



HAL
open science

STOCHASTIC MIRROR DESCENT FOR NONPARAMETRIC ADAPTIVE IMPORTANCE SAMPLING

Pascal Bianchi, Bernard Delyon, François Portier, Victor Priser

► **To cite this version:**

Pascal Bianchi, Bernard Delyon, François Portier, Victor Priser. STOCHASTIC MIRROR DESCENT FOR NONPARAMETRIC ADAPTIVE IMPORTANCE SAMPLING. 2024. hal-04702293

HAL Id: hal-04702293

<https://hal.science/hal-04702293v1>

Preprint submitted on 19 Sep 2024

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

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

Stochastic mirror descent for nonparametric adaptive importance sampling

Bianchi P., Delyon B., Portier F. and Priser V.

September 19, 2024

Abstract

This paper addresses the problem of approximating an unknown probability distribution with density f - which can only be evaluated up to an unknown scaling factor - with the help of a sequential algorithm that produces at each iteration $n \geq 1$ an estimated density q_n . The proposed method optimizes the Kullback-Leibler divergence using a mirror descent (MD) algorithm directly on the space of density functions, while a stochastic approximation technique helps to manage between algorithm complexity and variability. One of the key innovations of this work is the theoretical guarantee that is provided for an algorithm with a fixed MD learning rate $\eta \in (0, 1)$. The main result is that the sequence q_n converges almost surely to the target density f uniformly on compact sets. Through numerical experiments, we show that fixing the learning rate $\eta \in (0, 1)$ significantly improves the algorithm's performance, particularly in the context of multi-modal target distributions where a small value of η allows to increase the chance of finding all modes. Additionally, we propose a particle subsampling method to enhance computational efficiency and compare our method against other approaches through numerical experiments.

1 Introduction

Consider the problem of approximating an unknown probability distribution with density $f : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ using a sequential algorithm that produces an estimated density q_n . The index $n \in \mathbb{N}$ here stands for the number of point-wise evaluations of f - no evaluation of the gradient ∇f is needed - and the knowledge of the normalization constant of f should not be necessary, i.e., the algorithm remains the same when, for an arbitrary constant $c > 0$, cf is used instead of f . This framework is useful in many applications such as Bayesian inference [EKPS23, HCEP23] or variational inference [BKM17], reinforcement learning [MPFR18, HNS19] or stochastic optimization [NWS14, GCE24] to name a few among the statistical learning literature.

The *adaptive importance sampling* method [OB92, OZ00, BEM⁺17] or the *sequential Monte Carlo* approach [DMDJ06], are based on generating random variables according to a certain sampling distribution that evolves during the algorithm and using some re-weighting allows to obtain unbiased estimators. Depending on the problem of interest, the sampling distribution might be chosen by minimizing some discrepancy with respect to the target measure and many such different approaches have been investigated in [EC22, BBD⁺19, LT16]. A leading approach, coming from the variational inference literature [BKM17], consists in minimizing the Kullback-Leibler (KL) divergence, defined as

$$\text{KL}(q||f) := \int \log(q/f)q,$$

with respect to q chosen out of a parametric family of density functions. Throughout the paper, $\int h$ shall be used as a shortcut for $\int h(x)dx$. The optimization framework attached to the variational inference approach is attractive because of the recent development in stochastic optimization and related methods (e.g., stochastic gradient descent, variance reduction). This, in particular, allows to handle large scale problems as promoted in [HBWP13].

In [DHDS16, KP22, CCK23b], the *mirror descent* (MD) algorithm [BT03] is employed to optimize the *Kullback-Leibler* divergence directly on the space of density functions and thereby avoiding parametric misspecification issues as in standard stochastic variational inference [HBWP13]. When applied to $\text{KL}(q||f)$, the MD algorithm gives the following iteration, for $n \geq 0$,

$$q_{n+1} \propto q_n^{1-\eta} f^\eta \tag{1}$$

where $\eta \in (0, 1]$ is called the *learning rate*. The above iteration cannot be implemented within our framework simply because f is unknown (so is the normalizing constant in the above).

The approach taken in this paper is to rely on *stochastic approximation* [RM51] whose main idea is to resort to a sequential algorithm in which a computationally cheap stochastic update is conducted at each step. Suppose that X_{n+1} is generated from q_n and define the importance weights $w_{n+1} = f(X_{n+1})/q_n(X_{n+1})$. Let $K_b : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be a probability density with a mean zero and covariance $b^2 I_d$, where I_d is the identity matrix of size $d \times d$, and $b > 0$ is a small positive parameter known as the *bandwidth*. Define the random map

$$M_{n+1} : x \in \mathbb{R}^d \mapsto w_{n+1}^\eta K_b(x - X_{n+1}),$$

and note that $\mathbb{E}(M_{n+1}) = \{q_n^{1-\eta} f^\eta\} * K_b$, where $*$ denotes the convolution product between functions, i.e., $f * g(x) = \int f(y)g(x-y)dy$ when f, g are real-valued Lebesgue-integrable functions, which from well-known results from approximation theory should be near the MD iteration expressed in (1), $q_n^{1-\eta} f^\eta$, when b is small. Having this in mind, the proposed algorithm follows from the functional iteration

$$g_{n+1} = (1 - \gamma_{n+1})g_n + \gamma_{n+1} M_{n+1},$$

where γ_{n+1} is a positive step size converging to 0 as n tends to infinity. The final step is given by $q_{n+1} = (1 - \lambda_{n+1})g_{n+1} / \int g_{n+1} + \lambda_{n+1}q_0$ where $\int g_{n+1}$ can be easily determined through the algorithm and q_0 is a heavy-tailed distribution that ensures a comprehensive exploration of the space \mathbb{R}^d as suggested by [DP21]. The distribution q_n writes as a mixture between q_0 and a weighted sum of n kernels anchored at the particles locations. Such a sum of kernels is encountered in the wellknown context of *kernel density estimation* [Par62].

In the previous algorithm, the choice $\eta = 1$ might be attractive at first glance because of the bias term which is easy to analyze [DP21]. However, when the dimension d is large or when the function f is complex, the importance weights often collapse [BBL⁺08], i.e., only a few weights carry-out the whole probability mass. This implies that the weights exhibit significant variance and this characteristic severely hampers the algorithm’s efficiency, causing it to stagnate at excessively high weights. As noted by [KP22],

$$\text{Var}(w_n^\eta) \leq \text{Var}(w_n)$$

for $\eta \in (0, 1)$, suggesting that choosing small η might help to reduce the variability and avoid the degeneracy of the importance weights. In addition, as observed in practice, and similarly to several stochastic gradient descent optimization algorithm choosing a fixed stepsize allows the algorithm to explore the space of interest and thereby avoid local minima.

Related algorithms. The proposed method has connections with several well-known approaches within the adaptive importance sampling, sequential Monte Carlo and variational inference literature. The idea of using a stochastic approximation of the MD iteration (1) using an adaptive importance sampling approach was first investigated in [DHDS16] and further studied in [KP22]. In both previous papers, the importance weights w_n are “tempered” using the MD power transform w_n^η but their results are significantly different from the one established in this paper in that [DHDS16] considers η converging to 0, and [KP22] deals with learning rate η going to 1, while here we study the case where η remains fixed during the algorithm. Same weights transformations, referred to as *tempering*, as well as other similar transformations e.g., *clipping*, that implies a different behavior than standard weights, have been used in several sequential Monte Carlo algorithms [Nea01, KM15, APB22] (without relying specifically on kernel smoothing); see also [DMDJ06, Section 2.3.1]. For more detail on the connection between *tempering* and MD, we refer to [CCK23b].

The proposed algorithm bears resemblance to the ones of [Wes93, GR96, Zha96, DP21] even though the MD iteration was not considered in the previous work (i.e., $\eta = 1$). In all previous references, a kernel smoothing estimator is employed to estimate f . This is also done by several popular sequential Monte Carlo samplers as described in [Cho04, DMDJ06] where generating from normalized g_n , i.e., $X_{n+1} \sim g_n / \int g_n$, is often described using two steps: (i) the *selection step* chooses at random, using multinomial sampling, one particular particle X_i among the existing ones; (ii) the *mutation step* generates X_{n+1} around X_i using (for example) kernel $K_b(\cdot - X_i)$.

Note that the proposed algorithm only requires evaluation of the target density up to a normalization constant. This differs significantly from approaches that rely on gradient evaluations as for

instance the *Markov chain Monte Carlo* methods proposed in [WT11, SBCR16] or the variational inference approaches detailed in [LW16, KSA⁺20, LLK⁺23].

