



Parallel sequential Monte Carlo for stochastic gradient-free nonconvex optimization

Ömer Deniz Akyildiz^{1,2} · Dan Crisan³ · Joaquín Míguez^{4,5}

Received: 3 June 2019 / Accepted: 8 July 2020 / Published online: 29 July 2020
© The Author(s) 2020

Abstract

We introduce and analyze a parallel sequential Monte Carlo methodology for the numerical solution of optimization problems that involve the minimization of a cost function that consists of the sum of many individual components. The proposed scheme is a stochastic zeroth-order optimization algorithm which demands only the capability to evaluate small subsets of components of the cost function. It can be depicted as a bank of samplers that generate particle approximations of several sequences of probability measures. These measures are constructed in such a way that they have associated probability density functions whose global maxima coincide with the global minima of the original cost function. The algorithm selects the best performing sampler and uses it to approximate a global minimum of the cost function. We prove analytically that the resulting estimator converges to a global minimum of the cost function almost surely and provide explicit convergence rates in terms of the number of generated Monte Carlo samples and the dimension of the search space. We show, by way of numerical examples, that the algorithm can tackle cost functions with multiple minima or with broad “flat” regions which are hard to minimize using gradient-based techniques.

Keywords Sequential Monte Carlo · Stochastic optimization · Nonconvex optimization · Gradient-free optimization · Sampling

An important part of this work was carried out when Ö. D. A. was visiting Department of Mathematics, Imperial College London. This work was partially supported by *Agencia Estatal de Investigación* of Spain (RTI2018-099655-B-I00 CLARA), and the regional government of Madrid (program CASICAM-CM S2013/ICE-2845). The work of the second author has been partially supported by a UC3M-Santander Chair of Excellence grant held at the Universidad Carlos III de Madrid.

✉ Ömer Deniz Akyildiz
omer.akyildiz@warwick.ac.uk

Dan Crisan
d.crisan@imperial.ac.uk

Joaquín Míguez
joaquin.miguez@uc3m.es

¹ The Alan Turing Institute, London, UK

² University of Warwick, Coventry, UK

³ Imperial College London, London, England

⁴ Instituto de Investigación Sanitaria Gregorio Marañón, Madrid, Spain

⁵ Instituto de Investigación Sanitaria Gregorio Marañón, Universidad Carlos III de Madrid, Leganes, Madrid, Spain

1 Introduction

In signal processing and machine learning, optimization problems of the form

$$\min_{\theta \in \Theta} f(\theta) = \sum_{i=1}^n f_i(\theta), \quad (1.1)$$

where $\Theta \subset \mathbb{R}^d$ is the d -dimensional compact search space, have attracted significant attention in recent years for problems where n is very large. Such problems often arise in big data settings, e.g., when one needs to estimate parameters given a large number of observations (Bottou et al. 2018).

Because of their efficiency, the optimization community has focused mainly on stochastic gradient-based methods (Robbins and Monro 1951; Duchi et al. 2011; Kingma and Ba 2014) (see Bottou et al. (2018) for a recent review of the field) where an estimate of the gradient is obtained using a randomly selected subsample of the gradients of the component functions [the f_i 's in Eq. (1.1)] at each iteration. The resulting estimate is then used to perform a stochastic descent step. The majority of these stochastic gradient methods con-

struct the subsamples using sampling with replacement to obtain unbiased estimates of the gradient. The latter can then be seen as a noisy gradient estimate with additive, zero-mean noise. In practice, however, there are schemes that subsample the data set without replacement (hence producing biased gradient estimators) and it has been argued that such methods can attain better numerical performance (Gürbüzbalaban et al. 2015; Shamir 2016).

The gradient information may not be always available, however, due to different reasons. For example, in an engineering application, the system to be optimized might be a black-box, e.g., a piece of closed software code with free parameters, which can be *evaluated* but cannot be differentiated (Nesterov and Spokoiny 2011). In these cases, one needs to use a *gradient-free* optimization scheme, meaning that the scheme must rely only on function evaluations, rather than any sort of actual gradient information. *Classical* gradient-free optimization methods have attracted significant interest over the past decades (Appel et al. 2004; Spall 2005; Mariño and Míguez 2007; Conn et al. 2009). These methods proceed either by a random search [which is based on evaluating the cost function at random points and update the parameter whenever a descent in the function evaluation is achieved (Spall 2005)], or by constructing a numerical (finite-difference type) approximation of the gradient that can be used to take a descent step (Nesterov and Spokoiny 2011).

Such methods are not applicable, however, if one can only obtain noisy function evaluations or one can only evaluate certain subsets of component functions in a problem like (1.1). In this case, since the function evaluations are not exact, direct random search methods cannot be used reliably and it is only recently that some authors have described how to compute finite-difference approximations of the gradient (Wibisono et al. 2012; Ghadimi and Lan 2013; Chen and Wild 2015; Bach and Perchet 2016). Also in recent years, evolutionary methods, based on the mutation, recombination and selection of samples, have been suggested for the approximation of gradients. The resulting optimization algorithms, termed evolutionary strategies (ES) have been applied within reinforcement learning schemes (Salimans et al. 2017; Wierstra et al. 2014; Hansen and Ostermeier 2001; Morse and Stanley 2016).

However, when the cost function has multiple minima or has some regions where the gradient vanishes, gradient-based methods may suffer from poor numerical performance. In particular, the optimizer can get stuck in a local minimum easily, due to its reliance on gradient approximations. Moreover, when the gradient contains little information about any minimum (e.g., in flat regions), gradient-free stochastic optimizers (as well as perfect gradient schemes) can suffer from slow convergence.

Model-based random-search methods (Hu et al. 2012), which use probabilistic models of various types in order to

speed up the search procedure, have been investigated in order to address problems where gradients cannot be approximated or simply turn out ineffective. The latter include classical algorithms such as simulated annealing (SA) (Kirkpatrick et al. 1983), Monte Carlo expectation maximization (EM) (Robert and Casella 2004) and other Markov chain Monte Carlo (MCMC) based methods (Pereyra et al. 2015). The class of model-based random search schemes also encompasses sequential Monte Carlo (SMC) techniques, e.g., SMC implementations of SA (Zhou and Chen 2013) and several optimization algorithms that mimic standard particle filters (Zhou et al. 2013; Liu et al. 2016). Let us note that most of the latter MCMC- and SMC-based procedures can be cast within the class of SMC samplers described in Del Moral et al. (2006), albeit with a target distribution which is sometimes implicitly defined in order to satisfy certain properties related to the objective function (Zhou et al. 2013). Nevertheless, these optimization techniques are generally designed to be used in problems where the objective function can be evaluated exactly and their extension to stochastic optimization is not straightforward, neither from the point of view of practical performance nor in terms of theoretical convergence guarantees.

Some authors have also explored the duality between optimization and probability theory, in a way that potentially enables the use of general computational inference algorithms for solving optimization problems. While in model-based optimization the emphasis is put on the algorithms (e.g., how to use MCMC methods in Pereyra et al. (2015) or particle filters in Liu et al. (2016), for optimization), in this line of research the emphasis is in converting the optimization problem into an equivalent inference problem, which can then be tackled with *any* suitable inference algorithm. A rigorous mathematical treatment of the topic can be found in Del Moral and Doisy (1999), while Ikonen et al. (2005) and Míguez et al. (2013) address the problem from a methodological viewpoint. Again, these contributions deal with problems where the objective function can be computed deterministically and exactly, though.

The stochastic setting, where it is only possible to compute noisy evaluations of $f(\theta)$, is harder and the bibliography is limited in comparison with the deterministic setup. The recent survey in Homem-de Mello and Bayraksan (2014) covers various gradient-based Monte Carlo procedures, however it addresses a different class of stochastic optimization problems where the cost function itself is defined as an expectation, rather than a finite-sum as in (1.1). Existing model-based random methods for stochastic optimization include MCMC-based samplers which target a probability density function (pdf) matched to the objective function in (1.1) (meaning that the maxima of the pdf coincide with the minima of $f(\theta)$) (Welling and Teh 2011; Chen et al. 2016). Such schemes, however, also rely on the computation of noisy

gradients. Other MCMC-based methods (see, e.g., Alquier et al. (2016) which employs noisy Metropolis steps) do not require gradients, yet these techniques have been primarily designed and investigated as sampling algorithms, rather than optimization methods. Similarly, an adaptive importance sampler for a target pdf matched to $f(\theta)$ is reported in Akyildiz et al. (2017). This method uses subsampling to compute noisy weights, but the technique lacks any theoretical guarantees and does not address the problem optimization directly either. A particle filtering algorithm for stochastic global optimization has been proposed by Stinis (2012). The method is intuitive, simple to implement and has been shown to work efficiently in some simple examples, however the contribution of Stinis (2012) is strictly methodological: There is no analysis of performance and no theoretical guarantees.

In this paper, we propose a *parallel sequential Monte Carlo optimizer* (PSMCO) to minimize cost functions with the finite-sum structure of problem (1.1). The PSMCO is a zeroth-order stochastic optimization algorithm, in the sense that it only uses evaluations of small batches of individual components $f_i(\theta)$ in (1.1). In particular, it does not require the computation or approximation of gradients. The proposed scheme proceeds by constructing parallel samplers, each of which aims at minimizing the same cost function $f(\theta)$. Each sampler performs subsampling without replacement to obtain its mini-batches of individual components and processes each component only once. Using these mini-batches, the PSMCO constructs potential functions, propagates samples via a jittering scheme (Crisan and Míguez 2018) and selects samples by applying a weighting-resampling procedure. The communication between parallel samplers is only necessary when a joint estimate of the minimum is required. In this case, the best performing sampler is selected and the minimum is estimated.

We analytically prove that the estimate provided by each sampler converges almost surely to a global minimum of the cost function and provide explicit convergence rates in terms of the number of Monte Carlo samples generated by the algorithm. This type of analysis goes beyond standard results for particle filters: It tackles the problem of stochastic optimization directly and it yields stronger theoretical guarantees compared to other stochastic optimization methods in the literature. In particular, we obtain error bounds for the solution of problem (1.1) that hold almost surely (a.s.) and vanish at a rate $\mathcal{O}\left(N^{-\frac{1}{2(d+1)}}\right)$, where N is the number of Monte Carlo samples and d is the dimension of the search space Θ . This is in contrast to the usual results for random search methods in the literature, which are purely asymptotic and do not provide any rates (Appel et al. 2004; Míguez 2010; Hu et al. 2012; Zhou and Chen 2013; Zhou et al. 2013). Let us also remark the difference between the proposed scheme

and the SMC-based schemes in Míguez et al. (2013) where the authors partitioned the parameter vector and modeled it as a dynamical system, an approach that cannot be used in the more general setup of (1.1) because each individual function f_i depends on the complete vector θ . The PSMCO algorithm, in turn, is explicitly designed to provide an estimate of the full parameter θ at each iteration.

The main contribution of this paper includes the theoretical analysis of the proposed PSMCO scheme and its numerical demonstration on three problems where classical stochastic optimization methods (especially gradient-based algorithms) struggle to perform. The paper is organized as follows. After a brief survey of the relevant notation (below), we lay out the relationship between Bayesian inference and optimization in Sect. 2. Then, we develop a sequential Monte Carlo scheme in Sect. 3. In Sect. 4, we analyze this scheme and investigate its theoretical properties. We present some numerical results in Sect. 5 and make some concluding remarks in Sect. 6.

Notation

For $n \in \mathbb{N}$, we denote $[n] = \{1, \dots, n\}$. The space of bounded functions on the parameter space $\Theta \subset \mathbb{R}^d$ is denoted as $B(\Theta)$. The set of continuous and bounded real functions on Θ is denoted $C_b(\Theta)$. The family of Borel subsets of Θ is denoted with $\mathcal{B}(\Theta)$. The set of probability measures on the measurable space $(\Theta, \mathcal{B}(\Theta))$ is denoted $\mathcal{P}(\Theta)$. Given $\varphi \in B(\Theta)$ and $\pi \in \mathcal{P}(\Theta)$, the integral of φ with respect to (w.r.t.) π is written as

$$(\varphi, \pi) = \int_{\Theta} \varphi(\theta)\pi(d\theta).$$

Given a Markov kernel $\kappa : \mathcal{B}(\Theta) \times \Theta \mapsto [0, 1]$, we denote $\kappa\pi(d\theta) = \int \kappa(d\theta|\theta')\pi(d\theta')$. If $\varphi \in B(\Theta)$, then $\|\varphi\|_{\infty} = \sup_{\theta \in \Theta} |\varphi(\theta)| < \infty$.

