

Original citation:

Agapiou, Sergios, Papaspiliopoulos, Omiros, Sanz-Alonso, Daniel and Stuart, A. M. (2017) Importance sampling : intrinsic dimension and computational cost. *Statistical Science*, 32 (3). pp. 405-431.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/96050>

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:

<http://dx.doi.org/10.1214/17-STS611>

A note on versions:

The version presented in WRAP is the published version or, version of record, and may be cited as it appears here.

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

Importance Sampling: Intrinsic Dimension and Computational Cost

S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso and A. M. Stuart

Abstract. The basic idea of importance sampling is to use independent samples from a proposal measure in order to approximate expectations with respect to a target measure. It is key to understand how many samples are required in order to guarantee accurate approximations. Intuitively, some notion of distance between the target and the proposal should determine the computational cost of the method. A major challenge is to quantify this distance in terms of parameters or statistics that are pertinent for the practitioner. The subject has attracted substantial interest from within a variety of communities. The objective of this paper is to overview and unify the resulting literature by creating an overarching framework. A general theory is presented, with a focus on the use of importance sampling in Bayesian inverse problems and filtering.

Key words and phrases: Importance sampling, notions of dimension, small noise, absolute continuity, inverse problems, filtering.

1. INTRODUCTION

1.1 Our Purpose

Our purpose in this paper is to overview various ways of measuring the computational cost of importance sampling, to link them to one another through transparent mathematical reasoning, and to create cohesion in the vast published literature on this subject. In addressing these issues, we will study importance sampling in a general abstract setting, and then in the particular cases of Bayesian inversion and filtering. These

two application settings are particularly important as there are many pressing scientific, technological and societal problems which can be formulated via inversion or filtering. An example of such an inverse problem is the determination of subsurface properties of the Earth from surface measurements; an example of a filtering problem is the assimilation of atmospheric measurements into numerical weather forecasts. We now proceed to overview the subject of importance sampling, and the perspective on it that is our focus. In Section 1.2, we describe the organization of the paper and our main contributions. Section 1.3 then collects all the references linked to the material in the introduction, as well as other general references on importance sampling. Each subsequent section of the paper contains its own literature review subsection providing further elaboration of the literature, and linking it to the details of the material that we present in that section.

The general abstract setting in which we work is as follows. We let μ and π be two probability measures on a measurable space $(\mathcal{X}, \mathcal{F})$ related via the expression

$$(1.1) \quad \frac{d\mu}{d\pi}(u) := g(u) / \int_{\mathcal{X}} g(u)\pi(du).$$

Here, g is the unnormalised *density* (or *Radon–Nikodym derivative*) of μ with respect to π . Note that the

S. Agapiou is Lecturer, Department of Mathematics and Statistics, University of Cyprus, 1 University Avenue, 2109 Nicosia, Cyprus (e-mail: agapiou.sergios@ucy.ac.cy).

O. Papaspiliopoulos is ICREA Research Professor, ICREA & Department of Economics, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain (e-mail: omiros.papaspiliopoulos@upf.edu).

D. Sanz-Alonso is Postdoctoral Research Associate, Division of Applied Mathematics, Brown University, Providence, Rhode Island 02906, USA (e-mail: daniel_sanz-alonso1@brown.edu).

A. M. Stuart is Bren Professor of Computing and Mathematical Sciences, Computing and Mathematical Sciences, California Institute of Technology, Pasadena, California 91125, USA (e-mail: astuart@caltech.edu).

very existence of the density implies that the target is *absolutely continuous* with respect to the proposal; absolute continuity will play an important role in our subsequent developments.

Importance sampling is a method for using independent samples from the *proposal* π to approximately compute expectations with respect to the *target* μ . The way importance sampling (and more generally Monte Carlo integration methods) is used within Bayesian statistics and Bayesian inverse problems is as an approximation of the target measure μ by a random probability measure using weighted samples that are generated from π . (This perspective differs from that arising in other disciplines, e.g., in certain applications in mathematical finance, such as option pricing.) Our perspective is dictated by the need to use the samples to estimate expectations and quantiles of a wide range of functions defined on the state space, for example, functions of a single variable or pairs of variables, or marginal likelihood quantities. The resulting approximation is typically called a particle approximation. Our perspective on importance sampling as a probability measure approximation dictates in turn the tools for studying its performance. Its computational cost is measured by the number of samples required to control the worst error made when approximating expectations within a class of test functions. In this article, and following existing foundational work, we primarily focus on a total variation metric between random measures for assessing the particle approximation error. Intuitively, the size of the error is related to how far the target measure is from the proposal measure. We make this intuition precise, and connect the particle approximation error to a key quantity, the second moment of $d\mu/d\pi$ under the proposal, which we denote by ρ :

$$\rho = \pi(g^2)/\pi(g)^2.$$

As detailed below, ρ is essentially the χ^2 divergence between the target and the proposal.

The first application of this setting that we study is the linear inverse problem to determine $u \in \mathcal{X}$ from y where

$$(1.2) \quad y = Ku + \eta, \quad \eta \sim N(0, \Gamma).$$

We adopt a Bayesian approach in which we place a prior $u \sim \mathbb{P}_u = N(0, \Sigma)$, assume that η is independent of u and seek the posterior $u|y \sim \mathbb{P}_{u|y}$. We study importance sampling with $\mathbb{P}_{u|y}$ being the target μ and \mathbb{P}_u being the proposal π .

The second application is the linear filtering problem of sequentially updating the distribution of $v_j \in \mathcal{X}$ given $\{y_i\}_{i=1}^j$ where

$$(1.3) \quad \begin{aligned} v_{j+1} &= Mv_j + \xi_j, \\ \xi_j &\sim N(0, Q), \quad j \geq 0, \\ y_{j+1} &= Hv_{j+1} + \zeta_{j+1}, \\ \zeta_{j+1} &\sim N(0, R), \quad j \geq 0. \end{aligned}$$

We assume that the problem has a Markov structure. We study the approximation of one step of the filtering update by means of particles, building on the study of importance sampling for the linear inverse problem. To this end, it is expedient to work on the product space $\mathcal{X} \times \mathcal{X}$, and consider importance sampling for $(v_j, v_{j+1}) \in \mathcal{X} \times \mathcal{X}$. It then transpires that, for two different proposals, which are commonly termed the *standard proposal* and the *optimal proposal*, the cost of one step of particle filtering may be understood by the study of a linear inverse problem on \mathcal{X} ; we show this for both proposals, and then use the link to an inverse problem to derive results about the cost of particle filters based on these two proposals.

The linear Gaussian models that we study can—and typically should—be treated by direct analytic calculations or efficient simulation of Gaussians. However, it is possible to analytically study the dependence of ρ on key parameters within these model classes, and furthermore they are flexible enough to incorporate formulations on function spaces, and their finite dimensional approximations. Thus, they are an excellent framework for obtaining insight into the performance of importance sampling for inverse problems and filtering.

For the abstract importance sampling problem, we will relate ρ to a number of other natural quantities. These include the *effective sample size* ess , used heuristically in many application domains, and a variety of *distance metrics* between π and μ . Since the existence of a density between target and proposal plays an important role in this discussion, we will also investigate what happens as this absolute continuity property breaks down. We study this first in *high dimensional problems*, and second in *singular parameter limits* (by which we mean limits of important parameters defining the problem). The ideas behind these two different ways of breaking absolute continuity are presented in the general framework, and then substantially developed in the inverse problem and filtering settings. The motivation for studying these limits can be appreciated by considering the two examples mentioned at

the start of this introduction: inverse problems from the Earth's subsurface, and filtering for numerical weather prediction. In both cases, the unknown which we are trying to determine from data is best thought of as a spatially varying field for subsurface properties such as permeability, or atmospheric properties, such as temperature. In practice, the field will be discretized and represented as a high dimensional vector, for computational purposes, but for these types of application the state dimension can be of order 10^9 . Furthermore, as computer power advances there is pressure to resolve more physics, and hence for the state dimension to increase. Thus, it is important to understand infinite dimensional problems, and sequences of approximating finite dimensional problems which approach the infinite dimensional limit. A motivation for studying singular parameter limits arises, for example, from problems in which the noise is small and the relevant log-likelihoods scale inversely with the noise variance.

This paper aims in particular to contribute towards a better understanding of the recurrent claim that importance sampling suffers from the curse of dimensionality. Whilst there is some empirical truth in this statement, there is a great deal of confusion in the literature about what exactly makes importance sampling hard. In fact, such a statement about the role of dimension is vacuous unless “dimension” is defined precisely. We will substantially clarify these issues in the contexts of inverse problems and filtering. Throughout this paper, we use the following conventions:

- State space dimension is the dimension of the measurable space where the measures μ and π are defined. We will be mostly interested in the case where the measurable space \mathcal{X} is a separable Hilbert space, in which case the state space dimension is the cardinality of an orthonormal basis of the space. In the context of inverse problems and filtering, the state space dimension is the dimension of the unknown.
- Data space dimension is the dimension of the space where the data lives.
- Nominal dimension is the minimum of the state space dimension and the data state dimension.
- Intrinsic dimension: we will use two notions of intrinsic dimension for linear Gaussian inverse problems, denoted by efd and τ . These combine state/data dimension and small noise parameters. They can be interpreted as a measure of how informative the data is relative to the prior.

We show that the intrinsic dimensions are natural when studying the computational cost of importance

sampling for inverse problems. In particular, we show how these intrinsic dimensions relate to the parameter ρ introduced above, a parameter that we show to be central to the computational cost, and to the breakdown of absolute continuity. Finally, we apply our understanding of linear inverse problems to particle filters, translating the results from one to the other via an interesting correspondence between the two problems, for both standard and optimal proposals, that we describe here. In studying these quantities, and their inter-relations, we aim to achieve the purpose set out at the start of this introduction.

1.2 Organization of the Paper and Main Contributions

Section 2 describes importance sampling in abstract form. In Sections 3 and 4, the linear Gaussian inverse problem and the linear Gaussian filtering problem are studied. Our aim is to provide a digestible narrative, and hence all proofs—and all technical matters related to studying measures in infinite dimensional spaces—are left to the Supplementary Material [4].

Further to providing a unified narrative of the existing literature, this paper contains some original contributions that shed new light on the use of importance sampling for inverse problems and filtering. Our main new results are:

- Theorem 2.1 bounds the error of importance sampling for bounded test functions. The main appeal of this theorem is its nonasymptotic nature, together with its clean interpretation in terms of: (i) the key quantity ρ ; (ii) effective sample size; (iii) metrics between probability measures; (iv) existing asymptotic results. According to the perspective on importance sampling as an approximation of one probability measure by another, the metric used in Theorem 2.1 is natural and it has already been used in important theoretical developments in the field as we discuss in Section 2.5. On the other hand, the result is less useful for quantifying the error for a specific test function of interest, such as linear, bilinear or quadratic functions, typically used for computing moments and covariances. We discuss extensions and generalizations in Section 2.
- Theorem 3.8 studies importance sampling for inverse problems. It is formulated in the linear Gaussian setting to allow a clear and full development of the connections that it makes between heretofore disparate notions. In particular, we highlight the following. (i) It provides the first clear

connection between finite intrinsic dimension and absolute continuity between posterior and prior. (ii) It demonstrates the relevance of the intrinsic dimension—rather than the state space or the nominal dimension—in the performance of importance sampling, by linking the intrinsic dimension and the parameter ρ ; thus, it shows the combined effect of the prior, the forward map and the noise model in the efficacy of the method. (iii) It provides theoretical support for the use of algorithms based on importance sampling for posterior inference in function space, provided that the intrinsic dimension is finite and the value of ρ is moderate.

- Theorems 4.2 and 4.3 are proved by studying the inverse problem at the heart of importance sampling based particle filters. These theorems, together with Theorem 4.5 and Example 4.6, provide an improved understanding of the advantages of the optimal proposal over the standard proposal in the context of filtering.

1.3 Literature Review

In this subsection, we provide a historical review of the literature in importance sampling. Each of the following Sections 2, 3 and 4 will contain a further literature review subsection providing detailed references linked explicitly to the theory as outlined in those sections.

Early developments of importance sampling as a method to reduce the variance in Monte Carlo estimation date back to the early 1950s [47, 48]. In particular, the paper [48] demonstrates how to optimally choose the proposal density for given test function ϕ and target density. Standard text book references for importance sampling include [33] and [71]. Important methodological improvements were introduced in [66, 72, 82] and [96]. A modern view of importance sampling in the general framework (1.1) is given in [22]. A comprehensive description of Bayesian inverse problems in finite state/data space dimensions can be found in [49], and its formulation in infinite dimensional spaces in [30, 60–62, 95]. Text books overviewing the subject of filtering and particle filters include [6, 31], and the article [27] provides a readable introduction to the area. For an up-to-date and in-depth survey of nonlinear filtering, see [28]. The linear Gaussian inverse problem and the linear Gaussian filtering problem have been extensively studied because they arise naturally in many applications, lead to considerable algorithmic tractability, and provide theoretical insight. For references concerning linear Gaussian inverse problems, see [39, 53,

64, 74]. The linear Gaussian filter—the Kalman filter—was introduced in [51]; see [59] for further analysis. The inverse problem of determining subsurface properties of the Earth from surface measurements is discussed in [81], while the filtering problem of assimilating atmospheric measurements for numerical weather prediction is discussed in [52].

The key role of ρ , the second moment of the Radon–Nikodym derivative between the target and the proposal, has long been acknowledged [70, 83]. The crucial question of how to choose a proposal measure that leads to small value of ρ has been widely studied, and we refer to [67] and references therein. In this vein, our theory in Sections 3 and 4 shows precise conditions that guarantee $\rho < \infty$ in inverse problems and filtering settings, in terms of well-defined basic concepts such as absolute continuity of the target with respect to the proposal. Our study of importance sampling for inverse problems in Section 3 is limited to the choice of prior as proposal, which is of central theoretical relevance. In practice, however, more sophisticated proposals are often used, potentially leading to reduced parameter ρ ; two novel ideas include the implicit sampling method described in [78], and the use of proposals based on the ensemble Kalman filter suggested in [65]. The value of ρ is known to be asymptotically linked to the effective sample size [56, 57, 70]. Recent justification for the use of the effective sample size within particle filters is given in [101]. We provide a further nonasymptotic justification of the relevance of ρ through its appearance in error bounds on the error in importance sampling; a relevant related paper is [26] which proved nonasymptotic bounds on the error in the importance-sampling based particle filter algorithm. In this paper, we will also bound the importance sampling error in terms of different notions of distance between the target and the proposal measures. Our theory is based on the χ^2 divergence—as in [20]—while the recent complementary analysis of importance sampling in [19] highlights the advantages of the Kullback–Leibler divergence; a useful overview of the subject of distances between probability measures is [43].

We formulate problems in both finite dimensional and infinite dimensional state spaces. We refer to [50] for a modern presentation of probability appropriate for understanding the material in this article. Some of our results are built on the rich area of Gaussian measures in Hilbert space; we include all the required background in the Supplementary Material, and references are included there. However, we emphasize that the

presentation in the main body of the text is designed to keep technical material to a minimum and to be accessible to readers who are not versed in the theory of probability in infinite dimensional spaces. Absolute continuity of the target with respect to the proposal—or the existence of a density of the target with respect to the proposal—is central to our developments. This concept also plays a pivotal role in the understanding of Markov chain Monte Carlo (MCMC) methods in high and infinite dimensional spaces [97]. A key idea in MCMC is that breakdown of absolute continuity on sequences of problems of increasing state space dimension is responsible for poor algorithmic performance with respect to increasing dimension; this should be avoided if possible, such as for problems with a well-defined infinite dimensional limit [25]. Similar ideas will come into play in this paper.