Contribution and related results. The main result of the paper is to establish the almost sure convergence (uniformly on compact sets) of the sequence $(q_n)_{n \geq 0}$ to the target density f , under the assumption that the bandwidth tends to zero as n tends to infinity *i.e.*, $b = b_n$ vanishes, and that γ_n satisfies typical Robbins-Monro conditions [RM51], while η might be constant during the algorithm. One important consequence of the previous is a central limit theorem, with rescaling factor \sqrt{n} , for the estimation of $\int hf$, for compactly supported test functions $h : \mathbb{R}^d \rightarrow \mathbb{R}$, using empirical weighted average of $h(X_n)$ with weights w_n .

Existing theoretical results on the convergence of sequence $(q_n)_{n \geq 0}$ when $\eta = 1$ might be found in [Wes93, GR96, Zha96, DP21]. For instance, the almost sure uniform convergence to the target density with a convergence rate and some central limit theorem for the integral estimation problem are obtained in [DP21]. In [DHDS16] some results concerning the weak convergence of q_n to f and the convergence of the *Kullback-Leibler* objective are given when η goes to 0. In [KP22], the almost-sure convergence is established when η converges to 1.

To the best of our knowledge, the results of the present paper are the first theoretical guarantees about the convergence of q_n to f when η is fixed during the algorithm. The new parasitic stationary point at 0 when $\eta < 1$ complicates the proof as we need to establish that the algorithm is not trapped in the vicinity of this spurious equilibrium. In contrast, forcing η to converge to 1 as in [KP22] makes the algorithm behave asymptotically like the case $\eta = 1$ as studied in [DP21], which eases the proof.

As mentioned previously, similar types of algorithms, that involves a power transformation of the weights w_n^η , often referred to as tempering, have been studied within the sequential Monte Carlo literature [Cho04, DMDJ06, DM07]. To our knowledge, the results obtained are different as the evolution of the sequence $(q_n)_{n \geq 0}$ is stopped while allowing the number of particles to go to infinity [Cho04]. This constraint is heavy because in practice one might allow the policy to change in time without constraint.

Another line of work is the nonparametric recursive estimation problem in which data is used sequentially to update the estimators [DW80, MPS09, BCD19]. Note that the recursive estimation framework relies on Robbins-Monro type procedure, just as we do, but the context is different because our framework requires the variational policy $(q_n)_{n \geq 0}$ to be updated and then used to draw points whereas in the recursive framework the data is distributed according to a fixed density.

Outline. In Section 2, we the mathematical framework and the main algorithm along with several practical remarks. In Section 3, we state our main result and provide a sketch of proof. In Section 4, we consider several practical variations of the proposed method while in Section 5, we compare them to other approaches using numerical experiments. All the proofs are provided in the Appendix.

2 The algorithm

We consider a probability density function $f : \mathbb{R}^d \rightarrow [0, +\infty)$, referred to as the target. Let $f_u : \mathbb{R}^d \rightarrow [0, +\infty)$ be a Lebesgue integrable function, representing an unnormalized version of f . That is, there is a constant $c > 0$ such that $f_u = cf$.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Consider a sequence $(X_n)_{n \geq 1}$ of random variable on \mathbb{R}^d . Denote by $(\mathcal{F}_n)_{n \geq 0}$ the natural filtration associated to the sequence $(X_n)_{n \geq 1}$. That is, $\mathcal{F}_n = \sigma(X_1, \dots, X_n)$ for $n \geq 1$, and $\mathcal{F}_0 = \{\emptyset, \Omega\}$. The sequence $(X_n)_{n \geq 1}$ is specified by its *policy* $(q_n)_{n \geq 0}$ defined as follows.

Definition 1. The sequence of random variable $(q_n)_{n \geq 0}$ is said to be a *policy* of $(X_n)_{n \geq 1}$, if it is adapted to $(\mathcal{F}_n)_{n \geq 0}$ and if, for every $n \geq 0$, q_n is a conditional probability distribution function of X_{n+1} given \mathcal{F}_n , that is, $\mathbb{E}(h(X_{n+1})|\mathcal{F}_n) = \int h q_n$ for every bounded continuous function h on \mathbb{R}^d .

Define the *importance weights* by

$$w_{n+1} := \frac{f_u(X_{n+1})}{q_n(X_{n+1})}, \quad n \geq 1. \quad (2)$$

These weights play an important role in the importance sampling framework as they allow to shift the distribution from q_n toward the target distribution f . As such, it allows to estimate without

bias integrals with respect to the unnormalized target distribution as, whenever $q_n > 0$ implies $f_u > 0$, it holds

$$\mathbb{E}(w_{n+1}h(X_{n+1})|\mathcal{F}_n) = \int hf_u.$$

We are now in a position to introduce our algorithm characterizing sequentially the policy $(q_n)_{n \geq 0}$ based on a sequence of unnormalized density $(g_n)_{n \geq 0}$. Let $K : \mathbb{R}^d \rightarrow [0, \infty)$ be a probability density function and define $K_b(x) = b^{-d}K(x/b)$ for any $b > 0$, the corresponding density with variance b^2I_d . At each step $n \geq 0$, the random variable X_{n+1} is drawn from q_n and the distribution g_n is updated into g_{n+1} as follows:

$$\begin{aligned} X_{n+1} &\sim q_n \\ g_{n+1}(x) &= (1 - \gamma_{n+1})g_n(x) + \gamma_{n+1}w_{n+1}^\eta K_{b_{n+1}}(x - X_{n+1}), \quad \forall x \in \mathbb{R}^d, \end{aligned} \quad (3)$$

where $(\gamma_n)_{n \geq 1}$ and $(b_n)_{n \geq 1}$ are positive sequences, respectively referred to as the step size and the bandwidth sequences. We set $g_0 = 0$. The next step is therefore to define q_{n+1} from g_{n+1} , by:

$$q_{n+1}(x) = (1 - \lambda_{n+1})\frac{g_{n+1}(x)}{\int g_{n+1}} + \lambda_{n+1}q_0(x), \quad \forall x \in \mathbb{R}^d, \quad (4)$$

where q_0 is a fixed probability density function, and $(\lambda_n)_{n \geq 1}$ is a positive sequence. Thus, we do not define q_{n+1} as the normalized version of g_{n+1} , but as a mixture between the latter and a fixed distribution q_0 . This mixture step will be revealed essential in our proofs, in order to ensure sufficient exploration, and thus prevent q_n to converge to a spurious stationary point. The parameter λ_{n+1} determines the tradeoff between the exploration and the adaptation to f .

The next proposition is given without proof as it is an easy consequence of the normalization stage in (4).

Proposition 1. *The policy $(q_n)_{n \geq 0}$ obtained from (4) is invariant with respect the normalization constant $c > 0$.*

The previous property is attractive because it implies that even when f_u is attached to a small normalization constant $c > 0$, it has no effect on the algorithm (even in the first iterations). This fact, in our proof, will be useful as it will allow us to work directly with the true target density f .


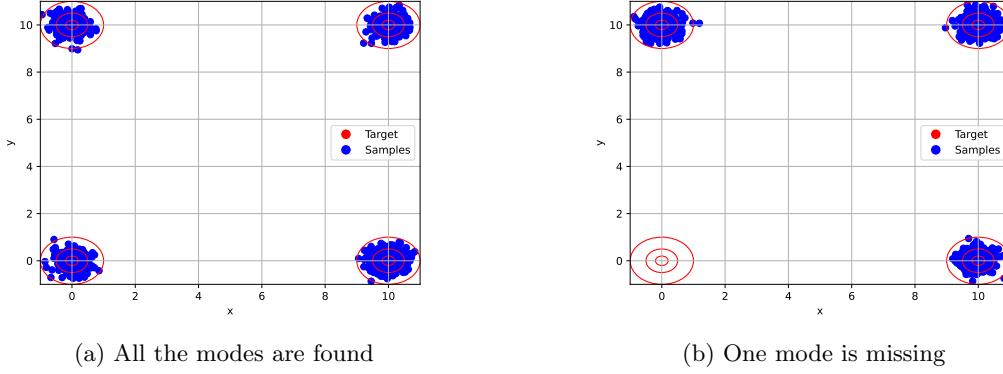
The role of η has already been depicted as balancing between bias and variance in [KP22] as it was shown that $\text{Var}(w_n^\eta) \leq \text{Var}(w_n)$ while $\mathbb{E}[w_n^\eta] \leq 1$. One related point is it can also enable to visit extensively the domain of interest. To illustrate this claim, we now consider the extreme case where $\eta = 0$, and we remark that in this case, the policy $(q_n)_{n \geq 0}$ obtained from (4) does not depend on f . The choice $\eta = 0$ is not expected to be efficient as it does not use the evaluations of f . It is nonetheless informative as it stresses that, as soon as η is small, there is no preference between the weights. As a consequence, a new particle tends to be drawn equally from any previous particle. This is particularly attractive as in some cases the algorithm can be trapped in a small part of the domain while missing other parts.