Let $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}^* \times \dots \times \mathbb{N}^*$, where $\mathbb{N}^* = \mathbb{N} \cup \{0\}$, be a multi-index. We define the partial derivative operator D^{α} as

$$D^{\alpha}h = \frac{\partial^{\alpha_1} \dots \partial^{\alpha_d} h}{\partial \theta_1^{\alpha_1} \dots \partial \theta_d^{\alpha_d}}$$

for a sufficiently differentiable function $h : \mathbb{R}^d \rightarrow \mathbb{R}$. We use $|\alpha| = \sum_{i=1}^d \alpha_i$ to denote the order of the derivative. Finally, the notation $\lfloor x \rfloor$ indicates the floor function for a real number x , which returns the biggest integer $k \leq x$.

2 Stochastic optimization as inference

In this section, we describe how to construct a sequence of probability distributions that can be linked to the solution of problem (1.1). Let $\pi_0 \in \mathcal{P}(\Theta)$ be the initial element of the sequence. We construct the rest of the sequence recursively as

$$\pi_t(d\theta) = \pi_{t-1}(d\theta) \frac{G_t(\theta)}{\int_{\Theta} G_t(\theta)\pi_{t-1}(d\theta)}, \quad \text{for } t \geq 1, \quad (2.1)$$

where the maps $G_t : \Theta \mapsto \mathbb{R}_+$ are termed *potential functions* (Del Moral 2004). The key idea is to associate these potentials $(G_t)_{t \geq 1}$ with mini-batches of individual components of the cost function (subsets of the f_i 's) in order to construct a sequence of measures $\pi_0, \pi_1, \dots, \pi_T$ such that (for a prescribed value of T) the global maxima of the density of π_T match the global minima of $f(\theta)$. We remark that the measures π_1, \dots, π_T are all absolutely continuous w.r.t π_0 if the potential functions $G_t, t = 1, \dots, T$, are bounded.

To construct the potentials, we use mini-batches consisting of K individual functions f_i for each iteration t . To be specific, we randomly select subsets of indices $\mathcal{I}_t, t = 1, \dots, T$, by drawing uniformly from $\{1, \dots, n\}$ without replacement. Each subset has $|\mathcal{I}_t| = K$ elements, in such a way that we obtain T subsets satisfying $\bigcup_{i=1}^T \mathcal{I}_i = [n]$ and $\mathcal{I}_i \cap \mathcal{I}_j = \emptyset$ when $i \neq j$. Finally, we define the potential functions $(G_t)_{t \geq 1}$ as

$$G_t(\theta) = \exp \left(- \sum_{i \in \mathcal{I}_t} f_i(\theta) \right), \quad t = 1, \dots, T. \quad (2.2)$$

In the sequel, we provide a result that establishes a precise connection between the optimization problem in (1.1) and the sequence of probability measures defined in (2.1), provided that Assumption 1 below is satisfied.

Assumption 1 The functions in the sequence $(G_t)_{t \geq 1}$ are positive and bounded, i.e.,

$$G_t(\theta) > 0 \quad \forall \theta \in \Theta \quad \text{and} \quad G_t \in \mathcal{B}(\Theta).$$

Next, we show the relationship between the minima of $f(\theta)$ and the maxima of $\frac{d\pi_T}{d\pi_0}$.

Proposition 1 Assume that the potentials are selected as in (2.2) for $1 \leq t \leq T$, with $\mathcal{I}_i \cap \mathcal{I}_j = \emptyset$ and $\bigcup_i \mathcal{I}_i = [n]$. Let π_T be the T -th probability measure constructed by means of recursion (2.1). If Assumption 1 holds and $\pi_0 \in \mathcal{P}(\Theta)$, then

$$\operatorname{argmax}_{\theta \in \Theta} \frac{d\pi_T}{d\pi_0}(\theta) = \operatorname{argmin}_{\theta \in \Theta} \sum_{i=1}^n f_i(\theta),$$

where $\frac{d\pi_T}{d\pi_0}(\theta) : \Theta \rightarrow \mathbb{R}_+$ denotes the Radon–Nikodym derivative of π_T w.r.t. the prior measure π_0 .

Proof See Appendix A.1. □

For conciseness, we abuse the notation and use $\pi(\theta), \theta \in \Theta$, to indicate the pdf associated to a probability measure $\pi(d\theta)$. The two objects are distinguished clearly by the context (e.g., for an integral (φ, π) , π necessarily is a measure) but also by their arguments. The probability measure $\pi(\cdot)$ takes arguments $d\theta$ or $A \in \mathcal{B}(\Theta)$, while the pdf $\pi(\theta)$ is a function $\Theta \rightarrow [0, \infty)$.

Remark 1 Notice that when π_0 is a uniform probability measure on Θ , we simply have

$$\pi_T(\theta) \propto \exp \left(- \sum_{i=1}^n f_i(\theta) \right), \quad \theta \in \Theta.$$

where $\pi_T(\theta)$ denotes the pdf (w.r.t. Lebesgue measure) of the measure $\pi_T(d\theta)$. □

Remark 2 Moreover, if we choose

$$\pi_0(\theta) \propto \exp(-f_1(\theta)) \quad (2.3)$$

and select index subsets such that $\bigcup_{i=1}^T \mathcal{I}_i = \{2, \dots, n\}$ then we also obtain

$$\pi_T(\theta) \propto \exp \left(- \sum_{i=1}^n f_i(\theta) \right), \quad \text{for } \theta \in \Theta.$$

When a Monte Carlo scheme is used to realize recursion (2.1), the use of a prior of the form (2.3) requires the ability to sample from it. □

In summary, if we can construct the sequence described by (2.1), then we can replace the minimization problem of $f(\theta)$ in (1.1) by the maximization of a pdf. This relationship was exploited in a Gaussian setting in Akyildiz et al. (2018), i.e., the special case of a Gaussian prior π_0 and log-quadratic potentials $(G_t)_{t \geq 1}$ (Gaussian likelihoods), which makes it possible to implement recursion (2.1) analytically. The solution of this special case can be shown to match a well-known stochastic optimization algorithm, called the incremental proximal method (Bertsekas 2011), with a variable-metric. However, for general priors and potentials, it is not possible to analytically construct (2.1) and maximize $\pi_T(\theta)$. For this reason, we propose a simulation method to approximate the recursion (2.1) and solve $\operatorname{argmax}_{\theta \in \Theta} \frac{d\pi_T}{d\pi_0}(\theta)$.

3 The algorithm

In this section, we first describe a sampler to simulate from the distributions defined by recursion (2.1). We then describe an algorithm which runs these samplers in parallel. The parallelization here is not primarily motivated by the computational gain (although it can be substantial). We have empirically found that non-interacting parallel samplers are able to keep track of multiple minima better than a single “big” sampler. For this reason, we will not focus on demonstrating computational gains in the experimental section. Rather, we will discuss what parallelization brings in terms of providing better estimates.

We consider M workers (corresponding to M samplers). Specifically, each worker sees a different configuration of the dataset, i.e., the m -th worker constructs a distinct sequence of index sets $(\mathcal{I}_t^{(m)})_{t \geq 1}$ which determine the mini-batches sampled from the full set of individual components. Having obtained different mini-batches which are randomly constructed, each worker then constructs different potentials $(G_t^{(m)})_{t \geq 1}$, where $G_t^{(m)}(\cdot) = \exp\{-\sum_{i \in \mathcal{I}_t} f_i(\cdot)\}$, as described in the previous section.

The m -th worker, therefore, aims at estimating a specific sequence of probability measures $\pi_t^{(m)}$, for $m \in \{1, \dots, M\}$. We denote the particle approximation of the posterior $\pi_t^{(m)}$ at time t as

$$\pi_t^{(m),N}(\mathrm{d}\theta) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta^{(i,m)}}(\mathrm{d}\theta),$$

where $\delta_{\theta'}(\mathrm{d}\theta)$ is the unit delta measure located at $\theta' \in \Theta$. Overall, the algorithm retains M probability distributions. Note that these distributions are different for each $t < T$, as they depend on different potentials, but $\pi_T^{(m)} = \pi_T$ for all workers because $\bigcup_{i=1}^T \mathcal{I}_i^{(m)} = [n]$ for every m .

One iteration of the algorithm on a local worker m can be described as follows. Assume the worker has computed the probability measure $\pi_{t-1}^{(m),N}$ using the particle system $\{\theta_{t-1}^{(m,i)}\}_{i=1}^N$. First, we use a jittering kernel $\kappa(\mathrm{d}\theta|\theta_{t-1})$ (a Markov kernel on Θ) to modify the particles (Crisan and Miguez 2018) (see Sect. 3.1 for the precise definition of $\kappa(\cdot|\cdot)$). The idea is to jitter a subset of the particles in order to modify and propagate them into better regions of Θ with higher probability density and lower cost. The particles are jittered by sampling,

$$\hat{\theta}_t^{(i,m)} \sim \kappa(\cdot|\theta_{t-1}^{(i,m)}) \quad \text{for } i = 1, \dots, N.$$

Algorithm 1 Sampler on a local node m

1: Sample $\theta_0^{(i,m)} \sim \pi_0$ for $i = 1, \dots, N$.

2: **for** $t \geq 1$ **do**

3: Jitter by generating samples

$$\hat{\theta}_t^{(i,m)} \sim \kappa(\mathrm{d}\theta|\theta_{t-1}^{(i,m)}) \quad \text{for } i = 1, \dots, N.$$

4: Compute normalized weights,

$$w_t^{(i,m)} = \frac{G_t^{(m)}(\hat{\theta}_t^{(i,m)})}{\sum_{i=1}^N G_t^{(m)}(\hat{\theta}_t^{(i,m)})} \quad \text{for } i = 1, \dots, N.$$

5: Resample by drawing N i.i.d. samples,

$$\theta_t^{(i,m)} \sim \hat{\pi}_t^{(m),N}(\mathrm{d}\theta) := \sum_{i=1}^N w_t^{(i,m)} \delta_{\hat{\theta}_t^{(i,m)}}(\mathrm{d}\theta),$$

for $i = 1, \dots, N$.

6: **end for**

Note that the jittering kernel may be designed so that it only modifies a subset of particles (again, see Sect. 3.1 for details). Next, we compute weights for the new set of particles $\{\hat{\theta}_t^{(i,m)}\}_{i=1}^N$ according to the t -th potential, namely

$$w_t^{(i,m)} = \frac{G_t^{(m)}(\hat{\theta}_t^{(i,m)})}{\sum_{i=1}^N G_t^{(m)}(\hat{\theta}_t^{(i,m)})} \quad \text{for } i = 1, \dots, N.$$

Remark 3 The particle weights can be made proportional to the potentials alone, i.e., $w_t^{(i,m)} \propto G_t^{(m)}(\hat{\theta}_t^{(i,m)})$, as long as the jittering kernels satisfy Assumption 2 in Sect. 3.1. Under mild assumptions, Algorithm 1 converges with standard error rates $\mathcal{O}(N^{-\frac{1}{2}})$, as proved in Sect. 4.

After obtaining weights, each worker performs a resampling step where for $i = 1, \dots, N$, we set $\theta_t^{(i,m)} = \hat{\theta}_t^{(i,k)}$ for $k \in \{1, \dots, N\}$ with probability $w_t^{(i,m)}$. The procedure just described corresponds to a simple multinomial resampling scheme, but other standard methods can be applied as well (Douc et al. 2005). We denote the resulting probability measure constructed at the t -th iteration of the m -th worker as

$$\pi_t^{(m),N}(\mathrm{d}\theta) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_t^{(i,m)}}(\mathrm{d}\theta).$$

The full procedure for the m -th worker is outlined in Algorithm 1. In Sect. 3.1, we elaborate on the selection of the jittering kernels and in Sect. 3.2, we detail the scheme for estimating a global minimum of $f(\theta)$ from the set of random measures $\{\pi_t^{(m),N}\}_{m=1}^M$.

3.1 Jittering kernel

The jittering kernel constitutes one of the key design choices of the proposed algorithm. Following Crisan and Miguez (2018), we put the following assumption on the kernel κ .

Assumption 2 The Markov kernel κ satisfies

$$\sup_{\theta' \in \Theta} \int_{\Theta} |\varphi(\theta) - \varphi(\theta')| \kappa(d\theta|\theta') \leq \frac{c_{\kappa} \|\varphi\|_{\infty}}{\sqrt{N}}$$

for any $\varphi \in B(\Theta)$ and some constant $c_{\kappa} < \infty$ independent of N .

In this paper, we use kernels of form

$$\kappa(d\theta|\theta') = (1 - \epsilon_N)\delta_{\theta'}(d\theta) + \epsilon_N\tau(d\theta|\theta'), \tag{3.1}$$

where $\epsilon_N \leq \frac{1}{\sqrt{N}}$, which satisfy Assumption 2 (Crisan and Miguez 2018). The kernel τ can be rather simple, such as a multivariate Gaussian or multivariate-t distribution centered around $\theta' \in \Theta$. Other choices of τ are possible as well.