As well as the breakdown of absolute continuity through increase in dimension, small noise limits can also lead to sequences of proposal/target measures which are increasingly close to mutually singular and for which absolute continuity breaks down. Small noise regimes are of theoretical and computational interest for both inverse problems and filtering. For instance, in inverse problems there is a growing interest in the study of the concentration rate of the posterior in the small observational noise limit; see [2, 5, 53–55, 84, 100]. In filtering and multiscale diffusions, the analysis and development of improved proposals in small noise limits is an active research area [37, 78, 94, 99, 104].

In order to quantify the computational cost of a problem, a recurrent concept is that of intrinsic dimension. Several notions of intrinsic dimension have been used in different fields, including dimension of learning problems [12, 102, 103], of statistical inverse problems [73], of functions in the context of quasi Monte Carlo (QMC) integration in finance applications [16, 58, 79] and of data assimilation problems [23]. The underlying theme is that in many application areas where models are formulated in high dimensional state spaces, there is often a small subspace which captures most of the features of the system. It is the dimension of this subspace that effects the cost of the problem. The recent subject of *active subspaces* shows promise in finding such low dimensional subspace of interest in certain applications [24]. In the context of inverse problems, the paper [8] proposed a notion of intrinsic dimension that was shown to have a direct connection with the performance of importance sampling. We introduce a further notion of intrinsic dimension

for Bayesian inverse problems which agrees with the notion of effective number of parameters used in machine learning and statistics [12]. We also establish that this notion of dimension and the one in [8] are finite, or otherwise, at the same time. Both intrinsic dimensions account for three key features of the cost of the inverse problem: the nominal dimension (i.e., the minimum of the dimension of the state space and the data), the size of the observational noise and the regularity of the prior relative to the observation noise. Varying the parameters related to these three features may cause a breakdown of absolute continuity. The deterioration of importance sampling in large nominal dimensional limits has been widely investigated [8, 11, 88–91]. In particular, the key role of the intrinsic dimension, rather than the nominal one, in explaining this deterioration was studied in [8]. Here, we study the different behaviour of importance sampling as absolute continuity is broken in the three regimes above, and we investigate whether, in all these regimes, the deterioration of importance sampling may be quantified by the various intrinsic dimensions that we introduce.

We emphasize that, whilst the theory and discussion in Section 2 is quite general, the applications to Bayesian inverse problems (Section 3) and filtering (Section 4) are in the case of linear problems with additive Gaussian noise. This linear Gaussian setting allows substantial explicit calculations and yields considerable insight. However, empirical evidence related to the behaviour of filters and Monte Carlo based methods when applied to nonlinear problems and non-Gaussian target measures suggests that similar ideas may apply in those situations; see [14, 24, 25, 29, 90]. Quantifying this empirical experience more precisely is an interesting and challenging direction for future study. We note in particular that extensions of the intrinsic dimension quantity that we employ have been provided in the literature for Bayesian hierarchical non-Gaussian models, more specifically within the so-called deviance information criterion of [93]; see Section 3.5.3 for more discussion.

1.4 Notation

Given a probability measure ν on a measurable space $(\mathcal{X}, \mathcal{F})$ expectations of a measurable function $\phi : \mathcal{X} \rightarrow \mathbb{R}$ with respect to ν will be written as both $\nu(\phi)$ and $\mathbb{E}_\nu[\phi]$. When it is clear which measure is being used we may drop the suffix ν and write simply $\mathbb{E}[\phi]$. Similarly, the variance will be written as $\text{Var}_\nu(\phi)$ and again we may drop the suffix when no confusion arises from

doing so. All test functions ϕ appearing in the paper are assumed to be measurable.

We will be interested in sequences of measures indexed by time or by the state space dimension. These are denoted with a subscript, for example, ν_t, ν_d . Anything to do with samples from a measure is denoted with a superscript: N for the number of samples, and n for the indices of the samples. The i th coordinate of a vector u is denoted by $u(i)$. Thus, $u_t^n(i)$ denotes i th coordinate of the n th sample from the measure of interest at time t . Finally, the law of a random variable v will be denoted by \mathbb{P}_v .

2. IMPORTANCE SAMPLING

In Section 2.1, we define importance sampling and in Section 2.2 we demonstrate the role of the second moment of the target-proposal density, ρ ; we prove two nonasymptotic theorems showing $\mathcal{O}((\rho/N)^{\frac{1}{2}})$ convergence rate of importance sampling with respect to the number N of particles. Then in Section 2.3.2 we show how ρ relates to the effective sample size ess as often defined by practitioners, whilst in Section 2.3.3 we link ρ to various distances between probability measures. In Section 2.4.1, we highlight the role of the breakdown of absolute continuity in the growth of ρ , as the dimension of the space \mathcal{X} grows. Section 2.4.2 follows with a similar discussion relating to singular limits of the density between target and proposal. Section 2.5 contains a literature review and, in particular, sources for all the material in this section.

2.1 General Setting

We consider target μ and proposal π , both probability measures on the measurable space $(\mathcal{X}, \mathcal{F})$, related by (1.1). In many statistical applications, interest lies in estimating expectations under μ , for a collection of test functions, using samples from π . For a test function $\phi : \mathcal{X} \rightarrow \mathbb{R}$ such that $\mu(|\phi|) < \infty$, the identity

$$\mu(\phi) = \frac{\pi(\phi g)}{\pi(g)},$$

leads to the *autonormalized importance sampling* estimator:

$$\begin{aligned} (2.1) \quad \mu^N(\phi) &:= \frac{\frac{1}{N} \sum_{n=1}^N \phi(u^n)g(u^n)}{\frac{1}{N} \sum_{m=1}^N g(u^m)}, \quad u^n \sim \pi \text{ i.i.d.} \\ &= \sum_{n=1}^N w^n \phi(u^n), \quad w^n := \frac{g(u^n)}{\sum_{m=1}^N g(u^m)}; \end{aligned}$$

here the w^n 's are called the *normalized weights*. As suggested by the notation, it is useful to view (2.1) as integrating a function ϕ with respect to the random probability measure $\mu^N := \sum_{n=1}^N w^n \delta_{u^n}$. Under this perspective, importance sampling consists of approximating the target μ by the measure μ^N , which is typically called the *particle approximation* of μ . Note that, while μ^N depends on the proposal π , we suppress this dependence for economy of notation. Our aim is to understand the quality of the approximation μ^N of μ . In particular, we would like to know how large to choose N in order to obtain small error. This will quantify the computational cost of importance sampling.

2.2 A Nonasymptotic Bound on Particle Approximation Error

A fundamental quantity in addressing this issue is ρ , defined by

$$(2.2) \quad \rho := \frac{\pi(g^2)}{\pi(g)^2}.$$

Thus, ρ is the second moment of the Radon–Nikodym derivative of the target with respect to the proposal. The Cauchy–Schwarz inequality shows that $\pi(g^2) \leq \pi(g)^2$ and hence that $\rho \geq 1$. Our first nonasymptotic result shows that, for bounded test functions ϕ , both the bias and the mean square error (MSE) of the autonormalized importance sampling estimator are $\mathcal{O}(N^{-1})$ with constant of proportionality linear in ρ .

THEOREM 2.1. *Assume that μ is absolutely continuous with respect to π , with square-integrable density g , that is, $\pi(g^2) < \infty$. The bias and MSE of importance sampling over bounded test functions may be characterized as follows:*

$$\sup_{|\phi| \leq 1} |\mathbb{E}[\mu^N(\phi) - \mu(\phi)]| \leq \frac{12}{N} \rho,$$

and

$$\sup_{|\phi| \leq 1} \mathbb{E}[(\mu^N(\phi) - \mu(\phi))^2] \leq \frac{4}{N} \rho.$$

REMARK 2.2. For a bounded test function $|\phi| \leq 1$, we trivially get $|\mu^N(\phi) - \mu(\phi)| \leq 2$; hence the bounds on bias and MSE provided in Theorem 2.1 are useful only when they are smaller than 2 and 4, respectively.

The upper bounds stated in this result suggest that it is good practice to keep ρ/N small in order to obtain good importance sampling approximations. This heuristic dominates the developments in the remainder

of the paper, and in particular our wish to study the behaviour of ρ in various limits. The result trivially extends to provide bounds on the mean square error for functions bounded by any other known bound different from 1. For practical purposes, the theorem is directly applicable to instances where importance sampling is used to estimate probabilities, such as in rare event simulation. However, its primary role is in providing a bound on the particle approximation error, which is naturally defined over bounded functions, as is common with weak convergence results. It is also important to realise that such a result will not hold without more assumptions on the weights for unbounded test functions; for example when g has third moment but not fourth under π , then $\mu(g^2) < \infty$, $\pi(g^2) < \infty$ but the importance sampling estimator of $\mu(g^2)$ has infinite variance. We return to extensions of the theorem for unbounded test functions in Section 2.3 below.

2.3 Connections, Interpretations and Extensions

Theorem 2.1 clearly demonstrates the role of ρ , the second moment of the target density with respect to the proposal, in determining the number of samples required to effectively approximate expectations. Here, we link ρ to other quantities used in analysis and monitoring of importance sampling algorithms, and we discuss some limitations of thinking entirely in terms of ρ .

2.3.1 *Asymptotic consistency.* It is interesting to contrast Theorem 2.1 to a well-known elementary asymptotic result. First, note that

$$\mu^N(\phi) - \mu(\phi) = \frac{N^{-1} \sum_{n=1}^N \frac{g(u^n)}{\pi(g)} [\phi(u^n) - \mu(\phi)]}{N^{-1} \sum_{n=1}^N \frac{g(u^n)}{\pi(g)}}.$$

Therefore, under the condition $\pi(g^2) < \infty$, and provided additionally that $\pi(g^2\phi^2) < \infty$, an application of the Slutsky lemmas gives that

$$(2.3) \quad \sqrt{N}(\mu^N(\phi) - \mu(\phi)) \implies N\left(0, \frac{\pi(g^2\bar{\phi}^2)}{\pi(g)^2}\right),$$

$$\text{where } \bar{\phi} := \phi - \mu(\phi).$$

For bounded $|\phi| \leq 1$, the only condition needed for appealing to the asymptotic result is $\pi(g^2) < \infty$. Then (2.3) gives that, for large N and since $|\bar{\phi}| \leq 2$,

$$\mathbb{E}[(\mu^N(\phi) - \mu(\phi))^2] \lesssim \frac{4}{N}\rho,$$

which is in precise agreement with Theorem 2.1.

2.3.2 *Effective sample size.* Many practitioners define the *effective sample size* by the formula

$$\begin{aligned} \text{ess} &:= \left(\sum_{n=1}^N (w^n)^2\right)^{-1} = \frac{(\sum_{n=1}^N g(u^n))^2}{\sum_{n=1}^N g(u^n)^2} \\ &= N \frac{\pi_{\text{MC}}^N(g)^2}{\pi_{\text{MC}}^N(g^2)}, \end{aligned}$$

where π_{MC}^N is the empirical Monte Carlo random measure

$$\pi_{\text{MC}}^N := \frac{1}{N} \sum_{n=1}^N \delta_{u^n}, \quad u^n \sim \pi.$$

By the Cauchy–Schwarz inequality, it follows that $\text{ess} \leq N$. Furthermore, since the weights lie in $[0, 1]$, we have

$$\sum_{n=1}^N (w^n)^2 \leq \sum_{n=1}^N w^n = 1$$

so that $\text{ess} \geq 1$. These upper and lower bounds may be attained as follows. If all the weights are equal, and hence take value N^{-1} , then $\text{ess} = N$, the optimal situation. On the other hand, if exactly k weights take the same value, with the remainder then zero, $\text{ess} = k$; in particular the lower bound of 1 is attained if precisely one weight takes the value 1 and all others are zero.

For large enough N , and provided $\pi(g^2) < \infty$, the strong law of large numbers gives

$$\text{ess} \approx N/\rho.$$

Recalling that $\rho \geq 1$, we see that ρ^{-1} quantifies the proportion of particles that effectively characterize the sample size, in the large particle size asymptotic. Furthermore, by Theorem 2.1, we have that, for large N ,

$$\sup_{|\phi| \leq 1} \mathbb{E}[(\mu^N(\phi) - \mu(\phi))^2] \lesssim \frac{4}{\text{ess}}.$$

This provides a further justification for the use of ess as an effective sample size, in the large N asymptotic regime.

2.3.3 *Probability metrics.* Intuition tells us that importance sampling will perform well when the distance between proposal π and target μ is not too large. Furthermore, we have shown the role of ρ in measuring the rate of convergence of importance sampling. It is hence of interest to explicitly link ρ to distance metrics between π and μ . In fact, we consider asymmetric divergences as distance measures; these are not strictly

metrics, but certainly represent useful distance measures in many contexts in probability. First, consider the χ^2 divergence, which satisfies

$$(2.4) \quad D_{\chi^2}(\mu \parallel \pi) := \pi \left(\left[\frac{g}{\pi(g)} - 1 \right]^2 \right) = \rho - 1.$$

The Kullback–Leibler divergence is given by

$$D_{\text{KL}}(\mu \parallel \pi) := \pi \left(\frac{g}{\pi(g)} \log \frac{g}{\pi(g)} \right),$$

and may be shown to satisfy

$$(2.5) \quad \rho \geq e^{D_{\text{KL}}(\mu \parallel \pi)}.$$

Thus, Theorem 2.1 suggests that the number of particles required for accurate importance sampling scales exponentially with the Kullback–Leibler divergence between proposal and target and linearly with the χ^2 divergence.

2.3.4 Beyond bounded test functions. In contrast to Theorem 2.1, the asymptotic result (2.3), establishes the convergence rate $N^{-1/2}$ (asymptotically) under the weaker moment assumption on the test function $\pi(g^2\phi^2) < \infty$. It is thus of interest to derive nonasymptotic bounds on the MSE and bias for much larger classes of test functions. This can be achieved at the expense of more assumptions on the importance weights. The next theorem addresses the issue of relaxing the class of test functions, whilst still deriving nonasymptotic bounds. By including the result, we also highlight the fact that, whilst ρ plays an important role in quantifying the difficulty of importance sampling, other quantities may be relevant in the analysis of importance sampling for unbounded test functions. Nonetheless, the sufficiency and necessity of scaling the number of samples with ρ is understood in certain settings, as will be discussed in the bibliography at the end of this section.

To simplify the statement, we first introduce the following notation. We write $m_t[h]$ for the t th central moment with respect to π of a function $h : \mathcal{X} \rightarrow \mathbb{R}$. That is,

$$m_t[h] := \pi(|h(u) - \pi(h)|^t).$$

We also define, as above, $\bar{\phi} := \phi - \mu(\phi)$.

THEOREM 2.3. *Suppose that ϕ and g are such that C_{MSE} defined below is finite:*

$$C_{\text{MSE}} := \frac{3}{\pi(g)^2} m_2[\phi g] + \frac{3}{\pi(g)^4} \pi(|\phi g|^{2d})^{\frac{1}{d}} C_{2e}^{\frac{1}{e}} m_{2e}[g]^{\frac{1}{e}}$$

$$+ \frac{3}{\pi(g)^{2(1+\frac{1}{p})}} \pi(|\phi|^{2p})^{\frac{1}{p}} \cdot C_{2q(1+\frac{1}{p})}^{\frac{1}{q}} m_{2q(1+\frac{1}{p})}[g]^{\frac{1}{q}}.$$

Then the bias and MSE of importance sampling when applied to approximate $\mu(\phi)$ may be characterized as follows:

$$|\mathbb{E}[\mu^N(\phi) - \mu(\phi)]| \leq \frac{1}{N} \left(\frac{2}{\pi(g)^2} m_2[g]^{\frac{1}{2}} m_2[\bar{\phi}g]^{\frac{1}{2}} + 2C_{\text{MSE}}^{\frac{1}{2}} \frac{\pi(g^2)^{\frac{1}{2}}}{\pi(g)} \right)$$

and

$$\mathbb{E}[(\mu^N(\phi) - \mu(\phi))^2] \leq \frac{1}{N} C_{\text{MSE}}.$$

The constants $C_t > 0$, $t \geq 2$, satisfy $C_t^{\frac{1}{t}} \leq t - 1$ and the two pairs of parameters d, e , and p, q are conjugate pairs of indices satisfying $d, e, p, q \in (1, \infty)$ and $d^{-1} + e^{-1} = 1$, $p^{-1} + q^{-1} = 1$.