To illustrate the previous property and thereby the importance of the parameter η , we provide a toy example with a multi-modal target in dimension $d = 2$, where all modes are challenging to find. The target density is a mixture of four Gaussian distributions with means at $(0, 0)$, $(10, 0)$, $(0, 10)$, and $(10, 10)$, each having a variance of $0.1I_2$. The heavy-tailed density q_0 is a Student's t-distribution with a location parameter of $(5, 5)$ and a scale parameter of $10I_2$. The target distribution possesses four modes. In Figure 1, we present two illustrative runs of our algorithm. In Figure 1a, all four modes of the target distribution are recovered while in Figure 1b, one mode is missing. A run of the algorithm with one mode missing can occur randomly for any value of η , but by setting the parameter low enough, the algorithm is more likely to recover all the modes of the target distribution. To evaluate how much this impact the outcome of the algorithm, we compute the sliced Wasserstein distance, as defined in Section 5.1, between the target distribution and the weighted empirical measure of the particles $(X_n, w_n)_{n \geq 1}$. The average over 50 independent runs of the log Wasserstein distance is provided in Figure 1c.

3 Main results

3.1 Almost sure convergence

We start by giving the assumptions needed on the target f_u and the density q_0 .



(c) Average over 50 independent runs of the logarithm of sliced Wasserstein distance for different value of η

Figure 1: The target is the multi-modal distribution described in the text. Figure 1a represents samples generated from q_n ($\eta = 1$) where all modes are identified while Figure 1b has missed one mode. Figure 1c shows a comparison (based on the sliced Wasserstein distance) of different algorithms when varying η .

Assumption 1.

- i) The functions f_u , q_0 , are bounded, continuous, nonnegative, and integrable on \mathbb{R}^d .
- ii) There exists $r > 0$ and $C_f > 0$ such that $\int_{\|x\|>t} f_u(x)dx \leq C_f t^{-r}$ for all $t \geq 0$.
- iii) There exists $c > 0$ such that $cf_u \leq q_0$.

Assumption 1-ii holds for instance if the function $x \mapsto \|x\|^{r+d} f_u(x)$ is bounded in \mathbb{R}^d . Assumption 1-iii ensures that the support of the target distribution f_u is included in the support of q_0 , which will implies (assuming that λ_n is large enough) that f_u can be explored thoroughly. This condition is necessary for the definition of the weights $(w_n)_{n \geq 1}$. We also need to have some regularity and integrability conditions on the *kernel* function K .

Assumption 2.

- i) $K : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is a bounded and Lipschitz density function and $K(0) > 0$.

ii) There exist $r > 0$ and $C_K > 0$ such that, for all $x \in \mathbb{R}^d$, $(1 + \|x\|^{d+r})K(x) \leq C_K$.

The Gaussian kernel, $K(u) \propto \exp(-\|u\|^2/2)$, or the Epanechnikov Kernel, $K(u) \propto (1 - \|u\|^2)_+$ where $u_+ = u$ when $u \geq 0$ and 0 else, satisfy the above assumption. In the numerical experiments, the Gaussian kernel will be used. Finally we state the assumption needed on the step-size sequence $(\gamma_n)_{n \geq 1}$, bandwidth sequence $(b_n)_{n \geq 1}$ and mixture parameter $(\lambda_n)_{n \geq 1}$.

Assumption 3.

i) The sequence $(\gamma_n)_{n \geq 1}$, $(b_n)_{n \geq 1}$, $(\lambda_n)_{n \geq 1}$ are decreasing to 0.

ii) There exists $C_\gamma > 0$ and $n_\gamma \geq 1$ such that for all $n \geq 1$, $n^{-1} \leq C_\gamma \gamma_n$, and, for all $n \geq n_\gamma$, $\gamma_n - \gamma_{n+1} \leq \gamma_n \gamma_{n+1}$. Moreover,

$$\sum_{n \geq 1} \gamma_n^2 < \infty.$$

iii) When $\eta \geq 1/2$, it holds that

$$\frac{n\gamma_n^2 \log n}{\lambda_n b_n^d} \rightarrow 0.$$

When $\eta < 1/2$, we have

$$\frac{n\gamma_n^2 \log n}{\lambda_n^{2(1-\eta)} b_n^{2d(1-\eta)}} \rightarrow 0.$$

We remark, that Assumption 3-ii holds when $\gamma_n = Cn^{-\alpha}$, for $C > 0$ and $\alpha \in (1/2, 1)$, or $C > 1$ and $\alpha = 1$. When $\gamma_n = Cn^{-1}$ and $\eta > 1/2$, Assumption 3-iii on (b_n) implies a classical condition in the kernel smoothing estimation literature, that is,

$$\lim_{n \rightarrow \infty} nb_n^d = \infty \text{ and } \lim_{n \rightarrow \infty} b_n = 0.$$

In non parametric estimation when the function is at least 2-times continuously differentiable and the kernel has order 2 [Sto80], the optimal bandwidth is $h_n = n^{-\frac{1}{4+d}}$. This choice is made possible by the assumption of our main result.

Based on the previous set of assumptions, we are able to prove the almost sure convergence on compacts. This is the main result of the paper.

Theorem 1. *Consider the policy $(q_n)_{n \geq 0}$ given by (4). Let Assumptions 1, 2 and 3 hold true and let $A \subset \mathbb{R}^d$ be a compact set. Then, almost surely, we have:*

$$\lim_{n \rightarrow \infty} \sup_{x \in A} |q_n(x) - f(x)| = 0.$$

3.2 Sketch of the proof

By Proposition (1), we can consider $f = f_u$. We can rewrite (3) as follows

$$g_{n+1} = (1 - \gamma_{n+1})g_n + \gamma_{n+1}q_n^{1-\eta}f^\eta * K_{b_{n+1}} + \gamma_{n+1}\xi_{n+1},$$

where $\xi_{n+1} : \mathbb{R}^d \rightarrow \mathbb{R}$ is a martingale increment in that, for all $x \in \mathbb{R}^d$, $\mathbb{E}[\xi_{n+1}(x)|\mathcal{F}_n] = 0$. From the latter equation, we see that the algorithm has two equilibria: f and 0. Note that the case $\eta = 1$ does not present this issue, which is why our proof differs significantly from those in [DP21] and [KP22]. The point f is stable, while the point 0 is unstable. This means that, without the martingale term involving $(\xi_n)_{n \geq 1}$, $(g_n)_{n \geq 0}$ would converge to f , but the presence of the martingale term can potentially cause it to get trapped at 0.

Using that $q_n = g_n / (\int g_n) + \lambda_n q_0$, a convexity inequality implies that

$$\begin{aligned} g_{n+1} &\geq (1 - \gamma_{n+1})g_n + \gamma_{n+1} \frac{(1 - \lambda_n)^{1-\eta}}{2^\eta (\int g_n)^{1-\eta}} T(g_n) * K_{b_{n+1}} \\ &\quad + \gamma_{n+1} \xi_{n+1} + \gamma_{n+1} \frac{\lambda_n}{2^\eta} T(q_0), \end{aligned}$$

where $T(q) = q^{1-\eta}f^\eta$. By iterating the latter equation, we obtain a lower bound on g_n . Using a Freedman-type inequality, we show that, for sufficiently large n , the term involving the heavy

tailed density q_0 is greater than the absolute value of the martingale term involving $(\xi_n)_{n \geq 1}$ in the expression of this lower bound. Hence, cases where the martingale term could lead the algorithm into the trap cannot occur. The trap is therefore avoided. With this in hand, the analysis of convergence becomes more straightforward.

Note that in [KP22, PD18], a Freedman-type concentration inequality is also used for the analysis of the martingale term. When $\eta \in [1/2, 1)$, we recover the same condition as in [KP22, PD18] while when $\eta \in (0, 1/2)$, the condition in Assumption 3-iii is stronger. This is due to the variance of the noise that does not scale the same way.

3.3 Convergence in total variation and weak convergence.

An application of Scheffé’s lemma allows to extend the uniform convergence on compact sets to L_1 -convergence.