Remark 4 The design of the kernel as a centered Gaussian or a multivariate-t distribution around θ' may not guarantee the propagation of samples into better (lower cost) regions. In this case, the weighting-and-resampling procedure naturally tends to keep and replicate the particles that attain a lower cost. However, the jittering kernel can also be designed to accelerate the optimization process. In particular, our setup allows for the use of gradient estimators [such as finite-difference schemes (Nesterov and Spokoiny 2011) or nudging steps (Akyildiz and Míguez 2020)] in the jittering kernel to accelerate the propagation of samples into lower-cost regions.

3.2 Estimating the global minima of $f(\theta)$

In order to estimate the global minima of $f(\theta)$, we first assess the performance of the samplers run by each worker. A typical performance measure is the *marginal likelihood estimate* resulting from $\pi_t^{(m),N}$. After choosing the worker which has attained the highest marginal likelihood (say the m_0 -th worker), we estimate a minimum of $f(\theta)$ by selecting the particle $\theta_t^{(i,m)}$ that yields the highest density $\pi_t^{(m_0)}(\theta_t^{(i,m_0)})$.

To be precise, let us start by denoting the *incremental* marginal likelihood associated to $\pi_t^{(m)}$ and its estimate $\pi_t^{(m),N}$ as $Z_{1:t}^{(m)}$ and $Z_{1:t}^{(m),N}$, respectively. They can be explicitly obtained by first computing

$$\begin{aligned} Z_t^{(m)} &= \int G_t^{(m)}(\theta)\hat{\pi}_t^{(m)}(d\theta) \\ &\approx \frac{1}{N} \sum_{i=1}^N G_t^{(m)}(\hat{\theta}_t^{(i,m)}) := Z_t^{(m),N} \end{aligned}$$

Algorithm 2 PSMCO

- 1: Sample $\theta_0^{(i,m)} \sim \pi_0$ for $i = 1, \dots, N$.
 - 2: **for** $t \geq 1$ **do**
 - 3: **for** $m = 1, \dots, M$ **do**
 - 4: Jitter N local particles (step 3 of Algorithm 1).
 - 5: Update the marginal likelihood

 - $Z_{1:t}^{(m),N} = Z_{1:t-1}^{(m),N} \times Z_t^{(m),N}$
 - where
 - $Z_t^{(m),N} = \frac{1}{N} \sum_{i=1}^N G_t^{(m)}(\hat{\theta}_t^{(i,m)})$.

 - 6: Compute weights (step 4 of Algorithm 1).
 - 7: Resample (step 5 of Algorithm 1).
 - 8: **end for**
 - 9: **if** an estimate of the solution of problem (1.1) is needed at time t **then**
 - 10: Choose
 - $m_t^* = \operatorname{argmax}_{m \in \{1, \dots, M\}} Z_{1:t}^{(m),N}$
 - 11: Estimate
 - $\theta_t^{*,N} = \operatorname{argmax}_{i \in \{1, \dots, N\}} p_t^{(m_t^*),N}(\theta_t^{(i,m_t^*)})$.
 - 12: **end if**
 - 13: **end for**
-

and then updating the running products

$$Z_{1:t}^{(m)} = Z_t^{(m)} Z_{1:t-1}^{(m)} = \prod_{k=1}^t Z_k^{(m)}$$

and

$$Z_{1:t}^{(m),N} = Z_t^{(m),N} Z_{1:t-1}^{(m),N} = \prod_{k=1}^t Z_k^{(m),N}.$$

The quantity $Z_{1:t}^{(m)}$ is a local performance index that keeps track of the “quality” of the m -th particle system $\{\theta_t^{(i,m)}\}_{i=1}^N$ (Elvira et al. 2017) and, hence, we use $\{Z_{1:t}^{(m),N}\}_{m=1}^M$ to determine the best performing worker¹. Given the index of the best performing sampler, which is given by

$$m_t^* = \operatorname{argmax}_{m \in \{1, \dots, M\}} Z_{1:t}^{(m),N},$$

¹ If we interpret each sequence of index sets $(Z_t^{(m)})_{t \geq 1}$ as a different model (since different indices yield different potentials) then $Z_{1:t}^{(m)}$ is the Bayesian evidence in favour of model m . Let us note, however, that $Z_{1:t}^{(m)}$ is not a direct indicator of the performance of worker m as an optimizer. The fact that $Z_{1:t}^{(m_1)} > Z_{1:t}^{(m_2)}$ does not necessarily imply that the estimate of θ^* computed from worker m_1 is quantifiably better than the estimate computed from worker m_2 .

we obtain a maximum-a-posteriori (MAP) estimator,

$$\theta_t^{*,N} = \operatorname{argmax}_{i \in \{1, \dots, N\}} p_t^{(m_i^*),N}(\theta^{(i,m_i^*)}), \tag{3.2}$$

where $p_t^{(m_i^*),N}(\theta)$ is the kernel density estimator (Silverman 1998; Wand and Jones 1994) described in Remark 5. Note that we do *not* construct the entire density estimator and maximize it. Since this operation is performed locally on the particles from the best performing sampler, it involves $\mathcal{O}(N^2)$ operations, where N is the number of particles on a single worker, which is much smaller than the total number MN . The full procedure is outlined in Algorithm 2.

Remark 5 Let $k : \Theta \rightarrow (0, \infty)$ be a bounded pdf with zero mean and finite second-order moment, i.e., we have $\int_{\Theta} \|\theta\|_2^2 k(\theta) d\theta < \infty$. We can use the particle system $\{\theta_t^{(i,m)}\}_{i=1}^N$ and the pdf $k(\cdot)$ to construct the kernel density estimator (KDE) of $\pi_t^{(m)}(\theta)$ as

$$\begin{aligned} p_t^{(m),N}(\theta) &= \frac{1}{N} \sum_{i=1}^N k(\theta - \theta_t^{(i,m)}) \\ &= (k^\theta, \pi_t^{(m),N}), \end{aligned} \tag{3.3}$$

where $k^\theta(\theta') = k(\theta - \theta')$. Note that $p_t^{(m),N}(\theta)$ is not a standard KDE because the particles $\{\theta_t^{(i,m)}\}_{i=1}^N$ are not i.i.d. samples from $\pi_t^{(m)}(\theta)$. Equation (3.3), however, suggests that the estimator, $p_t^{(m),N}(\theta)$ converges when the approximate measure $\pi_t^{(m),N}$ does. See Crisan and Míguez (2014) for an analysis of particle KDE’s. \square

4 Analysis

In this section, we provide some basic theoretical guarantees for Algorithm 2. In particular, we prove results regarding a sampler on a single worker m . To ease the notation, we skip the superscript (m) in the rest of this section and simply note that results presented below hold for every $m \in \{1, \dots, M\}$. All proofs are deferred to the Appendix.

When constrained to a single worker m , the approximation π_t^N is provably convergent. In particular, we have the following results that hold for every worker $m = 1, \dots, M$.

Theorem 1 *If the sequence $(G_t)_{t \geq 1}$ satisfies Assumption 1 and the jittering kernels satisfy Assumption 2, then, for any $\varphi \in B(\Theta)$, we have*

$$\left\| (\varphi, \pi_t) - (\varphi, \pi_t^N) \right\|_p \leq \frac{c_{t,p} \|\varphi\|_\infty}{\sqrt{N}}$$

for every $t = 1, \dots, T$ and for any $p \geq 1$, where $c_{t,p} > 0$ is a constant independent of N .

Proof See Appendix A.2. \square

Theorem 1 states that the samplers on local workers converge to their correct probability measures (for each m) with rate $\mathcal{O}(1/\sqrt{N})$, which is standard for Monte Carlo methods. Next we provide an upper bound for the random error $|(\varphi, \pi_t) - (\varphi, \pi_t^N)|$.

Corollary 1 *Under the assumptions of Theorem 1, for every $\varphi \in B(\Theta)$, we have*

$$\left| (\varphi, \pi_t^N) - (\varphi, \pi_t) \right| \leq \frac{U_{t,\delta}}{N^{\frac{1}{2}-\delta}}, \quad \text{and } 1 \leq t \leq T,$$

where $U_{t,\delta}$ is an a.s. finite random variable and $0 < \delta < \frac{1}{2}$ is an arbitrary constant independent of N . In particular,

$$\lim_{N \rightarrow \infty} (\varphi, \pi_t^N) = (\varphi, \pi_t) \quad \text{a.s.} \tag{4.1}$$

for any $\varphi \in B(\Theta)$.

Proof See Appendix A.3. \square

This result ensures that the *random error* made by the estimators vanishes as $N \rightarrow \infty$. Moreover, it provides us with a rate $\mathcal{O}(1/\sqrt{N})$ since the constant $\delta > 0$ can be chosen arbitrarily small.

These results are important because they enable us to analyze the properties of the kernel density estimators constructed using the samples at each worker. In order to be able to do so, we need to impose regularity conditions on the sequence of densities $\pi_t(\theta)$ and the kernels we use to approximate them.

Assumption 3 For every $\theta \in \Theta$, the derivatives $D^\alpha \pi_t(\theta)$ exist and they are Lipschitz continuous, i.e., there is a constant $L_{\alpha,t} > 0$ such that

$$|D^\alpha \pi_t(\theta) - D^\alpha \pi_t(\theta')| \leq L_{\alpha,t} \|\theta - \theta'\|$$

for all $\theta, \theta' \in \Theta$, $t = 1, \dots, T$ and for all $\alpha = (\alpha_1, \dots, \alpha_d) \in \{0, 1\}^d$.

Note that for $\alpha = (0, \dots, 0)$ it is not hard to relate Assumption 3 directly to the cost function as we do in the following proposition.

Proposition 2 *Assume that we define the incremental cost functions*

$$F_t(\theta) = \sum_{i \in \bigcup_{k=1}^t \mathcal{I}_k} f_i(\theta)$$

and there exists some ℓ_t such that

$$|F_t(\theta) - F_t(\theta')| \leq \ell_t \|\theta - \theta'\|,$$

i.e., F_t is Lipschitz. Assume there exists $F_t^* = \min_{\theta \in \Theta} F_t(\theta)$

such that $|F_t^*| < \infty$ and recall that $\pi_t(\theta) \propto \exp(-F_t(\theta))$. Then we have the following inequality,

$$|\pi_t(\theta) - \pi_t(\theta')| \leq \frac{\ell_t \exp(-F_t^*)}{Z_{\pi_t}} \|\theta - \theta'\|_2$$

where $Z_{\pi_t} = \int_{\Theta} \exp(-F_t(\theta)) d\theta$.

Proof See Appendix A.4. □

Next, we state assumptions on the kernel k . We first note that the kernels in practice are defined with a bandwidth parameter $h \in \mathbb{R}_+$. In particular, given a kernel k , we can define scaled kernels k_h as

$$k_h(\theta) = h^{-d} k(h^{-1}\theta), \quad h > 0,$$

where, we recall, d is the dimension of the parameter vector θ . Hence, given k we define a family of kernels $\{k_h, h \in \mathbb{R}_+\}$.

Assumption 4 The kernel $k : \Theta \rightarrow (0, \infty)$ is a zero-mean bounded pdf, i.e., $k(\theta) \geq 0 \forall \theta \in \Theta$ and $\int k(\theta) d\theta = 1$. The second moment of this density is bounded, i.e., $\int_{\Theta} \|\theta\|^2 k(\theta) d\theta < \infty$. Finally, $D^\alpha k \in C_b(\Theta)$, i.e., $\|D^\alpha k\|_\infty < \infty$ for any $\alpha \in \{0, 1\}^d$.

Remark 6 We note that Assumption 4 implies that $D^\alpha k_h \in C_b(\Theta)$ and we have $\|D^\alpha k_h\|_\infty = \frac{1}{h^{d+|\alpha|}} \|D^\alpha k\|_\infty$ for any $h > 0$ and $\alpha \in \{0, 1\}^d$. □

We denote the kernel density estimator defined using a scaled kernel k_h and the empirical measure π_t^N as $\mathbf{p}_t^{h,N}(\theta)$. In particular, given a normalized kernel (a pdf) $k : \Theta \rightarrow (0, \infty)$, satisfying the assumptions in Assumption 4, we can construct the KDE

$$\mathbf{p}_t^{h,N}(\theta) = (k_h^\theta, \pi_t^N).$$

where $k_h^\theta(\theta') = k_h(\theta - \theta')$ (see Remark 5). Now, we are ready to state the main results regarding the kernel density estimators, adapted from Crisan and Míguez (2014).