REMARK 2.4. In Bayesian inverse problems, $\pi(g) < \infty$ often implies that $\pi(g^s) < \infty$ for any positive s ; we will demonstrate this in a particular case in Section 3. In such a case, Theorem 2.3 combined with Hölder’s inequality shows that importance sampling converges at rate N^{-1} for any test function ϕ satisfying $\pi(|\phi|^{2+\epsilon}) < \infty$ for some $\epsilon > 0$. Note, however, that the constant in the $\mathcal{O}(N^{-1})$ error bound is not readily interpretable simply in terms of ρ ; in particular the expression necessarily involves moments of g with exponent greater than two.

2.4 Behaviour of the Second Moment ρ

Having demonstrated the importance of ρ , the second moment of the target-proposal density, we now show how it behaves in high dimensional problems and in problems where there are measure concentration phenomena due to a small parameter in the likelihood. These two limits will be of importance to us in subsequent sections of the paper, where the small parameter measure concentration effect will arise due to high quality data.

2.4.1 High state space dimension and absolute continuity. The preceding three subsections have demonstrated how, when the target is absolutely continuous with respect to the proposal, importance sampling converges as the square root of ρ/N . It is thus natural to ask if, and how, this desirable convergence breaks down for sequences of target and proposal measures

which become increasingly close to singular. To this end, suppose that the underlying space is the Cartesian product \mathbb{R}^d equipped with the corresponding product σ -algebra, the proposal is a product measure and the un-normalized weight function also has a product form, as follows:

$$\begin{aligned} \pi_d(du) &= \prod_{i=1}^d \pi_1(du(i)), \\ \mu_d(du) &= \prod_{i=1}^d \mu_1(du(i)), \\ g_d(u) &= \exp\left\{-\sum_{i=1}^d h(u(i))\right\}, \end{aligned}$$

for probability measures π_1, μ_1 on \mathbb{R} and $h : \mathbb{R} \rightarrow \mathbb{R}^+$ (and we assume it is not constant to remove the trivial case $\mu_1 = \pi_1$). We index the proposal, target, density and ρ with respect to d since interest here lies in the limiting behaviour as d increases. In the setting of (1.1), we now have

$$(2.6) \quad \mu_d(du) \propto g_d(u)\pi_d(du).$$

By construction, g_d has all polynomial moments under π_d and importance sampling for each d has the good properties developed in the previous sections. It is also fairly straightforward to see that μ_∞ and π_∞ are mutually singular when h is not constant: one way to see this is to note that

$$\frac{1}{d} \sum_{i=1}^d u(i)$$

has a different almost sure limit under μ_∞ and π_∞ . Two measures cannot be absolutely continuous unless they share the same almost sure properties. Therefore, μ_∞ is not absolutely continuous with respect to π_∞ and importance sampling is undefined in the limit $d = \infty$. As a consequence, we should expect to see a degradation in its performance for large state space dimension d .

To illustrate this degradation note that under the product structure (2.6), we have $\rho_d = (\rho_1)^d$. Furthermore, $\rho_1 > 1$ (since h is not constant). Thus, ρ_d grows exponentially with the state space dimension suggesting, when combined with Theorem 2.1, that exponentially many particles are required, with respect to dimension, to make importance sampling accurate.

It is important to realise that it is not the product structure *per se* that leads to the collapse, rather the lack of absolute continuity in the limit of infinite state

space dimension. Thinking about the role of high dimensions in this way is very instructive in our understanding of high dimensional problems, but is very much related to the setting in which all the coordinates of the problem play a similar role. This does not happen in many application areas. Often there is a diminishing response of the likelihood to perturbations in growing coordinate index. When this is the case, increasing the state space dimension has only a mild effect in the cost of the problem, and it is possible to have well-behaved infinite dimensional limits; we will see this perspective in Sections 3.1, 3.2 and 3.3 for inverse problems, and Sections 4.1, 4.2 and 4.3 for filtering.

2.4.2 *Singular limits.* In the previous subsection, we saw an example where for high dimensional state spaces the target and proposal became increasingly close to being mutually singular, resulting in ρ which grows exponentially with the state space dimension. In this subsection, we observe that mutual singularity can also occur because of small parameters in the unnormalized density g appearing in (1.1), even in problems of fixed dimension; this will lead to ρ which grows algebraically with respect to the small parameter. To understand this situation, let $\mathcal{X} = \mathbb{R}$ and consider (1.1) in the setting where

$$g_\epsilon(u) = \exp(-\epsilon^{-1}h(u)),$$

where $h : \mathbb{R} \rightarrow \mathbb{R}^+$. We will write g_ϵ and ρ_ϵ to highlight the dependence of these quantities on ϵ . Furthermore assume, for simplicity, that h is twice differentiable and has a unique minimum at u^* , and that $h''(u^*) > 0$. Assume, in addition, that π has a Lebesgue density with bounded first derivative. Then the Laplace method shows that

$$\mathbb{E} \exp(-2\epsilon^{-1}h(u)) \approx \exp(-2\epsilon^{-1}h(u^*)) \sqrt{\frac{2\pi\epsilon}{2h''(u^*)}}$$

and that

$$\mathbb{E} \exp(-\epsilon^{-1}h(u)) \approx \exp(-\epsilon^{-1}h(u^*)) \sqrt{\frac{2\pi\epsilon}{h''(u^*)}}.$$

It follows that

$$\rho_\epsilon \approx \sqrt{\frac{h''(u^*)}{4\pi\epsilon}}.$$

Thus, Theorem 2.1 indicates that the number of particles required for importance sampling to be accurate should grow at least as fast as $\epsilon^{-\frac{1}{2}}$.

2.5 Discussion and Connection to Literature

2.5.1 Metrics between random probability measures. In Section 2.1, we introduced the importance sampling approximation of a target μ using a proposal π , both related by (1.1). The resulting particle approximation measure μ^N is random because it is based on samples from π . Hence, $\mu^N(\phi)$ is a *random* estimator of $\mu(\phi)$. This estimator is in general biased and, therefore, a reasonable metric for its quality is the MSE

$$\mathbb{E}[(\mu^N(\phi) - \mu(\phi))^2],$$

where the expectation is with respect to the randomness in the measure μ^N . We bound the MSE over the class of bounded test functions in Theorem 2.1. In fact, we may view this theorem as giving a bound on a distance between the measure μ and its approximation μ^N . To this end, let ν and μ denote mappings from an underlying probability space (which for us will be that associated with π) into the space of probability measures on $(\mathcal{X}, \mathcal{F})$; in the following, expectation \mathbb{E} is with respect to this underlying probability space. In [85], a distance $d(\cdot, \cdot)$ between such random measures is defined by

$$(2.7) \quad d(\nu, \mu)^2 = \sup_{|\phi| \leq 1} \mathbb{E}[(\nu(\phi) - \mu(\phi))^2].$$

The paper [85] used this distance to study the convergence of particle filters. Note that if the measures are not random the distance reduces to total variation. Using this distance, together with the discussion in Section 2.3.3 linking ρ to the χ^2 divergence, we see that Theorem 2.1 states that

$$d(\mu^N, \mu)^2 \leq \frac{4}{N} (1 + D_{\chi^2}(\mu \parallel \pi)).$$

In Section 2.3.3, we also link ρ to the Kullback–Leibler divergence; the bound (2.5) can be found in Theorem 4.19 of [13]. As was already noted, this suggests the need to increase the number of particles linearly with $D_{\chi^2}(\mu \parallel \pi)$ or exponentially with $D_{\text{KL}}(\mu \parallel \pi)$.

2.5.2 Complementary analyses of importance sampling error. Provided that $\log(\frac{g(u)}{\pi(g)})$, $u \sim \mu$, is concentrated around its expected value, as often happens in large dimensional and singular limits, it has recently been shown [19] that using a sample size of approximately $\exp(D_{\text{KL}}(\mu \parallel \pi))$ is both necessary and sufficient in order to control the L^1 error $\mathbb{E}|\mu^N(\phi) - \mu(\phi)|$ of the importance sampling estimator $\mu^N(\phi)$. Theorem 2.1 is similar to [31], Theorem 7.4.3. However, the later result uses a metric defined over subclasses

of bounded functions. The resulting constants in their bounds rely on covering numbers, which are often intractable. In contrast, the constant ρ in Theorem 2.1 is more amenable to analysis and has several meaningful interpretations as we highlight in this paper. The central limit result in equation (2.3) shows that for large N the upper bound in Theorem 2.1 is sharp. Equation (2.3) can be seen as a trivial application of deeper central limit theorems for particle filters; see [21].

This discussion serves to illustrate the fact that a universal analysis of importance sampling in terms of ρ alone is not possible. Indeed Theorem 2.3 shows that the expression for the error constant in useful error bounds may be quite complex when considering test functions which are not bounded. The constants $C_t > 0$, $t \geq 2$ in Theorem 2.3 are determined by the Marcinkiewicz–Zygmund inequality [86]. The proof follows the approach of [35] for evaluating moments of ratios. Despite the complicated dependence of error constants on the problem at hand, there is further evidence for the centrality of the second moment ρ in the paper [87]. There it is shown (see Remark 4) that, when ρ is finite, a *necessary* condition for accuracy within the class of functions with bounded second moment under the proposal, is that the sample size N is of the order of the χ^2 divergence, and hence of the order of ρ .

Further importance sampling results have been proved within the study of convergence properties of various versions of the particle filter as a numerical method for the approximation of the true filtering/smoothing distribution. These results are often formulated in finite dimensional state spaces, under bounded likelihood assumptions and for bounded test functions; see [1, 26, 27, 32, 77]. Generalizations for continuous time filtering can be found in [6] and [45].

2.5.3 Effective sample size, and the case of infinite second moment. The effective sample size ess , introduced in Section 2.3.2, is a standard statistic used to assess and monitor particle approximation errors in importance sampling [56, 57]. The effective sample size ess does not depend on any specific test function, but is rather a particular function of the normalized weights which quantifies their variability. So does ρ , and as we show in Section 2.3.2 there is an asymptotic connection between both. Our discussion of ess relies on the condition $\pi(g^2) < \infty$. Intuitively, the particle approximation will be rather poor when this condition is not met. Extreme value theory provides some clues about the asymptotic particle approximation error. First, it

may be shown that, regardless of whether $\pi(g^2)$ is finite or not, but simply on the basis that $\pi(g) < \infty$, the largest normalised weight, $w^{(N)}$, will converge to 0 as $N \rightarrow \infty$; see, for example, Section 3 of [36] for a review of related results. On the other hand, [76] shows that, for large N ,

$$\mathbb{E}\left[\frac{N}{\text{ess}}\right] \approx \int_0^N \gamma S(\gamma) d\gamma,$$

where $S(\gamma)$ is the survival function of the distribution of the un-normalized weights, $\gamma := g(u)$ for $u \sim \pi$. For instance, if the weights have density proportional to γ^{-a-1} , for $1 < a < 2$, then $\pi(g^2) = \infty$ and, for large enough N and constant C ,

$$\mathbb{E}\left[\frac{N}{\text{ess}}\right] \approx CN^{-a+2}.$$

Thus, in contrast to the situation where $\pi(g^2) < \infty$, in this setting the effective sample size does not grow linearly with N .

2.5.4 Large state dimension, and singular limits. In Section 2.4.1, we studied high dimensional problems with a product structure that enables analytical calculations. The use of such product structure was pioneered for MCMC methods in [42]. It has then been recently employed in the analysis of importance sampling in high nominal dimensions, starting with the seminal paper [8], and leading on to others such as [9–11, 88–90] and [91].

In [8], Section 3.2, it is shown that the maximum normalised importance sampling weight can be approximately written as

$$w^{(N)} \approx \frac{1}{1 + \sum_{n>1} \exp\{-\sqrt{d}c(z^{(n)} - z^{(1)})\}},$$

where $\{z^n\}_{n=1}^N$ are samples from $N(0, 1)$ and the $z^{(n)}$ are the ordered statistics. In [11], a direct but nontrivial calculation shows that if N does not grow exponentially with d , the sum in the denominator converges to 0 in probability and as a result the maximum weight to 1. Of course, this means that all other weights are converging to zero, and that the effective sample size is 1. It chimes with the heuristic derived in Section 2.4.1 where we show that ρ grows exponentially with d and that choosing N to grow exponentially is thus necessary to keep the upper bound in Theorem 2.1 small. The phenomenon is an instance of what is sometimes termed *collapse of importance sampling* in high dimensions. This type of behaviour can be obtained for other classes of targets and proposals; see [8, 90]. Attempts to alleviate this behaviour include the use of

tempering [9] or combining importance sampling with Kalman-based algorithms [40]. However, the range of applicability of these ideas is still to be studied. In Section 2.4.2, we use the Laplace method. This is a classical methodology for approximating integrals and can be found in many text books; see, for instance, [7].

3. IMPORTANCE SAMPLING AND INVERSE PROBLEMS

The previous section showed that the distance between the proposal and the target is key in understanding the computational cost of importance sampling and the central role played by ρ . In this section, we study the computational cost of importance sampling applied in the context of Bayesian inverse problems, where the target will be the posterior and the proposal the prior. To make the analysis tractable, we consider linear Gaussian inverse problems, but our ideas extend beyond this setting. Section 3.1 describes the setting and necessary background on inverse problems. Then Section 3.2 introduces various notions of “intrinsic dimension” for linear Gaussian inverse problems; a key point to appreciate in the sequel is that this dimension can be finite even when the inverse problem is posed in an infinite dimensional Hilbert space. The analysis of importance sampling starts in Section 3.3. The main result is Theorem 3.8, that shows the equivalence between (i) finite *intrinsic* dimension, (ii) absolute continuity of the posterior (target) with respect to the prior (proposal), and (iii) the central quantity ρ being finite. The section closes with a thorough study of singular limits in Section 3.4 and a literature review in Section 3.5.

3.1 General Setting

We study the inverse problem of finding u from y where

$$(3.1) \quad y = Ku + \eta.$$

In particular, we work in the setting where u is an element of the (potentially infinite dimensional) separable Hilbert space $(\mathcal{H}, \langle \cdot, \cdot \rangle, \|\cdot\|)$. Two cases will help guide the reader.

EXAMPLE 3.1 (Linear Regression Model). In the context of the linear regression model, $u \in \mathbb{R}^{d_u}$ is the regression parameter vector, $y \in \mathbb{R}^{d_y}$ is a vector of training outputs and $K \in \mathbb{R}^{d_y \times d_u}$ is the so-called design matrix whose column space is used to construct a linear predictor for the scalar output. In this setting, $d_u, d_y < \infty$, although in modern applications both might be very large, and the case $d_u \gg d_y$ is the so-called “large p (here d_u) small N (here d_y)” problem.

