

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/164754>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Rethinking the Effective Sample Size

Víctor Elvira^{1,2} | Luca Martino³ | Christian P. Robert⁴

¹University of Edinburgh, United Kingdom

²The Alan Turing Institute, United Kingdom

³King Juan Carlos University of Madrid, Spain

⁴Université Paris Dauphine, France

Correspondence

School of Mathematics

University of Edinburgh

James Clerk Maxwell Building

The King's Buildings

Peter Guthrie Tait Road

Edinburgh EH9 3FD United Kingdom

Email: victor.elvira@ed.ac.uk

The effective sample size (ESS) is widely used in sample-based simulation methods for assessing the quality of a Monte Carlo approximation of a given distribution and of related integrals. In this paper, we revisit the approximation of the ESS in the specific context of importance sampling (IS). The derivation of this approximation, that we will denote as $\widehat{\text{ESS}}$, is partially available in [27]. This approximation has been widely used in the last 25 years due to its simplicity as a practical rule of thumb in a wide variety of importance sampling methods. However, we show that the multiple assumptions and approximations in the derivation of $\widehat{\text{ESS}}$, makes it difficult to be considered even as a reasonable approximation of the ESS. We extend the discussion of the $\widehat{\text{ESS}}$ in the multiple importance sampling (MIS) setting, we display numerical examples, and we discuss several avenues for developing alternative metrics. This paper does not cover the use of ESS for MCMC algorithms.

Keywords — Bayesian inference, effective sample size, importance sampling, Monte Carlo methods.

1 | INTRODUCTION

A usual problem in statistics, signal processing, and machine learning consists in approximating an intractable integral with respect to (w.r.t.) a targeted distribution. In the simplest Monte Carlo algorithm, the integral is approximated by random samples simulated from the targeted distribution itself [e.g., 42, Section 3.2]. However, in many problems it is not possible to simulate from the targeted distribution or it is not efficient to do so. Importance sampling (IS) is a

Abbreviations: ESS, effective sampling size; IS, importance sampling; MC, Monte Carlo; MCMC, Markov chain Monte Carlo.

well-established Monte Carlo technique where samples are drawn from a different distribution (called *proposal* distribution), and an importance weight is assigned to each sample to take into account the mismatch between proposal and targeted distributions [e.g., 42, Section 3.3].

Monte Carlo methods are usually evaluated by the mean square error (MSE) of the estimators that approximate the integral of interest. In many cases, these estimators are unbiased (as, e.g., for the plain Monte Carlo estimator) or the bias quickly vanishes when the number of samples increases (as, e.g., for the self-normalized IS (SNIS) estimator). Therefore, Monte Carlo estimators are most often characterized and compared through their variance [38, Chapter 9]. This is the case of the effective sample size (ESS), which involves the ratio of the variances of the plain Monte Carlo and SNIS estimators. However, computing these variances usually requires solving integrals of a similar complexity as the (intractable) integral at the source of the analysis. Hence, the problem to solve is two-fold. First, the integral of interest is intractable. Second, characterizing the method that approximates the integral best is also impossible, because of a second intractable integral. An alternative that bypasses this deadlock consists in a relative comparison of the performances of different Monte Carlo estimators. In other words, instead of characterizing in absolute terms a specific estimator, one can compare two or more Monte Carlo estimators. However, a naïve comparison of these via a second Monte Carlo experiment is almost always too costly to be considered.

Conceptually, the ESS can be and often is interpreted as the number of particles/samples that would need to be drawn from the target (*direct* or *raw* Monte Carlo) in order to yield the same variance as that from the self-normalized IS estimator under study [27, 31]. However, this relative measure is itself impossible to obtain in a closed form in most cases. Indeed, it is proportional to the ratio of the variances of the two aforementioned estimators, variances that are impossible to compute. For this reason, an approximation of the ESS has been used instead for the last two decades in the Monte Carlo literature. Its simple expression is given by

$$\widehat{\text{ESS}} = \frac{1}{\sum_{n=1}^N \bar{w}_n^2}, \quad (1)$$

where \bar{w}_n represent the normalized importance weights associated with the N samples. The approximation from ESS to $\widehat{\text{ESS}}$ is described in the technical report of [27], although the final part is not fully explicit (see more details in next section). Since the technical report [27] and the follow-up paper [28], the $\widehat{\text{ESS}}$ metric has been widely used (and sometimes misused) as an approximate relative measure of performance between the self-normalized IS estimator and the vanilla Monte Carlo estimators.¹ We highlight that the $\widehat{\text{ESS}}$ was derived in 1992 and has been applied in IS-based methodology that did not exist at that time (e.g., particle filters were proposed later in [20] and also the first adaptive IS samplers [5]). While it is hard to identify the number of publications where the ESS is used, the authors estimate it to be in the thousands. It is therefore of clear interest to understand the derivation of $\widehat{\text{ESS}}$, including its approximations, assumptions, and resulting limitations.

In this paper, we first review in detail the derivation of $\widehat{\text{ESS}}$, revisiting [27]. We clarify all steps, exposing all assumptions and approximations that are introduced with the goal of simplifying intractable terms and producing a computable expression. Then, we summarize the drawbacks of these approximations, focusing on situations that lead to serious problems and misleading conclusions. The main goal of this paper is to raise consciousness around the $\widehat{\text{ESS}}$, so practitioners are aware of its flaws and can be careful when they decide to use it. Note that in this work, we focus on the definition of ESS in IS, and not in the ESS used in Markov Chain Monte Carlo (MCMC) which was first defined in [45].

The rest of the paper is organized as follows. In Section 2, we briefly describe IS and the ESS. In Section 3,

¹The technical report [27] can be alternatively found in <https://victorelvira.github.io/papers/kong92.pdf>.

we completely derive the approximated $\widehat{\text{ESS}}$ from ESS, and the assumptions, approximations, and consequences are clearly listed. We analyze other derivations and connections of the $\widehat{\text{ESS}}$ in Section 4, along with its extension in MIS. We also review existing alternative metrics, and we include three numerical examples. In Section 5, we discuss new directions to develop novel metrics with the purpose of overcoming undesired limitations of the $\widehat{\text{ESS}}$. We conclude the paper with a discussion and some possible future lines in Section 6.

2 | EFFECTIVE SAMPLE SIZE IN IMPORTANCE SAMPLING

2.1 | Importance sampling

Let us consider the problem of estimating the integral

$$I = \int h(\mathbf{x})\bar{\pi}(\mathbf{x})d\mathbf{x}, \quad (2)$$

where $\bar{\pi}$ is a targeted pdf that can be evaluated up to a normalizing constant, and h is an integrable function w.r.t. $\bar{\pi}$. We thus assume that $\bar{\pi}$ is known to an unknown normalizing constant Z ,

$$\bar{\pi}(\mathbf{x}) = \frac{\pi(\mathbf{x})}{Z},$$

where obviously $Z = \int \pi(\mathbf{x})d\mathbf{x}$, meaning that only the non-negative function π can be evaluated. In the simplest version of IS, $\bar{\pi}(\mathbf{x})$ is approximated with a set of N weighted samples as follows. First, N samples $\mathbf{X}_1, \dots, \mathbf{X}_N$ are drawn from a proposal density $q(\mathbf{x})$. Second, each sample is assigned with an importance weights as $W_n = \frac{\pi(\mathbf{X}_n)}{q(\mathbf{X}_n)}$. The unbiased estimator of I , often called *unnormalized IS* (UIS) estimator, is generally used if the normalizing constant Z is known [14]. Moreover, in the general case where Z is unknown, the self-normalized IS (SNIS) estimator can always be used:

$$\tilde{I} = \sum_{n=1}^N \bar{W}_n h(\mathbf{X}_n), \quad \mathbf{X}_n \sim q(\mathbf{x}), \quad (3)$$

where $\bar{W}_n = \frac{W_n}{\sum_{i=1}^N W_i}$ are the normalized weights.

In the case of the UIS estimator, its variance is proportional to the χ^2 -divergence between a normalized version of the function $\bar{\pi}(\mathbf{x})|h(\mathbf{x})|$ and the proposal $q(\mathbf{x})$ (see e.g., [42, 25]). The optimal proposal for the SNIS estimator, \tilde{I} , is $q^*(\mathbf{x}) \propto |h(\mathbf{x}) - I|\pi(\mathbf{x})$ [21, 38, 14]. A possible interpretation is that the optimal proposal in SNIS trades-off minimizing both the aforementioned χ^2 -divergence and also the χ^2 -divergence between the target pdf (i.e., without the function h) and the proposal. It is therefore a legitimate question to examine the efficiency of the set of weighted samples simulated from a specific proposal q .

2.2 | The effective sample size

The effective sample size (ESS) is an indicator of the worth of the IS estimator defined as the number of samples simulated from the target pdf $\bar{\pi}$ that would provide an estimator with a *performance* equal to the performance of the IS estimator based on N samples (drawn from the proposal pdf). More precisely, if it were possible to simulate samples

from the target $\tilde{\pi}$, the integral I could also be approximated by

$$\tilde{I} = \frac{1}{N} \sum_{n=1}^N h(\mathbf{X}_n), \quad \mathbf{X}_n \sim \tilde{\pi}(\mathbf{x}). \quad (4)$$

In [27], the ESS is defined as the ratio between the variance estimator with N samples from the target, and the variance of the IS estimator. In particular, the ESS of the self-normalized estimator is given by

$$\text{ESS} = N \frac{\text{Var}[\tilde{I}]}{\text{Var}[\tilde{I}]}, \quad (5)$$

which is approximated in [27] and applied in [28] for a missing data problem. Note that ESS depends on a specific integrand $h(\cdot)$, which means that the same set of weighted samples can be suitable for a specific function $h_1(\cdot)$ but disastrous for another function $h_2(\cdot)$.