Theorem 2 Choose

$$h = \left[N^{\frac{1}{2(d+1)}} \right]^{-1} \tag{4.2}$$

and denote $\mathbf{p}_t^N(\theta) = \mathbf{p}_t^{h,N}(\theta)$ (since $h = h(N)$). If Assumptions 1, 2, 3 and 4 hold, and Θ is compact, then

$$\sup_{\theta \in \Theta} |\mathbf{p}_t^N(\theta) - \pi_t(\theta)| \leq \frac{V_\varepsilon}{N^{\frac{1-\varepsilon}{2(d+1)}}} \tag{4.3}$$

where $V_\varepsilon \geq 0$ is an a.s. finite random variable and $0 < \varepsilon < 1$ is a constant, both of which are independent of N and θ . In particular,

$$\lim_{N \rightarrow \infty} \sup_{\theta \in \Theta} |\mathbf{p}_t^N(\theta) - \pi_t(\theta)| = 0 \quad a.s. \tag{4.4}$$

Proof It follows from the proof of Theorem 4.2 and Corollary 4.1 in Crisan and Míguez (2014). See Appendix A.5 for an outline. □

This theorem is a uniform convergence result, i.e., it holds uniformly in a compact parameter space Θ . We note that Theorem 2 specifies the dependence of the bandwidth h on the number of Monte Carlo samples N for convergence to be attained at that rate. Based on this result, we can relate the empirical maxima to the true maxima.

Theorem 3 Let $\theta_t^{*,N} \in \operatorname{argmax}_{i \in \{1, \dots, N\}} \mathbf{p}_t^N(\theta_t^{(i)})$ be an estimate of a global maximum of π_t and let $\theta_t^* \in \operatorname{argmax}_{\theta \in \Theta} \pi_t(\theta)$ be an actual global maximum. If Θ is compact, π_t is continuous at θ_t^* and Assumptions 1, 2, 3 and 4 hold, then for N sufficiently large

$$\pi_t(\theta_t^*) - \pi_t(\theta_t^{*,N}) \leq \frac{W_{t,d,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}}, \quad 1 \leq t \leq T,$$

where $\varepsilon \in (0, 1)$ is an arbitrarily small constant and $W_{t,d,\varepsilon}$ is an a.s. finite random variable, both independent of N .

Proof See Appendix A.6. □

Remark 7 By choosing $t = T$, Theorem 3 provides a convergence rate for the MAP estimator θ_T^* , which is also the approximate solution of problem (1.1). □

Theorem 3 also yields a convergence rate for the error $f(\theta_T^{*,N}) - f(\theta^*)$, where $f(\cdot)$ is the original cost function in problem (1.1), provided that the prior is chosen so that $\pi_T(\theta) \propto \exp(-f(\theta))$ (see Remark 1).

Corollary 2 Choose any

$$\theta^* \in \operatorname{argmin}_{\theta \in \Theta} f(\theta) \quad \text{and} \quad \theta_T^{*,N} \in \operatorname{argmax}_{i \in \{1, \dots, N\}} \mathbf{p}_T^N(\theta_T^{(i)}).$$

Under the same assumptions as in Theorem 3, if $\|f\|_\infty < \infty$ then we have

$$0 \leq f(\theta_T^{*,N}) - f(\theta^*) \leq \frac{\tilde{W}_{T,d,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}},$$

where $\tilde{W}_{T,d,\varepsilon}$ is an a.s. finite random variable.

Proof See Appendix A.7. □

Finally, we obtain a convergence rate for the expected error.

Corollary 3 Choose any

$$\theta^* \in \operatorname{argmin}_{\theta \in \Theta} f(\theta) \quad \text{and} \quad \theta_T^{*,N} \in \operatorname{argmax}_{i \in \{1, \dots, N\}} \mathbf{p}_T^N(\theta_T^{(i)}).$$

Under the same assumptions as in Theorem 3, if $\|f\|_\infty < \infty$ then we have

$$0 \leq \mathbb{E}[f(\theta_T^{*,N})] - f(\theta^*) \leq \frac{C_{T,d,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}},$$

where $C_{T,d,\varepsilon} = \mathbb{E}[\tilde{W}_{T,d,\varepsilon}] < \infty$ is a constant independent of N .

Proof The proof follows from Corollary 2, since $\tilde{W}_{T,d,\varepsilon}$ is an a.s. finite random variable. \square

4.1 Discussion

Theorem 3 and Corollaries 2 and 3 go beyond standard results on the convergence of SMC methods. While the latter refer to the approximation of integrals (in the vein of Theorem 1 and Corollary 1), Corollaries 2 and 3 directly address the convergence of the sequence of optimizers $\theta_T^{*,N}$ and state that the proposed algorithm yields, with probability 1, an asymptotically optimal solution to problem (1.1) even if $f(\theta)$ is non-convex and presents multiple local and/or global minima. These results also provide explicit convergence rates that depend on the computational cost (the number of particles N) and the dimension d of the search space.

Note that the analyses available in the literature for most Monte Carlo optimization algorithms are purely asymptotical (see Appel et al. 2004; Ikonen et al. 2005; Míguez 2010; Hu et al. 2012; Zhou et al. 2013, i.e., they do not provide explicit convergence rates. Moreover, they often rely on restrictive assumptions. For example, Hu et al. (2012) and Zhou et al. (2013) require that the objective function present a unique global minimum. More detailed analyses are carried out by Zhou and Chen (2013) and Míguez et al. (2013). However, the former falls short of providing explicit error rates for the sequence of optimizers (bounds are given for the total variation distance between the Boltzmann distributions and their SMC approximations in a SA scheme) and the latter relies on a sequential decomposition of the cost function that is not satisfied by $f(\theta)$ in problem (1.1). Moreover, all the analytical results in these papers (Appel et al. 2004; Ikonen et al. 2005; Míguez 2010; Hu et al. 2012; Zhou et al. 2013; Zhou and Chen 2013; Míguez et al. 2013) are obtained for deterministic optimization problems where the objective function can be evaluated exactly, while Theorem 3 and Corollaries 2 and 3 hold for a more general stochastic optimization framework where $f(\theta)$ can only be estimated using mini-batches of data.

5 Numerical results

In this section, we show numerical results for three optimization problems which are hard to solve with conventional

methods. In the first example, we focus on minimizing a function with multiple global minima. The aim of this experiment is to show that, when the cost function has several global minima, the PSMCO algorithm can successfully populate with Monte Carlo samples the regions of Θ that contain these minima. In the second example, we tackle the minimization of a challenging cost function, with broad flat regions, for which standard stochastic gradient optimizers struggle. The third example involves a non-convex, non-smooth cost function and we use it to compare the proposed PSMCO scheme with a similar SMC-based optimization method proposed in Stinis (2012).

5.1 Minimization of a function with multiple global minima

In this experiment, we tackle the problem

$$\min_{\theta \in \mathbb{R}^2} f(\theta), \text{ where } f(\theta) = \sum_{i=1}^n f_i(\theta)$$

and

$$f_i(\theta) = -\frac{1}{\lambda} \log \left(\sum_{k=1}^4 \mathcal{N}(\theta; m_{i,k}, R) \right),$$

with $\lambda = 10$ and $R = r I_2$, with I_2 denoting the 2×2 identity matrix and $r = 0.2$. We choose the means $m_{i,k}$ randomly, namely $m_{i,k} \sim \mathcal{N}(m_{i,k}; m_k, \sigma^2)$ where,

$$m_1 = [4, 4]^\top, \quad m_2 = [-4, -4]^\top, \\ m_3 = [-4, 4]^\top, \quad m_4 = [4, -4]^\top,$$

and $\sigma^2 = 0.5$. This selection results in a cost function with four global minima. Such functions arise in many machine learning problems, see, e.g., Mei et al. (2018). In this experiment, we have chosen $n = 1,000$. Although a small number for stochastic optimization problems, we note that each $f_i(\theta)$ represents a mini-batch in this scenario and we set $K = 1$ in the PSMCO algorithm.

In order to run the algorithm, we choose a uniform prior measure $\pi_0(\theta) = \mathcal{U}([-a, a] \times [-a, a])$ with $a = 50$. It follows from Proposition 1 that the pdf that matches the cost function $f(\theta)$ can be written as

$$\pi_T(\theta) \propto \exp(-f(\theta)),$$

and it has four global maxima. This pdf is displayed in Fig. 1a. We run $M = 100$ samplers, with $N = 50$ particles each, yielding a total number of $MN = 5,000$ particles. We choose a Gaussian jittering scheme; specifically, the jittering kernel is defined as

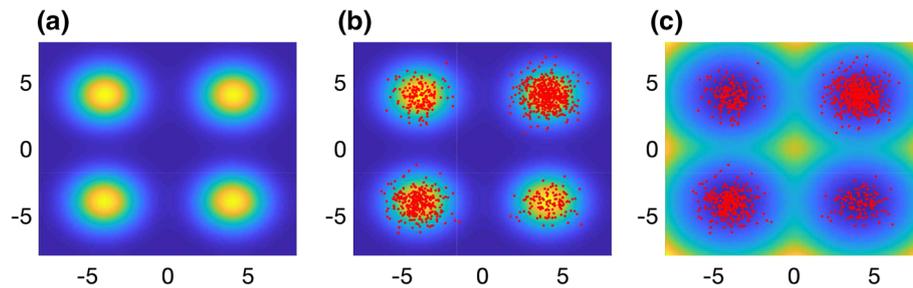


Fig. 1 An illustration of the performance of the proposed algorithm for a cost function with four global minima. **a** The plot of $\pi_T(\theta) \propto \exp(-f(\theta))$. The blue regions indicate low values. It can be seen that

there are four global maxima. **b** Samples drawn by the PSMCO at a single time instant. **c** The plot of the samples together with the actual cost function $f(\theta)$

$$\kappa(d\theta|\theta') = (1 - \epsilon_N)\delta_{\theta'}(d\theta) + \epsilon_N \mathcal{N}(\theta; \theta', \sigma_j^2)d\theta, \quad (5.1)$$

where $\epsilon_N = 1/\sqrt{N}$ and $\sigma_j^2 = 0.5$.

Some illustrative results can be seen from Fig. 1. To be specific, we have run independent samplers and plot all samples for this experiment (instead of estimating a minimum with the best performing sampler). From Fig. 1b, it can be seen that the algorithm populates the regions surrounding all maxima with samples. Finally, Fig. 1c shows the location of the samples relative to the actual cost function $f(\theta)$. These plots illustrate how the algorithm “locates” multiple, distinct global maxima with independent samplers. Note different samplers can converge to different global maxima in practice—which is in agreement with the analysis provided in Sect. 4.

5.2 Minimization of the sigmoid function

In this experiment, we address the problem,

$$\min_{\theta \in \mathbb{R}^2} f(\theta) := \sum_{i=1}^n (y_i - g_i(\theta))^2, \quad (5.2)$$

where

$$g_i(\theta) = \frac{1}{1 + \exp(-\theta_1 - \theta_2 x_i)},$$

with $x_i \in \mathbb{R}$, $f_i(\theta) = (y_i - g_i(\theta))^2$ and $\theta = [\theta_1, \theta_2]^\top$. The function g_i is called as the sigmoid function. Cost functions of the form in eq. (5.2) are widely used in nonlinear regression with neural networks in machine learning (Bishop 2006).

In this experiment, we have $n = 100,000$. We choose $M = 25$ and $MN = 1,000$, leading to $N = 40$ particles for every sampler. The mini-batch size is $K = 100$. The jittering kernel κ is defined in the same way as in (5.1), where the Gaussian pdf has a variance chosen as the ratio of the dataset size L to the mini-batch size K , i.e., $\sigma_j^2 = n/K$, which

yields a rather large variance² $\sigma_j^2 = 1000$. To compute the maximum as described in Eq. (3.2), we use a Gaussian kernel density with bandwidth $h = \lfloor N^{1/6} \rfloor^{-1}$.

In Fig. 2, we compare the PSMCO algorithm with a parallel stochastic gradient descent (PSGD) scheme (Zinkevich et al. 2010) using M optimizers. We note that, given a particular realization³ of $(x_i)_{i=1}^n$, searching for a minimum of $f(\theta)$ may be a hard task. Figure 2a shows one such case, where the cost function has broad flat regions which make it difficult to find its maxima using gradient-based methods unless their initialization is sufficiently good. Accordingly, we have run two instances of PSGD with “bad” and “good” initializations.

