1}} = \widehat{Z}_d. \quad (24)
\end{aligned}
$$

Therefore, in SIS, $\widehat{Z}_d$ in Eq. (22) and $\overline{Z}_d$ in Eq. (24) are equivalent formulations of the same estimator of Z_d [15].

4.1. Estimators of the marginal likelihood in SIR

Sequential Importance Resampling (SIR) [11, 16] combines the SIS approach with the application of the resampling procedure described in Section 3. Considering the proper importance weight of a resampled particle given in Definition 1 and

recalling that $w_d^{(n)}$ the weight at the d-th iteration, we obtain the following recursion, $w_1^{(n)} = \beta_1^{(n)}$ and

$$w_d^{(n)} = w_{d-1}^{(n)} \beta_d^{(n)}, \tag{25}$$

for $d = 2, \ldots, D$, where

$$w_{d-1}^{(n)} = \begin{cases} w_{d-1}^{(n)}, & \text{without resampling at } (d-1)\text{-th iter.,} \\ \widehat{Z}_{d-1}, & \text{with resampling at } (d-1)\text{-th iter.,} \end{cases} \tag{26}$$

i.e., if a resampling is applied at $(d-1)$-th iteration then $\xi_{d-1}^{(n)} = \widehat{Z}_{d-1}, \forall n = 1, \ldots, N$.

Remark 2. *Using the Definition 1 and the recursive definition of the weights $w_d^{(n)}$ in Eqs. (25)-(26), $\widehat{Z}_d$ and $\overline{Z}_d$ are both consistent and equivalent estimators of the marginal likelihood, also in SIR. Namely, the two estimators are*

$$\widehat{Z}_d = \frac{1}{N} \sum_{n=1}^{N} w_d^{(n)}, \quad \overline{Z}_d = \prod_{j=1}^{d} \left[\sum_{n=1}^{N} \bar{w}_{j-1}^{(n)} \beta_j^{(n)} \right] \tag{27}$$

are equivalent, $\widehat{Z}_d \equiv \overline{Z}_d$. For instance, if the resampling is applied at each iteration, they become

$$\overline{Z}_d = \prod_{j=1}^{d} \left[\frac{1}{N} \sum_{n=1}^{N} \beta_j^{(n)} \right], \tag{28}$$

$$\widehat{Z}_d = \widehat{Z}_{d-1} \left[\frac{1}{N} \sum_{n=1}^{N} \beta_d^{(n)} \right] = \prod_{j=1}^{d} \left[\frac{1}{N} \sum_{n=1}^{N} \beta_j^{(n)} \right], \tag{29}$$

and clearly coincide.

Remark 3. *Let us focus on the marginal likelihood estimators at the final iteration, i.e., $\widehat{Z} = \widehat{Z}_D$ and $\overline{Z} = \overline{Z}_D$. Underline{Without} using the Definition 1 and the recursive definition of the weights $w_d^{(n)}$ in Eqs. (25)-(26), the only estimator of the marginal likelihood that can be properly computed in SIR is $\overline{Z}$, that involves only the computation of the normalized weights $\bar{w}_d^{(n)}$ (omitting the values of the corresponding unnormalized ones).*

5. PARTIAL RESAMPLING

The core of Sequential Monte Carlo methods is the SIR approach. Namely, the weights are constructed recursively as in (20) and resampling steps, *involving all the particles*, are applied at some iterations. The combination of both, SIS and resampling schemes, is possible in the standard SIR approach *only* if the entire set of particles is employed in the resampling [6, 7], so that the assumption

$$w(\widetilde{\mathbf{x}}_1) = w(\widetilde{\mathbf{x}}_2) = \ldots = w(\widetilde{\mathbf{x}}_N), \tag{30}$$

is enough for computing $\overline{I}$ and $\overline{Z}$, since $\bar{w}(\widetilde{\mathbf{x}}_n) = \frac{1}{N}$ for all $n \in \{1, \ldots, N\}$. If Definition 1 is used and then the recursive expression (25)-(26) is applied, we can define a resampling procedure *involving only a subset of particles*, as described in the following. We consider to apply a partial resampling scheme at the d-th iteration:

1. Choose randomly without replacement a subset of $M \leq N$ samples,

   $$\mathcal{R} = \{x_d^{(j_1)}, \ldots, x_d^{(j_M)}\},$$

   contained within the set of N particles $\{x_d^{(1)}, \ldots, x_d^{(N)}\}$. Let us denote also the set

   $$\mathcal{N} = \{x_d^{(1)}, \ldots, x_d^{(N)}\} \backslash \mathcal{R},$$

   of the particles which do not take part in the resampling.