EXAMPLE 3.2 (Deconvolution Problem). In the context of signal deconvolution, $u \in L^2(0, 1)$ is a square integrable unknown signal on the unit interval, $K : L^2(0, 1) \rightarrow L^2(0, 1)$ is a convolution operator $Ku(x) = (\phi \star u)(x) = \int_0^1 \phi(x - z)u(z) dz$, and $y = Ku + \eta$ is the noisy observation of the convoluted signal where η is observational noise. The convolution kernel ϕ might be, for example, a Gaussian kernel $\phi(x) = e^{-\delta x^2}$. Note also that discretization of the deconvolution problem will lead to a family of instances of the preceding linear regression model, parametrized by the dimension of the discretization space.

The infinite dimensional setting does require some technical background, and this is outlined in the Supplementary Material. Nevertheless, the reader versed only in finite dimensional Gaussian concepts will readily make sense of the notions of intrinsic dimension described in Section 3.2 simply by thinking of (potentially infinite dimensional) matrix representations of covariances.

In equation (3.1), the data y is comprised of the image of the unknown u under a linear map K , with added observational noise η . Here, K can be formally thought of as being a bounded linear operator in \mathcal{H} , which is ill-posed in the sense that if we attempt to invert the data using the (generalized) inverse of K , we get amplification of small errors η in the observation to large errors in the reconstruction of u . In such situations, we need to use regularization techniques in order to stably reconstruct the unknown u from the noisy data y .

We assume Gaussian observation noise $\eta \sim \mathbb{P}_\eta := N(0, \Gamma)$ and adopt a Bayesian approach by putting a prior on the unknown $u \sim \mathbb{P}_u = N(0, \Sigma)$. Here and throughout, $\Gamma : \mathcal{H} \rightarrow \mathcal{H}$ and $\Sigma : \mathcal{H} \rightarrow \mathcal{H}$ are bounded, self-adjoint, positive-definite linear operators. Note that we do not assume that Γ and Σ are trace class, which introduces some technical difficulties since η and u do not necessarily live in \mathcal{H} . This is discussed in the Supplementary Material.

The Bayesian solution is the posterior distribution $u|y \sim \mathbb{P}_{u|y}$. In the finite dimensional setting, the prior \mathbb{P}_u and the posterior $\mathbb{P}_{u|y}$ are Gaussian conjugate and $\mathbb{P}_{u|y} = N(m, C)$, with mean and covariance given by

$$(3.2) \quad m = \Sigma K^*(K \Sigma K^* + \Gamma)^{-1} y,$$

$$(3.3) \quad C = \Sigma - \Sigma K^*(K \Sigma K^* + \Gamma)^{-1} K \Sigma.$$

A simple way to derive the expressions above is by working with precision matrices. Indeed, using Bayes'

rule and completion of the square gives

$$(3.4) \quad C^{-1} = \Sigma^{-1} + K^* \Gamma^{-1} K,$$

$$(3.5) \quad C^{-1} m = K^* \Gamma^{-1} y.$$

An application of Schur complement then yields (3.2) and (3.3).

REMARK 3.3. Under appropriate conditions—see the references in the literature review Section 3.5 and the Supplementary Material—formulae (3.2)–(3.5) can be established in the infinite dimensional setting. From now on and whenever necessary, we *assume* that these expressions are available in the general Hilbert space setting that we work in. In particular, Proposition 3.5 makes use of the formula (3.4) for the posterior precision.

Under the prior and noise models, we may write $u = \Sigma^{\frac{1}{2}} u_0$ and $\eta = \Gamma^{\frac{1}{2}} \eta_0$ where u_0 and η_0 are independent centred Gaussians with identity covariance operators (white noises). Thus, we can write (3.1), for $y_0 = \Gamma^{-\frac{1}{2}} y$, as

$$(3.6) \quad y_0 = S u_0 + \eta_0, \quad S = \Gamma^{-\frac{1}{2}} K \Sigma^{\frac{1}{2}}.$$

Therefore, all results may be derived for this inverse problem, and translated back to the original setting. This intuition demonstrates the centrality of the operator S linking K , Σ and Γ . The following assumption will be in place in the remainder of the paper.

ASSUMPTION 3.4. Define $S = \Gamma^{-\frac{1}{2}} K \Sigma^{\frac{1}{2}}$, $A = S^* S$ and assume that A , viewed as a linear operator in \mathcal{H} , is bounded. Furthermore, assume that the spectrum of A consists of a countable number of eigenvalues, sorted without loss of generality in a nonincreasing way:

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_j \geq \dots \geq 0.$$

In Section 3.5, we give further intuition on the centrality of the operator S , and hence A , and discuss the role of the assumption in the context of inverse problems.

3.2 Intrinsic Dimension

Section 2 demonstrates the importance of the distance between the target (here the posterior) and the proposal (here the prior) in the performance of importance sampling. In the Gaussian setting considered in this section, any such distance is characterized in terms of means and covariances. We now show that the “size” of the operator A defined in Assumption 3.4 can be

used to quantify the distance between the prior and the posterior covariances, Σ and C . In Sections 3.3 and 3.4, we will see that, although A does not contain explicit information on the prior and posterior means, its size largely determines the computational cost of importance sampling.

PROPOSITION 3.5. *Under the general setting of Section 3.1, the following identities hold:*

$$\begin{aligned} \text{Tr}((C^{-1} - \Sigma^{-1})\Sigma) &= \text{Tr}(A), \\ \text{Tr}((\Sigma - C)\Sigma^{-1}) &= \text{Tr}((I + A)^{-1}A). \end{aligned}$$

Thus, the traces of A and of $(I + A)^{-1}A$ measure the differences between the posterior and prior precision and covariance operators, respectively, relative to their prior values. For this reason, they provide useful measures of the computational cost of importance sampling, motivating the following definitions:

$$(3.7) \quad \tau := \text{Tr}(A), \quad \text{efd} := \text{Tr}((I + A)^{-1}A).$$

Note that the trace calculates the sum of the eigenvalues and is well-defined, although may be infinite, in the Hilbert space setting. We refer to efd as effective dimension; both τ and efd are measures of the intrinsic dimension of the inverse problem at hand. The next result shows that the intrinsic dimension efd has the appealing property of being bounded above by the nominal dimension.

PROPOSITION 3.6. *Let S and A be defined as in Assumption 3.4, and consider the finite dimensional setting with the notation introduced in Example 3.1:*

1. *The matrices $\Gamma^{1/2}S(I + A)^{-1}S^*\Gamma^{-1/2} \in \mathbb{R}^{d_y \times d_y}$, $S(I + A)^{-1}S^* \in \mathbb{R}^{d_y \times d_y}$ and $(I + A)^{-1}A \in \mathbb{R}^{d_u \times d_u}$ have the same nonzero eigenvalues, and hence the same trace.*
2. *If $\lambda_i > 0$ is a nonzero eigenvalue of A then these three matrices have corresponding eigenvalue $\lambda_i(1 + \lambda_i)^{-1} < 1$, and*

$$\text{efd} = \sum_i \frac{\lambda_i}{1 + \lambda_i} \leq d = \min\{d_u, d_y\}.$$

Here, recall, $d = \min\{d_u, d_y\}$ is the nominal dimension of the problem. Part 2 of the preceding result demonstrates the connection between efd and the physical dimensions of the unknown and observation spaces, whilst part 1 demonstrates the equivalence between the traces of a variety of operators, all of which are used in the literature; this is discussed in greater detail in Section 3.5. In the Hilbert space setting, recall,

the intrinsic dimensions efd and τ can be infinite. It is important to note, however, that this cannot happen if the rank of K is finite. That is, the intrinsic dimension efd (and, as we now show, also τ) is finite whenever the unknown u or the data y live in a finite dimensional subspace of \mathcal{H} . The following result relates efd and τ . It shows in particular that in the infinite dimensional setting they are finite, or otherwise, at the same time.

LEMMA 3.7. *Under the general setting of Section 3.1, the operator A is trace class if and only if $(I + A)^{-1}A$ is trace class. Moreover, the following inequalities hold:*

$$\frac{1}{\|I + A\|} \text{Tr}(A) \leq \text{Tr}((I + A)^{-1}A) \leq \text{Tr}(A).$$

As a consequence,

$$(3.8) \quad \frac{1}{\|I + A\|} \tau \leq \text{efd} \leq \tau.$$

We are now ready to study the performance of importance sampling with posterior as target and prior as proposal. In Section 3.3, we identify conditions under which we can guarantee that ρ in Theorem 2.1 is finite and absolute continuity holds. In Section 3.4, we then study the growth of ρ as mutual singularity is approached in different regimes. The intrinsic dimensions τ and efd will be woven into these developments.

3.3 Absolute Continuity

In the finite dimensional setting, the Gaussian proposal and target distributions have densities with respect to the Lebesgue measure. They are hence mutually absolutely continuous and it is hence straightforward to find the Radon–Nikodym derivative of the target with respect to the proposal by taking the ratio of the respective Lebesgue densities once the posterior is identified via Bayes’ theorem; this gives

$$(3.9) \quad \begin{aligned} \frac{d\mu}{d\pi}(u) &= \frac{d\mathbb{P}_{u|y}}{d\mathbb{P}_u}(u; y) \\ &\propto \exp\left(-\frac{1}{2}u^*K^*\Gamma^{-1}Ku + u^*K^*\Gamma^{-1}y\right) \\ &=: g(u; y). \end{aligned}$$

Direct calculation shows that, for $d_u, d_y < \infty$ the ratio ρ defined in (2.2) is finite, and indeed that g admits all polynomial moments, all of which are positive. In this subsection, we study ρ in the Hilbert space setting. In general, there is no guarantee that the posterior

is absolutely continuous with respect to the prior; when it is not, g , and hence ρ , are not defined. We thus seek conditions under which such absolute continuity may be established.

To this end, we define the likelihood measure $y|u \sim \mathbb{P}_{y|u} := N(Ku, \Gamma)$, and the joint distribution of (u, y) under the model $\nu(du, dy) := \mathbb{P}_{y|u}(dy|u)\mathbb{P}_u(du)$, recalling that $\mathbb{P}_u = N(0, \Sigma)$. We also define the marginal distribution of the data under the joint distribution, $\nu_y(dy) = \mathbb{P}_y(dy)$. We have the following result.

THEOREM 3.8. *Let Assumption 3.4 hold and let $\mu = \mathbb{P}_{u|y}$ and $\pi = \mathbb{P}_u$. The following are equivalent:*

- (i) $\text{efd} < \infty$;
- (ii) $\tau < \infty$;
- (iii) $\Gamma^{-1/2}Ku \in \mathcal{H}$, π -almost surely;
- (iv) for ν_y -almost all y , the posterior μ is well-defined as a measure in any space of full prior measure and is absolutely continuous with respect to the prior with

$$\begin{aligned}
 \frac{d\mu}{d\pi}(u) & \propto \exp\left(-\frac{1}{2}\|\Gamma^{-1/2}Ku\|^2\right. \\
 (3.10) \quad & \left. + \frac{1}{2}(\Gamma^{-1/2}y, \Gamma^{-1/2}Ku)\right) \\
 & =: g(u; y),
 \end{aligned}$$

where $0 < \pi(g(\cdot; y)) < \infty$.

REMARK 3.9. Due to the exponential structure of g , we have that assertion (iv) of the last theorem is immediately equivalent to g being ν -almost surely positive and finite and for ν_y -almost all y the second moment of the target-proposal density is finite:

$$\rho = \frac{\pi(g(\cdot; y)^2)}{\pi(g(\cdot; y))^2} < \infty.$$

Item (iii) is a requirement on the regularity of the forward image of draws from the prior, relative to the regularity of the noise. This regularity condition heavily constrains the space of possible reconstructions and is thus related to the intrinsic dimension of the inverse problem, as we establish here. For a discussion on the regularity of draws from Gaussian measures in Hilbert spaces, see the Supplementary Material.

We have established something very interesting: there are meaningful notions of intrinsic dimension for inverse problems formulated in infinite state/data state dimensions and, when the intrinsic dimension is finite,

importance sampling may be possible as there is absolute continuity; moreover, in such a situation ρ is finite. Thus, under any of the equivalent conditions (i)–(iv), Theorem 2.1 can be used to provide bounds on the effective sample size ess , defined in Section 2.3.2; indeed the effective sample size is then proportional to N .

It is now of interest to understand how ρ , and the intrinsic dimensions τ and efd , depend on various parameters, such as small observational noise or the dimension of finite dimensional approximations of the inverse problem. Such questions are studied in the next subsection.

3.4 Large Nominal Dimension and Singular Parameter Limits

The parameter ρ is a complicated nonlinear function of the eigenvalues of A and the data y . However, there are some situations in which we can lower bound ρ in terms of the intrinsic dimensions τ , efd and the size of the eigenvalues of A . We present two classes of examples of this type. The first is a simple but insightful example in which the eigenvalues cluster into a finite dimensional set of large eigenvalues and a set of small remaining eigenvalues. The second involves asymptotic considerations in a simultaneously diagonalizable setting.

3.4.1 Spectral jump. Consider the setting where u and y both live in finite dimensional spaces of dimensions d_u and d_y , respectively. Suppose that A has eigenvalues $\{\lambda_i\}_{i=1}^{d_u}$ with $\lambda_i = C \gg 1$ for $1 \leq i \leq k$, and $\lambda_i \ll 1$ for $k+1 \leq i \leq d_u$; indeed, we assume that

$$\sum_{i=k+1}^{d_u} \lambda_i \ll 1.$$

Then $\tau(A) \approx Ck$, whilst the effective dimension satisfies $\text{efd} \approx k$. Using the identity,

$$\begin{aligned}
 2D_{\text{KL}}(\mathbb{P}_{u|y} \|\mathbb{P}_u) & = \log(\det(I + A)) \\
 & \quad - \text{Tr}((I + A)^{-1}A) + m^* \Sigma^{-1}m
 \end{aligned}$$

and studying the asymptotics for fixed m , with k and C large, we obtain

$$D_{\text{KL}}(\mathbb{P}_{u|y} \|\mathbb{P}_u) \approx \frac{\text{efd}}{2} \log(C).$$

Therefore, using (2.5),

$$\rho \gtrsim C^{\frac{\text{efd}}{2}}.$$

This suggests that ρ grows exponentially with the number of large eigenvalues, whereas it has an algebraic

dependence on the *size* of the eigenvalues. Theorem 2.1 then suggests that the number of particles required for accurate importance sampling will grow exponentially with the number of large eigenvalues, and algebraically with the size of the eigenvalues. A similar distinction may be found by comparing the behaviour of ρ in large state space dimension in Section 2.4.1 (exponential) and with respect to small scaling parameter in Section 2.4.2 (algebraic).

3.4.2 *Spectral cascade.* We now introduce a three-parameter family of inverse problems, defined through the eigenvalues of A . These three parameters represent the regularity of the prior and the forward map, the size of the observational noise and the number of positive eigenvalues of A , which corresponds to the nominal dimension. We are interested in investigating the performance of importance sampling, as quantified by ρ , in different regimes for these parameters. We work in the framework of Assumption 3.4, and under the following additional assumption.

ASSUMPTION 3.10. Within the framework of Assumption 3.4, we assume that $\Gamma = \gamma I$ and that A has eigenvalues $\{\frac{j^{-\beta}}{\gamma}\}_{j=1}^\infty$ with $\gamma > 0$, and $\beta \geq 0$. We consider a truncated sequence of problems with $A(\beta, \gamma, d)$, with eigenvalues $\{\frac{j^{-\beta}}{\gamma}\}_{j=1}^d$, $d \in \mathbb{N} \cup \{\infty\}$. Finally, we assume that the data is generated from a fixed underlying infinite dimensional truth u^\dagger ,

$$(3.11) \quad y = Ku^\dagger + \eta, \quad Ku^\dagger \in \mathcal{H},$$

and for the truncated problems the data is given by projecting y onto the first d eigenfunctions of A .