The bad initial point for PSGD can be seen from Fig 2a, at $[-190, 0]^\top$ (the blue dot). We initialize M parallel SGD optimizers around $[-190, 0]^\top$, each with a small zero-mean Gaussian perturbation with variance 10^{-8} . This is a poor initialization because gradients are nearly zero in this region (yellow area in Fig. 2a). We refer to the PSGD algorithm starting from this point as PSGD with B/I, which refers to bad initialization. We also initialize the PSMCO from this region, with Gaussian perturbations around $[-190, 0]^\top$, with the same small variance $\sigma_{\text{init}}^2 = 10^{-8}$.

The “good” initialization for the PSGD is selected from a better region, namely around the point $[0, -100]^\top$, where gradient values actually contain useful information about the minimum. We refer to the PSGD algorithm starting from this point as PSGD with G/I.

The results can be seen in Fig. 2b. We observe that the PSGD with good initialization (G/I) moves towards a better region, however, it gets stuck because the gradient becomes nearly zero. On the other hand, PSGD with B/I is unable

² Note that this is for efficient exploration of the global minima, which are hard to find for this example. A large jittering variance may not be adequate in practice when there are multiple minima close to each other, see, e.g., Sect. 5.1.

³ For this experiment, we generate i.i.d. uniform realizations, $x_k \sim \mathcal{U}([-2.5, 2.5])$ for $k = 1, \dots, n$.

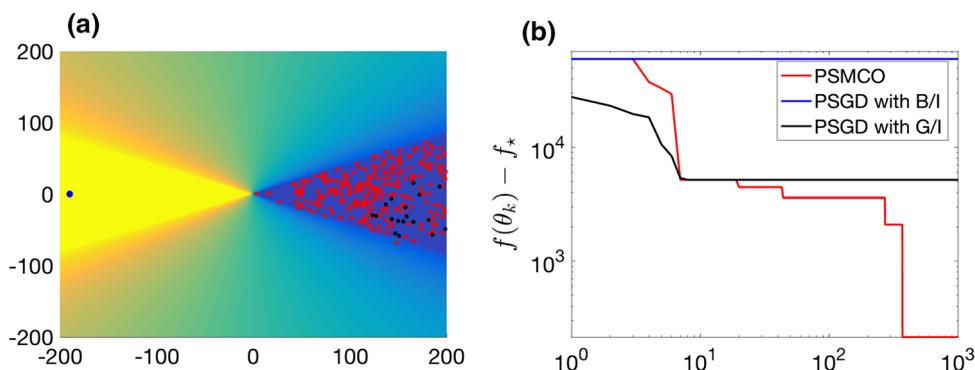


Fig. 2 **a** The cost function and a snapshot of samples from the 50th iteration of the PSMCO, PSGD with bad initialization (blue dot on the yellow area) and PSGD with good initialization (black dots on the blue area). **b** Performance of each algorithm: it can be seen that PSMCO first converges to the wide region with low values (blue region) and then

jumps to the minimum. This is because the marginal likelihood estimate of the sampler close to the minimum dominates after a while. There is effectively full communication among samplers only to determine the minimizer

to move at all, since it is initialized in a region where all gradients are negligible (which is true even for the mini-batch observations). The PSMCO method, on the other hand, searches the space effectively to find the global minimum, as depicted in Fig. 2b.

5.3 Constrained nonsmooth nonconvex optimization

In this section, we compare the proposed PSMCO scheme to the method of Stinis (2012), labeled here as ‘particle filtering for stochastic global optimization’ (PFSGO), and the stochastic evolution strategies (SES) algorithm in Salimans et al. (2017) for a high-dimensional non-smooth and non-convex optimization problem. In particular, we apply this algorithms to numerically solve the problem

$$\min_{\theta \in \Theta} \frac{1}{2} \|y - X^T \theta\|^2 + \frac{\rho}{2} \sum_{i=1}^d P_{\lambda, \gamma}(\theta_i), \tag{5.3}$$

where $y \in \mathbb{R}^n$, $X \in \mathbb{R}^{d \times n}$, $\Theta = [-5, 5]^d$, the dimension d is set to different values (see below), and $P_{\lambda, \gamma} : \mathbb{R} \mapsto \mathbb{R}$ is given by

$$P_{\lambda, \gamma}(x) = \begin{cases} \lambda|x| & \text{if } |x| \leq \lambda, \\ \frac{2\gamma\lambda|x| - x^2 - \lambda^2}{2(\gamma-1)} & \text{if } \lambda < |x| < \gamma\lambda, \\ \frac{\lambda^2(\gamma+1)}{2} & \text{if } |x| \geq \gamma\lambda, \end{cases}$$

where $\lambda > 2$ and $\gamma > 0$. This problem formulation is useful for variable selection, see, e.g., Fan and Li (2001) or Lan and Yang (2019). It is easy to see that problem (5.3) can be written as

$$\min_{\theta \in \Theta} \frac{1}{2} \sum_{i=1}^n (y_i - x_i^T \theta)^2 + \frac{\rho}{2} \sum_{i=1}^d P_{\lambda, \gamma}(\theta_i), \tag{5.4}$$

where $y_i \in \mathbb{R}$, and $x_i \in \mathbb{R}^d$. This, in turn, makes the problem an instance of (1.1), with

$$f_i(\theta) = \frac{1}{2} (y_i - x_i^T \theta)^2 + \frac{\tilde{\rho}}{2} \sum_{i=1}^d P_{\lambda, \gamma}(\theta_i),$$

and $\tilde{\rho} = \rho/n$.

In this problem, we also test the single-worker version of the proposed optimization scheme. We refer to this algorithm simply as SMCO and it is obtained as the particular case of PSMCO with $M = 1$. We use the usual jittering kernel of the form (3.1)

$$\kappa(d\theta|\theta') = (1 - \epsilon_N)\delta_{\theta'}(d\theta) + \epsilon_N\tau(d\theta|\theta'),$$

where τ is a Gaussian kernel with covariance $C = \sigma^2 I_d$ for both methods. We also use the same Gaussian transition kernel for the PFSGO. Let us remark, though, that (unlike SMCO and PSMCO) the PFSGO scheme modifies all particles at every iteration, i.e., it uses $\tau(\cdot|\theta')$ instead of $\kappa(\cdot|\theta')$ for sampling. The SES scheme also uses τ in order to estimate the gradients.

We choose $\sigma^2 = 10^{-2}$ and $N = 100$. The mini-batch size is taken as $K = 1$ and the number of components is $n = 1,000$. For the PSMCO, we chose $M = 5$, so it essentially runs 5 samplers with 20 particles each while the SMCO scheme runs a single sampler with $N = 100$ particles. For the regularization parameters, we choose $\tilde{\rho} = 1$, $\lambda = 10^{-3}$, and $\gamma = 2.01$. For the SES, we choose a small step size of $\alpha = 10^{-7}$ as larger values cause it to diverge. We simulate the data using a sparse parameter θ^* , where only three values

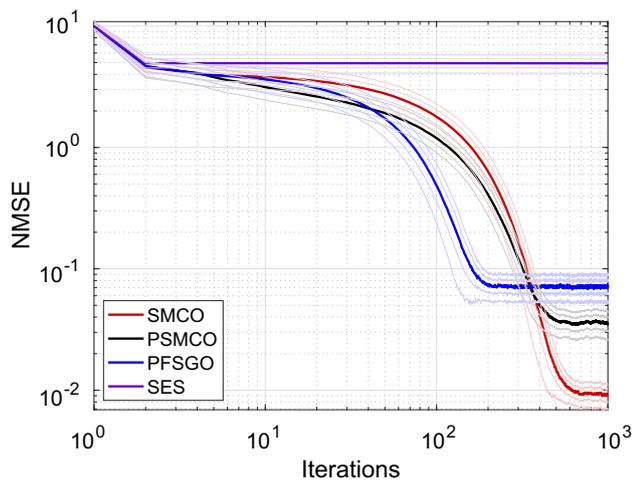


Fig. 3 Comparison of algorithms for problem (5.4) with $d = 10$, $N = 100$, and $n = 1,000$. It can be seen that the SMCO is the most efficient method for this problem and the PSMCO ($M = 5$) is the second best. Although PFSGO converges faster, the steady error that it attains is higher. The results are averaged over 1,000 Monte Carlo runs

are nonzero. We simulate the entries of the matrix X as i.i.d. variates from $\mathcal{N}(0, 1)$ and compute $y = X^\top \theta^*$. In order to compute the error for an iterate θ_k produced by any method, we compute

$$\text{NMSE}(k) = \frac{\|\theta_k - \theta^*\|^2}{\|\theta^*\|^2}.$$

The results can be seen in Fig. 3. We also plot the 0.5σ curves around the error curves which are averaged over 1,000 Monte Carlo runs. It can be seen that, for this particular example, the SMCO performs the best, while the PSMCO still outperforms the PFSGO. The SES basically is very slow due to the inefficiency of the gradient estimators for this problem.

To gain further insight, we also compare PSMCO ($M = 25$) and the PFSGO on a problem that is higher-dimensional, namely $d = 30$, and with more data points, $n = 10,000$. We set $\sigma^2 = 10^{-3}$ and leaving other parameters same as in the example with $d = 10$.

Figure 4 displays the results for this example. It can be seen that again the PSMCO algorithm converges to a point which has lower NMSE than the PFSGO. We believe that this is mainly due to the difference in the transition kernels. The PFSGO uses a full transition kernel where every particle is modified whereas jittering enables us to induce slower and more careful changes and also gives us a chance to keep a particle unmodified if it is in a good location.

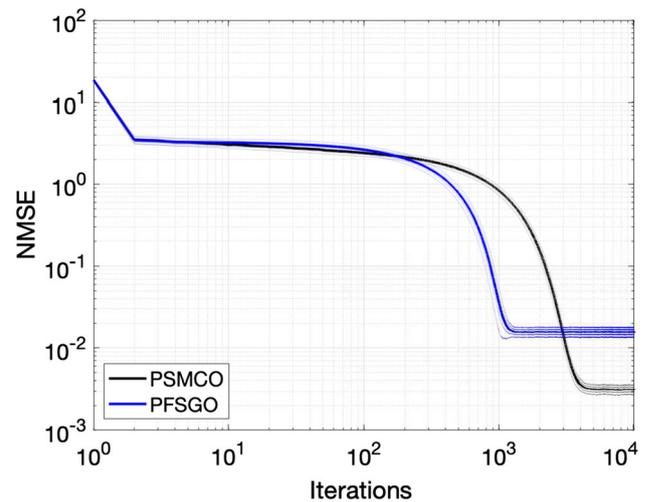


Fig. 4 Comparison of algorithms for problem (5.4), with $d = 30$, $N = 1,000$, and $n = 10,000$ —only for the PSMCO ($M = 25$) and PFSGO schemes. The results are averaged over 1,000 Monte Carlo runs

6 Conclusions

We have proposed a parallel sequential Monte Carlo optimization algorithm which does not require the computation (either exact or approximate) of gradients and, therefore, can be applied to the minimization of challenging cost functions, e.g., with multiple global minima or with broad “flat” regions. The proposed method uses jittering kernels to propagate samples (Crisan and Míguez 2018) and particle kernel density estimators to find the minima (Crisan and Míguez 2014), within a stochastic optimization setup. We have provided a detailed analysis of the proposed scheme. In particular, we have proved that it yields asymptotically optimal solutions to the stochastic optimization problem (1.1) (as the number of samples N is increased) and we have computed explicit convergence rates for the resulting optimizers that depend on N and the dimension of the search space, d . These results are new and improve on classical asymptotic analyses for Monte Carlo optimization methods, which typically lack convergence rates.

From a practical perspective, we argue that the parallel setting where each sampler uses a different configuration of the same dataset can be useful to improve the practical behaviour of the algorithm. To illustrate this point, we have studied the numerical performance of the PSMCO algorithm in scenarios where gradient-based methods struggle to converge. In this work, we have focused on challenging but relatively low-dimensional cost functions. We leave the potential applications of our scheme to high-dimensional optimization problems as a future work. Also the design of an interacting extension of our method similar to particle

islands (Vergé et al. 2015) can be potentially useful in more challenging settings.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Appendix

A.1 Proof of Proposition 1

We prove this result by induction. For $t = 1$, let

$$\pi_1(d\theta) = \pi_0(d\theta) \frac{G_1(\theta)}{\int_{\Theta} G_1(\theta)\pi_0(d\theta)} = \pi_0(d\theta) \frac{G_1(\theta)}{(G_1, \pi_0)}.$$

Since $G_1 \in B(\Theta)$ it follows that

$$\sup_{\theta \in \Theta} \left| \frac{G_1(\theta)}{(G_1, \pi_0)} \right| = \frac{\sup_{\theta \in \Theta} G_1(\theta)}{(G_1, \pi_0)} < \infty$$

because of Assumption 1. Hence $\pi_1 \ll \pi_0$ is a proper measure. Assume next, as an induction hypothesis, that $\pi_{T-1} \ll \pi_0$. Then

$$\pi_T(d\theta) = \pi_{T-1}(d\theta) \frac{G_T(\theta)}{(G_T, \pi_{T-1})}$$

and Assumption 1 implies (again) that

$$\frac{\sup_{\theta \in \Theta} G_T(\theta)}{(G_T, \pi_{T-1})} < \infty,$$

hence π_T is proper and $\pi_T \ll \pi_0$. Therefore, the Radon-Nikodym derivative of the final measure π_T w.r.t. the prior π_0 is

$$\frac{d\pi_T}{d\pi_0}(\theta) \propto \prod_{t=1}^T G_t(\theta) = \exp\left(-\sum_{i=1}^n f_i(\theta)\right).$$

From here, it readily follows that maximizing this Radon-Nikodym derivative is equivalent to solving problem (1.1). \square

A.2 Proof of Theorem 1

We proceed by an induction argument. At time $t = 0$, the bound

$$\|(\varphi, \pi_0^N) - (\varphi, \pi_0)\|_p \leq \frac{c_{0,p} \|\varphi\|_{\infty}}{\sqrt{N}}$$

is a straightforward consequence of the Marcinkiewicz–Zygmund inequality (Shiryaev 1996) because the particles $\{\theta_0^{(i)}\}_{i=1}^N$ are i.i.d samples from π_0 .