2. Give the set $\mathcal{R}$, resample with replacement M particles according to the normalized weights $\bar{w}_d^{(j_m)} = \frac{w_d^{(j_m)}}{\sum_{k=1}^{M} w_d^{(j_k)}}$, for $m = 1, \ldots, M$, obtaining $\widetilde{\mathcal{R}} = \{\widetilde{x}_d^{(1)}, \ldots, \widetilde{x}_d^{(M)}\}$. Clearly, $\widetilde{\mathcal{R}} \subseteq \mathcal{R}$.

3. For all the resampled particles in $\widetilde{\mathcal{R}}$ set

   $$w_d^{(m)} = \frac{1}{M} \sum_{k=1}^{m} w_d^{(j_k)}, \tag{31}$$

   for $m = 1, \ldots, M$, whereas the unnormalized importance weights of the particles in $\mathcal{N}$ remains invariant.

4. Go forward to the iteration $d+1$ of the SIR method, using the recursive formula (20).

The procedure above is valid since yields proper weighted samples by Liu's definition. If $M = N$, it coincides with the traditional resampling procedure. This approach can reduce the loss of diversity due to the application of the resampling [6, 7, 11].

6. CONCLUSIONS

In this work, we have introduced a proper choice of the unnormalized weight assigned to a resampled particle. This choice entails several theoretical and practical consequences. We have described two of them, regarding (1) the estimation of the marginal likelihood and, (2) the application of a partial resampling involving only a subset of the cloud of particles, within SIR techniques. Other novel algorithms (based on the partial resampling perspective) and theoretical consequences (also affecting well-known the MCMC techniques, such as the particle Metropolis-Hastings method [1, 14], and parallel SMC implementations) will be highlighted out in an extended version of this work. Several numerical simulations will be also included.

References

[1] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *J. R. Statist. Soc. B*, 72(3):269–342, 2010.

[2] M. S. Arulumpalam, S. Maskell, N. Gordon, and T. Klapp. A tutorial on particle filters for on-line nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions Signal Processing*, 50(2):174–188, February 2002.

[3] M. Bolić, P. M. Djurić, and S. Hong. Resampling algorithms for particle filters: A computational complexity perspective. *EURASIP Journal on Advances in Signal Processing*, 2004(15):2267–2277, November 2004.

[4] M. F. Bugallo, L. Martino, and J. Corander. Adaptive importance sampling in signal processing. *Digital Signal Processing*, (47):36–49, 2015.

[5] O. Cappé, A. Guillin, J. M. Marin, and C. P. Robert. Population Monte Carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929, 2004.

[6] P. M. Djurić, J. H. Kotecha, J. Zhang, Y. Huang, T. Ghirmai, M. F. Bugallo, and J. Míguez. Particle filtering. *IEEE Signal Processing Magazine*, 20(5):19–38, September 2003.

[7] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York (USA), 2001.

[8] A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: fifteen years later. *technical report*, 2008.

[9] V. Elvira, L. Martino, D. Luengo, and M. F. Bugallo. Generalized multiple importance sampling. *arXiv:1511.03095*, 2015.

[10] 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 Signal Processing*, 50(2):425–437, February 2002.

[11] J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, 2004.

[12] L. Martino, V. Elvira, D. Luengo, and J. Corander. An adaptive population importance sampler: Learning from the uncertainty. *IEEE Transactions on Signal Processing*, 63(16):4422–4437, 2015.

[13] L. Martino, V. Elvira, D. Luengo, and J. Corander. Layered adaptive importance sampling. *arXiv:1505.04732*, 2015.

[14] L. Martino, F. Leisen, and J. Corander. On multiple try schemes and the particle metropolis-hastings algorithm. *viXra:1409.0051*, 2014.

[15] L. Martino, J. Read, V. Elvira, and F. Louzada. Cooperative parallel particle filters for on-line model selection and applications to urban mobility. *viXra:1512.0420*, 2015.

[16] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2004.