REMARK 3.11. Since $Ku^\dagger \in \mathcal{H}$, using the Gaussian theory provided in the Supplementary Material one can check that the distribution of the data in (3.11) is equivalent to the marginal probability measure of the data under the model, $\nu_y(dy)$. Hence, the conclusions of Theorem 3.8 and Remark 3.9 which are formulated for ν_y -almost all y , also hold for almost all y of the form of (3.11).

Note that d in the previous assumption is the data space dimension, which agrees here with the nominal dimension. The setting of the previous assumption arises, for example, when d is finite, from discretizing the data of an inverse problem formulated in an infinite dimensional state space. Provided that the forward map K and the prior covariance Σ commute, our analysis extends to the case where both the unknown and the data are discretized in the common eigenbasis. In

all these cases, interest lies in understanding how the cost of importance sampling depends on the level of the discretizations. The parameter γ may arise as an observational noise scaling, and it is hence of interest to study the cost of importance sampling when γ is small. And finally, the parameter β reflects regularity of the problem, as determined by the prior and noise covariances, and the forward map; critical phase transitions occur in computational cost as this parameter is varied, as we will show.

EXAMPLE 3.12 (Example 3.2 revisited). We revisit the deconvolution problem in the unit interval. In particular, we consider the problem of deconvolution of a periodic signal which is blurred by a periodic kernel and polluted by Gaussian white noise $N(0, \gamma I)$. This problem is diagonalized by the discrete Fourier transform, giving rise to a countable number of decoupled equations in frequency space of the form:

$$y_j = K_j u_j + \eta_j, \quad j \in \mathbb{N}.$$

Here, u_j are the Fourier coefficients of the unknown signal u , K_j the Fourier coefficients of the blurring kernel ϕ which is assumed to be known, and $\eta_j \stackrel{\text{i.i.d.}}{\sim} N(0, \gamma)$ the Fourier coefficients of the observational noise η . Consider the case in which $K_j \asymp j^{-t}$, $t \geq 0$; the case $t = 0$ corresponds to the direct observation case while the bigger t is the more severe the blurring. We put a Gaussian prior on u , $u \sim N(0, (-\Delta)^{-s})$, $s \geq 0$, where Δ is the Laplacian with periodic boundary conditions on $(0, 1)$, so that by the Karhunen–Loeve expansion $u_j = \sqrt{\kappa_j} \zeta_j$, with $\kappa_j \asymp j^{-2s}$ and $\zeta_j \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$. The larger s is the higher the regularity of draws from the prior. In this case, the operator A has eigenvalues $\{c \frac{j^{-2t-2s}}{\gamma}\}_{j=1}^\infty$, where c is independent of j, γ . For this example, the value of β in Assumption 3.10 is $\beta = 2t + 2s$ and large values of β correspond to more severe blurring and/or higher regularity of the prior. A natural way of discretizing this problem is to truncate the infinite sequence of 1-dimensional problems to d terms, resulting in truncation of the sequence of eigenvalues of A . The limit $\gamma \rightarrow 0$ corresponds to vanishing noise in the observation of the blurred signal.

The intrinsic dimensions $\tau = \tau(\beta, \gamma, d)$ and $\text{efd} = \text{efd}(\beta, \gamma, d)$ read

$$(3.12) \quad \tau = \frac{1}{\gamma} \sum_{j=1}^d j^{-\beta}, \quad \text{efd} = \sum_{j=1}^d \frac{j^{-\beta}}{\gamma + j^{-\beta}}.$$

TABLE 1

The third and fourth columns show the scaling of the intrinsic dimensions with model parameters for the spectra cascade example of Section 3.4.2. The fourth one gives a lower bound on the growth of ρ , suggesting that the number of particles should be increased at least as indicated by this column in terms of the model parameters. This lower bound holds for all realizations of the data y when $\gamma \rightarrow 0$, and in probability for those regimes where γ is fixed. ϵ can be chosen arbitrarily small

Regime	Parameters	efd	τ	ρ
Small noise	$\gamma \rightarrow 0, d < \infty$	d	γ^{-1}	$\gamma^{-d/2}$
	$\gamma \rightarrow 0, d = \infty, \beta > 1$	$\gamma^{-1/\beta}$	γ^{-1}	$\gamma^{-\frac{\epsilon\beta}{2}}(\gamma^{-1/\beta-\epsilon})$
Large d	$d \rightarrow \infty, \beta < 1$	$d^{1-\beta}$	$d^{1-\beta}$	$\exp(d^{1-\beta})$
Small noise and large d	$\gamma = d^{-\alpha}, d \rightarrow \infty, \beta > 1, \alpha > \beta$	d	d^α	$d^{(\alpha-\beta)d}$
	$\gamma = d^{-\alpha}, d \rightarrow \infty, \beta > 1, \alpha < \beta$	$d^{\alpha/\beta}$	d^α	$d^{\epsilon d^{\alpha/\beta-\epsilon}}$
	$\gamma = d^{-\alpha}, d \rightarrow \infty, \beta < 1, \alpha > \beta$	d	$d^{1+\alpha-\beta}$	$d^{(\alpha-\beta)d}$
	$\gamma = d^{-\alpha}, d \rightarrow \infty, \beta < 1, \alpha < \beta$	$d^{1+\alpha-\beta}$	$d^{1+\alpha-\beta}$	$d^{\epsilon d^{\alpha/\beta-\epsilon}}$
Regularity	$d = \infty, \beta \searrow 1$	$\frac{1}{\beta-1}$	$\frac{1}{\beta-1}$	$\exp(\frac{1}{\beta-1})$

Table 1 shows the scalings of the effective dimensions efd and τ with the model parameters. It also shows how ρ behaves under these scalings, and hence gives, by Theorem 2.1, an indication of the number of particles required for accurate importance sampling in a given regime. In all the scaling limits where ρ grows to infinity, the posterior and prior are approaching mutual singularity; we can then apply Theorem 2.1 to get an indication of how importance sampling deteriorates in these limits.

Note that by Theorem 3.8 we have $\tau(\beta, \gamma, d) < \infty$ if and only if $\text{efd}(\beta, \gamma, d) < \infty$. It is clear from (3.12) that $\tau = \infty$ if and only if $\{d = \infty, \beta \leq 1\}$. By Theorem 3.8 again, this implies, in particular, that absolute continuity is lost in the limit as $d \rightarrow \infty$ when $\beta \leq 1$, and as $\beta \searrow 1$ when $d = \infty$. Absolute continuity is also lost in the limit $\gamma \rightarrow 0$, in which the posterior is fully concentrated around the data (at least in those directions in which the data live). In this limit, we always have $\tau = \infty$, whereas $\text{efd} < \infty$ in the case where $d < \infty$ and $\text{efd} = \infty$ when $d = \infty$. Note that in the limit $\gamma = 0$ Assumption 3.4 does not hold, which explains why τ and efd are not finite simultaneously. Indeed, as was noted before, efd is always bounded by the nominal dimension d irrespective of the size γ of the noise.

Some important remarks on Table 1 are:

- ρ grows algebraically in the small noise limit ($\gamma \rightarrow 0$) if the nominal dimension d is finite.
- ρ grows exponentially in τ or efd as the nominal dimension grows ($d \rightarrow \infty$) if $\beta < 1$, and as the prior becomes rougher ($\beta \searrow 1$) if $d = \infty$.

- ρ grows factorially in the small noise limit ($\gamma \rightarrow 0$) if $d = \infty$, and in the joint limit $\gamma = d^{-\alpha}, d \rightarrow \infty$. The exponent in the rates relates naturally to efd.

The scalings of τ and efd can be readily deduced by comparing the sums defining τ and efd with integrals. The analysis of the sensitivity of ρ to the model parameters relies on an explicit expression for this quantity. Details are given in the Supplementary Material.

3.5 Discussion and Connection to Literature

3.5.1 *Examples and Hilbert space formulation of inverse problems.* Further examples of linear inverse problems in both finite and infinite dimensions include the Radon transform inversion used for X-ray imaging, the determination of the initial temperature from later measurements and the inversion of the Laplace transform. Many case studies and more elaborate nonlinear inverse problems can be found, for example, in [49, 95], which adopt a Bayesian approach to their solution, and [38, 80], which adopt a classical approach. The periodic deconvolution problem considered in Example 3.12 is discussed, for instance, in [18], Section 5, where an example of a convolution operator with algebraically decaying spectrum is also provided. The Bayesian approach we undertake, in the example of linear regression (Example 3.1) becomes the Gaussian conjugate Bayesian analysis of linear regression models, as in [69]. This paper also derives formulae (3.4), (3.5) for the mean and covariance expressed via precisions in the finite dimensional setting. For the infinite dimensional counterpart, see [2], Section 5. Formulae (3.2), (3.3) in the infinite dimensional setting are derived in [64, 74]; in the specific case of inverting for the

initial condition in the heat equation they were derived in [39]. The Supplementary Material has a discussion of Gaussian measures in Hilbert spaces and contains further background references.

3.5.2 The operator A : Centrality and assumptions. The assumption that the spectrum of A introduced in Assumption 3.4 consists of a countable number of eigenvalues, means that the operator A can be thought of as an infinitely large diagonal matrix. It holds if A is compact [63], Theorem 3, Chapter 28, but is in fact more general since it covers, for example, the noncompact case $A = I$.

We note here that the inverse problem

$$(3.13) \quad y_0 = w_0 + \eta_0$$

with η_0 a white noise and $w_0 \sim N(0, SS^*)$ is equivalent to (3.6), but formulated in terms of unknown $w_0 = Au_0$, rather than unknown u_0 . In this picture, the key operator is SS^* rather than $A = S^*S$. Note that by Lemma S.M.1.5 in the Supplementary Material $\text{Tr}(S^*S) = \text{Tr}(SS^*)$. Furthermore, if S is compact the operators SS^* and S^*S have the same nonzero eigenvalues [38], Section 2.2, thus $\text{Tr}((I + SS^*)^{-1}SS^*) = \text{Tr}((I + S^*S)^{-1}S^*S)$. The last equality holds even if S is noncompact, since then Lemma S.M.1.5 together with Lemma 3.7 imply that both sides are infinite. Combining, we see that the intrinsic dimension (τ or efd) is the same regardless of whether we view w_0 or u_0 as the unknown. In particular, the assumption that A is bounded is equivalent to assuming that the operators S, S^* or SS^* are bounded [63], Theorem 14, Chapter 19. For the equivalent formulation (3.13), the posterior mean equation (3.2) is

$$m = SS^*(SS^* + I)^{-1}y.$$

If SS^* is compact, that is, if its nonzero eigenvalues λ_i go to 0, then m is a regularized approximation of w_0 , since the components of the data corresponding to small eigenvalues λ_i are shrunk towards zero. On the other hand, if SS^* is unbounded, that is, if its nonzero eigenvalues λ_i go to infinity, then there is no regularization and high frequency components in the data remain almost unaffected by SS^* in m . Therefore, the case SS^* is bounded is the borderline case determining whether the prior has a regularizing effect in the inversion of the data.

The operator A has played an important role in the study of linear inverse problems. First, it has been used for obtaining posterior contraction rates in the small noise limit; see the operator B^*B in [3, 68]. Its use was motivated by techniques for analyzing classical

regularization methods, in particular regularization in Hilbert scales; see [38], Chapter 8. Furthermore, its eigenvalues and eigendirections can be used to determine (optimal) low-rank approximations of the posterior covariance [15], [92], Theorem 2.3. The analogue of A in nonlinear Bayesian inverse problems is the so-called prior-preconditioned data-misfit Hessian, which has been used in [75] to design Metropolis–Hastings proposals. In more realistic settings, the spectrum of A may not be analytically available and needs to be numerically approximated; for example, see [15], Section 6.7, in the context of linearized global seismic inversion.

3.5.3 Notions of dimension and interpretations. In Section 3.2, we study notions of dimension for Bayesian inverse problems. In the Bayesian setting, the prior imparts information and correlations on the components of the unknown u , reducing the number of parameters that are estimated. In the context of Bayesian or penalized likelihood frameworks, this has led to the notion of *effective number of parameters*, defined as

$$\text{Tr}(\Gamma^{1/2}S(I + S^*S)^{-1}S^*\Gamma^{-1/2}).$$

This quantity agrees with efd by Proposition 3.6 and has been used extensively in statistics and machine learning; see, for example, the deviance information criterion in [93] (which generalises this notion to more general Bayesian hierarchical models), and Section 3.5.3 of [12] and references therein. One motivation for this definition is based on a Bayesian version of the ‘‘hat matrix’’; see, for example, [93]. In this article, we provide a different motivation that is more relevant to our aims: rather than as an effective number of parameters, we interpret efd as the effective dimension of the Bayesian linear model. Similar forms of effective dimension have been used for learning problems in [17, 102, 103] and for statistical inverse problems in [73]. In all of these contexts, the size of the operator A quantifies how informative the data is; see the discussion below. The paper [11] introduced the notion of $\tau = \text{Tr}(A)$ as an effective dimension for importance sampling within linear inverse problems and filtering. In that paper, several transformations of the inverse problem are performed before doing the analysis. We undo these transformations here. The role of τ in the performance of the ensemble Kalman filter had been previously studied in [41].

Proposition 3.6 shows that efd is at most as large as the nominal dimension. The difference between both is a measure of the effect the prior has on the inference relative to the maximum likelihood solution. Proposition 3.5 shows that efd quantifies how far the poste-

rior is from the prior, measured in terms of how distant their covariances are in units of the prior; and similarly for τ , but expressed in terms of precisions and again in units of the prior. By the cyclic property of the trace, Lemma S.M.1.5(ii) in the Supplementary Material, and by Proposition 3.5, τ and efd may also be characterized as follows:

$$\begin{aligned}\tau &= \text{Tr}((C^{-1} - \Sigma^{-1})\Sigma) = \text{Tr}((\Sigma - C)C^{-1}), \\ \text{efd} &= \text{Tr}((\Sigma - C)\Sigma^{-1}) = \text{Tr}((C^{-1} - \Sigma^{-1})C).\end{aligned}$$

Thus, we may also view efd as measuring the change in the precision, measured in units given by the posterior precision; whilst τ measures the change in the covariance, measured in units given by the posterior covariance.

4. IMPORTANCE SAMPLING AND FILTERING

This section studies importance sampling in the context of filtering. In particular, we study two different choices of proposals that play an important role in the subject of filtering. The analysis relies on the relationship between Bayesian inversion and filtering mentioned in the introductory section, and detailed here. In Section 4.1, we set-up the problem and derive a link between importance sampling based particle filters and the inverse problem. In Sections 4.2 and 4.3, we use this connection to study, respectively, the intrinsic dimension of filtering and the connection to absolute continuity between the two proposals considered and the target. Section 4.4 contains some explicit computations which enable comparison of the cost of the two proposals in various singular limits relating to high dimension or small observational noise. We conclude with the literature review Section 4.5.

The component of particle filtering which we analyse in this section is only that related to sequential importance sampling; we do not discuss the interaction between the simulated particles which arises via resampling schemes. Such interaction would not typically be very relevant in the two time-unit dynamical systems we study here, but would be necessary to get reasonable numerical schemes when assimilating data over many time units. We comment further on this, and the choice of the assimilation problem we study, in the literature review.