When devising IS-based methods, an assessment of the ESS is of considerable interest, since it provides a direct measure of the efficiency of a potential method. Unfortunately, obtaining analytically this ESS is not possible in most of the scenarios where IS methods are applied; similar integrals to that of Eq. (2) need to be computed, e.g., the variance of \tilde{I} is given by

$$\text{Var}_{\tilde{\pi}}[\tilde{I}] = \frac{1}{N} \left(\mathbb{E}_{\tilde{\pi}}[H^2] - (\mathbb{E}_{\tilde{\pi}}[H])^2 \right), \quad (6)$$

where we denote $H \equiv h(\mathbf{X}_n)$ to alleviate notations.² Unfortunately, neither $\mathbb{E}_{\tilde{\pi}}[H^2] = \int h^2(\mathbf{x})\tilde{\pi}(\mathbf{x})d\mathbf{x}$ nor $\mathbb{E}_{\tilde{\pi}}[H] = \int h(\mathbf{x})\tilde{\pi}(\mathbf{x})d\mathbf{x}$ can be computed in most problems of interest. Moreover, the exact variance of the SNIS estimator, \tilde{I} , has almost never a closed form. Due to the intractability of the ESS, the ratio in Eq. (5) is most often approximated in the Monte Carlo literature by the expression

$$\widehat{\text{ESS}} = \frac{1}{\sum_{n=1}^N \tilde{w}_n^2}. \quad (7)$$

In the following we describe the path leading from (5) to the approximation in (7).

3 | APPROXIMATING THE EFFECTIVE SAMPLE SIZE

In the following, we first detail all the steps in the derivation of the approximation $\widehat{\text{ESS}}$. Then, we summarize the different assumptions and approximations behind $\widehat{\text{ESS}}$. Finally, we analyze the consequences of these assumptions and approximations.

3.1 | From ESS to $\widehat{\text{ESS}}$

First, note that \tilde{I} is a ratio of two dependent r.v.'s and a biased estimator of I . The bias of \tilde{I} can be considered negligible w.r.t. its variance when N grows, since the bias is typically $O(N^{-1})$, so that the bias squared is $O(N^{-2})$, while the variance is $O(N^{-1})$ [42]. Nevertheless, for low or intermediate N , the bias might be far from negligible (see,

²Note that, here, all samples are i.i.d. so we can drop the subindex n from the summation.

e.g., [38]). A more relevant expression for a relative comparison of performance might be

$$ESS^* = N \frac{\text{Var}[\tilde{I}]}{\text{MSE}[\tilde{I}]} \tag{8}$$

Otherwise, using IS with N simulations does not exactly produce the same MSE as using direct simulation with ESS samples from the target. Therefore, a first implicit approximation involved in the handling of ESS is

$$\text{MSE}_q[\tilde{I}] \approx \text{Var}_q[\tilde{I}]. \tag{9}$$

Note that \tilde{I} in Eq. (3) is the ratio between N realizations of the r.v.'s WH and W , and thus both numerator and denominator are dependent. Then, the variance of \tilde{I} in Eq. (3) cannot be computed in a closed form and must be approximated. It is possible to use the delta method, which is a second order Taylor expansion used to approximate expectations of functions involving two dependent r.v.'s [40]. In Appendix A, we provide the full derivation of the delta method for generic functions, and here we apply the resulting Eq. (46) for the case where the function is a ratio, $f(T, Y) = \frac{T}{Y}$, with $T = WH$ and $Y = W$. In this case, the expansion is done around the expectation of each r.v., i.e., $(T_0, Y_0) = (IZ, Z)$. The first derivatives are $f'_T(T, Y) = \frac{1}{Y} = \frac{1}{W}$ and $f'_Y(T, Y) = -\frac{T}{Y^2} = -\frac{H}{W}$. This yields

$$\begin{aligned} \text{Var}_q[\tilde{I}] &\approx \frac{1}{N} \left(\left(\frac{\text{E}_q[HW]}{\text{E}_q[W]} \right)^2 \text{Var}_q[W] + \text{Var}_q[HW] - 2 \left(\frac{\text{E}_q[HW]}{\text{E}_q[W]} \right) \text{Cov}_q[W, HW] \right) \\ &= \frac{1}{N} \left(I^2 \text{Var}_q[W] + \text{Var}_q[HW] - 2I \text{Cov}_q[W, HW] \right), \end{aligned} \tag{10}$$

where we use the identities $\text{E}_q[HW] = ZI$ and $\text{E}_q[W] = Z$. In the following, we assume without loss of generality that the target is normalized, i.e., $Z = 1$, as it is done in [27]. It is easy to show that the derivation holds for any Z by realizing that the normalized weights in \tilde{I} do not depend on Z . Let us expand the expressions

$$\begin{aligned} \text{Cov}_q[W, HW] &= \text{E}_q[HW^2] - \text{E}_q[W]\text{E}_q[WH] \\ &= \text{E}_q[HW^2] - I \\ &= \text{E}_{\tilde{\pi}}[HW] - I \\ &= \text{Cov}_{\tilde{\pi}}[W, H] + I\text{E}_{\tilde{\pi}}[W] - I, \end{aligned} \tag{11}$$

and

$$\begin{aligned} \text{Var}_q[HW] &= \text{E}_q[W^2H^2] - \text{E}_q^2[WH] \\ &= \text{E}_{\tilde{\pi}}[WH^2] - I^2. \end{aligned} \tag{12}$$

The delta method is again applied to approximate the expectation $\text{E}_{\tilde{\pi}}[WH^2]$ using the first two moments of W and H , i.e., as a second order Taylor approximation of this expectation. In this case, we apply the generic approximation for expectations in Eq. (44) in Appendix A, particularizing for function $f(T, Y) = TY^2$, where $T = W$ and $Y = W$:

$$\begin{aligned}
E_{\bar{\pi}}[WH^2] &\approx E_{\bar{\pi}}[W]E_{\bar{\pi}}^2[H] + \frac{1}{2}\text{Var}_{\bar{\pi}}[H](2E_{\bar{\pi}}[W]) + \text{Cov}_{\bar{\pi}}[W, H](2E_{\bar{\pi}}[H]) \\
&= I^2E_{\bar{\pi}}[W] + \text{Var}_{\bar{\pi}}[H]E_{\bar{\pi}}[W] + 2I\text{Cov}_{\bar{\pi}}[W, H].
\end{aligned} \tag{13}$$

The remainder term in this approximation is the expectation of a polynomial order three, namely $E_{\bar{\pi}}[(W - E_{\bar{\pi}}[W])(H - E_{\bar{\pi}}[H])^2] = E_{\bar{\pi}}[(W - E_{\bar{\pi}}[W])(H - I)^2]$. This is easy to show by generalizing Eq. (43) in Appendix A to a third order Taylor expansion.

Then, plugging Eqs. (11), (12), and (13) in Eq. (10), we obtain

$$\begin{aligned}
\text{Var}_q[\bar{I}] &\approx \frac{1}{N} \left(\text{Var}_{\bar{\pi}}[H]E_{\bar{\pi}}[W] + I^2(1 + \text{Var}_q[W] - E_{\bar{\pi}}[W]) \right) \\
&= \text{Var}_{\bar{\pi}}[\bar{I}](1 + \text{Var}_q[W]),
\end{aligned} \tag{14}$$

where in the last equality, we use $\frac{1}{N}\text{Var}_{\bar{\pi}}[H] = \text{Var}_{\bar{\pi}}[\bar{I}]$ (see estimator of Eq. (4)) and $E_{\bar{\pi}}[W] = E_q[W^2] = \text{Var}_q[W] + 1$ (we recall we assume $E_q[W] = Z = 1$ without loss of generality). Therefore, using Eq. (14), the ESS is approximated as

$$\text{ESS} = N \frac{\text{Var}_{\bar{\pi}}[\bar{I}]}{\text{MSE}_q[\bar{I}]} \approx N \frac{\text{Var}_{\bar{\pi}}[\bar{I}]}{\text{Var}_q[\bar{I}]} \approx \frac{N}{1 + \text{Var}_q[W]}. \tag{15}$$

which implies that ESS is always *less than* N , a drawback discussed below. Note that the derivation of [27] ends at this point. In [28, Section 4.1], the authors state that “*although there is no guarantee that the remainder term in this approximation is ignorable, what is nice about Eq. (15) is that it does not involve h* ”. First, the fact that the expression does not depend on h indicates that it is necessarily a loose approximation, since the quality of the IS approximation depends on the mismatch of the proposal $q(\mathbf{x})$ w.r.t. $|h(\mathbf{x})|\pi(\mathbf{x})$ [e.g., 42, 25] (see also the example in Section 3.4.2).

Second, other approximations can be derived from the simulated sample, including unbiased estimates of $E_q[W^2H^2]$ and $E_q^2[WH]$, obtained at the same cost as the original estimates, as well as bootstrap versions. At last, we point out the most detrimental aspect of (15), namely that it always increases the variance from the i.i.d. simulation reference, when it is well-known that there always exist IS versions that on the contrary decrease the variance (see Section 3.3 for more details).

The further approximation from Eq. (15) to $\widehat{\text{ESS}} = 1/\sum_{n=1}^N \bar{w}_n^2$ does not appear in [27]. Note also that the derivation assumes twice a normalized target for the approximation of Eq. (15). This is without loss of generality, since the SNIS estimator in (3) can be rewritten as

$$\bar{I} = \sum_{n=1}^N \frac{W_n}{\sum_{i=1}^N W_i} h(\mathbf{X}_n), \quad \mathbf{X}_n \sim q(\mathbf{x}). \tag{16}$$

Then, it becomes evident that the estimator remains the same (and hence its variance) if all weights are multiplied by a constant. Hence, one can assume in the previous formulation that the weights are $\widetilde{W}_n = \frac{W_n}{Z}$ instead W_n , which yields $E_q[\widetilde{W}] = 1$ true. In that case, (14) turns

$$\text{ESS} \approx \frac{N}{1 + \text{Var}_q[\widetilde{W}]} \tag{17}$$

Then, if Z is known, Eq. (17) can be adapted as

$$\text{ESS} \approx \frac{N}{1 + \frac{\text{Var}_q[W]}{Z^2}} \quad (18)$$

$$= N \frac{Z^2}{E_q[W^2]}, \quad (19)$$

where we have used $\text{Var}_q[W] = E_q[W^2] - E_q^2[W] = E_q[W^2] - Z^2$.³ In practical scenarios, neither Z or $E_q[W^2]$ is known. A particle approximation of the expectations can be used to produce the estimator

$$\widehat{\text{ESS}} = N \frac{\left(\frac{1}{N} \sum_{n=1}^N w_n\right)^2}{\frac{1}{N} \sum_{n=1}^N w_n^2} \quad (20)$$

$$= \frac{1}{\sum_{n=1}^N \tilde{w}_n^2}, \quad (21)$$

where we use $Z \approx \frac{1}{N} \sum_{n=1}^N w_n$ and $E_q[W^2] \approx \frac{1}{N} \sum_{n=1}^N w_n^2$. We remark that, although increasing N improves the quality of the particle approximation in the last step, this is not a guarantee of $\widehat{\text{ESS}}$ getting closer to ESS, as we show in the example of Section 3.4.1.