Assume now that, after iteration $t - 1$, we have a particle set $\{\theta_{t-1}^{(i)}\}_{i=1}^N$ and the empirical measure $\pi_{t-1}^N(d\theta_{t-1}) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_{t-1}^{(i)}}(d\theta_{t-1})$, which satisfies

$$\|(\varphi, \pi_{t-1}) - (\varphi, \pi_{t-1}^N)\|_p \leq \frac{c_{t-1,p} \|\varphi\|_{\infty}}{\sqrt{N}}. \tag{A.1}$$

We first analyze the error in the jittering step. To this end, we construct the *jittered* random measure

$$\hat{\pi}_t^N(d\theta) = \frac{1}{N} \sum_{i=1}^N \delta_{\hat{\theta}_t^{(i)}}(d\theta)$$

and iterate the triangle inequality to obtain

$$\begin{aligned} \|(\varphi, \pi_{t-1}) - (\varphi, \hat{\pi}_t^N)\|_p &\leq \|(\varphi, \pi_{t-1}) - (\varphi, \pi_{t-1}^N)\|_p \\ &\quad + \|(\varphi, \pi_{t-1}^N) - (\varphi, \kappa\pi_{t-1}^N)\|_p \\ &\quad + \|(\varphi, \kappa\pi_{t-1}^N) - (\varphi, \hat{\pi}_t^N)\|_p, \end{aligned} \tag{A.2}$$

where

$$\kappa\pi_{t-1}^N = \int \kappa(d\theta|\theta_{t-1})\pi_{t-1}^N(d\theta_{t-1}) = \frac{1}{N} \sum_{i=1}^N \kappa(d\theta|\theta_{t-1}^{(i)}).$$

The first term on the right-hand side (rhs) of (A.2) is bounded by the induction hypothesis (A.1). For the second term, we note that,

$$\begin{aligned} &\left| (\varphi, \pi_{t-1}^N) - (\varphi, \kappa\pi_{t-1}^N) \right| \\ &= \left| \frac{1}{N} \sum_{i=1}^N \varphi(\theta_{t-1}^{(i)}) - \frac{1}{N} \sum_{i=1}^N \int \varphi(\theta)\kappa(d\theta|\theta_{t-1}^{(i)}) \right| \\ &= \left| \frac{1}{N} \sum_{i=1}^N \int (\varphi(\theta_{t-1}^{(i)}) - \varphi(\theta)) \kappa(d\theta|\theta_{t-1}^{(i)}) \right| \\ &\leq \frac{1}{N} \sum_{i=1}^N \int |\varphi(\theta_{t-1}^{(i)}) - \varphi(\theta)| \kappa(d\theta|\theta_{t-1}^{(i)}) \\ &\leq \frac{c_{\kappa} \|\varphi\|_{\infty}}{\sqrt{N}}, \end{aligned} \tag{A.3}$$

where the last inequality follows from Assumption 2. The upper bound in (A.3) is deterministic, so the inequality readily implies that

$$\|(\varphi, \pi_{t-1}^N) - (\varphi, \kappa\pi_{t-1}^N)\|_p \leq \frac{c_\kappa \|\varphi\|_\infty}{\sqrt{N}}. \quad (\text{A.4})$$

For the last term on the right-hand side of (A.2), we let \mathcal{F}_{t-1} be the σ -algebra generated by the random sequence $\{\theta_{0:t-1}^{(i)}, \hat{\theta}_{1:t-1}^{(i)}\}_{i=1}^N$. Let us first note that

$$\begin{aligned} \mathbb{E}[(\varphi, \hat{\pi}_t) | \mathcal{F}_{t-1}] &= \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\varphi(\hat{\theta}_t^{(i)}) | \mathcal{F}_{t-1}] \\ &= \frac{1}{N} \sum_{i=1}^N \int \varphi(\theta) \kappa(d\theta | \theta_{t-1}^{(i)}) = (\varphi, \kappa\pi_{t-1}^N). \end{aligned}$$

Therefore, the difference $(\varphi, \hat{\pi}_t^N) - (\varphi, \kappa\pi_{t-1}^N)$ takes the form

$$(\varphi, \hat{\pi}_t^N) - (\varphi, \kappa\pi_{t-1}^N) = \frac{1}{N} \sum_{i=1}^N S^{(i)},$$

where $S^{(i)} = \varphi(\hat{\theta}_t^{(i)}) - \mathbb{E}[\varphi(\hat{\theta}_t^{(i)}) | \mathcal{F}_{t-1}]$, $i = 1, \dots, N$, are zero-mean and conditionally independent random variables, with $|S^{(i)}| \leq 2\|\varphi\|_\infty$. Then, we readily obtain the bound

$$\begin{aligned} \mathbb{E}\left[|(\varphi, \hat{\pi}_t^N) - (\varphi, \kappa\pi_{t-1}^N)|^p | \mathcal{F}_{t-1}\right] &= \frac{1}{N^p} \mathbb{E}\left[\left|\sum_{i=1}^N S^{(i)}\right|^p | \mathcal{F}_{t-1}\right] \\ &\leq \frac{B_{t,p} N^{\frac{p}{2}} \|\varphi\|_\infty^p}{N^p}. \end{aligned} \quad (\text{A.5})$$

where the relation (A.5) follows from the Marcinkiewicz–Zygmund inequality (Shiryaev 1996) and $B_{t,p} < \infty$ is some constant independent of N . Taking unconditional expectations on both sides of (A.5) and then computing $(\cdot)^{\frac{1}{p}}$ yields

$$\|(\varphi, \hat{\pi}_t^N) - (\varphi, \kappa\pi_{t-1}^N)\|_p \leq \frac{\hat{c}_{t,p} \|\varphi\|_\infty}{\sqrt{N}}. \quad (\text{A.6})$$

where $\hat{c}_{t,p} = B_{t,p}^{\frac{1}{p}}$ is a finite constant independent of N . Therefore, taking together (A.1), (A.4) and (A.6) we have established that

$$\|(\varphi, \pi_{t-1}) - (\varphi, \hat{\pi}_t^N)\|_p \leq \frac{c_{1,t,p} \|\varphi\|_\infty}{\sqrt{N}}, \quad (\text{A.7})$$

where $c_{1,t,p} = c_{t-1,p} + c_\kappa + \hat{c}_{t,p} < \infty$ is a finite constant independent of N .

Next, we have to bound the error after the weighting step. We recall that

$$\pi_t(d\theta) = \pi_{t-1}(d\theta) \frac{G_t(\theta)}{(G_t, \pi_{t-1})}$$

and define

$$\tilde{\pi}_t^N(d\theta) = \hat{\pi}_t^N(d\theta) \frac{G_t(\theta)}{(G_t, \hat{\pi}_t^N)}$$

where $\tilde{\pi}_t^N$ denotes the *weighted* measure. We first note that

$$\begin{aligned} |(\varphi, \pi_t) - (\varphi, \tilde{\pi}_t^N)| &= \left| \frac{(\varphi G_t, \pi_{t-1})}{(G_t, \pi_{t-1})} - \frac{(\varphi G_t, \hat{\pi}_t^N)}{(G_t, \hat{\pi}_t^N)} \pm \frac{(\varphi G_t, \hat{\pi}_t^N)}{(G_t, \pi_{t-1})} \right| \\ &\leq \frac{|(\varphi G_t, \pi_{t-1}) - (\varphi G_t, \hat{\pi}_t^N)| + \|\varphi\|_\infty |(G_t, \hat{\pi}_t^N) - (G_t, \pi_{t-1})|}{(G_t, \pi_{t-1})}. \end{aligned} \quad (\text{A.8})$$

Using Minkowski's inequality together with (A.7) and (A.8) yields

$$\begin{aligned} \|(\varphi, \pi_t) - (\varphi, \tilde{\pi}_t^N)\|_p &\leq \frac{c_{1,t,p} \|\varphi G_t\|_\infty + c_{1,t,p} \|\varphi\|_\infty \|G_t\|_\infty}{(G_t, \pi_{t-1}) \sqrt{N}}, \\ &\leq \frac{2c_{1,t,p} \|\varphi\|_\infty \|G_t\|_\infty}{(G_t, \pi_{t-1}) \sqrt{N}} \end{aligned}$$

where the second inequality follows from $\|\varphi G_t\|_\infty \leq \|\varphi\|_\infty \|G_t\|_\infty$. More concisely, we have

$$\|(\varphi, \pi_t) - (\varphi, \tilde{\pi}_t^N)\|_p \leq \frac{c_{2,t,p} \|\varphi\|_\infty}{\sqrt{N}} \quad (\text{A.9})$$

where the constant

$$c_{2,t,p} = \frac{2c_{1,t,p} \|G_t\|_\infty}{(G_t, \pi_{t-1})} < \infty$$

is independent of N . Note that the assumptions on $(G_t)_{t \geq 1}$ imply that $(G_t, \pi_{t-1}) > 0$.

Finally, we bound the resampling step. Note that the resampling step consists of drawing N i.i.d samples from $\tilde{\pi}_t^N$, i.e., $\theta_t^{(i)} \sim \tilde{\pi}_t^N$ i.i.d for $i = 1, \dots, N$, and then constructing

$$\pi_t^N(d\theta) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_t^{(i)}}(d\theta).$$

Since samples are i.i.d, as in the base case, we have,

$$\|(\varphi, \tilde{\pi}_t^N) - (\varphi, \pi_t^N)\|_p \leq \frac{\tilde{c}_p \|\varphi\|_\infty}{\sqrt{N}}, \quad (\text{A.10})$$

for some constant $\tilde{c}_p < \infty$ independent of N . Now combining (A.9) and (A.10), we have the desired result,

$$\|(\varphi, \pi_t) - (\varphi, \pi_t^N)\|_p \leq \frac{c_t \|\varphi\|_\infty}{\sqrt{N}}$$

where $c_t = c_{2,t,p} + \tilde{c}_p$ is a finite constant independent of N . \square

A.3 Proof of Corollary 1

From Theorem 1, we obtain

$$\|(\varphi, \pi_t) - (\varphi, \pi_t^N)\|_p \leq \frac{c_t \|\varphi\|_\infty}{\sqrt{N}},$$

where $c_t < \infty$ is a constant independent of N . Let us choose $p \geq 4$ and $0 < \epsilon < 1$. We construct the nonnegative random variable

$$U_{t,\epsilon}^p = \sum_{N=1}^\infty N^{\frac{p}{2}-1-\epsilon} |(\varphi, \pi_t) - (\varphi, \pi_t^N)|^p.$$

and use Fatou’s lemma to obtain

$$\begin{aligned} \mathbb{E}[U_{t,\epsilon}^p] &\leq \sum_{N=1}^\infty N^{\frac{p}{2}-1-\epsilon} \mathbb{E} \left[|(\varphi, \pi_t) - (\varphi, \pi_t^N)|^p \right], \\ &\leq c^p \|\varphi\|_\infty^p \sum_{N=1}^\infty N^{-1-\epsilon} < \infty, \end{aligned} \tag{A.11}$$

where the second inequality follows from Theorem 1. The relationship (A.11) implies that the r.v. $U_{t,\epsilon}^p$ is a.s. finite.