Corollary 1. *Under the assumptions of Theorem 1, we have, almost surely,*

$$\lim_{n \rightarrow \infty} \int |q_n - f| = 0.$$

Now we can turn our attention to weak convergence type of results for the estimation of integrals. This property has some practical interest in regards of the Bayesian application where one is often interested in computing the mean with respect to posterior distribution. With the help of Algorithm 1, integral of the form $\mu(h) = \int h f$, for a given integrable function h , can be easily estimated using the normalized quantity

$$\hat{\mu}_n(h) = \frac{\sum_{i=1}^n w_i h(X_i)}{\sum_{i=1}^n w_i}.$$

The asymptotic normality is established in the next proposition.

Corollary 2. *Let $h : \mathbb{R}^d \mapsto \mathbb{R}$ with compact support A such that $\int f h^2 < \infty$ and suppose that $\inf_{x \in A} f(x) > 0$. Under the assumptions of Theorem 1, $\sqrt{n}(\hat{\mu}_n(h) - \mu(h)) \rightsquigarrow \mathcal{N}(0, \sigma^2(h))$ with $\sigma^2(h) = \mu(h^2) - \mu(h)^2$.*

The proof of this result, which is given in the Appendix, follows from an application of the Lindeberg central limit theorem [HH14] with a careful use of the convergence of $(q_n)_{n \geq 1}$ in order to check each of the conditions leading to the right asymptotic variance. We note that the expression of the asymptotic variance is the same as the one of the oracle Monte Carlo estimate $(1/n) \sum_{i=1}^n h(X_i)$ where $(X_i)_{i \geq 1}$ is an independent sequence of random variables with common distribution f . This equicontinuity property is reminiscent of Corollary 1 stated in [PD18] where a high-level condition is given on q_n to satisfy such a central limit theorem.

4 Practical considerations

In this section, we provide a complete description of the considered algorithms including subsampling and minibatching variants.

4.1 Initial algorithm

To present a concise description of the studied algorithm, let us start with some algebra expanding (3) and (4) which together provide the incremental expression of the algorithm, as a damped stochastic version of mirror descent. Assuming $g_0 = 0$ and taking $\lambda_0 = 1$, the policy $(q_n)_{n \geq 0}$ writes:

$$q_n(x) = (1 - \lambda_n) \left(\frac{\sum_{i=1}^n W_{i,n} K_{b_i}(x - X_i)}{\sum_{i=1}^n W_{i,n}} \right) + \lambda_n q_0(x), \quad (5)$$

where for $1 \leq i \leq n$,

$$W_{i,n} := w_i^n \gamma_i \prod_{j=i}^n (1 - \gamma_j). \quad (6)$$

Algorithm 1: Mirror Descent for Adaptive Sampling (MIDAS)

Input: Budget $N \geq 1$, step sizes $(\gamma_n)_{1 \leq n \leq N}$, bandwidths $(b_n)_{1 \leq n \leq N}$, mixture weights $(\lambda_n)_{1 \leq n \leq N}$, learning rate $\eta \in (0, 1]$, initial distribution q_0 , kernel K

Output: Weighted particles $(X_n, W_{n,N})_{1 \leq n \leq N}$

```
1 for  $n \leftarrow 0$  to  $N - 1$  do
2   generate  $X_{n+1}$  from the mixture  $q_n$  defined in (5);
3    $W_{n+1,n+1} \leftarrow \left( \frac{f_u(X_{n+1})}{q_n(X_{n+1})} \right)^\eta \gamma_{n+1}$ ;
4   if  $n \geq 1$ , for all  $1 \leq i \leq n$ ,  $W_{i,n+1} \leftarrow (1 - \gamma_{n+1})W_{i,n}$ ;
end
```

The practical implementation of (5) and (6) is detailed in Algorithm 1, referred to as *mirror descent for adaptive sampling* (MIDAS), for a budget $N \geq 1$ corresponding to the number of evaluation of f .

Thanks to (5), it is easy to see that MIDAS is invariant with respect to the choice of f_u among all possible scaled versions of the density f . As a result the algorithm does not require the knowledge of the normalizing constant but also the behavior of the algorithm is not sensible to the value of this constant. This makes easier the choice of the hyperparameter such as the learning rate η or the step-size γ_n . An interesting special case is obtained when setting $\gamma_n = 1/n$. In this case, we obtain the simplification: $W_{i,n} = w_i^\eta/n$, for all $1 \leq i \leq n$. If, moreover, $\eta = 1$, the algorithm is closely related to the algorithm of [DP21] even though we consider here a slightly different choice of the kernel's bandwidth sequence.

Sampling from the mixture q_n in (5) is achieved by drawing a random index in $\{1, \dots, n\}$ with a probability equal to the weights $W_{i,n}$ for $1 \leq i \leq n$. The cost of generating the index is $O(\log n)$. Next, a random variable is generated according to the kernel density K , which yields the final particle X_{n+1} (up to shifting and rescaling). Of course, the kernel K is chosen to make the latter step computationally effective.

As justified in the previous paragraph, we shall neglect the (logarithmic) cost of drawing a particle from the current distribution q_n . Hence the main computational cost of the algorithm is carried out by the evaluation, at each iteration, of the importance weight $W_{n+1,n+1}$ and in particular to the computation of $f_u(X_{n+1})$ and $q_n(X_{n+1})$. We denote by c_u the cost of evaluating f_u . The evaluation of $q_n(X_{n+1})$ requires n evaluations of K . Denoting by c_K the cost for evaluating K at a given point, the n -th iteration of the algorithm requires an order of $c_u + nc_K$. This leads to an overall computing cost of $Nc_u + N^2c_K$. Even if in some practical situations (e.g., complex Bayesian model or when f_u is the result of a heavy simulation program), c_K is might be smaller than c_u , the complexity is dominated by the quadratic term N^2c_K when N is large. In this case, it is interesting to consider a variant of our algorithm, in which the complexity is reduced. This is the purpose of the next section.

4.2 Subsampling variant

In order to decrease the quadratic in N computing cost of MIDAS, we propose a subsampling version of MIDAS, which is inspired from [GSS93]. The aim is to restrict the iteration cost of the algorithm to approximately $O(\ell_n)$ operations, where $(\ell_n)_{n \geq 0}$ is a sequence of integers chosen by the user herself, and such that $\ell_n \ll n$. One may for instance consider $\ell_n \sim n^\delta$, with $0 < \delta < 1$.

At each iteration n , the main idea is to draw ℓ_n indices $(u(i, n))_{1 \leq i \leq \ell_n}$ according to the following weighted empirical distribution

$$\mathbb{P}_n := \frac{\sum_{i=1}^{\ell_n} W_{i,n} \delta_i}{\sum_{i=1}^{\ell_n} W_{i,n}},$$

where we recall the definition of $W_{i,n}$ in (6). In other words, at the n th iteration of the algorithm, ℓ_n particles with its bandwidth are drawn with replacement among the n particles generated so far by the algorithm. Then, we define

$$q_n^*(x) := (1 - \lambda_n) \sum_{i=1}^{\ell_n} K_{b_{u(i,n)}}(x - X_{u(i,n)}) + \lambda_n q_0(\cdot), \quad (7)$$

The algorithm, which will be referred to as SubMIDAS for subsampling MIDAS, is described in Algorithm 2.

Algorithm 2: SubMIDAS

Input: Budget $N \geq 1$, step sizes $(\gamma_n)_{1 \leq n \leq N}$, bandwidths $(b_n)_{1 \leq n \leq N}$, mixture weights $(\lambda_n)_{1 \leq n \leq N}$, learning rate η , bootstrap sample sizes $(\ell_n)_{1 \leq n \leq N}$, initial distribution q_0 , kernel K .

Output: Weighted particles $(X_n, W_{n,N})_{1 \leq n \leq N}$

- 1 Generate X_1 from q_0 and set $W_{1,1} \leftarrow \left(\frac{f_u(X_1)}{q_0(X_1)} \right)^\eta \gamma_1$;
- 2 **for** $n \leftarrow 1$ **to** $N - 1$ **do**
- 3 generate independent indices $(u(i, n))_{1 \leq i \leq \ell_n}$ from $\mathbb{P}_n = \frac{\sum_{i=1}^n W_{i,n} \delta_i}{\sum_{i=1}^n W_{i,n}}$;
- 4 generate X_{n+1} from the mixture q_n^* defined in (7);
- 5 $W_{n+1,n+1} \leftarrow \left(\frac{f_u(X_{n+1})}{q_n^*(X_{n+1})} \right)^\eta \gamma_{n+1}$;
- 6 for all $1 \leq i \leq n$, $W_{i,n+1} \leftarrow (1 - \gamma_{n+1})W_{i,n}$;

end

The computing cost of the n -th iteration is then of order $\ell_n \log(n)$ instead of $\ell_n \log(n)$ in Algorithm 1, which, when neglecting the operation to update the past weights (line 6 in Algorithm 2 compared to line 5), leads to an overall computing time $c_K \sum_{n=1}^N \ell_n \log(n) + Nc_f$. The assumption that line 6 is negligible compared to line 5 is observed in practice when using the Gaussian kernel. More importantly, in the case when $\gamma_i = 1/i$, line 6 is not necessary anymore and therefore leading to the mentioned computing time.