3.2 | Summary of assumptions and approximations

The question at this stage is whether or not $\widehat{\text{ESS}}$ is a good mathematical approximation of the effective sample size, but also whether or not it keeps some of its desirable properties. Let us summarize the assumptions and approximations behind the derivation of ESS:

1. The ESS is conceptually defined as the ratio between the performance of two estimators, the direct Monte Carlo estimator, \bar{I} , where samples are drawn from the target $\bar{\pi}$, and the SNIS estimator \tilde{I} . The expression $\text{ESS} = N\text{Var}_{\bar{\pi}}[\bar{I}]/\text{Var}_q[\tilde{I}]$ in Eq. (15) does not take in account the bias of \tilde{I} (which can be significant for small N). Then, the ratio of variances overestimates the theoretical value $\text{ESS} = N\text{Var}_{\bar{\pi}}[\bar{I}]/\text{MSE}_q[\tilde{I}]$, especially for low N . This can be particularly dangerous since the approximated ESS can report a good performance in a scenario with large bias. However, the error of this first approximation vanishes as N grows.
2. In the derivation of [27], all the samples are considered to be i.i.d. from a unique proposal q , i.e., $\mathbf{X}_n \sim q(\mathbf{x})$, for $n = 1, \dots, N$. Nevertheless, $\widehat{\text{ESS}}$ is used mostly in algorithms with multiple proposals, as in adaptive importance sampling (AIS) or sequential Monte Carlo methods (SMC) methods [9, 12, 7, 5, 11, 16, 20, 35, 13]. It is worth noting that controlling the $\widehat{\text{ESS}}$ has been proved to guarantee the long term stability of sequential Monte Carlo methods under certain assumptions [47] (see Section 4.1.4). Extensions to dependent samples as in MCMC are even more delicate.
3. A first delta method is applied in order to approximate $\text{Var}_q[\tilde{I}]$ in Eq. (10). Then, the delta method is applied a second time to approximate $E_{\bar{\pi}}[WH^2]$ in Eq. (13).

³The authors have unsuccessfully tried to find a published article with the remaining derivation from (15) to (21). The first and only reference from this final part can be found in the blog post of one of the authors [41].

3.3 | Some (undesirable) consequences

As a consequence of the unrealistic assumptions and approximations (mostly the Taylor expansions in the delta methods) in the derivation of $\widehat{\text{ESS}}$, this approximated ESS exhibits some flawed features that parallel the flaws in its derivation.

1. The approximation is bounded as

$$1 \leq \widehat{\text{ESS}} \leq N.$$

We provide a proof in Section 4.1.1. This leads to two undesired consequences:

- Even when the sample approximation to the integral I is very poor, we have that $\widehat{\text{ESS}} \geq 1$. Note that the lower bound of the (exact) ESS is 0, corresponding to the case when the IS estimator has infinity variance and is completely inappropriate. That $\widehat{\text{ESS}}$ cannot detect infinite variance estimator appears to us as a major defect of this indicator.
 - The bound of $\widehat{\text{ESS}} \leq N$ is equally limiting. In many settings, $\text{Var}_q(\tilde{I}) \leq \text{Var}_{\tilde{\pi}}[\tilde{I}]$ for an adequate choice of the proposal q . For instance, in connection with several advanced IS techniques such as PMC [5], AIS methods [3], a smart choice of MIS schemes [17], or variance reduction techniques [38, Section 8], importance sampling estimates can significantly lower the variance of the i.i.d. solution. In such scenarios, ESS is *larger* than N , but $\widehat{\text{ESS}}$ cannot capture this variance reduction. Note that in the extreme case where the theoretically optimal proposal is used, the UIS estimator can achieve zero variance, but also the SNIS estimator can largely reduce variance w.r.t. standard Monte Carlo. Hence, a good approximation of ESS should never be upper bounded.
2. $\widehat{\text{ESS}}$ does not depend on the function h . While this is an advantage for practical reasons, the ESS conveys the efficiency or lack thereof of the approximation of the moment in h of $\tilde{\pi}$, but this dependence is completely lost in the approximation. For the same target distribution, a particle approximation can provide a low-variance estimator of some moment h_1 and be disastrous for another moment h_2 . See Section 3.4.2 for a toy example of estimation of rare events.
 3. $\widehat{\text{ESS}}$ does not depend on the set samples $\{\mathbf{X}_n\}_{n=1}^N$, but only in their associated (normalized) weights $\{\tilde{w}_n\}_{n=1}^N$. In other words, the diversity of the sample cannot be directly evaluated through the corresponding weights, which are a mapping from the dimension of X to dimension one. While this might seem to be a feature, $\widehat{\text{ESS}}$ clearly lacks the ability to assess the diversity of the sample approximation, which is key to evaluate its worth. It is well-known that having diversity in the samples is essential to improve the performance of the estimators, e.g., negative correlation is induced to reduce their variance in antithetic and stratified sampling [38, Chapter 8]. This problem is also evident when resampling steps are needed, e.g., in particle filtering and other iterative or adaptive IS-based methods [10, 29]. Note that, despite a potential inadequacy of the sampling approximation, the ESS is always equal to N after a resampling step since all the weights are then set equal, even if the same particle was replicated N times (potentially, in the worst scenarios, in an area of the target with low mass or where h can be zero).

3.4 | Examples

In the following, we provide two examples that allow us to evaluate the behavior of the $\widehat{\text{ESS}}$, and a third example with multiple proposals is provided in Section 4.3. In the first example, we study how the disparity between ESS and $\widehat{\text{ESS}}$

grows when the mismatch between target and proposal densities is increased. The second example focuses on the estimation of rare events. More precisely, h is an indicator function that takes value one in a subset of the support where the target density has low probability mass. This example highlights the disadvantage of not including function h in $\widehat{\text{ESS}}$.

3.4.1 | ESS and $\widehat{\text{ESS}}$ depending on the mismatch between target and proposal

Let us consider the target $\tilde{\pi}(x) = \mathcal{N}(x; 0, 1)$. In Figure 1, we show both the ESS and $\widehat{\text{ESS}}$ of the SNIS estimator of the target mean, i.e. $h(x) = x$, with $N \in \{4, 16, 256\}$ samples, when $q(x) = \mathcal{N}(x; \mu_q, 1)$, varying μ_q in the interval $[0, 3]$. The values of ESS and $\widehat{\text{ESS}}$ are normalized over N , so all cases can be compared. We note that ESS and $\widehat{\text{ESS}}$ are only available in the case without mismatch. In this example, we approximate both metrics through 10,000 independent Monte Carlo runs. In the right subplot, we display the ratio between $\widehat{\text{ESS}}$ and ESS (i.e., when the line is above 1, the $\widehat{\text{ESS}}$ overestimates the true ESS). Note that the efficiency deteriorates for large N . In all cases, $\widehat{\text{ESS}}$ overestimates the ESS, with a larger gap when the mismatch between target and proposal increases. In this example, all estimators have finite variance.

Figure 2 shows the results of the same experiment when there is a mismatch between the variances of target and proposal. In particular, the proposal is now $q(x) = \mathcal{N}(x; 0, \sigma_q^2)$, and we vary σ_q in the interval $[0.6, 3.6]$ (note that too small a proposal variance can yield IS estimators with infinite variance). Note that again there is a gap between ESS and $\widehat{\text{ESS}}$. Interestingly, for high values of N , and for high values of σ_q , the $\widehat{\text{ESS}}$ now underestimates the ESS (contrary to the example with mean mismatch). This toy example contradicts the broad belief that states that if the $\widehat{\text{ESS}}$ is high, it is unclear if the quality of the approximation is good but if the $\widehat{\text{ESS}}$ is low, then necessarily the IS method is failing. One of the reasons for this contradiction is that in this example, the optimal proposal is broader than the target. As discussed above, the $\widehat{\text{ESS}}$ is blind to $h(x)$ and also to cases where $\text{ESS} > N$. Note that in the left subplot of Fig. 2, for $N \in \{16, 256\}$ and in the range $\sigma_q \in [1, 2.5]$, the ESS is larger than N (larger than 1 in the normalized version displayed in the figure).⁴

We now repeat the experiment to estimate the second moment of the target $\tilde{\pi}(x)$, i.e., with $h(x) = x^2$. Figures 3 and 4 show, for the same setup, the ESS and $\widehat{\text{ESS}}$ (left plot) and the the ratio between $\widehat{\text{ESS}}$ and ESS (right plot). We note that the $\widehat{\text{ESS}}$ is the same as in the previous setup, since the metric is not aware of $h(x)$, as thoroughly discussed in this paper. This explains why the mismatch between ESS and $\widehat{\text{ESS}}$ is even more evident than in the previous case (with $h(x) = x$).

3.4.2 | Estimation of rare events

We now focus on a rare-event estimation problem [2]. In particular, we consider a simplified version of the setup in [39, 19], targeting the estimation of the integral in Eq. (2) with $h_\alpha = \mathbb{1}_{x \in \mathcal{S}_\alpha}(x)$, i.e., the indicator function that takes value one if $x \in \mathcal{S}_\alpha$ and zero otherwise. We consider the case with no mismatch between target and proposal, with $\tilde{\pi}(x) = q(x) = \mathcal{N}(0, \mathbf{I}_{d_x})$, where d_x is the dimension of \mathbf{x} , and \mathbf{I}_{d_x} is the identity matrix of such dimension. We choose $\mathcal{S}_\alpha = \mathbb{R}^{d_x} \setminus [-\alpha, \alpha]^{d_x}$. For instance, for $d_x = 1$, the region is $\mathcal{S}_\alpha = \{x : |x| > \alpha\}$.