4.1 General Setting

We simplify the notation by setting $j = 0$ in (1.3) to obtain

$$(4.1) \quad \begin{aligned}v_1 &= Mv_0 + \xi, & v_0 &\sim N(0, P), \xi \sim N(0, Q), \\ y_1 &= Hv_1 + \zeta, & \zeta &\sim N(0, R).\end{aligned}$$

Note that we have also imposed a Gaussian assumption on v_0 . Because of the Markov assumption on the dynamics for $\{v_j\}$, we have that v_0 and ξ are independent. As in Section 3, we set-up the problem in a separable Hilbert space \mathcal{H} , although the reader versed only in finite dimensional Gaussian measures should have no trouble following the developments, simply by thinking of the covariance operators as (possibly infinite) matrices. We assume throughout that the covariance operators $P, Q, R : \mathcal{H} \rightarrow \mathcal{H}$ are bounded, self-adjoint, positive linear operators, but not necessarily trace-class (see the discussion on this trace-class issue in Section 3). We also assume that the operators $M, H : \mathcal{H} \rightarrow \mathcal{H}$ that describe, respectively, the unconditioned signal dynamics and the observation operator, can be extended to larger spaces if necessary; see the Supplementary Material for further details on these technical issues.

Our goal in this section is to study the cost of importance sampling within the context of both the standard and optimal proposals for particle filtering. For both these proposals, we show that there is an inverse problem embedded within the particle filtering method, and compute the proposal covariance, the observation operator and the observational noise covariance. We may then use the material from the previous section, concerning inverse problems, to make direct conclusions about the cost of importance sampling for particle filters.

The aim of one step of filtering may be expressed as sampling from the target $\mathbb{P}_{v_1, v_0 | y_1}$. Particle filters do this by importance sampling, with this measure on the product space $\mathcal{X} \times \mathcal{X}$ as the target. We wish to compare two ways of doing this, one by using the proposal distribution $\mathbb{P}_{v_1 | v_0} \mathbb{P}_{v_0}$ and the second by using as proposal distribution $\mathbb{P}_{v_1 | v_0, y_1} \mathbb{P}_{v_0}$. The first is known as the *standard proposal*, and the second as the *optimal proposal*. We now connect each of these proposals to a different inverse problem.

4.1.1 *Standard proposal.* For the standard proposal, we note that, using Bayes' theorem, conditioning, and that the observation y_1 does not depend on v_0 explicitly,

$$\begin{aligned}\mathbb{P}_{v_1, v_0 | y_1} &\propto \mathbb{P}_{y_1 | v_1, v_0} \mathbb{P}_{v_1, v_0} \\ &= \mathbb{P}_{y_1 | v_1, v_0} \mathbb{P}_{v_1 | v_0} \mathbb{P}_{v_0} \\ &= \mathbb{P}_{y_1 | v_1} \mathbb{P}_{v_1 | v_0} \mathbb{P}_{v_0}.\end{aligned}$$

Thus, the density of the target $\mathbb{P}_{v_1, v_0 | y_1}$ with respect to the proposal $\mathbb{P}_{v_1 | v_0} \mathbb{P}_{v_0}$ is proportional to $\mathbb{P}_{y_1 | v_1}$. Although this density concerns a proposal on the joint




FIG. 1. Filtering step decomposed in two different ways. The upper path first pushes forward the measure \mathbb{P}_{v_0} using the signal dynamics, and then incorporates the observation y_1 . The lower path assimilates the observation y_1 first, and then propagates the conditioned measure using the signal dynamics. The standard proposal corresponds to the upper decomposition and the optimal one to the lower decomposition.

space of (v_0, v_1) , since it involves only v_1 we may consider the related inverse problem of finding v_1 , given y_1 , and ignore v_0 .

In this picture, filtering via the standard proposal proceeds as follows:

$$\mathbb{P}_{v_0} \mapsto \mathbb{P}_{v_1} \mapsto \mathbb{P}_{v_1|y_1}.$$

Here, the first step involves propagation of probability measures under the dynamics. This provides the proposal $\pi = \mathbb{P}_{v_1}$ used for importance sampling to determine the target $\mu = \mathbb{P}_{v_1|y_1}$. The situation is illustrated in the upper branch of Figure 1. Since

$$\mathbb{E}(v_1 v_1^*) = \mathbb{E}(Mv_0 + \xi)(Mv_0 + \xi)^*,$$

and v_0 and ξ are independent under the Markov assumption, the proposal distribution is readily seen to be a centred Gaussian with covariance $\Sigma = MPM^* + Q$. The observation operator is $K = H$ and the noise covariance $\Gamma = R$. We have established a direct connection between the particle filter, with standard proposal and the inverse problem of the previous section. We will use this connection to study the cost of the particle filter, with standard proposal, in what follows.

4.1.2 *Optimal proposal.* For the optimal proposal, we note that, by conditioning on v_0 ,

$$\begin{aligned} \mathbb{P}_{v_1, v_0|y_1} &= \mathbb{P}_{v_1|v_0, y_1} \mathbb{P}_{v_0|y_1} \\ &= \mathbb{P}_{v_1|v_0, y_1} \mathbb{P}_{v_0} \frac{\mathbb{P}_{v_0|y_1}}{\mathbb{P}_{v_0}}. \end{aligned}$$

Thus, the density of the target $\mathbb{P}_{v_1, v_0|y_1}$ with respect to the proposal $\mathbb{P}_{v_1|v_0, y_1} \mathbb{P}_{v_0}$ is the same as the density of $\mathbb{P}_{v_0|y_1}$ with respect to \mathbb{P}_{v_0} . As a consequence, although this density concerns a proposal on the joint space of (v_0, v_1) , it is equivalent to an inverse problem involving only v_0 . We may thus consider the related inverse problem of finding v_0 given y_1 , and ignore v_1 .

In this picture, filtering via the optimal proposal proceeds as follows:

$$\mathbb{P}_{v_0} \mapsto \mathbb{P}_{v_0|y_1} \mapsto \mathbb{P}_{v_1|y_1}.$$

Here, the first step involves importance sampling with proposal $\pi = \mathbb{P}_{v_0}$ and target $\mu = \mathbb{P}_{v_0|y_1}$. This target measure is then propagated under the conditioned dynamics to find $\mathbb{P}_{v_1|y_1}$; the underlying assumption of the optimal proposal is that $\mathbb{P}_{v_1|v_0, y_1}$ can be sampled so that this conditioned dynamics can be implemented particle by particle. The situation is illustrated in the lower branch of Figure 1. Since

$$y_1 = HMv_0 + H\xi + \zeta$$

the proposal distribution is readily seen to be a centred Gaussian with covariance $\Sigma = P$, the observation operator $K = HM$ and the noise covariance given by the covariance of $H\xi + \zeta$, namely $\Gamma = HQH^* + R$. Again we have established a direct connection between the particle filter, with optimal proposal, and the inverse problem of the previous section. We will use this connection to study the cost of the particle filter, with optimal proposal, in what follows.

A key assumption of the optimal proposal is the second step: the ability to sample from the conditioned dynamics $\mathbb{P}_{v_1|v_0, y_1}$ and we make a few comments on this before returning to our main purpose, namely to study cost of particle filtering via the connection to an inverse problem. The first comment is to note that since we are in a purely Gaussian setting, this conditioned dynamics is itself determined by a Gaussian and so may in principle be performed in a straightforward fashion. In fact, the conditioned dynamics remains Gaussian even if the forward model Mv_0 is replaced by a nonlinear map $f(v_0)$, so that the optimal proposal has wider applicability than might at first be appreciated. Second,

TABLE 2
Comparison of the standard and optimal proposals

	Standard proposal	Optimal proposal
Proposal	$\mathbb{P}_{v_0}(dv_0)\mathbb{P}_{v_1 v_0}(dv_1)$	$\mathbb{P}_{v_0}(dv_0)\mathbb{P}_{v_1 v_0,y_1}(dv_1)$
BIP	$y_1 = H v_1 + \eta_{st}$	$y_1 = H M v_0 + \eta_{op}$
Prior Cov.	$M P M^* + Q$	P
Data Cov.	R	$R + H Q H^*$
$\log g(u; y_1)$	$-\frac{1}{2} \ H v_1\ _R^2 + \langle y_1, H v_1 \rangle_R$	$-\frac{1}{2} \ H M v_0\ _{R+H Q H^*}^2 + \langle y_1, H M v_0 \rangle_{R+H Q H^*}$

we comment that the Gaussian arising in the conditioned dynamics has mean m and variance Ξ given by the formulae

$$\Xi = Q - Q H^*(H Q H^* + R)^{-1} H Q,$$

$$m = M v_0 + Q H^*(H Q H^* + R)^{-1} (y_1 - H M v_0).$$

It is a tacit assumption in what follows that the operators defining the filtering problem are such that $\Xi : \mathcal{H} \rightarrow \mathcal{H}$ is well-defined and that $m \in \mathcal{H}$ is well-defined. More can be said about these points, but doing so will add further technicalities without contributing to the main goals of this paper.

4.2 Intrinsic Dimension

Using the inverse problems that arise for the standard proposal and for the optimal proposal, and employing them within the definition of A from Assumption 3.4, we find the two operators A arising for these two different proposals:

$$A := A_{st}$$

$$:= (M P M^* + Q)^{1/2} H^* R^{-1} H (M P M^* + Q)^{1/2}$$

for the standard proposal, and

$$A := A_{op} := P^{1/2} M^* H^* (R + H Q H^*)^{-1} H M P^{1/2}$$

for the optimal proposal. Again here it is assumed that these operators are bounded in \mathcal{H} :

ASSUMPTION 4.1. The operators A_{st} and A_{op} , viewed as linear operators in \mathcal{H} , are bounded. Furthermore, the spectra of both A_{st} and A_{op} consist of a countable number of eigenvalues.

Using these definitions of A_{st} and A_{op} , we may define, from (3.7), the intrinsic dimensions τ_{st} , efd_{st} for the standard proposal and τ_{op} , efd_{op} for the optimal one in the following way:

$$\tau_{st} = \text{Tr}(A_{st}), \quad \text{efd}_{st} = \text{Tr}((I + A_{st})^{-1} A_{st})$$

and

$$\tau_{op} = \text{Tr}(A_{op}), \quad \text{efd}_{op} = \text{Tr}((I + A_{op})^{-1} A_{op}).$$

4.3 Absolute Continuity

The following two theorems are a straightforward application of Theorem 3.8, using the connections between filtering and inverse problems made above. The contents of the two theorems are summarized in Table 2.

THEOREM 4.2. Consider one-step of particle filtering for (4.1). Let $\mu = \mathbb{P}_{v_1|y_1}$ and $\pi = \mathbb{P}_{v_1} = N(0, Q + M P M^*)$. Then the following are equivalent:

- (i) $\text{efd}_{st} < \infty$;
- (ii) $\tau_{st} < \infty$;
- (iii) $R^{-1/2} H v_1 \in \mathcal{H}$, π -almost surely;
- (iv) for v_y -almost all y , the target distribution μ is well-defined as a measure in \mathcal{X} and is absolutely continuous with respect to the proposal with

$$(4.2) \quad \frac{d\mu}{d\pi}(v_1)$$

$$\propto \exp\left(-\frac{1}{2} \|R^{-1/2} H v_1\|^2 + \frac{1}{2} \langle R^{-1/2} y_1, R^{-1/2} H v_1 \rangle\right)$$

$$=: g_{st}(v_1; y_1),$$

where $0 < \pi(g_{st}(\cdot; y_1)) < \infty$.

THEOREM 4.3. Consider one-step of particle filtering for (4.1). Let $\mu = \mathbb{P}_{v_0|y_1}$ and $\pi = \mathbb{P}_{v_0} = N(0, Q)$. Then, for $R_{op} = R + H Q H^*$, the following are equivalent:

- (i) $\text{efd}_{op} < \infty$;
- (ii) $\tau_{op} < \infty$;
- (iii) $R_{op}^{-1/2} H M v_0 \in \mathcal{H}$, π -almost surely;
- (iv) for v_y -almost all y , the target distribution μ is well defined as a measure in \mathcal{X} and is absolutely

continuous with respect to the proposal with

$$\begin{aligned}
 & \frac{d\mu}{d\pi}(v_0) \\
 (4.3) \quad & \propto \exp\left(-\frac{1}{2}\|R_{\text{op}}^{-1/2}HMv_0\|^2\right. \\
 & \left. + \frac{1}{2}(R_{\text{op}}^{-1/2}y_1, R_{\text{op}}^{-1/2}HMv_0)\right) \\
 & =: g_{\text{op}}(v_0; y_1),
 \end{aligned}$$

where $0 < \pi(g_{\text{op}}(\cdot; y_1)) < \infty$.

REMARK 4.4. Because of the exponential structure of g_{st} and g_{op} , the assertion (iv) in the preceding two theorems is equivalent to g_{st} and g_{op} being ν -almost surely positive and finite and for almost all y_1 the second moment of the target-proposal density being finite. This second moment is given, for the standard and optimal proposals, by

$$\rho_{\text{st}} = \frac{\pi(g_{\text{st}}(\cdot; y)^2)}{\pi(g_{\text{st}}(\cdot; y))^2} < \infty$$

and

$$\rho_{\text{op}} = \frac{\pi(g_{\text{op}}(\cdot; y)^2)}{\pi(g_{\text{op}}(\cdot; y))^2} < \infty$$

respectively. The relative sizes of ρ_{st} and ρ_{op} determine the relative efficiency of the standard and optimal proposal versions of filtering.

The following theorem shows that there is loss of absolute continuity for the standard proposal whenever there is for the optimal one. The result is formulated in terms of the intrinsic dimension τ , and we show that $\tau_{\text{op}} = \infty$ implies $\tau_{\text{st}} = \infty$; by Theorem 3.8, this implies the result concerning absolute continuity. Recalling that poor behaviour of importance sampling is intimately related to such breakdown, this suggests that the optimal proposal is always at least as good as the standard one. The following theorem also gives a condition on the operators H , Q and R under which collapse for both proposals occurs at the same time, irrespective of the regularity of the operators M and P . Roughly speaking, this simultaneous collapse result states that if R is large compared to Q then absolute continuity for both proposals is equivalent; and hence collapse of importance sampling happens under one proposal if and only if it happens under the other. Intuitively, the advantages of the optimal proposal stem from the noise in the dynamics; they disappear completely if the dynamics is deterministic. The theorem quantifies this idea. Finally, an example demonstrates that there are

situations where τ_{op} is finite, so that optimal proposal based importance sampling works well for finite dimensional approximations of an infinite dimensional problem, whilst τ_{st} is infinite, so that standard proposal based importance sampling works poorly for finite dimensional approximations.

THEOREM 4.5. Suppose that Assumption 4.1 holds. Then

$$(4.4) \quad \tau_{\text{op}} \leq \tau_{\text{st}}.$$

Moreover, if $\text{Tr}(HQH^*R^{-1}) < \infty$, then

$$\tau_{\text{st}} < \infty \iff \tau_{\text{op}} < \infty.$$

We remark that, under additional simplifying assumptions, we can obtain bounds of the form (4.4) for efd and ρ . We chose to formulate the result in terms of τ since we can prove the bound (4.4) in full generality. Moreover, by Theorem 3.8 the bound in terms of τ suffices in order to understand the different collapse properties of both proposals.

The following example demonstrates that it is possible that $\tau_{\text{op}} < \infty$ while $\tau_{\text{st}} = \infty$; in this situation, filtering via the optimal proposal is well-defined, whilst using the standard proposal it is not. Loosely speaking, this happens if y_1 provides more information on v_1 than v_0 .