4.3 Mini-batching variant

In the context of stochastic algorithms, the utilization of mini-batches, each consisting of m particles, is a common methodological refinement. At each iteration $n+1$, instead of generating a single particle X_{n+1} according to the density q_n , the algorithm generates m particles $X_{n+1,1}, \dots, X_{n+1,m}$, each sampled from the same density q_n .

The update formula for the unnormalized density, originally delineated by (3), is accordingly modified to accommodate this batch processing strategy:

$$g_{n+1}(x) = (1 - \gamma_{n+1})g_n(x) + \gamma_{n+1} \frac{1}{m} \sum_{k=1}^m w_{n+1,k}^\eta K_{b_{n+1}}(x - X_{n+1,k}), \quad \forall x \in \mathbb{R}^d,$$

where the weights $w_{n+1,k}$ are computed as $w_{n+1,k} = f(X_{n+1,k})/q_n(X_{n+1,k})$.

This mini-batch approach offers the advantage of preserving the total number of particles, while reducing the computational time as it can generate multiple particles at each step through parallelization. However, it is worth noting that the fewer number of updates potentially impacts the adaptivity of the sampling mechanism. Thus, the choice of m represents a trade-off between computational efficiency and adaptive capability.

It is noteworthy that the convergence properties of our algorithm remain intact even under mini-batch adaptations. Indeed, the conditional expectation $\mathbb{E}[g_{n+1}|\mathcal{F}_n]$ remains invariant regardless of whether mini-batching is employed. Consequently, we can extend the proof to the mini-batched context. Moreover, as m is a predetermined constant in our mini-batching setup, the martingale increment term $\mathbb{E}[g_{n+1}|\mathcal{F}_n] - g_{n+1}$ retains its asymptotic characteristics. This further substantiates that the introduction of mini-batching does not perturb the convergence behavior of the algorithm.

5 Numerical experiments

This section is dedicated to the practical evaluation of the MIDAS algorithm (in particular the subsampling version) based on several synthetic examples as well as a real data Bayesian estimation problem. Two recent competitors, from Markov chain Monte Carlo and Sequential Monte Carlo literature, shall be considered for the sake of comparison.

5.1 Competitors

We use the subMIDAS algorithm as described in Algorithm 2, with $\ell_n = \sqrt{n}$ and mini-batches of size $m = 300$ as described in Section 4.3. The bandwidths, mixture weights and step sizes are given by

$$b_n = \frac{0.4}{\sqrt{d}} \left(\frac{mn}{10000} + 1 \right)^{-1/(4+d)}, \quad \lambda_n = \frac{1}{\log(mn + 10)}, \quad \gamma_n = \frac{1}{(n + 10)}.$$

To allow reasonable initialization, we implement a burn-in phase as follows: at $n = 1$, we set an initial batch size of $m_0 = 2000$ and for the first ten steps ($n \leq 10$), we set $\lambda_n = 0.5$.

A key feature of our algorithm is that it generates random variables with a density known up to a normalization constant, and it does not require the gradient of the density. We hence compare with two other algorithms sharing the same specifications: Annealed Importance Sampling (AIS) [Nea01] and Kernel Adaptive Metropolis-Hastings (KAMH) [SSG⁺14].

For AIS, we use a batch size of 300 with 20 Metropolis updates. Therefore, for a given number K of intermediate distributions, we evaluate the unnormalized target density function $K \times 300 \times 20$ times. We use intermediate distributions in the form described by [Nea01], with a geometrically spaced schedule. We run AIS independently for various numbers of intermediate distributions K , and in the final iteration, we plot the distance to the target distribution with respect to the number of evaluations of the target, $K \times 300 \times 20$.

For KAMH, in each run, the starting particle is generated from q_0 . We use a Gaussian kernel with a covariance matrix $\sigma^2 I_d$ where $\sigma = 5$, and we set the scaling parameters to $\nu = 2.38/\sqrt{d}$ and $\gamma = 0.2$. In this case, one step corresponds to one evaluation of the unnormalized target density function.

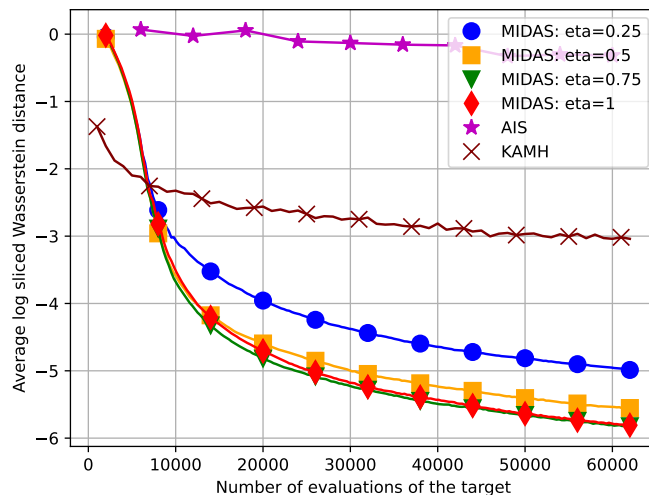


Figure 2: Cold start

5.2 Evaluation

For each competing method, we evaluate the distance to the target distribution with respect to the budget, i.e., the number of evaluation of f_u . After a given number of evaluations of f_u , each method (subMIDAS, AIS, KAMH) gives an empirical distribution described with the help weighted particles. To evaluate the performance, we compute the sliced Wasserstein distance between the empirical distribution and the target distribution. Recall that for two probability measures μ and ν on \mathbb{R}^d with finite second moments,

$$SW_2(\mu, \nu) := \mathbb{E}_{\theta \sim \mathcal{U}(\mathbb{S}^{d-1})} [W_2(\theta_{\#}\mu, \theta_{\#}\nu)^2].$$

Here, W_2 denotes the Wasserstein distance on probability measures on \mathbb{R} with finite second moments, \mathbb{S}^{d-1} denotes the unit sphere $\{\theta \in \mathbb{R}^d : \|\theta\| = 1\}$, $\mathcal{U}(\mathbb{S}^{d-1})$ denotes the uniform distribution on \mathbb{S}^{d-1} , and $\theta_{\#}\mu$ denotes the pushforward of the measure μ through the map $x \in \mathbb{R}^d \mapsto \langle \theta, x \rangle$.

With the help of Monte Carlo simulation from $\mathcal{U}(\mathbb{S}^{d-1})$, this distance is easily approximated since the Wasserstein distance W_2 on measures on \mathbb{R} admits a closed form. In the presented graphs, we average the sliced Wasserstein distance over 50 runs.

5.3 Toy examples

Cold start

The "cold start" scenario occurs when the target density is far from the initial density. In this case, our target distribution is a Gaussian with a mean of $5/\sqrt{d}1_d$ and a variance of $(0.4)^2/dI_d$. The initial distribution is a Gaussian with a mean of 0 and a covariance matrix: $5/dI_d$.

The results are given in Figure 2. In such a simple example of a target with a single mode, $\eta = 1$ seems to be a reasonable choice. While taking $\eta = 3/4$ makes the algorithm slightly faster at the start, it appears that taking a small value for η might slow the convergence of MIDAS. AIS performs poorly, as the nature of the algorithm leads to high variance in the weights when the starting distribution is far from the target.

Gaussian mixture

In this case, the target distribution is a mixture of two Gaussians with equal weights and mean equal to $1/(2\sqrt{d})1_d$ and $-1/(2\sqrt{d})1_d$. The covariance matrix of the two Gaussians is $0.4^2/dI_d$. The initial distribution is a Student's distribution with mean 0 and scale parameter $5/dI_d$.

In this multi-modal target example, we clearly see the importance of η . In this case, the algorithm performs better for the lowest value of η . Furthermore, for all four values of η , MIDAS outperforms our competitors.