Finally, since (trivially) $N^{\frac{p}{2}-1-\epsilon} |(\varphi, \pi_t) - (\varphi, \pi_t^N)|^p \leq U_{t,\epsilon}^p$, we have

$$|(\varphi, \pi_t) - (\varphi, \pi_t^N)| \leq \frac{U_{t,\delta}}{N^{\frac{1}{2}-\delta}}, \tag{A.12}$$

where $\delta = \frac{1+\epsilon}{p}$ and $U_{t,\delta} = (U_{t,\epsilon}^p)^{\frac{1}{p}}$. Since $p \geq 4$ and $0 < \epsilon < 1$, it follows that $0 < \delta < \frac{1}{2}$. The almost sure convergence follows from (A.12). Taking $N \rightarrow \infty$ yields

$$\lim_{N \rightarrow \infty} |(\varphi, \pi_t) - (\varphi, \pi_t^N)| = 0 \quad \text{a.s.}$$

□

A.4 Proof of Proposition 2

Recall the assumption

$$|F_t(\theta) - F_t(\theta')| \leq \ell_t \|\theta - \theta'\|.$$

We write $F_t^* = \min_{\theta \in \Theta} F_t(\theta)$, which is assumed to be finite, but not necessarily nonnegative. We first prove that $\exp(-F_t(\theta))$ is also Lipschitz continuous. Note that we trivially have $\exp(-F_t(\theta)) \leq \exp(-F_t^*)$ for all θ since $F_t(\theta) \geq F_t^*$ for all θ . Now consider any $(\theta, \theta') \in \Theta \times \Theta$. We first consider the case where $F_t(\theta) \leq F_t(\theta')$. We obtain

$$\begin{aligned} 0 < e^{-F_t(\theta)} - e^{-F_t(\theta')} &= e^{-F_t(\theta)} \left(1 - e^{F_t(\theta) - F_t(\theta')} \right), \\ &\leq e^{-F_t(\theta)} \left(1 - (1 + F_t(\theta) - F_t(\theta')) \right), \end{aligned} \tag{A.13}$$

where we have used the inequality $e^a \geq 1 + a$. Therefore, we readily obtain from (A.13)

$$\begin{aligned} 0 < e^{-F_t(\theta)} - e^{-F_t(\theta')} &\leq e^{-F_t(\theta)} (F_t(\theta') - F_t(\theta)), \\ &\leq e^{-F_t^*} (F_t(\theta') - F_t(\theta)) \end{aligned} \tag{A.14}$$

$$= e^{-F_t^*} |F_t(\theta') - F_t(\theta)|, \tag{A.15}$$

since $F_t(\theta) \leq F_t(\theta')$. Next, assume otherwise, i.e., $F_t(\theta) \geq F_t(\theta')$. In this case, we can also show using the same line of reasoning that

$$e^{-F_t(\theta')} - e^{-F_t(\theta)} \leq e^{-F_t^*} (F_t(\theta) - F_t(\theta')) \tag{A.16}$$

$$= e^{-F_t^*} |F_t(\theta') - F_t(\theta)|, \tag{A.17}$$

since $F_t(\theta) \geq F_t(\theta')$. Therefore, we can conclude (combining (A.15) and (A.16)) that

$$|e^{-F_t(\theta)} - e^{-F_t(\theta')}| \leq e^{-F_t^*} |F_t(\theta') - F_t(\theta)| \leq e^{-F_t^*} \ell_t \|\theta - \theta'\|,$$

where the last inequality holds because F_t is Lipschitz. Finally recall that

$$\pi_t(\theta) = \frac{e^{-F_t(\theta)}}{Z_{\pi_t}},$$

where we denote $Z_{\pi_t} = \int_{\Theta} e^{-F_t(\theta)} d\theta$. We straightforwardly obtain

$$|\pi_t(\theta) - \pi_t(\theta')| \leq \frac{1}{Z_{\pi_t}} e^{-F_t^*} \ell_t \|\theta - \theta'\|.$$

□

A.5 Proof of Theorem 2

Using the proof of Theorem 4.2 and Corollary 4.1 in Crisan and Míguez (2014), we obtain

$$\sup_{\theta \in \Theta} |p_t^N(\theta) - \pi_t(\theta)| \leq \frac{V_{1,\epsilon}}{\left[N^{\frac{1}{2(d+1)}} \right]^{1-\epsilon}},$$

where $V_{1,\epsilon}$ is an a.s. finite random variable. Noting that

$$\sup_{a \geq 1} \frac{a}{[a]} = 2,$$

we obtain

$$\sup_{\theta \in \Theta} |\mathbf{p}_t^N(\theta) - \pi_t(\theta)| \leq \frac{V_\varepsilon}{N^{\frac{1-\varepsilon}{2(d+1)}}},$$

where $V_\varepsilon = 2V_{1,\varepsilon}$ is an almost surely finite random variable. \square

A.6 Proof of Theorem 3

Recall that $\pi_t(\theta)$ is a probability density w.r.t. the Lebesgue measure. Choose $\theta_t^* \in \arg \max_{\theta \in \Theta} \pi_t(\theta)$ and construct the ball

$$B_{t,n}^* := B\left(\theta_t^*, \frac{1}{n}\right) \subset \Theta$$

where $n \geq 1$ is a positive integer. We assume, without loss of generality, that $B_{t,1}^* \subseteq \Theta$ and denote

$$\pi_t(B_{t,n}^*) = \int_{B_{t,n}^*} \pi_t(\theta) d\theta \quad \text{and} \quad \pi_t^N(B_{t,n}^*) = \int_{B_{t,n}^*} \pi_t^N(d\theta).$$

Also recall that the grid of points generated by the SMC sampler at time t is $\{\theta_t^{(i)}\}_{1 \leq i \leq N} \subset \Theta$ and the estimate of θ_t^* obtained from the grid is denoted

$$\theta_t^{*,N} \in \arg \max_{\theta \in \{\theta_t^{(i)}\}_{1 \leq i \leq N}} \mathbf{p}_t^N(\theta), \tag{A.18}$$

where $\mathbf{p}_t^N(\theta)$ is the kernel density estimator of π_t . Our argument to prove Theorem 3 proceeds in two steps:

1. We show that, for any given $n \geq 1$, one can a.s. find N sufficiently large to ensure that $\{\theta_t^{(i)}\}_{1 \leq i \leq N} \cap B_{t,n}^* \neq \emptyset$, i.e., that there are points of the grid contained in the ball $B_{t,n}^*$. Moreover, we deduce an inequality that relates the radius n^{-1} of the ball $B_{t,n}^*$ with the number of necessary particles N .
2. From the existence of at least one particle $\theta_t^{(i)}$ inside $B_{t,n}^*$ and the assumption that $\pi_t(\theta)$ is Lipschitz, we deduce bounds for the differences $|\pi_t(\theta_t^*) - \pi_t(\theta_t^{(i)})|$ and $|\pi_t(\theta_t^{*,N}) - \pi_t(\theta_t^{(i)})|$, and, as a consequence, for the approximation error $|\pi_t(\theta_t^{*,N}) - \pi_t(\theta_t^*)|$.

A.6.1 The ball $B_{t,n}^*$ is a.s. non-empty

Since $\pi_t(\theta)$ is assumed continuous at every $\theta_t^* \in \arg \max_{\theta \in \Theta} \pi_t(\theta)$, we have $\pi_t(B_{t,n}^*) > 0$. Therefore, for every $n < \infty$, Theorem 2 ensures that there exists N_n (a.s. finite) such that for all $N \geq N_n$,

$$\left| \pi_t^N(B_{t,n}^*) - \pi_t(B_{t,n}^*) \right| < \frac{U_{t,\delta}}{N^{\frac{1}{2}-\delta}} < \frac{\pi_t(B_{t,n}^*)}{2}, \tag{A.19}$$

where $U_{t,\delta}$ is an a.s. finite random variable and $\delta \in (0, \frac{1}{2})$ is an arbitrarily small constant (both independent of N). Moreover, the second inequality in (A.19) implies that

$$\pi_t^N(B_{t,n}^*) > \frac{\pi_t(B_{t,n}^*)}{2} > 0. \tag{A.20}$$

Therefore, for all $N > N_n$ there exists at least one integer $i_b \in \{1, \dots, N\}$ such that $\theta_t^{(i_b)} \in B_{t,n}^*$.

To be specific, since $\pi_t(\theta)$ is a density w.r.t. the Lebesgue measure, we can readily find a lower bound for the integral $\pi_t(B_{t,n}^*)$, namely

$$\frac{\pi_t(B_{t,n}^*)}{2} > \frac{1}{2} \text{Leb}(B_{t,n}^*) \times \inf_{\theta \in B_{t,n}^*} \pi_t(\theta) > c_{t,d} n^{-d}$$

where $\text{Leb}(B_{t,n}^*) = \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)n^d}$ is the Lebesgue measure of the d -dimensional ball with radius n^{-1} , $\Gamma(\cdot)$ is Euler's gamma function and

$$c_{t,d} := \frac{\pi^{\frac{d}{2}}}{2\Gamma(\frac{d}{2}+1)n^d} \times \inf_{\theta \in B_{t,n}^*} \pi_t(\theta) > 0.$$

Therefore, for any given $n < \infty$, if we choose N such that $0 < \frac{U_{t,\delta}}{N^{\frac{1}{2}-\delta}} < c_{t,d} n^{-d}$, i.e.,

$$N \geq N_n := \left(\frac{U_{t,\delta}}{c_{t,d}}\right)^{\frac{2}{1-2\delta}} n^{\frac{2d}{1-2\delta}} \tag{A.21}$$

then the inequalities (A.19) and (A.20) hold a.s. (note that $N_n < \infty$ a.s. because $n < \infty$ and $U_{t,\delta} < \infty$ a.s.).

A.6.2 Error bounds

Choose $i_b \in \{1, \dots, N\}$ such that $\theta_t^{(i_b)} \in B_{t,n}^*$. Such index exists a.s. whenever N satisfies the inequality (A.21). Let us recall the construction of the estimate $\theta_t^{*,N}$ from expression (A.18) and denote

$$\hat{\theta}_t^{*,N} \in \arg \max_{\theta \in \Theta} \mathbf{p}_t^N(\theta).$$

Let $L_t < \infty$ be the Lipschitz constant of the pdf $\pi_t(\theta)$. Since $\theta_t^{(i_b)} \in B_{t,n}^*$, we readily obtain the upper bound

$$\pi_t(\theta_t^*) - \pi_t(\theta_t^{(i_b)}) < L_t n^{-1}$$

and, therefore,

$$\pi_t(\theta_t^*) - L_t n^{-1} < \pi_t(\theta_t^{(i_b)}). \tag{A.22}$$

However, using Theorem 2 we obtain

$$\left| \pi_t(\theta_t^{(ib)}) - \mathbf{p}_t^N(\theta_t^{(ib)}) \right| < \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}}, \tag{A.23}$$

where $\varepsilon \in (0, 1)$ is an arbitrarily small constant and $V_{t,\varepsilon}$ is an a.s. finite random variable, both independent of N . Combining (A.22) and (A.23) yields

$$\mathbf{p}_t^N(\theta_t^{(ib)}) > \pi_t(\theta_t^{(ib)}) - \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}} > \pi_t(\theta_t^*) - L_t n^{-1} - \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}}$$

and, as a consequence,

$$\mathbf{p}_t^N(\theta_t^{*,N}) \geq \mathbf{p}_t^N(\theta_t^{(ib)}) > \pi_t(\theta_t^*) - L_t n^{-1} - \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}}. \tag{A.24}$$

Moreover, using Theorem 2 again, we find that

$$\left| \pi_t(\hat{\theta}_t^{*,N}) - \mathbf{p}_t^N(\hat{\theta}_t^{*,N}) \right| < \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}}, \tag{A.25}$$

with the same constant $\varepsilon \in (0, 1)$ and a.s. finite random variable $V_{t,\varepsilon}$ as in (A.24). Since $\pi_t(\hat{\theta}_t^{*,N}) \leq \pi_t(\theta_t^*)$, the inequality (A.25) implies that

$$\mathbf{p}_t^N(\hat{\theta}_t^{*,N}) < \pi_t(\theta_t^*) + \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}}$$

and, since $\mathbf{p}_t^N(\theta_t^{*,N}) \leq \mathbf{p}_t^N(\hat{\theta}_t^{*,N})$, we arrive at

$$\mathbf{p}_t^N(\theta_t^{*,N}) < \pi_t(\theta_t^*) + \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}}. \tag{A.26}$$

Taking the inequalities (A.24) and (A.26) together, we readily obtain the uniform bound (for $\theta \in \Theta$)

$$\left| \pi_t(\theta_t^*) - \mathbf{p}_t^N(\theta_t^{*,N}) \right| < \frac{V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}} + L_t n^{-1} \tag{A.27}$$

and a simple triangle inequality then yields

$$\begin{aligned} \left| \pi_t(\theta_t^{*,N}) - \pi_t(\theta_t^*) \right| &\leq \left| \pi_t(\theta_t^{*,N}) - \mathbf{p}_t^N(\theta_t^{*,N}) \right| \\ &\quad + \left| \mathbf{p}_t^N(\theta_t^{*,N}) - \pi_t(\theta_t^*) \right| \\ &< \frac{2V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}} + L_t n^{-1}, \end{aligned} \tag{A.28}$$

where the second inequality follows from (A.27) and yet another application of Theorem 2.