EXAMPLE 4.6. Suppose that

$$H = Q = R = M = I, \quad \text{Tr}(P) < \infty.$$

Then it is straightforward from the definitions that $A_{\text{st}} = P + I$ and $A_{\text{op}} = P/2$. In an infinite dimensional Hilbert space setting, the identity operator has infinite trace, $\text{Tr}(I) = \infty$, and so

$$\tau_{\text{st}} = \text{Tr}(A_{\text{st}}) = \text{Tr}(P + I) = \infty,$$

$$\tau_{\text{op}} = \text{Tr}(A_{\text{op}}) = \text{Tr}(P/2) < \infty.$$

We have thus established an example of a filtering model for which $\tau_{\text{st}} = \infty$ and $\tau_{\text{op}} < \infty$. We note that by Theorem 4.5, any such example satisfies the condition $\text{Tr}(HQH^*R^{-1}) = \infty$. When this condition is met, automatically $\tau_{\text{st}} = \infty$. However, τ_{op} can still be finite. Indeed, within the proof of that theorem in the Supplementary Material we show that the inequality

$$\tau_{\text{op}} \leq \text{Tr}(R^{-1}HMPM^*H^*)$$

always holds. The right-hand side may be finite provided that the eigenvalues of P decay fast enough. A simple example of this situation is where HM is a bounded operator and all the relevant operators have

eigenvalues. In this case, the Rayleigh–Courant–Fisher theorem (see the Supplementary Material) guarantees that the eigenvalues of $HMPM^*H^*$ can be bounded in terms of those of P . Again by the Rayleigh–Courant–Fisher theorem, since we are always assuming that the covariance R is bounded, it is possible to bound the eigenvalues of $R^{-1}HMPM^*H^*$ in terms of those of $HMPM^*H^*$. This provides a wider range of examples where $\tau_{st} = \infty$ while $\tau_{op} < \infty$.

4.4 Large Nominal Dimension and Singular Parameter Limits

We noted in Remark 4.4 that the values of the second moment of the target-proposal density, ρ_{st} and ρ_{op} , can be used to characterize the performance of particle filters based on the standard and optimal proposals, respectively. By comparing the values of ρ_{st} and ρ_{op} , we can ascertain situations in which the optimal proposal has significant advantage over the standard proposal. We also recall, from Section 3, the role of the intrinsic dimensions in determining the scaling of the second moment of the target-proposal density.

The following example will illustrate a number of interesting phenomena in this regard. In the setting of fixed finite state/data state dimension, it will illustrate how the scalings of the various covariances entering the problem effect computational cost. In the setting of increasing nominal dimension d , when the limiting target is singular with respect to the proposal, it will illustrate how computational cost scales with d . And finally, we will contrast the cost of the filters in two differing initialization scenarios: (i) from an arbitrary initial covariance P , and from a steady state covariance P_∞ . Such a steady state covariance is a fixed point of the covariance update map for the Kalman filter defined by (1.3).

EXAMPLE 4.7. Suppose that $M = H = I \in \mathbb{R}^{d \times d}$, and $R = rI$, $Q = qI$, with $r, q > 0$. A simple calculation shows that the steady state covariance is given

by

$$P_\infty = \frac{\sqrt{q^2 + 4qr} - q}{2} I,$$

and that the operators A_{st} and A_{op} when $P = P_\infty$ are

$$A_{st} = \frac{\sqrt{q^2 + 4qr} + q}{2r} I, \quad A_{op} = \frac{\sqrt{q^2 + 4qr} - q}{2(q+r)} I.$$

Note that A_{st} and A_{op} are a function of q/r only, whereas P_∞ is not.

If the filtering step is initialized outside stationarity at $P = pI$, with $p > 0$, then

$$A_{st} = \frac{p+q}{r} I, \quad A_{op} = \frac{p}{q+r} I.$$

Both the size and number of the eigenvalues of A_{op}/A_{st} play a role in determining the size of ρ , the second moment of the target-proposal variance. It is thus interesting to study how ρ scales in both the small observational noise regime $r \ll 1$ and the high dimensional regime $d \gg 1$. The results are summarized in Tables 3 and 4. Some conclusions from these tables are:

- The standard proposal degenerates at an algebraic rate as $r \rightarrow 0$, for fixed dimension d , for both initializations of P .
- The optimal proposal is not sensitive to the small observation limit $r \rightarrow 0$ if the size of the signal noise, q , is fixed. If started outside stationarity, the optimal proposal degenerates algebraically if $q \propto r \rightarrow 0$. However, even in this situation the optimal proposal scales well if initialized in the stationary regime.
- In this example, the limiting problem with $d = \infty$ has infinite intrinsic dimension for both proposals, because the target and the proposal are mutually singular. As a result, ρ grows exponentially in the large d limit.
- Example 4.6 suggests that there are cases where ρ_{st} grows exponentially in the large dimensional limit

TABLE 3

Scalings of the standard and optimal proposals in small noise and large d regimes for one filter step initialized from stationarity ($P = P_\infty$). This table should be interpreted in the same way as Table 1

Regime	Param.	eig(A_{st})	eig(A_{op})	eig(P_∞)	ρ_{st}	ρ_{op}
Small obs. noise	$r \rightarrow 0$	r^{-1}	r	r	$r^{-d/2}$	1
	$r = q \rightarrow 0$	1	1	$r(=q)$	1	1
Large d	$d \rightarrow \infty$	1	1	1	$\exp(d)$	$\exp(d)$

TABLE 4

Scalings of the standard and optimal proposals in small noise and large d regimes for one filter step initialized from $P = pI$. This table should be interpreted in the same way as Table 1

Regime	Param.	$\text{eig}(A_{\text{st}})$	$\text{eig}(A_{\text{op}})$	ρ_{st}	ρ_{op}
Small obs. noise	$r \rightarrow 0$	r^{-1}	1	$r^{-d/2}$	1
	$r = q \rightarrow 0$	r^{-1}	r^{-1}	$r^{-d/2}$	$r^{-d/2}$
Large d	$d \rightarrow \infty$	1	1	$\exp(d)$	$\exp(d)$

$d \rightarrow \infty$ but ρ_{op} converges to a finite value. This may happen if $\text{Tr}(HQH^*R^{-1}) < \infty$, but the prior covariance P is sufficiently smooth.

4.5 Discussion and Connection to Literature

In Section 4.1, we follow [8, 11, 88–91] and consider one step of the filtering model (1.3). There are two main motivations for studying one step of the filter. First, if keeping the filter error small is prohibitively costly for one step, then there is no hope that an on-line particle filter will be successful [8]. Second, it can provide insight for filters initialized close to stationarity [23]. As in [88, 89, 91], we cast the analysis of importance sampling in joint space and consider as target $\mu := \mathbb{P}_{u|y_1}$, with $u := (v_0, v_1)$ and with the standard and optimal proposals defined in Section 4.1.

In general nonlinear, non-Gaussian problems, the optimal proposal is usually not implementable, since it is not possible to evaluate the corresponding weights, or to sample from the distribution $\mathbb{P}_{v_1|v_0, y_1}$. However, the optimal proposal is implementable in our framework (see, e.g., [34]) and understanding its behaviour is important in order to build and analyse improved and computable proposals which are informed by the data [44, 78, 98]. It is worth making the point that the so-called ‘‘optimal proposal’’ is really only *locally* optimal. In particular, this choice is optimal in minimizing the variance of the weights at the given step given that all previous proposals have been already chosen. This choice does not minimize the Monte Carlo variance for some time horizon for some family of test functions. A different optimality criterion is obtained by trying to simultaneously minimize the variance of weights at times $t \leq r \leq t + m$, for some $m \geq 1$, or minimize some function of these variances, say their sum or their maximum. Such look ahead procedures might not be feasible in practice. Surprisingly, examples exist where the standard proposal leads to smaller variance of weights some steps ahead relative to the locally optimally tuned particle filter; see, for example, Section 3 of [46], and

the discussion in [22], Chapter 10. Still, such examples are quite contrived and experience suggests that local adaptation is useful in practice.

Similarly as for inverse problems, the values of ρ_{st} and ρ_{op} determine the performance of importance sampling for the filtering model with the standard and optimal proposals. In Section 4.3, we show that the conditions of collapse for the standard and optimal proposals (found in [89] and [11], resp.) correspond to any of the equivalent conditions of finite intrinsic dimension or finite ρ in Theorems 4.2 and 4.3.

In Section 4.4, we study singular limits in the framework of [23]. Thus, we consider a diagonal filtering setting in the Euclidean space \mathbb{R}^d , and assume that all coordinates of the problem play the same role, which corresponds to the extreme case $\beta = 0$ in Section 3.4. The paper [23] introduced a notion of effective dimension for detectable and stabilizable linear Gaussian data assimilation problems as the Frobenius norm of the steady state covariance of the filtering distribution. It is well known that the detectability and stabilizability conditions ensure the existence of such steady state covariance [59]. This notion of dimension quantifies the success of data assimilation in having reduced uncertainty on the unknown once the data has been assimilated. Therefore, the definition of dimension given in [23] is at odds with both τ and efd : it does not quantify how much is learned from the data in one step, but instead how concentrated the filtering distribution is in the time asymptotic regime when the filter is in steady state. Our calculations demonstrate differences which can occur in the computational cost of filtering, depending on whether it is initialized in this statistical steady state, or at an arbitrary point. The paper [23] also highlights the importance of the size of the operator A in studying the performance of importance sampling, both for the standard and optimal proposals. Motivated by computational and physical intuition, the authors of [23] quantify the size of this operator by means of the Frobenius norm rather than the trace which we

employ here. The trace is more natural in the infinite dimensional limit, as demonstrated through large intrinsic dimension limits in [11], and through the connection with absolute continuity in Theorems 4.2 and 4.3 above. We remark that the analysis in [89] also relies on traces, but an unfortunate typo may trick the reader into believing that the Frobenius norm is being used. Note also that some authors [11] write the eigenvalues as squares which can cause confusion to the casual reader.

5. CONCLUSIONS

In this article, we have provided a framework which unifies the multitude of publications with bearing on importance sampling. We have aimed to give new insight into the potential use of importance sampling for inference in inverse problems and filtering settings in models that involve high and infinite state space and data dimensions. Our study has required revisiting the fundamental structure of importance sampling on general state spaces. We have derived nonasymptotic concentration inequalities for the particle approximation error and related what turns out to be the key parameter of performance, the second moment of the density between the target and proposal, to many different importance sampling input and output quantities.

As a compromise between mathematical tractability and practical relevance, we have focused on Bayesian linear models for regression and statistical inversion of ill-posed inverse problems. We have studied the efficiency of sampling-based posterior inference in these contexts carried out by importance sampling using the prior as proposal. We have demonstrated that performance is controlled by an intrinsic dimension, as opposed to the state space or data dimensions, and we have discussed and related two different notions of intrinsic dimension. It is important to emphasise that the intrinsic dimension quantifies the relative strength between the prior and the likelihood in forming the posterior, as opposed to quantifying the “degrees of freedom” in the prior. In other words, infinite-dimensional Bayesian linear models with finite intrinsic dimension are not identified with models for which the prior is concentrated on a finite-dimensional manifold of the infinite-dimensional state space.

A similar consideration of balancing tractability and practical relevance has dictated the choice not to study interacting particles typically used for filtering, but rather to focus on one-step filtering using importance sampling. For such problems, we introduce appropriate

notions of intrinsic dimension and compare the relative merits of popular alternative schemes.

The most pressing topic for future research stemming from this article is the development of concrete recommendations for algorithmic design within classes of Bayesian models used in practice. Within the model structure, we have studied here, practically relevant and important extensions include models with non-Gaussian priors on the unknown, nonlinear operators that link the unknown to the data, and unknown hyperparameters involved in the model specification. Linearisation of a nonlinear model around some reasonable value for the unknown (e.g., the posterior mean) is one way to extend our measures of intrinsic dimension in such frameworks. We can expect the subject area to see considerable development in the coming decade.

ACKNOWLEDGMENTS

The authors are thankful to Alexandre Chorin, Arnaud Doucet, Adam Johansen and Matthias Morzfeld for their generous feedback. SA and DSA are grateful to EPSRC for financial support. AMS is grateful to DARPA, EPSRC and ONR for financial support.

SUPPLEMENTARY MATERIAL

Supplement to “Importance sampling: Intrinsic dimension and computational cost”. (DOI: [10.1214/17-STS611SUPP](https://doi.org/10.1214/17-STS611SUPP); .pdf). The Supplementary Material contains the proofs of all our results. It also contains some background on Gaussian measures on Hilbert spaces, and related technical aspects arising from considering measures in infinite dimensional spaces.

REFERENCES

- [1] ACHUTEGUI, K., CRISAN, D., MIGUEZ, J. and RIOS, G. (2014). A simple scheme for the parallelization of particle filters and its application to the tracking of complex stochastic systems. Preprint. Available at [arXiv:1407.8071](https://arxiv.org/abs/1407.8071).
- [2] AGAPIOU, S., LARSSON, S. and STUART, A. M. (2013). Posterior contraction rates for the Bayesian approach to linear ill-posed inverse problems. *Stochastic Process. Appl.* **123** 3828–3860. [MR3084161](https://doi.org/10.1016/j.spl.2013.08.011)
- [3] AGAPIOU, S. and MATHÉ, P. (2014). Preconditioning the prior to overcome saturation in Bayesian inverse problems. Preprint. Available at [arXiv:1409.6496](https://arxiv.org/abs/1409.6496).
- [4] AGAPIOU, S., PAPASPILIOPOULOS, O., SANZ-ALONSO, D. and STUART, A. M. (2017). Supplement to “Importance sampling: Intrinsic dimension and computational cost.” DOI:[10.1214/17-STS611SUPP](https://doi.org/10.1214/17-STS611SUPP).