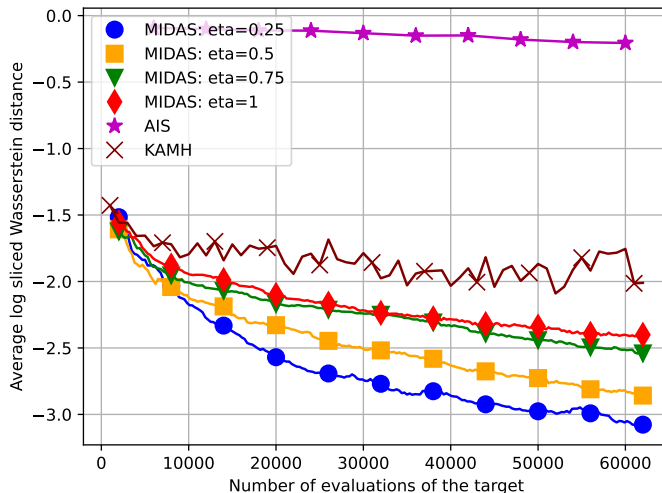


Figure 3: Gaussian Mixture

Anisotropic Gaussian mixture

In this case, the target distribution is a mixture of two Gaussians with mean equal to $1/(2\sqrt{d})1_d$ and $-1/(2\sqrt{d})1_d$. The covariance matrix of the two Gaussians is $0.4^2/d\text{Diag}(10, 1, \dots, 1)$. The initial distribution is a Student's distribution with a location parameter of 0 and a scale parameter of $5/dI_d$.

This case produces results that are similar to those of the Gaussian mixture case.

5.4 Bayesian logistic regression in real data set

We consider the Bayesian logistic regression setting of [GHB12], also considered in the recent Bayesian inference literature [LW16, KP22, LLK⁺23]. More precisely, we consider a data set

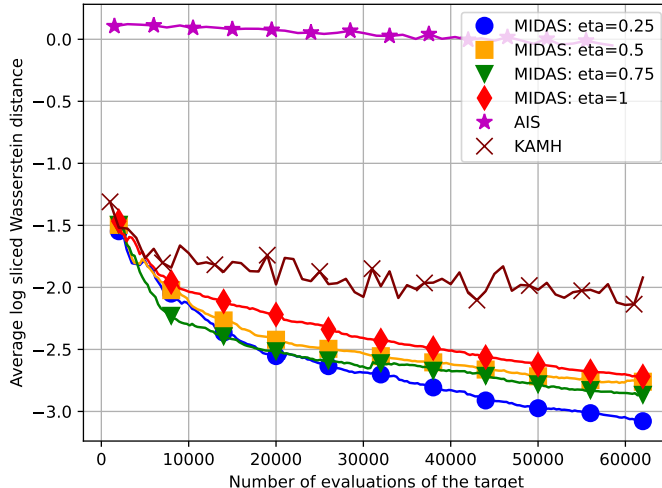


Figure 4: Anisotropic Gaussian Mixture

$\mathcal{D} = \{c_i, z_i\}_{i \in I}$ of points $z_i \in \mathbb{R}^{d-1}$ and binary class labels $c_i \in \{-1, 1\}$ for $i \in I$ where d is the dimension.

The model is parameterized by a parameter $\theta = (w, \beta) \in \mathbb{R}^{d-1} \times \mathbb{R}$ and the hyper parameters $a, b \in \mathbb{R}$. For each $i \in I$, we have

$$\mathbb{P}(c_i = 1 | \theta, z_i) = \frac{1}{1 + e^{-w^T z_i}}. \quad (8)$$

The parameter β follows a Gamma distribution of shape parameter $a = 1$ and rate parameter $b = 0.01$. The parameter w conditionally to β follows a Gaussian distribution of mean 0 and variance $(1/\beta)I_{d-1}$.

We are interested in estimating the posterior density $p(\theta | \mathcal{D}_{\text{train}})$ of the parameter θ according to a training data set $\mathcal{D}_{\text{train}}$ which is given by:

$$p(\theta | \mathcal{D}_{\text{train}}) = \frac{p(\mathcal{D}_{\text{train}} | \theta) p(\theta)}{\mathbb{P}(\mathcal{D}_{\text{train}})}$$

where the prior $p(\theta)$ is the density of the parameter θ given by the model, the likelihood $p(\mathcal{D}_{\text{train}} | \theta)$ is the probability of obtaining the data set $\mathcal{D}_{\text{train}}$ with parameter θ and is given by (8) and the marginal likelihood $\mathbb{P}(\mathcal{D}_{\text{train}})$ is the probability of obtaining the data set $\mathcal{D}_{\text{train}}$ which does not need to be computed here. Thus, the unnormalized target distribution is given by $p(\mathcal{D}_{\text{train}} | \theta) p(\theta)$.

Given a new data point z_{new} , we are interested in predicting the label c_{new} . Using the posterior density and (8), we have: $\mathbb{P}(c_{\text{new}} = 1 | z_{\text{new}}, \mathcal{D}_{\text{train}}) = \int \mathbb{P}(c_{\text{new}} = 1 | \theta, z_{\text{new}}) p(\theta | \mathcal{D}_{\text{train}}) d\theta$. And if the computed probability exceeds 1/2, we infer that $c_{\text{new}} = 1$. Otherwise, we assign $c_{\text{new}} = -1$.

We consider the dataset 'waveform' made of 5000 entries with dimension $d = 22$. For each competing method, the predictions are made based on the training set $\mathcal{D}_{\text{train}}$ of size 400 and the average accuracy is computed with the help of the test dataset made of the remaining 4600 points. This is displayed in Figure 5 where we observe that MIDAS with $\eta = 1/4$ outperforms all competitors. Moreover, in this case, the classical algorithm with $\eta = 1$ performs even worse than the competitors like AIS and KAMH.

6 Possible extensions

Several extensions of this work may be worth exploring further. We conjecture that the subsampling algorithm subMIDAS satisfies similar convergence properties as the one described in Theorem 1 depending on the choice of the parameters λ_n , γ_n and b_n with respect to ℓ_n . This is left for further research. In addition, we believe that the assumptions on the previous sequences, in case $\eta < 1/2$, might be alleviated at the price of a more thorough examination.

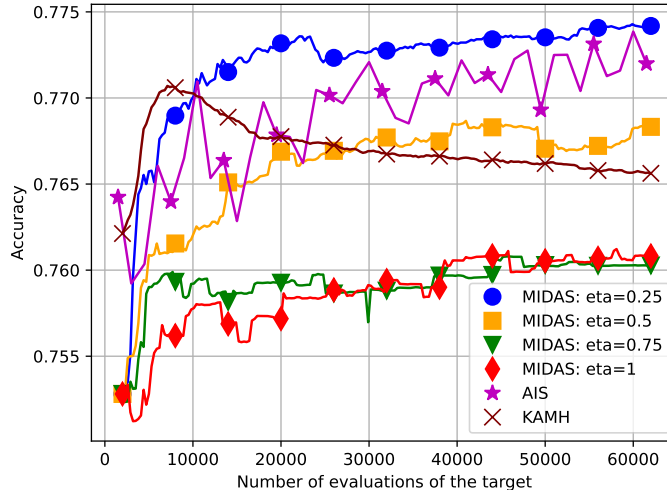


Figure 5: Bayesian Logistic Regression

The adaptive choice of the learning rate η might help to improve the convergence rate and the practical behavior of the algorithm. As noted in [CCK23a], our algorithm can be viewed as an annealed importance sampling (AIS) algorithm. Such an algorithm chooses a distribution path f_k of the form $f_k \propto q_0^{1-\beta_k} f^{\beta_k}$, where the annealing schedule satisfies $\beta_N = 1$ for the final iteration N . The annealing schedule corresponding to our algorithm is $\beta_n = 1 - (1 - \eta)^n$ as studied in [CCK23a], for instance. There are many choices of schedules that work well for different types of problems. The existence of an optimal schedule that minimizes the variance of the estimates is discussed in [SBCCD24]. Studying optimal choice of the parameter η in our context could have some interesting connections with the previously mentioned papers. We leave this investigation for future work.

References

- [APB22] Grégoire Aufort, Pierre Pudlo, and Denis Burgarella. Tempered, anti-truncated, multiple importance sampling, 2022.
- [BBD⁺19] Alessandro Barp, Francois-Xavier Briol, Andrew Duncan, Mark Girolami, and Lester Mackey. Minimum stein discrepancy estimators. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [BBL⁺08] Thomas Bengtsson, Peter Bickel, Bo Li, et al. 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*, pages 316–334. Institute of Mathematical Statistics, 2008.
- [BCD19] Bernard Bercu, Sami Capderou, and Gilles Durrieu. Nonparametric recursive estimation of the derivative of the regression function with application to sea shores water quality. *Statistical Inference for Stochastic Processes*, 22:17–40, 2019.
- [BEM⁺17] Monica F Bugallo, Victor Elvira, Luca Martino, David Luengo, Joaquin Miguez, and Petar M Djuric. Adaptive importance sampling: The past, the present, and the future. *IEEE Signal Processing Magazine*, 34(4):60–79, 2017.
- [BKM17] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.