This example is relevant because the efficiency of direct sampling decays when both α and d_x grow. First, we note that, since function h_α is not present in $\widehat{\text{ESS}}$, this metric cannot capture the efficiency loss that depends on α . Moreover, in this case, we can also question the adequacy of the ESS as an efficiency metric per se, since regardless

⁴When $\sigma_q < 1$, the variance of SNIS estimator can be infinite. This situation should be detected more easily with a different approach, while the $\widehat{\text{ESS}}$ is clearly ineffective in detecting the malfunctioning.

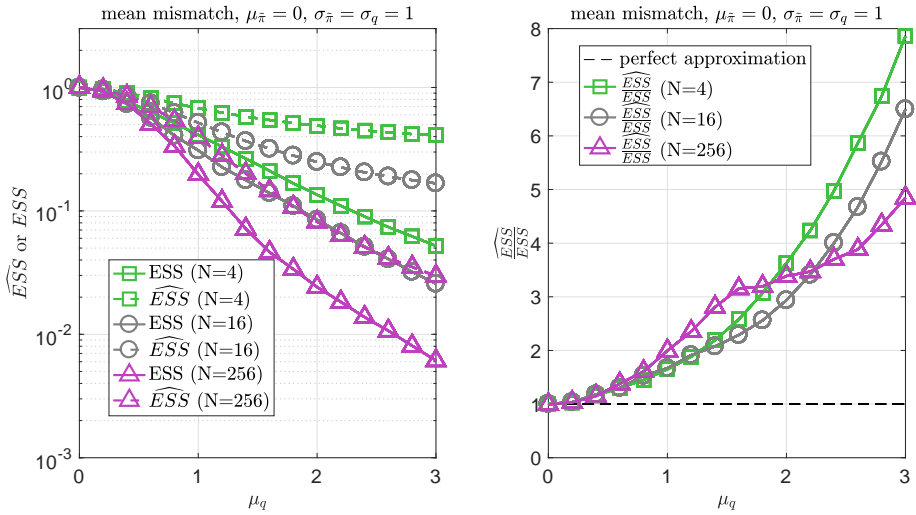


FIGURE 1 Left plot shows the ESS and the \widehat{ESS} for the SNIS estimator of the mean of the target $\tilde{\pi}(x) = \mathcal{N}(x; 0, 1)$, when the mean of the proposal, μ_q , is different from that of the target. The right plot shows the ratio between \widehat{ESS} and ESS. In all situations, the \widehat{ESS} overestimates ESS. Both ESS and \widehat{ESS} are normalized by the number of samples (divided by N), so they can be compared across different values of N .

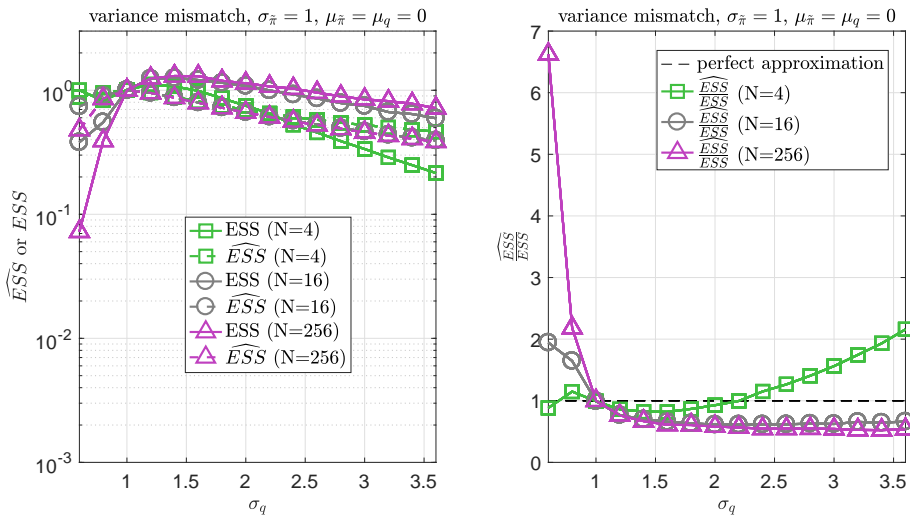


FIGURE 2 Left plot shows the ESS and the \widehat{ESS} for the SNIS estimator of the mean of the target $\tilde{\pi}(x) = \mathcal{N}(x; 0, 1)$, when the variance of the proposal, σ_q^2 , is different from that of the target. The right plot shows the ratio between \widehat{ESS} and ESS. Note that now the \widehat{ESS} sometimes underestimates and sometimes overestimates the true ESS. Both ESS and \widehat{ESS} are normalized by the number of samples (divided by N), so they can be compared across different values of N .




FIGURE 3 Left plot shows the ESS and the $\widehat{\text{ESS}}$ for the SNIS estimator of the second moment of the target $\tilde{\pi}(x) = \mathcal{N}(x; 0, 1)$, when the mean of the proposal, μ_q , is different from that of the target. The right plot shows the ratio between $\widehat{\text{ESS}}$ and ESS. In all situations, the $\widehat{\text{ESS}}$ overestimates ESS. Both ESS and $\widehat{\text{ESS}}$ are normalized by the number of samples (divided by N), so they can be compared across different values of N .

of N , all weights are always equal to 1, and hence $\widehat{\text{ESS}} = N$. The normalized weights are thus identical, and estimator \tilde{I} of Eq. (3) becomes the estimator \bar{I} of Eq. (4). Hence, the original effective sample size of (5) also yields $\text{ESS} = N$, regardless the choice of h .

The top plot of Fig. 5 shows the variance of the SNIS estimator in log-scale, for $d_x \in [1, 2, 5]$ and with $N = 10$ samples (we recall it is the same as the standard Monte Carlo estimator in this case). The bottom plot shows the relative root MSE (RRMSE) that normalizes the square root MSE by the true value of the integral I_α . It can be seen that, when α grows, the efficiency decays: the RRMSE grows to infinity with α , while $\text{ESS} = \widehat{\text{ESS}} = N$ in the whole setup.

We now consider instead the optimal SNIS proposal, $q^*(\mathbf{x}) \propto |h(\mathbf{x}) - I|\pi(\mathbf{x})$, and connect this experiment with [38, Example 9.4]. For simplicity, we particularize for $d_x = 1$ and $S_\alpha = \{x : x > \alpha\}$. Since $0 < I < 1$ for any $\alpha > 0$, it is easy to show that $q^*(x) = \frac{1}{2I(1-I)} ((1-I)\tilde{\pi}(x)\mathbb{1}_{x \in S_\alpha}(x) + I\tilde{\pi}(x)\mathbb{1}_{x \notin S_\alpha}(x))$. Regardless α , the optimal SNIS places half of the mass out of the region of interest S_α , which can be considered an inefficiency and explains why the UIS estimator is preferred for rare events estimation when it can be used (i.e., when Z is known) [39, 18]. In Fig. 6, we show the variance (top) and the RRMSE (middle) of the SNIS estimator in two cases: with the sub-optimal proposal $q(x) = \tilde{\pi}(x)$ (as in the previous setup) and with the optimal proposal $q^*(x)$. Moreover, we display the ESS and $\widehat{\text{ESS}}$ (bottom) for both estimators. The first remark is that, as expected, the optimal SNIS estimator obtains a better performance than the sub-optimal SNIS estimator. This is obviously translated into a higher ESS for the optimal SNIS. As discussed in Section (3.3), the ESS can be in general unbounded, and in this example, the ESS of the optimal SNIS estimator grows to infinity when α grows. However, the $\widehat{\text{ESS}}$ associated to the optimal proposal does not grow (we

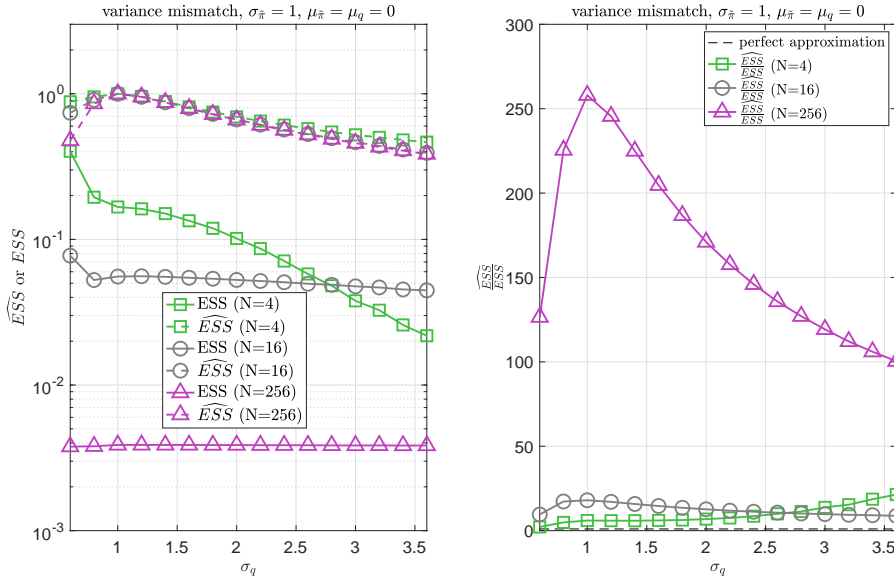


FIGURE 4 Left plot shows the ESS and the \widehat{ESS} for the SNIS estimator of the second moment of the target $\tilde{\pi}(x) = \mathcal{N}(x; 0, 1)$, when the variance of the proposal, σ_q^2 , is different from that of the target. The right plot shows the ratio between \widehat{ESS} and ESS. Note that now the \widehat{ESS} sometimes underestimates and sometimes overestimates the true ESS. Both ESS and \widehat{ESS} are normalized by the number of samples (divided by N), so they can be compared across different values of N .

recall that by construction is upper-bounded by N), and even more, it is inferior to the \widehat{ESS} of the sub-optimal SNIS. Moreover, the \widehat{ESS} of the optimal SNIS decays as α grows, while its ESS clearly grows.

As a summary, in this example the ESS metric and its approximation are unable to capture the loss of efficiency. This is because the ESS is a relative metric of the SNIS efficiency w.r.t. the standard Monte Carlo estimator, \bar{I} , which is considered as the gold standard in ESS, while it is known to be highly inefficient in rare events [39]. Moreover, more sophisticated rare-event estimation methods based on IS [39, 19] can achieve an ESS of orders of magnitude above N , while the \widehat{ESS} will never be above N by construction, as discussed in Section 3.3. We strongly discourage the use of \widehat{ESS} in rare-events problems.