The inequality (A.28) holds for any pair of integers (N, n) that satisfies the relationship (A.21). For any given N , sufficiently large for

$$n_N := \sup \left\{ m \in \mathbb{N} : m^{-1} > \left(\frac{U_{t,\delta}}{c_{t,d}} \right)^{\frac{1}{d}} \frac{1}{N^{\frac{1-2\delta}{2d}}} \right\}$$

to be well defined, the pair consisting of N and $n = n_N$ satisfies (A.21), while

$$n_N^{-1} \leq 2 \left(\frac{U_{t,\delta}}{c_{t,d}} \right)^{\frac{1}{d}} \frac{1}{N^{\frac{1-2\delta}{2d}}}. \tag{A.29}$$

Hence, if we substitute $n = n_N$ in the inequality (A.28) and then apply the inequality (A.29) we arrive at

$$\left| \pi_t(\theta_t^{*,N}) - \pi_t(\theta_t^*) \right| < \frac{2V_{t,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}} + 2 \left(\frac{U_{t,\delta}}{c_{t,d}} \right)^{\frac{1}{d}} \frac{L_t}{N^{\frac{1-2\delta}{2d}}}, \tag{A.30}$$

where $V_{t,\varepsilon}$ and $U_{t,\delta}$ are a.s. finite, and L_t and $c_{t,d}$ are finite. The constants $\varepsilon \in (0, 1)$ and $\delta \in (0, 1/2)$ can be chosen arbitrarily small. Hence, if we let $0 < \delta = \varepsilon/2 < \frac{1}{2}$, the r.h.s. of (A.30) can be upper bounded, which results in the bound

$$\left| \pi_t(\theta_t^{*,N}) - \pi_t(\theta_t^*) \right| < \frac{W_{t,d,\varepsilon}}{N^{\frac{1-\varepsilon}{2(d+1)}}},$$

where

$$W_{t,d,\varepsilon} = 2 \left[V_{t,\varepsilon} + \left(\frac{U_{t,\delta(\varepsilon)}}{c_{t,d}} \right)^{\frac{1}{d}} L_t \right] < \infty \text{ a.s.}$$

A.7 Proof of Corollary 2

Recall that

$$\|f\|_\infty = \sup_{\theta \in \Theta} |f(\theta)| < \infty.$$

Note that Theorem 3 implies that

$$0 \leq e^{-f(\theta^*)} - e^{-f(\theta_T^{*,N})} \leq \frac{W_{T,d,\varepsilon} Z_{\pi_T}}{N^{\frac{1-\varepsilon}{2(d+1)}}}, \tag{A.31}$$

where Z_{π_T} is the normalizing constant of π_T . Next, we lower bound the left-hand side of (A.31) as

$$\begin{aligned} e^{-f(\theta^*)} - e^{-f(\theta_T^{*,N})} &= e^{-f(\theta_T^{*,N})} \left(e^{f(\theta_T^{*,N}) - f(\theta^*)} - 1 \right) \\ &\geq e^{-\|f\|_\infty} (f(\theta_T^{*,N}) - f(\theta^*)) \end{aligned} \tag{A.32}$$

where the last inequality follows from the relationships

$$e^{-f(\theta_T^{*,N})} \geq e^{-\|f\|_\infty}$$

(since $f(\theta_T^{*,N}) \leq \|f\|_\infty$) and $e^a \geq a + 1$ for $a \in \mathbb{R}$. Combining (A.31) and (A.32), we obtain

$$f(\theta_T^{*,N}) - f(\theta^*) \leq \frac{\tilde{W}_{T,d,\varepsilon}}{N^{\frac{1}{2(d+1)}}}$$

where

$$\tilde{W}_{T,d,\varepsilon} = Z_{\pi_T} W_{T,d,\varepsilon} e^{\|f\|_\infty}$$

is a.s. finite.

References

- Akyildiz, Ö.D., Míguez, J.: Nudging the particle filter. *Stat. Comput.* **30**(2), 305–330 (2020)
- Akyildiz, O.D., Mariño, I.P., Míguez, J.: Adaptive noisy importance sampling for stochastic optimization. In: *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2017 IEEE 7th International Workshop on*, IEEE, pp 1–5 (2017)
- Akyildiz, O.D., Elvira, V., Miguez, J.: The Incremental Proximal Method: A Probabilistic Perspective. In: *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada (2018)
- Alquier, P., Friel, N., Everitt, R., Boland, A.: Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. *Stat. Comput.* **26**(1–2), 29–47 (2016)
- Appel, M., Labarre, R., Radulovic, D.: On accelerated random search. *SIAM J. Optim.* **14**(3), 708–731 (2004)
- Bach, F., Perchet, V.: Highly-smooth zero-th order online optimization. In: *Conference on Learning Theory*, pp 257–283 (2016)
- Bertsekas, D.P.: Incremental gradient, subgradient, and proximal methods for convex optimization: a survey. *Optim. Mach. Learn.* **2010**, 1–38 (2011)
- Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer, Secaucus (2006)
- Bottou, L., Curtis, F.E., Nocedal, J.: Optimization methods for large-scale machine learning. *Siam Rev.* **60**(2), 223–311 (2018)
- Chen, C., Carlson, D., Gan, Z., Li, C., Carin, L.: Bridging the gap between stochastic gradient MCMC and stochastic optimization. In: *Artificial Intelligence and Statistics*, pp 1051–1060 (2016)
- Chen, R., Wild, S.: Randomized derivative-free optimization of noisy convex functions. arXiv preprint [arXiv:1507.03332](https://arxiv.org/abs/1507.03332) (2015)
- Conn, A.R., Scheinberg, K., Vicente, L.N.: *Introduction to derivative-free optimization*, MPS-SIAM Series on Optimization, vol 8. SIAM (2009)
- Crisan, D., Míguez, J.: Particle-kernel estimation of the filter density in state-space models. *Bernoulli* **20**(4), 1879–1929 (2014)
- Crisan, D., Miguez, J.: Nested particle filters for online parameter estimation in discrete-time state-space markov models. *Bernoulli* **24**(4A), 3039–3086 (2018)
- Del Moral, P.: *Feynman–Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer, Berlin (2004)
- Del Moral, P., Doisy, M.: Maslov idempotent probability calculus, *I. Theory Probab Appl* **43**(4), 562–576 (1999)
- Del Moral, P., Doucet, A., Jasra, A.: Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **68**(3), 411–436 (2006)
- Douc, R., Cappé, O.: Comparison of resampling schemes for particle filtering. In: *Image and Signal Processing and Analysis, 2005. ISPA 2005. In: Proceedings of the 4th International Symposium on*, IEEE, pp 64–69 (2005)
- Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learning Res.* **12**, 2121–2159 (2011)
- Elvira, V., Míguez, J., Djurić, P.M.: Adapting the number of particles in sequential monte carlo methods through an online scheme for convergence assessment. *IEEE Trans. Signal Process.* **65**(7), 1781–1794 (2017)
- Fan, J., Li, R.: Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Am. Stat. Assoc.* **96**(456), 1348–1360 (2001)
- Ghadimi, S., Lan, G.: Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM J. Optim.* **23**(4), 2341–2368 (2013)
- Gürbüzbalaban, M., Ozdaglar, A., Parrilo, P.: Why random reshuffling beats stochastic gradient descent. arXiv preprint [arXiv:1510.08560](https://arxiv.org/abs/1510.08560) (2015)
- Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. *Evolut. Comput.* **9**(2), 159–195 (2001)
- Hu J, Wang Y, Zhou E, Fu MC, Marcus SI (2012) A survey of some model-based methods for global optimization. In: Hernández-Hernández, D., Minjárez, J.A. (eds.) *Optimization, Control, and Applications of Stochastic Systems*, Springer, pp 157–179
- Ikonen, E., Najim, K., Del Moral, P.: Application of genealogical decision trees for open-loop tracking control. *IFAC Proc. Vol.* **38**(1), 288–293 (2005)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
- Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220**(4598), 671–680 (1983)
- Lan, G., Yang, Y.: Accelerated stochastic algorithms for nonconvex finite-sum and multiblock optimization. *SIAM J. Optim.* **29**(4), 2753–2784 (2019)
- Liu B, Cheng S, Shi Y (2016) Particle filter optimization: A brief introduction. In: *International Conference on Swarm Intelligence*, Springer, pp 95–104
- Mariño, I.P., Míguez, J.: Monte Carlo method for multiparameter estimation in coupled chaotic systems. *Phys. Rev. E* **76**(5), 057203 (2007)
- Mei, S., Bai, Y., Montanari, A.: The landscape of empirical risk for nonconvex losses. *Ann. Stat.* **46**(6A), 2747–2774 (2018)
- Homem-de Mello, T., Bayraksan, G.: Monte carlo sampling-based methods for stochastic optimization. *Surv. Oper. Res. Manag. Sci.* **19**(1), 56–85 (2014)
- Míguez, J.: Analysis of a sequential Monte Carlo method for optimization in dynamical systems. *Signal Process.* **90**(5), 1609–1622 (2010)
- Míguez, J., Crisan, D., Djurić, P.M.: On the convergence of two sequential monte carlo methods for maximum a posteriori sequence estimation and stochastic global optimization. *Stat. Comput.* **23**(1), 91–107 (2013)
- Morse, G., Stanley, K.O.: Simple evolutionary optimization can rival stochastic gradient descent in neural networks. In: *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, ACM, pp 477–484 (2016)
- Nesterov, Y., Spokoiny, V.: *Random gradient-free minimization of convex functions*. Technical report, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE) (2011)
- Pereyra, M., Schniter, P., Chouzenoux, E., Pesquet, J.C., Tourneret, J.Y., Hero, A.O., McLaughlin, S.: A survey of stochastic simulation and

- optimization methods in signal processing. *IEEE J. Select. Top. Signal Process.* **10**(2), 224–241 (2015)
- Robbins, H., Monro, S.: A stochastic approximation method. *Ann. Math. Stat.* **22**, 400–407 (1951)
- Robert, C.P., Casella, G.: *Monte Carlo statistical methods*. Wiley, New York (2004)
- Salimans, T., Ho, J., Chen, X., Sidor, S., Sutskever, I.: Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint [arXiv:1703.03864](https://arxiv.org/abs/1703.03864)* (2017)
- Shamir, O.: Without-replacement sampling for stochastic gradient methods. In: *Advances in Neural Information Processing Systems*, pp 46–54 (2016)
- Shiryayev, A.N.: *Probability*. Springer, Berlin (1996)
- Silverman, B.W.: *Density Estimation for Statistics and Data Analysis*. Routledge, Abingdon (1998)
- Spall, J.C.: *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*, vol. 65. Wiley, New York (2005)
- Stinis, P.: Stochastic global optimization as a filtering problem. *J. Comput. Phys.* **231**(4), 2002–2014 (2012)
- Vergé, C., Dubarry, C., Del Moral, P., Moulines, E.: On parallel implementation of sequential monte carlo methods: the island particle model. *Stat. Comput.* **25**(2), 243–260 (2015)
- Wand, M.P., Jones, M.C.: *Kernel Smoothing*. Chapman and Hall/CRC, Boca Raton (1994)
- Welling, M., Teh, Y.W.: Bayesian learning via stochastic gradient langevin dynamics. In: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp 681–688 (2011)
- Wibisono, A., Wainwright, M.J., Jordan, M.I., Duchi, J.C.: Finite sample convergence rates of zero-order stochastic optimization methods. In: *Advances in Neural Information Processing Systems*, pp 1439–1447 (2012)
- Wierstra, D., Schaul, T., Glasmachers, T., Sun, Y., Peters, J., Schmidhuber, J.: Natural evolution strategies. *J. Mach. Learn. Res.* **15**(1), 949–980 (2014)
- Zhou, E., Chen, X.: Sequential monte carlo simulated annealing. *J. Global Optim.* **55**(1), 101–124 (2013)
- Zhou, E., Fu, M.C., Marcus, S.I.: Particle filtering framework for a class of randomized optimization algorithms. *IEEE Trans. Autom. Contr.* **59**(4), 1025–1030 (2013)
- Zinkevich, M., Weimer, M., Li, L., Smola, A.J.: Parallelized stochastic gradient descent. In: *Advances in neural information processing systems*, pp 2595–2603 (2010)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.