- [BT03] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [CCK23a] N. Chopin, F. Crucinio, and A. Korba. A connection between tempering and entropic mirror descent. In *Forty-first International Conference on Machine Learning, 2023*.
- [CCK23b] Nicolas Chopin, Francesca R Crucinio, and Anna Korba. A connection between tempering and entropic mirror descent, 2023.
- [Cho04] Nicolas Chopin. Central limit theorem for sequential monte carlo methods and its application to bayesian inference. *The Annals of Statistics*, 32(6):2385–2411, 2004.
- [DHDS16] Bo Dai, Niao He, Hanjun Dai, and Le Song. Provable bayesian inference via particle mirror descent. In *Artificial Intelligence and Statistics*, pages 985–994, 2016.
- [DM07] Randal Douc and Eric Moulines. Limit theorems for weighted samples with applications to sequential monte carlo methods. In *ESAIM: Proceedings*, volume 19, pages 101–107. EDP Sciences, 2007.
- [DMDJ06] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.
- [DP21] Bernard Delyon and François Portier. Safe adaptive importance sampling: A mixture approach. *The Annals of Statistics*, 49(2):885–917, 2021.
- [DW80] Luc P Devroye and TJ Wagner. On the L_1 convergence of kernel estimators of regression functions with applications in discrimination. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 51:15–25, 1980.
- [EC22] Víctor Elvira and Emilie Chouzenoux. Optimized population monte carlo. *IEEE Transactions on Signal Processing*, 70:2489–2501, 2022.
- [EKPS23] Michael Engel, Oindrila Kanjilal, Iason Papaioannou, and Daniel Straub. Bayesian updating and marginal likelihood estimation by cross entropy based importance sampling. *Journal of Computational Physics*, 473:111746, 2023.
- [GCE24] Thomas Guilmeau, Emilie Chouzenoux, and Víctor Elvira. A divergence-based condition to ensure quantile improvement in black-box global optimization, 2024.
- [GHB12] Samuel J Gershman, Matthew D Hoffman, and David M Blei. Nonparametric variational inference. In *Proceedings of the 29th International Conference on Machine Learning*, pages 235–242, 2012.
- [GR96] Geof H. Givens and Adrian E. Raftery. Local adaptive importance sampling for multivariate densities with strong nonlinear relationships. *Journal of the American Statistical Association*, 91(433):132–141, 1996.
- [GSS93] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proc. F Radar Signal Process. UK*, 140(2):107, 1993.
- [HBWP13] Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347, 2013.
- [HCEP23] Yunshi Huang, Emilie Chouzenoux, Víctor Elvira, and Jean-Christophe Pesquet. Efficient bayes inference in neural networks through adaptive importance sampling. *Journal of the Franklin Institute*, 360(16):12125–12149, 2023.
- [HH14] Peter Hall and Christopher C Heyde. *Martingale limit theory and its application*. Academic press, 2014.
- [HNS19] Josiah Hanna, Scott Niekum, and Peter Stone. Importance sampling policy evaluation with an estimated behavior policy. In *International Conference on Machine Learning*, pages 2605–2613. PMLR, 2019.

- [KM15] Eugenia Koblents and Joaquín Míguez. A population monte carlo scheme with transformed weights and its application to stochastic kinetic models. *Statistics and Computing*, 25(2):407–425, 2015.
- [KP22] Anna Korba and François Portier. Adaptive importance sampling meets mirror descent: a bias-variance tradeoff. In *International Conference on Artificial Intelligence and Statistics*, pages 11503–11527. PMLR, 2022.
- [KSA⁺20] Anna Korba, Adil Salim, Michael Arbel, Giulia Luise, and Arthur Gretton. A non-asymptotic analysis for stein variational gradient descent. *Advances in Neural Information Processing Systems*, 33:4672–4682, 2020.
- [LLK⁺23] Lingxiao Li, Qiang Liu, Anna Korba, Mikhail Yurochkin, and Justin Solomon. Sampling with mollified interaction energy descent. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.
- [LT16] Yingzhen Li and Richard E Turner. Rényi divergence variational inference. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [LW16] Qiang Liu and Dilin Wang. Stein variational gradient descent: a general purpose bayesian inference algorithm. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 2378–2386, 2016.
- [MPFR18] Alberto Maria Metelli, Matteo Papini, Francesco Faccio, and Marcello Restelli. Policy optimization via importance sampling. *Advances in Neural Information Processing Systems*, 31:5442—5454, 2018.
- [MPS09] Abdelkader Mokkadem, Mariane Pelletier, and Yousri Slaoui. The stochastic approximation method for the estimation of a multivariate probability density. *Journal of Statistical Planning and Inference*, 139(7):2459–2478, 2009.
- [Nea01] Radford M Neal. Annealed importance sampling. *Statistics and computing*, 11:125–139, 2001.
- [NWS14] Deanna Needell, Rachel Ward, and Nati Srebro. Stochastic gradient descent, weighted sampling, and the randomized kaczmarz algorithm. *Advances in neural information processing systems*, 27:1017–1025, 2014.
- [OB92] Man-Suk Oh and James O. Berger. Adaptive importance sampling in Monte Carlo integration. *J. Statist. Comput. Simulation*, 41(3-4):143–168, 1992.
- [OZ00] Art Owen and Yi Zhou. Safe and effective importance sampling. *J. Amer. Statist. Assoc.*, 95(449):135–143, 2000.
- [Par62] Emanuel Parzen. On Estimation of a Probability Density Function and Mode. *The Annals of Mathematical Statistics*, 33(3):1065 – 1076, 1962.
- [PD18] François Portier and Bernard Delyon. Asymptotic optimality of adaptive importance sampling. In *Advances in Neural Information Processing Systems*, 2018.
- [RM51] Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400 – 407, 1951.
- [SBCCD24] S. Syed, A. Bouchard-Côté, K. Chern, and A. Doucet. Optimised annealed sequential monte carlo samplers. *arXiv preprint arXiv:2408.12057*, 2024.
- [SBCR16] Umut Simsekli, Roland Badeau, Taylan Cemgil, and Gaël Richard. Stochastic quasi-newton langevin monte carlo. In *International Conference on Machine Learning*, pages 642–651. PMLR, 2016.
- [SSG⁺14] Dino Sejdinovic, Heiko Strathmann, Maria Lomeli Garcia, Christophe Andrieu, and Arthur Gretton. Kernel adaptive metropolis-hastings. In *International conference on machine learning*, pages 1665–1673. PMLR, 2014.

- [Sto80] Charles J. Stone. Optimal Rates of Convergence for Nonparametric Estimators. *The Annals of Statistics*, 8(6):1348 – 1360, 1980.
- [Wes93] Mike West. Approximating posterior distributions by mixtures. *Journal of the Royal Statistical Society: Series B (Methodological)*, 55(2):409–422, 1993.
- [WT11] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688. Citeseer, 2011.
- [Zha96] Ping Zhang. Nonparametric importance sampling. *Journal of the American Statistical Association*, 91(435):1245–1253, 1996.

A Proofs

A.1 Notation

Let $\mathcal{C}(\mathbb{R}^d)$ denote the set of continuous functions on \mathbb{R}^d . We denote by $B(x, \varepsilon)$ the d -dimensional open ball of radius ε and center x . We define $(x)^+ = \max(x, 0)$. We use the convention that $\inf \emptyset = +\infty$, and that $0^0 = 1$. For positive sequences $(a_n)_{n \geq 1}$, $(b_n)_{n \geq 1}$, the notation $a_n = O(b_n)$ means that there exists $C > 0$ such that $a_n \leq C b_n$ for all $n \geq 1$. For a sequence $(a_n)_{n \geq 0}$, we denote $\underline{\lim} a_n := \lim_{n \rightarrow \infty} \inf_{k \geq n} a_k$ and $\overline{\lim} a_n := \lim_{n \rightarrow \infty} \sup_{k \geq n} a_k$.