4 | BEYOND \widehat{ESS} FOR DIAGNOSTICS IN IMPORTANCE SAMPLING

In this section, we provide alternative derivations/interpretations of the \widehat{ESS} and we discuss its use when several proposals are used for simulating the samples.

4.1 | Alternative derivations of \widehat{ESS}

The broad applicability of the \widehat{ESS} in the literature of particle-based methods has motivated alternative derivations or connections with other metrics. In the following, we present four of them.




FIGURE 5 Estimation of rare event with $\tilde{\pi}(x) = q(x) = \mathcal{N}(0, 1)$, and $h_\alpha = \mathbb{1}_{x \in \mathcal{S}_\alpha}(x)$, with $\mathcal{S}_\alpha = \mathbb{R}^{d_x} \setminus [-\alpha, \alpha]^{d_x}$, with $d_x \in \{1, 2, 5\}$. **Top:** Variance of standard Monte Carlo (and SNIS estimator) in log-scale when α grows. Note that for high values of α , although the efficiency keeps deteriorating, the variance decays because the true value of the integral I_α becomes close to zero faster. **Bottom:** Relative root MSE (RRMSE) in log-scale of both estimators defined as $RRMSE = \frac{\text{Var}(\hat{I})}{I_\alpha}$. In this example, both ESS and $\widehat{\text{ESS}}$ are equal to N regardless α or d_x .

4.1.1 | Relation of $\widehat{\text{ESS}}$ with the discrepancy of the normalized weights

Let us consider the Euclidean distance L_2 between the uniform probability mass function (pmf) on $\mathcal{U}\{1, 2, \dots, N\}$, denoted as $\bar{\mathbf{w}}^* = [1/N, \dots, 1/N]$, and the pmf described by $\bar{\mathbf{w}} = [\bar{w}_1, \dots, \bar{w}_N]$ on the same support:

$$\begin{aligned} L_2 &= \|\bar{\mathbf{w}} - \bar{\mathbf{w}}^*\|_2 \\ &= \sqrt{\sum_{n=1}^N \left(\bar{w}_n - \frac{1}{N}\right)^2}. \end{aligned}$$

Developing the square and re-arranging terms:

$$\begin{aligned} L_2 &= \sqrt{\left(\sum_{n=1}^N \bar{w}_n^2\right) + N\left(\frac{1}{N^2}\right) - \frac{2}{N} \sum_{n=1}^N \bar{w}_n} \\ &= \sqrt{\left(\sum_{n=1}^N \bar{w}_n^2\right) - \frac{1}{N}} \\ &= \sqrt{\frac{1}{\widehat{\text{ESS}}} - \frac{1}{N}}. \end{aligned} \tag{22}$$




FIGURE 6 Estimation of rare event with SNIS estimators, with sub-optimal proposal $q(x) = \tilde{\pi}(x) = \mathcal{N}(0, 1)$ and optimal proposal $q^*(x)$. The function is $h_\alpha(x) = \mathbb{1}_{x > \alpha}(x)$. **Top:** Variance of standard Monte Carlo (and SNIS estimator) in log-scale when α grows. Note that, for high values of α , while the efficiency of SNIS keeps deteriorating, the variance decays because the true value of the integral I_α becomes close to zero faster. **Middle:** Relative root MSE (RRMSE) in log-scale of both estimators defined as $RRMSE = \frac{\text{Var}(\hat{I})}{I_\alpha}$. **Bottom:** ESS and $\widehat{\text{ESS}}$ for both SNIS estimators. The efficiency with the optimal proposal is superior to the case of the sub-optimal proposal, with an increasing gap when α grows. However, the averaged $\widehat{\text{ESS}}$ is inferior in the optimal case, which again shows a misleading behavior of this metric.

Alternatively, we can express this relation as

$$\widehat{\text{ESS}} = \frac{1}{L_2^2 + \frac{1}{N}}. \quad (23)$$

It is clear that decreasing the mismatch between the pmf described by the normalized weights, \bar{w}_n , and the uniform pmf is equivalent to increasing the ESS. Moreover, this connection allows to prove that $1 \leq \text{ESS} \leq N$ as follows. First, since $L_2 \geq 0$, the maximum value of the approximated effective sample size is $\widehat{\text{ESS}} = N$, when $L_2 = 0$. Second, since $\bar{\mathbf{w}}$ is in the $(N-1)$ -simplex, the vertices are the points that maximize the distance to the center, $\bar{\mathbf{w}}^*$. Their distance is $L_2 = \sqrt{1 - 1/N}$, which corresponds to the minimum approximated effective sample size, $\widehat{\text{ESS}} = 1$, when $\bar{\mathbf{w}}$ is any permutation of the vector $[1, 0, \dots, 0]$. This perspective of *approximated* ESS as a discrepancy measure of the normalized weights, is exploited in [23, 34] to propose novel discrepancy measures with similar properties to $\widehat{\text{ESS}}$ but with beneficial behavior in different situations. These alternative metrics are discussed in Section 4.2.

4.1.2 | $\widehat{\text{ESS}}$ as loss of efficiency in a convex combination of independent r.v.'s

In [38, Section 9.4], the following interesting toy example is proposed. Let us consider N i.i.d. r.v.'s Z_n , with $n = 1, \dots, N$ with variance σ_Z^2 and mean μ_Z . Let us assume a set of unnormalized weights $\mathbf{w} = \{w_n\}_{n=1}^N$ that are used to build the following linear combination

$$C_{\mathbf{w}} = \frac{\sum_{n=1}^N w_n Z_n}{\sum_{n=1}^N w_n}. \quad (24)$$

Note that if all weights were the same, denoting \mathbf{w}^* as a vector with equal non-negative entries, the variance of $C_{\mathbf{w}^*}$ would be $\sigma_{C_{\mathbf{w}^*}}^2 = \frac{\sigma_Z^2}{N}$. The question is, to how many N_{eff} i.i.d. samples (instead of N), the weighted combination $C_{\mathbf{w}}$ is equivalent to? The variance of $C_{\mathbf{w}}$ can be readily computed as

$$\begin{aligned} \sigma_{C_{\mathbf{w}}}^2 &= \frac{\sum_{n=1}^N w_n^2 \sigma_Z^2}{\left(\sum_{n=1}^N w_n\right)^2} \\ &= \frac{\sigma_Z^2}{\widehat{\text{ESS}}}, \end{aligned} \quad (25)$$

i.e., the variance of the combination $C_{\mathbf{w}}$ is equivalent to the variance of the equal-weights combination of $\widehat{\text{ESS}}$ samples. Note that this example is far from being realistic in the intricate case of IS. In the case of single-proposal IS, the weights are also r.v.'s and not deterministic as in the previous example. Moreover, the weights are also dependent of the r.v.'s they are weighting, i.e., $W_n = \frac{\pi(\mathbf{X}_n)}{q(\mathbf{X}_n)}$ and $h(\mathbf{X}_n)$ are usually dependent r.v.'s. Finally, in the previous example, the variance of Z_n is fixed, while in IS, depending if the samples are simulated from the proposal or the target, the r.v. $h(\mathbf{X}_n)$ has a different variance. However, this interesting and illustrative derivation highlights the lost of efficiency resulting from having a few weights dominate over the others.

4.1.3 | Relation of $\widehat{\text{ESS}}$ with the coefficient of variation

The derivation of [31] starts defining the coefficient of variation (CV) of the normalized weights as

$$\text{CV} = \sqrt{\frac{1}{N} \sum_{n=1}^N (N \bar{w}_n - 1)^2}. \quad (26)$$

Then, by basic manipulations, it can be shown that

$$\widehat{\text{ESS}} = \frac{N}{1 + \text{CV}^2} = \frac{1}{\sum_{n=1}^N \bar{w}_n^2}. \quad (27)$$

The authors of [31] discuss that CV^2 is a reasonable approximation of $\text{Var}_q[W]$, assuming that $E_q[W] = 1$ (see the list of assumptions and approximations in Section 3.2). Note the relation of this derivation with the minimum square distance of Section 4.1.1. In fact, the relation between CV and the Euclidean distance L_2 previously defined is $\text{CV} = N\sqrt{L_2}$.

4.1.4 | $\widehat{\text{ESS}}$ as approximation of the χ^2 divergence between the target and proposal

First, we show the connection between the second moment of the importance weights and the χ^2 divergence between target and proposal (see for instance [6]), that can be expressed as

$$\chi^2(\tilde{\pi}, q) = \int \frac{(\tilde{\pi}(\mathbf{x}) - q(\mathbf{x}))^2}{q(\mathbf{x})} d\mathbf{x} \quad (28)$$

$$= \int \frac{\tilde{\pi}^2(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x} - 1 \quad (29)$$

$$= E_q[W^2] - 1. \quad (30)$$

The quantity $\rho = E_q[W^2]$, which appears in the divergence above, has been studied for instance in [1] and [47], in both cases due to the connection to the $\chi^2(\tilde{\pi}, q)$ and for its relation with the $\widehat{\text{ESS}}$. This quantity plays an important role since it helps to bound the error of the SNIS estimator under certain conditions [1, Theorem 2.1]:

$$\sup_{|h| \leq 1} E \left[\left(\tilde{I} - I \right)^2 \right] \leq \frac{4}{N} \rho, \quad \forall N \geq 1. \quad (31)$$

By inspecting Eq. (20), we note that $\widehat{\text{ESS}}/N$ can be interpreted as a sample approximation of the quantity $\frac{E_q[W]^2}{E_q[W^2]}$. Therefore, for large enough N we have that

$$\sup_{|h| \leq 1} E \left[\left(\tilde{I} - I \right)^2 \right] \lesssim \frac{4}{\widehat{\text{ESS}}} \quad (32)$$

as shown in [1, Section 2.3.2]. Although the result is only asymptotic and limited to particular choices of h , it provides an interesting connection among the $\widehat{\text{ESS}}$, the χ^2 between target and proposal, and the efficiency of the SNIS estimator. In relation with this result, [47] show that by controlling the $\widehat{\text{ESS}}$, it is possible to guarantee the long term stability of sequential Monte Carlo methods.