- [5] AGAPIOU, S., STUART, A. M. and ZHANG, Y.-X. (2014). Bayesian posterior contraction rates for linear severely ill-posed inverse problems. *J. Inverse Ill-Posed Probl.* **22** 297–321. [MR3215928](#)
- [6] BAIN, A. and CRISAN, D. (2009). *Fundamentals of Stochastic Filtering* **3**. Springer, Berlin.
- [7] BENDER, C. M. and ORSZAG, S. A. (1999). *Advanced Mathematical Methods for Scientists and Engineers. I. Asymptotic Methods and Perturbation Theory*. Springer, New York. [MR1721985](#)
- [8] BENGTTSSON, T., BICKEL, P., LI, B. et al. (2008). Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. In *Probability and Statistics: Essays in Honor of David A. Freedman. Inst. Math. Stat. (IMS) Collect.* 316–334. IMS, Beachwood, OH.
- [9] BESKOS, A., CRISAN, D. and JASRA, A. (2014). On the stability of sequential Monte Carlo methods in high dimensions. *Ann. Appl. Probab.* **24** 1396–1445. [MR3211000](#)
- [10] BESKOS, A., CRISAN, D., JASRA, A., KAMATANI, K. and ZHOU, Y. (2014). A stable particle filter in high-dimensions. Preprint. Available at [1412.3501](#).
- [11] BICKEL, P., LI, B. and BENGTTSSON, T. (2008). Sharp failure rates for the bootstrap particle filter in high dimensions. In *Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh. Inst. Math. Stat. (IMS) Collect.* **3** 318–329. IMS, Beachwood, OH. [MR2459233](#)
- [12] BISHOP, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer, New York. [MR2247587](#)
- [13] BOUCHERON, S., LUGOSI, G. and MASSART, P. (2013). *Concentration Inequalities*. Oxford Univ. Press, Oxford. [MR3185193](#)
- [14] BUI-THANH, T. and GHATTAS, O. (2012). Analysis of the Hessian for inverse scattering problems: I. Inverse shape scattering of acoustic waves. *Inverse Problems* **28** 055001, 32. [MR2915289](#)
- [15] BUI-THANH, T., GHATTAS, O., MARTIN, J. and STADLER, G. (2013). A computational framework for infinite-dimensional Bayesian inverse problems part I: The linearized case, with application to global seismic inversion. *SIAM J. Sci. Comput.* **35** A2494–A2523.
- [16] CAFLISCH, R. E., MOROKOFF, W. J. and OWEN, A. B. (1997). *Valuation of Mortgage Backed Securities Using Brownian Bridges to Reduce Effective Dimension*. Dept. Mathematics, Univ. California, Los Angeles.
- [17] CAPONNETTO, A. and DE VITO, E. (2007). Optimal rates for the regularized least-squares algorithm. *Found. Comput. Math.* **7** 331–368.
- [18] CAVALIER, L. and TSYBAKOV, A. (2002). Sharp adaptation for inverse problems with random noise. *Probab. Theory Related Fields* **123** 323–354.
- [19] CHATTERJEE, S. and DIACONIS, P. (2015). The sample size required in importance sampling. Preprint. Available at [arXiv:1511.01437](#).
- [20] CHEN, Y. (2005). Another look at rejection sampling through importance sampling. *Statist. Probab. Lett.* **72** 277–283. [MR2153124](#)
- [21] CHOPIN, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann. Statist.* **32** 2385–2411. [MR2153989](#)
- [22] CHOPIN, N. and PAPASPILIOPOULOS, O. (2016). A Concise Introduction to Sequential Monte Carlo.
- [23] CHORIN, A. J. and MORZFELD, M. (2013). Conditions for successful data assimilation. *Journal of Geophysical Research: Atmospheres* **118** 11–522.
- [24] CONSTANTINE, P. G. (2015). *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies* **2**. SIAM Spotlights, Philadelphia, PA.
- [25] COTTER, S. L., ROBERTS, G. O., STUART, A. M. and WHITE, D. (2013). MCMC methods for functions: Modifying old algorithms to make them faster. *Statist. Sci.* **28** 424–446. [MR3135540](#)
- [26] CRISAN, D., DEL MORAL, P. and LYONS, T. (1999). Discrete filtering using branching and interacting particle systems. *Markov Process. Related Fields* **5** 293–318.
- [27] CRISAN, D. and DOUCET, A. (2002). A survey of convergence results on particle filtering methods for practitioners. *IEEE Trans. Signal Process.* **50** 736–746.
- [28] CRISAN, D. and ROZOVSKII, B., eds. (2011). *The Oxford Handbook of Nonlinear Filtering*. Oxford Univ. Press, Oxford. [MR2882749](#)
- [29] CUI, T., MARTIN, J., MARZOUK, Y. M., SOLONEN, A. and SPANTINI, A. (2014). Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems* **30** 114015, 28. [MR3274599](#)
- [30] DASHTI, M. and STUART, A. M. (2016). The Bayesian approach to inverse problems. In *Handbook of Uncertainty Quantification* (R. Ghanem, D. Higdon and H. Owahdi, eds.). <http://arxiv.org/abs/1302.6989>.
- [31] DEL MORAL, P. (2004). *Feynman–Kac Formulae*. Springer, New York. [MR2044973](#)
- [32] DEL MORAL, P. and MICLO, L. (2000). Branching and interacting particle systems approximations of Feynman–Kac formulae with applications to non-linear filtering. In *Séminaire de Probabilités, XXXIV. Lecture Notes in Math.* **1729** 1–145. Springer, Berlin. [MR1768060](#)
- [33] DOUCET, A., DE FREITAS, N. and GORDON, N. (2001). An introduction to sequential Monte Carlo methods. In *Sequential Monte Carlo Methods in Practice*. 3–14. Springer, New York. [MR1847784](#)
- [34] DOUCET, A., GODSILL, S. and ANDRIEU, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. Comput.* **10** 197–208.
- [35] DOUKHAN, P. and LANG, G. (2009). Evaluation for moments of a ratio with application to regression estimation. *Bernoulli* **15** 1259–1286.
- [36] DOWNEY, P. J. and WRIGHT, P. E. (2007). The ratio of the extreme to the sum in a random sequence. *Extremes* **10** 249–266.
- [37] DUPUIS, P., SPILIOPOULOS, K. and WANG, H. (2012). Importance sampling for multiscale diffusions. *Multiscale Model. Simul.* **10** 1–27. [MR2902596](#)
- [38] ENGL, H. W., HANKE, M. and NEUBAUER, A. (1996). *Regularization of Inverse Problems. Mathematics and Its Applications* **375**. Kluwer Academic, Dordrecht. [MR1408680](#)
- [39] FRANKLIN, J. N. (1970). Well-posed stochastic extensions of ill-posed linear problems. *J. Math. Anal. Appl.* **31** 682–716. [MR0267654](#)

- [40] FREI, M. and KÜNSCH, H. R. (2013). Bridging the ensemble Kalman and particle filters. *Biometrika* **100** 781–800.
- [41] FURRER, R. and BENGTSOON, T. (2007). Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *J. Multivariate Anal.* **98** 227–255.
- [42] GELMAN, A., ROBERTS, G. O. and GILKS, W. R. (1996). Efficient Metropolis jumping rules. In *Bayesian Statistics, 5 (Alicante, 1994)*. 599–607. Oxford Univ. Press, New York. MR1425429
- [43] GIBBS, A. L. and SU, F. E. (2002). On choosing and bounding probability metrics. *Int. Stat. Rev.* **70** 419–435.
- [44] GOODMAN, J., LIN, K. K. and MORZFELD, M. (2015). Small-noise analysis and symmetrization of implicit Monte Carlo samplers. *Comm. Pure Appl. Math.*
- [45] HAN, W. (2013). On the Numerical Solution of the Filtering Problem Ph.D. thesis, Dept. Mathematics, Imperial College, London.
- [46] JOHANSEN, A. M. and DOUCET, A. (2008). A note on auxiliary particle filters. *Statist. Probab. Lett.* **78** 1498–1504.
- [47] KAHN, H. (1955). *Use of Different Monte Carlo Sampling Techniques*. Rand Corporation, Santa Monica, CA.
- [48] KAHN, H. and MARSHALL, A. W. (1953). Methods of reducing sample size in Monte Carlo computations. *Journal of the Operations Research Society of America* **1** 263–278.
- [49] KAIPIO, J. and SOMERSALO, E. (2005). *Statistical and Computational Inverse Problems*. Applied Mathematical Sciences **160**. Springer, New York. MR2102218
- [50] KALLENBERG, O. (2002). *Foundations of Modern Probability. Probability and Its Applications*, 2nd ed. Springer, Berlin.
- [51] KALMAN, R. E. (1960). A new approach to linear filtering and prediction problems. *J. Fluids Eng.* **82** 35–45.
- [52] KALNAY, E. (2003). *Atmospheric Modeling, Data Assimilation, and Predictability*. Cambridge Univ. Press, Cambridge.
- [53] KEKKONEN, H., LASSAS, M. and SILTANEN, S. (2016). Posterior consistency and convergence rates for Bayesian inversion with hypoelliptic operators. *Inverse Probl.* **32** 085005, 31. MR3535664
- [54] KNAPIK, B. T., VAN DER VAART, A. W. and VAN ZANTEN, J. H. (2011). Bayesian inverse problems with Gaussian priors. *Ann. Statist.* **39** 2626–2657.
- [55] KNAPIK, B. T., VAN DER VAART, A. W. and VAN ZANTEN, J. H. (2013). Bayesian recovery of the initial condition for the heat equation. *Comm. Statist. Theory Methods* **42** 1294–1313.
- [56] KONG, A. (1992). A note on importance sampling using standardized weights. Technical Report 348, Dept. Statistics, Univ. Chicago, Chicago, IL.
- [57] KONG, A., LIU, J. S. and WONG, W. H. (1994). Sequential imputations and Bayesian missing data problems. *J. Amer. Statist. Assoc.* **89** 278–288.
- [58] KUO, F. Y. and SLOAN, I. H. (2005). Lifting the curse of dimensionality. *Notices Amer. Math. Soc.* **52** 1320–1329. MR2183869
- [59] LANCASTER, P. and RODMAN, L. (1995). *Algebraic Riccati Equations*. Oxford Univ. Press, London.
- [60] LASANEN, S. (2007). Measurements and infinite-dimensional statistical inverse theory. *PAMM* **7** 1080101–1080102.
- [61] LASANEN, S. (2012). Non-Gaussian statistical inverse problems. Part I: Posterior distributions. *Inverse Probl. Imaging* **6** 215–266. MR2942739
- [62] LASANEN, S. (2012). Non-Gaussian statistical inverse problems. Part II: Posterior convergence for approximated unknowns. *Inverse Probl. Imaging* **6** 267–287.
- [63] LAX, P. D. (2002). *Functional Analysis*. Wiley, New York.
- [64] LEHTINEN, M. S., PAIVARINTA, L. and SOMERSALO, E. (1989). Linear inverse problems for generalised random variables. *Inverse Probl.* **5** 599.
- [65] LE GLAND, F., MONBET, V. and TRAN, V. (2009). Large sample asymptotics for the ensemble Kalman filter Ph.D. thesis, INRIA.
- [66] LI, W., TAN, Z. and CHEN, R. (2013). Two-stage importance sampling with mixture proposals. *J. Amer. Statist. Assoc.* **108** 1350–1365.
- [67] LIANG, F., LIU, C. and CARROLL, R. J. (2007). Stochastic approximation in Monte Carlo computation. *J. Amer. Statist. Assoc.* **102** 305–320.
- [68] LIN, K., LU, S. and MATHÉ, P. (2015). Oracle-type posterior contraction rates in Bayesian inverse problems. *Inverse Probl. Imaging* **9**.
- [69] LINDLEY, D. V. and SMITH, A. F. M. (1972). Bayes estimates for the linear model. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **34** 1–41. MR0415861
- [70] LIU, J. S. (1996). Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Stat. Comput.* **6** 113–119.
- [71] LIU, J. S. (2008). *Monte Carlo Strategies in Scientific Computing*. Springer, New York. MR2401592
- [72] LIU, J. S. and CHEN, R. (1998). Sequential Monte Carlo methods for dynamic systems. *J. Amer. Statist. Assoc.* **93** 1032–1044. MR1649198
- [73] LU, S. and MATHÉ, P. (2014). Discrepancy based model selection in statistical inverse problems. *J. Complexity* **30** 290–308. MR3183335
- [74] MANDELBAUM, A. (1984). Linear estimators and measurable linear transformations on a Hilbert space. *Z. Wahrsch. Verw. Gebiete* **65** 385–397. MR0731228
- [75] MARTIN, J., WILCOX, L. C., BURSTEDDE, C. and GHATTAS, O. (2012). A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion. *SIAM J. Sci. Comput.* **34** A1460–A1487. MR2970260
- [76] MCLEISH, D. L. and O’BRIEN, G. L. (1982). The expected ratio of the sum of squares to the square of the sum. *Ann. Probab.* **10** 1019–1028. MR0672301
- [77] MÍGUEZ, J., CRISAN, D. and DJURIĆ, P. M. (2013). On the convergence of two sequential Monte Carlo methods for maximum a posteriori sequence estimation and stochastic global optimization. *Stat. Comput.* **23** 91–107.
- [78] MORZFELD, M., TU, X., WILKENING, J. and CHORIN, A. J. (2015). Parameter estimation by implicit sampling. *Commun. Appl. Math. Comput. Sci.* **10** 205–225.
- [79] MOSKOWITZ, B. and CAFLISCH, R. E. (1996). Smoothness and dimension reduction in quasi-Monte Carlo methods. *Math. Comput. Modelling* **23** 37–54. MR1398000
- [80] MUELLER, J. L. and SILTANEN, S. (2012). *Linear and Nonlinear Inverse Problems with Practical Applications*.

- Computational Science & Engineering* **10**. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA. [MR2986262](#)
- [81] OLIVER, D. S., REYNOLDS, A. C. and LIU, N. (2008). *Inverse Theory for Petroleum Reservoir Characterization and History Matching*. Cambridge Univ. Press, Cambridge.
- [82] OWEN, A. and ZHOU, Y. (2000). Safe and effective importance sampling. *J. Amer. Statist. Assoc.* **95** 135–143. [MR1803146](#)
- [83] PITT, M. K. and SHEPHARD, N. (1999). Filtering via simulation: Auxiliary particle filters. *J. Amer. Statist. Assoc.* **94** 590–599. [MR1702328](#)
- [84] RAY, K. (2013). Bayesian inverse problems with non-conjugate priors. *Electron. J. Stat.* **7** 2516–2549.
- [85] REBESCHINI, P. and VAN HANDEL, R. (2015). Can local particle filters beat the curse of dimensionality? *Ann. Appl. Probab.* **25** 2809–2866. [MR3375889](#)
- [86] REN, Y.-F. and LIANG, H.-Y. (2001). On the best constant in Marcinkiewicz–Zygmund inequality. *Statist. Probab. Lett.* **53** 227–233. [MR1841623](#)
- [87] SANZ-ALONSO, D. (2016). Importance sampling and necessary sample size: An information theory approach. Preprint. Available at [arXiv:1608.08814](#).
- [88] SLIVINSKI, L. and SNYDER, C. Practical estimates of the ensemble size necessary for particle filters.
- [89] SNYDER, C. (2011). Particle filters, the “optimal” proposal and high-dimensional systems. In *Proceedings of the ECMWF Seminar on Data Assimilation for Atmosphere and Ocean*.
- [90] SNYDER, C., BENGTSOON, T., BICKEL, P. and ANDERSON, J. (2008). Obstacles to high-dimensional particle filtering. *Mon. Weather Rev.* **136** 4629–4640.
- [91] SNYDER, C., BENGTSOON, T. and MORZFELD, M. (2015). Performance bounds for particle filters using the optimal proposal. *Mon. Weather Rev.* **143** 4750–4761.
- [92] SPANTINI, A., SOLONEN, A., CUI, T., MARTIN, J., TENORIO, L. and MARZOUK, Y. (2015). Optimal low-rank approximations of Bayesian linear inverse problems. *SIAM J. Sci. Comput.* **37** A2451–A2487.
- [93] SPIEGELHALTER, D. J., BEST, N. G., CARLIN, B. P. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 583–639. [MR1979380](#)
- [94] SPILIOPOULOS, K. (2013). Large deviations and importance sampling for systems of slow-fast motion. *Appl. Math. Optim.* **67** 123–161.
- [95] STUART, A. M. (2010). Inverse problems: A Bayesian perspective. *Acta Numer.* **19** 451–559. [MR2652785](#)
- [96] TAN, Z. (2004). On a likelihood approach for Monte Carlo integration. *J. Amer. Statist. Assoc.* **99** 1027–1036. [MR2109492](#)
- [97] TIERNEY, L. (1998). A note on Metropolis–Hastings kernels for general state spaces. *Ann. Appl. Probab.* **8** 1–9. [MR1620401](#)
- [98] VAN LEEUWEN, P. J. (2010). Nonlinear data assimilation in geosciences: An extremely efficient particle filter. *Q. J. R. Meteorol. Soc.* **136** 1991–1999.
- [99] VANDEN-EIJNDEN, E. and WEARE, J. (2012). Data assimilation in the low noise, accurate observation regime with application to the Kuroshio current. *Mon. Weather Rev.* **141**.
- [100] VOLLMER, S. J. (2013). Posterior consistency for Bayesian inverse problems through stability and regression results. *Inverse Probl.* **29** 125011.
- [101] WHITELEY, N., LEE, A. and HEINE, K. (2016). On the role of interaction in sequential Monte Carlo algorithms. *Bernoulli* **22** 494–529. [MR3449791](#)
- [102] ZHANG, T. (2002). Effective dimension and generalization of kernel learning. In *Advances in Neural Information Processing Systems* 454–461.
- [103] ZHANG, T. (2005). Learning bounds for kernel regression using effective data dimensionality. *Neural Comput.* **17** 2077–2098.
- [104] ZHANG, W., HARTMANN, C., WEBER, M. and SCHÜTTE, C. (2013). Importance sampling in path space for diffusion processes. *Multiscale Model. Simul.*