4.2 | Alternative metrics to $\widehat{\text{ESS}}$

Due to the different drawbacks of the $\widehat{\text{ESS}}$ expression in Eq. (7) highlighted above, several other approximations have been studied in literature [22, 34]. In general, those metrics do not aim at improving the approximation of the ESS in (5), but instead look for alternative metrics with similar properties as the $\widehat{\text{ESS}}$.

For instance, a metric called perplexity, involving the discrete entropy [8] of the normalized weights was originally proposed in [4] (see also [43, Chapter 4] and [12, Section 3.5]). Another alternative metric that has shown a good performance in a variety of applications is defined as $\frac{1}{\max[\tilde{w}_1, \dots, \tilde{w}_N]}$, i.e., the inverse of the maximum of the normalized weights \tilde{w}_n [34]. Here, we review a family of metrics, called *Huggins-Roy's family*, introduced in [22] and independently proposed, studied, and generalized in [34]. All these alternative metrics fulfill the required conditions described in the

framework proposed in [34]. The Huggins-Roy's family is defined as

$$\begin{aligned} \widehat{ESS} &= H_N^{(\beta)}(\bar{\mathbf{w}}) = \left(\frac{1}{\sum_{n=1}^N \bar{w}_n^\beta} \right)^{\frac{1}{\beta-1}}, \\ &= \left(\sum_{n=1}^N \bar{w}_n^\beta \right)^{\frac{1}{1-\beta}}, \quad \beta \geq 0. \end{aligned}$$

This family is related to the Rényi entropy of the probability mass function (pmf) defined by $\{\bar{w}_n\}_{n=1}^N$ [8]. Indeed, the Rényi entropy is defined as

$$R_N^{(\beta)}(\bar{\mathbf{w}}) = \frac{1}{1-\beta} \log \left[\sum_{n=1}^N \bar{w}_n^\beta \right], \quad \beta \geq 0$$

Then, it is straightforward to note that

$$H_N^{(\beta)}(\bar{\mathbf{w}}) = \exp \left(R_N^{(\beta)}(\bar{\mathbf{w}}) \right).$$

Relevant specific cases are shown in Table 1. However, more research is needed in order to understand theoretically these families of metrics, and to propose new alternatives to \widehat{ESS} that do not necessarily inherit some of its undesired properties.

TABLE 1 Special cases of ESS measures contained in the Huggins-Roy's family $H_N^{(\beta)}(\bar{\mathbf{w}})$.

$\beta = 0$	$\beta = 1/2$	$\beta = 1$	$\beta = 2$	$\beta = \infty$
$\ \bar{\mathbf{w}}\ _0$	$\left(\sum_{n=1}^N \sqrt{\bar{w}_n} \right)^2$	$\exp \left(- \sum_{n=1}^N \bar{w}_n \log \bar{w}_n \right)$	$\frac{1}{\sum_{n=1}^N \bar{w}_n^2}$	$\frac{1}{\max\{\bar{w}_1, \dots, \bar{w}_N\}}$
L_0 norm of $\bar{\mathbf{w}}$ (number of non-zero elements)		(perplexity) [4]	\widehat{ESS} [27]	[34]

4.3 | ESS in Multiple Importance Sampling

As discussed in Section 3.2, the derivation of \widehat{ESS} assumes i.i.d. samples from a single proposal. However, \widehat{ESS} is used in situations where this is clearly not the case. The extension of single IS to multiple IS (MIS) is not unique, and many possible weighting and sampling schemes are possible (see a thorough review in [17]). This means that the N samples and weights are r.v.'s that can follow a large set of possible distributions, and still build consistent estimators. Obviously, the assumptions and approximations from (8) to (7) are in this case even more unrealistic than in single-proposal IS. For instance, if N proposals are available and N samples must be simulated, one can decide to deterministically sample once per proposal as $\mathbf{x}_n \sim q_n(\mathbf{x})$, $n = 1, \dots, N$, among many other options [17, Section 3]. For this sampling scheme, several weighting schemes are also possible, e.g., the traditional interpretation of MIS (denoted as N1 in [17]) given by

$$w_n = \frac{\pi(\mathbf{x}_n)}{q_n(\mathbf{x}_n)}, \quad n = 1, \dots, N. \tag{33}$$

while the *deterministic mixture* scheme (N3) implements the weights as

$$w_n = \frac{\pi(\mathbf{x}_n)}{\frac{1}{N} \sum_{j=1}^N q_j(\mathbf{x}_n)}, \quad n = 1, \dots, N. \quad (34)$$

More details can be found in [17], where it is shown that the UIS of N3 always present a better performance than N1 in terms of variance of the UIS estimator. The effect in the $\widehat{\text{ESS}}$ of the different sampling and weighting schemes is yet to be understood as we show in the next example.

4.3.1 | Toy example in MIS

Let us consider the multimodal target distribution,

$$\pi(x) = \frac{1}{3} \mathcal{N}(x; v_1, c^2) + \frac{1}{3} \mathcal{N}(x; v_2, c^2) + \frac{1}{3} \mathcal{N}(x; v_3, c^2),$$

with means $v_1 = -3$, $v_2 = 0$, and $v_3 = 3$, and variance $c^2 = 1$. We implement MIS estimators with three proposals, $q_i(x) = \mathcal{N}(x; \mu_i, \sigma^2)$, $i = 1, 2, 3$.

We select three scenarios depending on the mismatch between the target and the mixture of proposals:

1. **Scenario 1: no mismatch.** The proposal means are $\mu_i = v_i$ and $i = 1, 2, 3$ and the variance is $\sigma^2 = c^2 = 1$, i.e., the proposal pdfs can be seen as a whole mixture that exactly replicates the target, i.e., $\pi(x) = \psi(x) = \frac{1}{3} q_1(x) + \frac{1}{3} q_2(x) + \frac{1}{3} q_3(x)$.
2. **Scenario 2: mild mismatch.** The proposal means are $\mu_1 = -3$, $\mu_2 = -1$, and $\mu_3 = 3$, and the variance is $\sigma^2 = 2$.
3. **Scenario 3: large mismatch.** The proposal means are $\mu_1 = -4$, $\mu_2 = -1$, and $\mu_3 = 1$, and the variance is $\sigma^2 = 2$.

In all situations, the goal is estimating the mean of the target pdf with the six MIS schemes proposed in [17, Section 5]. The schemes N1 and N3 are described above, and the R3 index refers to the schemes where samples are simulated from the mixture of proposals, $\psi(x)$, and all the weights are computed as $w_n = \frac{\pi(x)}{\psi(x)}$.

Figure 7 shows, for the case of no mismatch (Scenario 1), the ESS, the $\widehat{\text{ESS}}$, and the ratio of both quantities, for all MIS schemes and for different values of total number of samples N (from 3 to 3×2^9). Both ESS and $\widehat{\text{ESS}}$ are normalized (divided by N) so the behavior of the later can be analyzed. First, note that in this example, the scheme R3 corresponds to direct sampling from the target distribution and its ESS is 1 regardless of the number of total samples drawn, which explains why it is the only scenario and scheme where $\widehat{\text{ESS}}$ is accurate. Note that the $\widehat{\text{ESS}}$ of N3 is around 7 times smaller than the true ESS. The reason is that the N3 proposal in this scenario samples from the target with variance reduction (there are the same number of samples for each component of the mixture, rather than a random number as in an i.i.d. simulation from the mixture). Note that for most schemes, $\widehat{\text{ESS}}$ overestimates the ESS, when N grows, but this is not the case for low values of N .

It is very illustrative that $\widehat{\text{ESS}}$ is blind to the difference between N3 and R3. In both schemes, due to the perfect match between the target and the mixture of proposals, all weights are always equal to 1, and hence the $\widehat{\text{ESS}}$ is maximum. However, N3 simulates the samples with variance reduction, which cannot be captured by $\widehat{\text{ESS}}$ since the samples play no role in its calculation.

Figure 8 shows the same curves for the case of mild mismatch between the target and the mixture of proposals (Scenario 2). In this scenario, all MIS schemes loose efficiency due to the mismatch, but $\widehat{\text{ESS}}$ keeps underestimating the ESS in N3. Finally, Figure 9 shows the same curves when the mismatch between the target and the mixture of

proposals is large (Scenario 3). In this scenario, the \widehat{ESS} underestimates the ESS in all MIS schemes except for $N3$.




FIGURE 7 MIS example (Scenario 1). True ESS (left), \widehat{ESS} (middle), and ratio of both (right) of the self-normalized IS estimator of the target mean, for the different MIS schemes as the total number of samples N grows. Note that ESS is normalized, i.e. divided by the number of samples N .

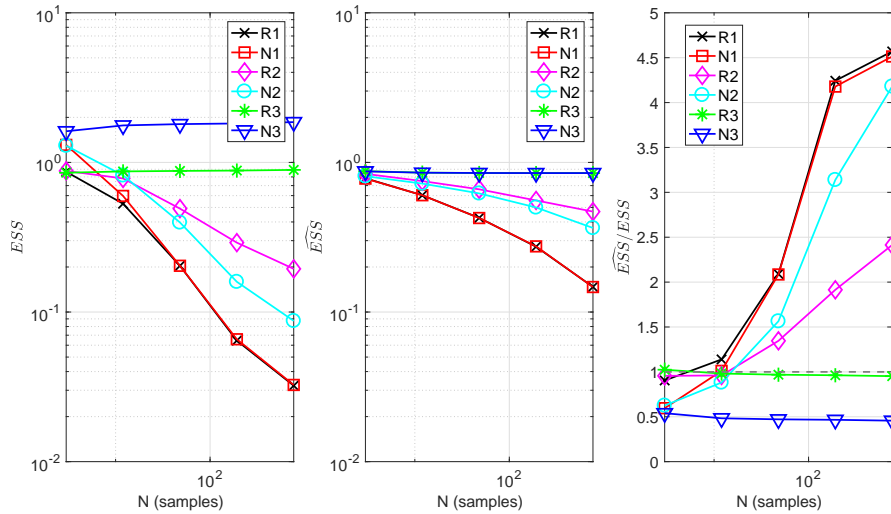


FIGURE 8 MIS example (Scenario 2). True ESS (left), \widehat{ESS} (middle), and ratio of both (right) of the self-normalized IS estimator of the target mean, for the different MIS schemes as the total number of samples N grows. Note that ESS is normalized, i.e. divided by the number of samples N .




FIGURE 9 MIS example (Scenario 3). True ESS (left), \widehat{ESS} (middle), and ratio of both (right) of the self-normalized IS estimator of the target mean, for the different MIS schemes as the total number of samples N grows. Note that ESS is normalized, i.e. divided by the number of samples N .

4.3.2 | An alternative \widehat{ESS} in MIS

An appropriate extension of Eq. (15) to MIS is not straightforward, since the samples are drawn from different proposal pdfs. Noting that in standard IS $\text{Var}_q[\widehat{Z}] = \frac{\text{Var}_q[w]}{N}$, we propose the following natural extension of the ESS to the MIS approach:

$$ESS_{\text{MIS}} \approx \frac{N}{1 + \text{Var}[\widehat{Z}]}, \quad (35)$$

where the variance of \widehat{Z} is now calculated taking into account the whole set of proposals. In this case, $\text{Var}[\widehat{Z}]$ needs to be still estimated. Note also that the approximation from the true ESS to Eq. (15) assumes a single proposal (and a normalized target).

5 | NEW DIRECTIONS FOR ALTERNATIVE METRICS

In Section 4.2, we have revisited some of the alternative metrics to the \widehat{ESS} that have been recently proposed. However, all those metrics inherit some of the undesired properties of the \widehat{ESS} , as we discussed in Section 3.3. We now explore two research lines that aim at dropping two of those properties, namely the lack of presence of both the function h and the set of samples $\{\mathbf{x}_n\}_{n=1}^N$.

5.1 | $h(\mathbf{x})$ -aware $\widehat{\text{ESS}}$

In most estimation problems, the function $h(\mathbf{x})$ plays a crucial role in the efficiency of the estimators, e.g., in the rare-event estimation problem of Section 3.4.2. Note that, for the optimal UIS proposal $q^*(\mathbf{x}) \propto |h(\mathbf{x})| \bar{\pi}(\mathbf{x})$, all evaluations $\frac{|h(\mathbf{x}_n)| \bar{\pi}(\mathbf{x}_n)}{q^*(\mathbf{x}_n)}$ are identical regardless of the value of \mathbf{x}_n . Hence, a natural extension of the $\widehat{\text{ESS}}$ is considering

$$\widehat{\text{ESS}}^{(h)} = \frac{1}{\sum_{n=1}^N \left(\bar{w}_n^{(h)} \right)^2} \quad (36)$$

where

$$\bar{w}_n^{(h)} = \frac{|h(\mathbf{x}_n)| w_n}{\sum_{j=1}^N |h(\mathbf{x}_j)| w_j}. \quad (37)$$

In this case, if the samples are simulated from the optimal UIS proposal $q^*(\mathbf{x})$, then $\widehat{\text{ESS}}^{(h)} = N$ always holds for non-negative functions h (i.e., it takes the maximum value). It is important to note that this modification goes beyond the definition of Eq. (5) (or the proposed variation with the MSE in the denominator in Eq. (8)). The original definition implicitly assumes that simulating from the target $\bar{\pi}(\mathbf{x})$ is the benchmark to compare with. However, it is well-known that this is not the case when attempting to reduce the variance of the IS estimator, which is clearly dependent of h .

Other alternative approximations could be considered in the derivation of Section 3.1. For instance, the natural variance estimate of SNIS estimator, proposed in [38, Sec. 9.1] and recently used in [37], could be plugged in the derivation. The difficulty of this approach is that the original derivation approximates $\text{Var}_q[\bar{I}]$ in such a way it depends on $\text{Var}_{\bar{\pi}}[\bar{I}]$, and therefore this term cancels out in the ratio of Eq. (15). This is particularly convenient, since $\text{Var}_{\bar{\pi}}[\bar{I}]$ is arguably the hardest term to estimate.

5.2 | Sample-aware $\widehat{\text{ESS}}$

The diversity of the samples is known to increase the performance of the estimators [38, Chapter 8]. The diversity in the set of samples is also key in IS-based methods that incorporate resampling steps (e.g., particle filtering [29] or adaptive IS [3]) to avoid degeneracy, even in the case where the $\widehat{\text{ESS}}$ is high. For these reasons, we would like the set of samples $\{\mathbf{x}_n\}_{n=1}^N$, to play a role in the approximated ESS. Intuitively, when the samples are farther apart this should yield larger ESS (for a fixed set of weights). An approximated ESS that would remain constant (in average) after a resampling step would be also desirable in order to avoid the well-known paradox that $\widehat{\text{ESS}} = N$ after a resampling step.

Here we follow the ideas of partitioning the space suggested in recent works in the literature [30, 33], which have shown to improve the quality of the resampling steps in sequential Monte Carlo schemes [32]. The main idea is to divide the domain \mathcal{X} of the random variable \mathbf{X} into M separate mutually exclusive regions. More specifically, let us consider an integer $M \in \mathbb{N}^+$, and a partition $\mathcal{P} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_M\}$ of the state space with M disjoint subsets,

$$\begin{aligned} \mathcal{X}_1 \cup \mathcal{X}_2 \cup \dots \cup \mathcal{X}_M &= \mathcal{X}, \\ \mathcal{X}_i \cap \mathcal{X}_k &= \emptyset, \quad i \neq k, \quad \forall i, j \in \{1, \dots, M\}. \end{aligned}$$

Let consider N weighted samples $\{\mathbf{x}_n, w_n\}_{n=1}^N$ such that $N \gg M$. We denote the subset of the set of indices

$\{1, \dots, N\}$,

$$\mathcal{J}_m = \{i = 1, \dots, N : \mathbf{x}_i \in \mathcal{X}_m\},$$

which are associated with the samples in the m -th sub-region \mathcal{X}_m . The cardinality $|\mathcal{J}_m|$ denotes the number of samples in \mathcal{X}_m , and we have $\sum_{m=1}^M |\mathcal{J}_m| = N$. We can associate a *summary weight* for each region \mathcal{X}_m , as

$$\bar{v}_m = \sum_{i \in \mathcal{J}_m} \bar{w}_i. \quad (38)$$

Note that $\sum_{m=1}^M \bar{v}_m = 1$, since $\bar{w}_n = \frac{w_n}{\sum_{i=1}^N w_i}$. We define a summary metric, alternatively to $\widehat{\text{ESS}}$, as

$$\widehat{\text{ESS}}^{(x)} = \frac{1}{\sum_{m=1}^M \bar{v}_m^2}, \quad 1 \leq \widehat{\text{ESS}}^{(x)} \leq M, \quad (39)$$

that measures the dissimilarity between the summary weight of each partition. We can also define new re-normalized weights (to sum up to one in each subset) as

$$\bar{w}_{n,m} = \frac{w_n}{\sum_{i \in \mathcal{J}_m} w_i} = \frac{w_n}{\bar{v}_m}, \quad (40)$$

and define a *local ESS* of each subset as

$$\widehat{\text{ESS}}_m = \frac{1}{\sum_{n \in \mathcal{J}_m} \bar{w}_{n,m}^2}, \quad 1 \leq \widehat{\text{ESS}}_m \leq |\mathcal{J}_m|. \quad (41)$$

It is easy to show that $M \leq \sum_{m=1}^M \widehat{\text{ESS}}_m \leq N$.

It is possible to find the standard $\widehat{\text{ESS}}$ as a function of the $M + 1$ metrics defined above (i.e., the M local metrics and the $\widehat{\text{ESS}}^{(x)}$) as:

$$\widehat{\text{ESS}} = \widehat{\text{ESS}}^{(x)} \left(\frac{1}{M} \sum_{m=1}^M \widehat{\text{ESS}}_m \right), \quad 1 \leq \widehat{\text{ESS}} \leq N. \quad (42)$$

We note that $\widehat{\text{ESS}}^{(x)}$ penalizes the concentration of the samples in the space. The $M + 1$ metrics give then more information about the diversity in the space of the set of samples, while the $\widehat{\text{ESS}}$ is again recovered by compressing this information as in Eq. (42).

This approach raises interesting questions for future research. For instance, it is possible to design the optimal choice of the partition. It is possible to study the behavior of $\widehat{\text{ESS}}$ if the partition is done empirically, e.g., with a clustering algorithm. Instead of using the same structure as in $\widehat{\text{ESS}}$, we could use other metrics in the literature (e.g., those revisited in Section 4.2) to build $\widehat{\text{ESS}}^{(x)}$ and the set of local metrics $\{\widehat{\text{ESS}}_m\}_{m=1}^M$.

Finally, there is a connection of these sample-aware metrics with the metric $\widehat{\text{ESS}}$ -IGH, proposed in [15, Eq. (23)] in the context novel quadrature methods derived under an importance sampling perspective. This metric, can be seen as a generalization of $\widehat{\text{ESS}}$ and penalizes situations when large weights correspond to samples that are farther apart from the mean of the proposal, i.e., it incorporates information from the sample space.

6 | CONCLUSION

Since the $\widehat{\text{ESS}}$ metric was derived in [27], it has been widely used in the literature of importance sampling. In this paper, we have revisited its derivation, highlighting the assumptions and approximations, and summarizing the undesired properties. Moreover, we have shown in representative numerical examples that the metric can produce misleading conclusions. The purpose of this paper is to raise awareness about the $\widehat{\text{ESS}}$, so it can be applied carefully.

We have also described alternative derivations and interesting connections with other metrics. We have revisited alternative metrics that have been recently proposed in the literature, and we have discussed potential improvements of the metric. Namely, we have discussed the generalization for the case of multiple proposals, the incorporation of the function $h(\mathbf{x})$ which is key in the efficiency of the SNIS estimator, and the inclusion of information that is present in the set of samples. Our preliminary exploration will allow for further research in those directions.

In the future, we would like to use the unnormalized IS weights to evaluate a sequence of particle approximations. Even if they do not have an absolute measure of performance (because Z is unknown), in adaptive IS schemes, one could exploit this information in iterative schemes. Finally, it would be desirable to have a more stable particle approximation. Note that the unnormalized IS weights often follow a right skewed distribution, which translates into having many zeros in the normalized weights. We could either use the *median* instead of the mean. We could also nonlinearly transform the unnormalized weights before normalizing them as in [24, 26, 46] (see a review in [36]).

7 | ACKNOWLEDGMENTS

The authors want to thank Matthew Ludkin and the referees for their useful comments.

A | DELTA METHOD FOR THE APPROXIMATION OF EXPECTATIONS

In this appendix, we provide a generic derivation of the delta method, which uses a second order Taylor expansion to approximate expectations of functions that involve dependent r.v.'s. In the main text, we particularize twice this generic formulation, making explicit the two delta methods involved in the derivation of $\widehat{\text{ESS}}$.

Let us define a bivariate function $f(T, Y)$. A second order expansion around a point (T_0, Y_0) is

$$f(T, Y) \approx f(T_0, Y_0) + f'_T(T_0, Y_0) (T - T_0) + f'_Y(T_0, Y_0) (Y - Y_0) + \frac{1}{2} \left(f''_{TT}(T_0, Y_0) (T - T_0)^2 + 2f''_{TY}(T_0, Y_0) (T - T_0) (Y - Y_0) + f''_{YY}(T_0, Y_0) (Y - Y_0)^2 \right). \quad (43)$$

Therefore, a good approximation to $E[f(T, Y)]$ with $f(T, Y)$ expanded around (T_0, Y_0) is

$$E[f(T, Y)] \approx f(T_0, Y_0) + \frac{1}{2} \left(f''_{TT}(T_0, Y_0) \text{Var}[T] + 2f''_{TY}(T_0, Y_0) \text{Cov}[T, Y] + f''_{YY}(T_0, Y_0) \text{Var}[Y] \right). \quad (44)$$

Now we derive an approximation of the variance of $f(T, Y)$, that by definition, is expressed as

$$\text{Var}[f(T, Y)] = E \left[(f(T, Y) - E[f(T, Y)])^2 \right]. \quad (45)$$

We take a rough approximation $E[f(T, Y)] \approx f(T_0, Y_0)$ from Eq. (43), as it is done for instance in [44]. Then, by taking

a first order expansion $f(T, Y) \approx f(T_0, Y_0) + f'_T(T_0, Y_0)(T - T_0) + f'_Y(T_0, Y_0)(Y - Y_0)$, Eq. (45) turns into

$$\begin{aligned}
 \text{Var}[f(T, Y)] &\approx \mathbb{E}\left[\left(f(T_0, Y_0) + f'_T(T_0, Y_0)(T - T_0) + f'_Y(T_0, Y_0)(Y - Y_0) - f(T_0, Y_0)\right)^2\right] \\
 &= \mathbb{E}\left[\left(f'_T(T_0, Y_0)(T - T_0) + f'_Y(T_0, Y_0)(Y - Y_0)\right)^2\right] \\
 &= \mathbb{E}\left[\left(f'_T(T_0, Y_0)(T - T_0)\right)^2 + \left(f'_Y(T_0, Y_0)(Y - Y_0)\right)^2\right. \\
 &\quad \left.+ 2f'_T(T_0, Y_0)(T - T_0)f'_Y(T_0, Y_0)(Y - Y_0)\right] \\
 &= f_T'^2(T_0, Y_0)\text{Var}[T] + f_Y'^2(T_0, Y_0)\text{Var}[Y] + 2f'_T(T_0, Y_0)f'_Y(T_0, Y_0)\text{Cov}[T, Y].
 \end{aligned} \tag{46}$$

References

- [1] S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso, and A.M. Stuart. Importance sampling: Intrinsic dimension and computational cost. *Statistical Science*, pages 405–431, 2017.
- [2] J. Bucklew. *Introduction to rare event simulation*. Springer Science & Business Media, 2004.
- [3] M. F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Míguez, and P. M. Djuric. Adaptive importance sampling: The past, the present, and the future. *IEEE Signal Processing Magazine*, 34(4):60–79, 2017.
- [4] O. Cappé, R. Douc, A. Guillin, J. M. Marin, and C. P. Robert. Adaptive importance sampling in general mixture classes. *Statistics and Computing*, 18:447–459, 2008.
- [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] Y. Chen. Another look at rejection sampling through importance sampling. *Statistics & probability letters*, 72(4):277–283, 2005.
- [7] Z. Chen. Bayesian filtering: From Kalman filters to particle filters, and beyond. *technical report*, pages 1–69, 2003.
- [8] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York (USA), 1991.
- [9] 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.
- [10] R. Douc, O. Cappé, and E. Moulines. Comparison of resampling schemes for particle filtering. In *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis*, pages 64–69, September 2005.
- [11] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001.
- [12] A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: fifteen years later. *technical report*, 2008.
- [13] V. Elvira and E. Chouzenoux. Langevin-based strategy for efficient proposal adaptation in population Monte Carlo. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5077–5081. IEEE, 2019.
- [14] V. Elvira and L. Martino. Advances in importance sampling. *Wiley StatsRef: Statistics Reference Online*, pages 1–22, 2021.
- [15] V. Elvira, L. Martino, and P. Closas. Importance gaussian quadrature. *IEEE Transactions on Signal Processing*, 69:474–488, 2020.

- [16] V. Elvira, L. Martino, D. Luengo, and M. F. Bugallo. Improving Population Monte Carlo: Alternative weighting and resampling schemes. *Signal Processing*, 131(12):77–91, 2017.
- [17] V. Elvira, L. Martino, D. Luengo, and M. F. Bugallo. Generalized multiple importance sampling. *Stat. Sci.*, 34(1):129–155, 2019.
- [18] V. Elvira and I. Santamaria. Multiple importance sampling for efficient symbol error rate estimation. *IEEE Signal Processing Letters*, 26(3):420–424, 2019.
- [19] V. Elvira and I. Santamaria. Multiple importance sampling for symbol error rate estimation of maximum-likelihood detectors in mimo channels. *IEEE Transactions on Signal Processing*, 69:1200–1212, 2021.
- [20] N. Gordon, D. Salmond, and A. F. M. Smith. Novel approach to nonlinear and non-Gaussian Bayesian state estimation. *IEE Proceedings-F Radar and Signal Processing*, 140:107–113, 1993.
- [21] T. C. Hesterberg. *Advances in importance sampling*. PhD thesis, Stanford University, 1988.
- [22] J. H Huggins and D. M Roy. Convergence of sequential Monte Carlo based sampling methods. *arXiv:1503.00966*, 2015.
- [23] J. H. Huggins and D. M. Roy. Sequential monte carlo as approximate sampling: bounds, adaptive resampling via ∞ -ess, and an application to particle gibbs. *Bernoulli (to appear in)*, *arXiv preprint arXiv:1503.00966*, 2018.
- [24] E. L. Ionides. Truncated importance sampling. *Journal of Computational and Graphical Statistics*, 17(2):295–311, 2008.
- [25] H. Kahn and A. Marshall. Methods of reducing sample size in monte carlo computations. *Journal of the Operations Research Society of America*, 1(5):263–278, 1953.
- [26] E. Koblents and J. Míguez. A population Monte Carlo scheme with transformed weights and its application to stochastic kinetic models. *Statistics and Computing*, pages 1–19, 2013.
- [27] A. Kong. A note on importance sampling using standardized weights. *University of Chicago, Dept. of Statistics, Tech. Rep.*, 348, 1992.
- [28] A. Kong, J. S. Liu, and W. H. Wong. Sequential imputations and Bayesian missing data problems. *Journal of the American Statistical Association*, 9:278–288, 1994.
- [29] T. Li, M. Bolic, and P. M. Djuric. Resampling methods for particle filtering: Classification, implementation, and strategies. *IEEE Signal Processing Magazine*, 32(3):70–86, 2015.
- [30] T. Li, T. P. Sattar, and S. Sun. Deterministic resampling: Unbiased sampling to avoid sample impoverishment in particle filters. *Signal Processing*, 92(7):1637–1645, 2012.
- [31] J. S. Liu and R. Chen. Blind deconvolution via sequential imputations. *Journal of the American Statistical Association*, 90(430):567–576, June 1995.
- [32] L. Martino and V. Elvira. Compressed Monte Carlo with application in particle filtering. *Information Sciences*, 553:331–352, 2021.
- [33] L. Martino, V. Elvira, and G. Camps-Valls. Group importance sampling for particle filtering and MCMC. *Digital Signal Processing*, 82:133–151, 2018.
- [34] L. Martino, V. Elvira, and F. Louzada. Effective sample size for importance sampling based on discrepancy measures. *Signal Processing*, 131:386 – 401, 2017.
- [35] L. Martino, V. Elvira, D. Luengo, and J. Corander. Layered adaptive importance sampling. *Statistics and Computing*, 27(3):599–623, 2017.

- [36] L. Martino, V. Elvira, J. Míguez, A. Artés-Rodríguez, and P.M. Djurić. A comparison of clipping strategies for importance sampling. In *2018 IEEE Statistical Signal Processing Workshop (SSP)*, pages 558–562. IEEE, 2018.
- [37] H. Nilakanta. *Output Analysis of Monte Carlo Methods with Applications to Networks and Functional Approximation*. PhD thesis, University of Minnesota, 2020.
- [38] A. Owen. *Monte Carlo theory, methods and examples*. <http://statweb.stanford.edu/~owen/mc/>, 2013.
- [39] A. Owen, Y. Maximov, and M. Chertkov. Importance sampling the union of rare events with an application to power systems analysis. *Electronic Journal of Statistics*, 13(1):231–254, 2019.
- [40] J. Rice. *Mathematical statistics and data analysis*. Nelson Education, 2006.
- [41] C. Robert. Effective Sample Size. *Christian P. Robert's blog*: <https://xianblog.wordpress.com/2010/09/24/effective-sample-size/>, 2010.
- [42] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2004.
- [43] C. P. Robert and G. Casella. *Introducing Monte Carlo Methods with R*. Springer, 2010.
- [44] H. Seltman. Approximations for mean and variance of a ratio. *unpublished note*, 2012.
- [45] A Sokal. Monte carlo methods in statistical mechanics: foundations and new algorithms. In *Functional integration*, pages 131–192. Springer, 1997.
- [46] A. Vehtari and A. Gelman. Pareto smoothed importance sampling. *arXiv preprint arXiv:1507.02646*, 2015.
- [47] N. Whiteley, A. Lee, and K. Heine. On the role of interaction in sequential Monte Carlo algorithms. *Bernoulli*, 22(1):494–529, 2016.