A.2 Proof of Theorem 1

We remark that the sequence $(q_n)_{n \geq 0}$ constructed with f_u is the same as the one using f . Indeed, for a target f normalized, the algorithm gives two sequences $(q_n)_{n \geq 0}$ and $(g_n)_{n \geq 0}$. Now we define $(\tilde{q}_n)_{n \geq 0}, (\tilde{g}_n)_{n \geq 0}$ the sequences associated to the algorithm with target $\tilde{f} = Cf$. We see by induction that for all $n \geq 0$, $\tilde{g}_n = C^n g_n$ and $\tilde{q}_n = q_n$. Then, in the rest of the proof we will consider $f_u = f$ so that for each $n \geq 1$,

$$w_n = \frac{f(X_n)}{q_{n-1}(X_n)}.$$

We define the operator T on the set of non-negative continuous functions on \mathbb{R}^d by:

$$Tg : x \mapsto f(x)^\eta g(x)^{1-\eta},$$

for all g . The iterates (g_n) given by (3) can be rewritten as:

$$g_{n+1} = (1 - \gamma_{n+1})g_n + \gamma_{n+1}Tq_n * K_{b_{n+1}} + \gamma_{n+1}\xi_{n+1}, \quad (9)$$

where ξ_{n+1} is a r.v. on $\mathcal{C}(\mathbb{R}^d)$ given by:

$$\xi_{n+1}(x) := w_{n+1}^\eta K_{b_{n+1}}(x - X_{n+1}) - Tq_n * K_{b_{n+1}}(x). \quad (10)$$

We remark that, for every $x \in \mathbb{R}^d$, $\mathbb{E}(\xi_{n+1}(x) | \mathcal{F}_n) = 0$. Iterating (9) with

$$\psi_{i+1}^n := \prod_{j=i+1}^n (1 - \gamma_j), \quad (11)$$

we obtain for every $n \geq 0$

$$g_n = \sum_{i=1}^n \psi_{i+1}^n \gamma_i Tq_{i-1} * K_{b_i} + M_n, \quad (12)$$

where the martingale term is defined as

$$M_n := \sum_{i=1}^n \psi_{i+1}^n \gamma_i \xi_i. \quad (13)$$

We define $Z_n := \int g_n$. By integrating (9), we have:

$$Z_{n+1} = (1 - \gamma_{n+1})Z_n + \gamma_{n+1} \int Tq_n + \gamma_{n+1} \int \xi_{n+1}.$$

We are now ready to proceed with the main steps of the proof.

Step 1: Control of the martingale term. We obtain the following bound on the martingale term M_n defined previously.

Proposition 2. *Let Assumptions 1, 2 and 3 hold true. There exists $\varepsilon > 0$ such that for any $p > 0$ the process M_n of equation (13) satisfies:*

$$\lim_{n \rightarrow \infty} \frac{\sup_{\|x\| \leq n^p} |M_n(x)|}{\lambda_n^{1-\eta}} = 0, \quad (14)$$

almost surely.

Step 2: An upper bound on Z_n . By applying Hölder's inequality, $\int Tq_n \leq 1$. Consequently, if we define the sequence $(\bar{Z}_n)_{n \geq 0}$ with $\bar{Z}_0 = Z_0$ as

$$\bar{Z}_{n+1} := (1 - \gamma_{n+1})\bar{Z}_n + \gamma_{n+1} + \gamma_{n+1} \int \xi_{n+1}, \quad (15)$$

we obtain $Z_n \leq \bar{Z}_n$ for every n a.s. .

We obtain the following result.

Proposition 3. *Suppose $f = f_u$, let Assumptions 1 and 3 hold true. $(\bar{Z}_n)_{n \geq 0}$ defined in (15) satisfies $\lim_{n \rightarrow \infty} \bar{Z}_n = 1$, a.s..*

Step 3: g_n is away from 0. By applying a convexity inequality ($1 - \eta < 1$) we obtain:

$$Tq_n(x) \geq \left(\frac{f(x)}{2} \right)^\eta \left((1 - \lambda_n)^{1-\eta} \left(\frac{g_n(x)}{Z_n} \right)^{1-\eta} + \lambda_n^{1-\eta} q_0(x)^{1-\eta} \right),$$

for every x . We define the operator \underline{T} on the set on non-negative continuous function on \mathbb{R}^d as

$$\underline{T}g : x \mapsto \left(\frac{f(x)}{2} \right)^\eta g(x)^{1-\eta},$$

for every g . Then, using that $\bar{Z}_n \geq Z_n$, we obtain

$$g_{n+1} \geq (1 - \gamma_{n+1})g_n + \gamma_{n+1} \frac{(1 - \lambda_n)^{1-\eta}}{\bar{Z}_n^{1-\eta}} \underline{T}g_n * K_{b_{n+1}} + \gamma_{n+1} \lambda_n^{1-\eta} \underline{T}q_0 * K_{b_{n+1}} + \gamma_{n+1} \xi_{n+1},$$

for every $n \in \mathbb{N}$. Iterating the latter equation, we obtain for every $n \geq 0$

$$g_n \geq v_n + M_n + \sum_{i=1}^n \psi_{i+1}^n \gamma_i \lambda_i^{1-\eta} \underline{T}q_0 * K_{b_i},$$

where

$$v_n := \sum_{i=1}^n \psi_{i+1}^n \gamma_i \frac{(1 - \lambda_i)^{1-\eta}}{\bar{Z}_n^{1-\eta}} \underline{T}g_{i-1} * K_{b_i}.$$

By the martingale control given by Proposition 2, we obtain:

Lemma 1. *Let Assumptions 1, 2 and 3 hold true. For any $\alpha > 0$ and $p > 0$, we have that almost surely, there exists $N_{\alpha,p} \in (0, \infty)$ such that for every $n \geq N_{\alpha,p}$*

$$\min_{f(x) \geq \alpha, \|x\| \leq n^p} (g_n - v_n)(x) \geq 0. \quad (16)$$

The previous result allows to show that g_n cannot reach 0 in places where f is positive.

Proposition 4. *Let Assumptions 1, 2 and 3 hold true. Then for every $\varepsilon > 0$, and every $x \in \mathbb{R}^d$ satisfying $\inf_{y \in B(x, \varepsilon)} f(y) > 0$, we obtain*

$$\begin{aligned} \liminf_{y \in B(x, \varepsilon)} v_n(y) &> 0, \\ \liminf_{y \in B(x, \varepsilon)} g_n(y) &> 0, \end{aligned}$$

almost surely.

Step 4: A lower bound on g_n

Proposition 5. *Let Assumptions 1, 2 and 3 hold true. Then for every compact set $A \subset \mathbb{R}^d$, almost surely, we obtain*

$$\liminf_{x \in A} g_n(x) - f(x) \geq 0.$$

According to Proposition 5, we obtain

$$\underline{\lim}_{x \in A} \inf (g_n(x) - f(x)) \geq 0, \quad (17)$$

for every compact set $A \subset \mathbb{R}^d$. Let us fix $\varepsilon > 0$, by Assumption 1, there exists a compact set A_ε such that:

$$\int_{A_\varepsilon} f(x) dx \geq 1 - \varepsilon.$$

Moreover,

$$\int_{A_\varepsilon} (g_n - f) \geq |A_\varepsilon| \inf_{x \in A_\varepsilon} (g_n(x) - f(x)).$$

By (17), we obtain

$$\underline{\lim} Z_n \geq \underline{\lim} \int_{A_\varepsilon} g_n \geq \int_{A_\varepsilon} f \geq 1 - \varepsilon,$$

for every $\varepsilon > 0$, which gives the result.

$$Z_n \xrightarrow{a.s.} 1, \quad (18)$$

under the assumption of Proposition 5.

Step 5: An upper bound on g_n

Proposition 6. *Suppose $f = f_u$, let Assumptions 1, 2, 3 hold true. Then, for every compact $A \subset \mathbb{R}^d$,*

$$\overline{\lim}_{x \in A} \sup (g_n(x) - f(x)) \leq 0.$$

And, using the Propositions 5 and 6, we obtain Theorem 1, leading to Corollary 2.

A.3 Proof of Proposition 2

In this subsection, we suppose $f = f_u$, and we let Assumptions 1, 2, and 3 hold true. The martingale increment $(\xi_i)_{i \geq 1}$ can also be rewritten as:

$$\xi_i(x) = w_i^\eta K_{b_i}(x - X_i) - \mathbb{E}[w_i^\eta K_{b_i}(x - X_i) | \mathcal{F}_{i-1}], \quad (19)$$

where $X_i \sim q_{i-1}$ and $w_i = \frac{f}{q_{i-1}}(X_i)$. By Hölder's inequality, we obtain a useful inequality

$$\mathbb{E}[w_{n+1}^\eta | \mathcal{F}_n] \leq 1. \quad (20)$$

The purpose of the next lemmas is to prove Proposition 2.

Lemma 2. *There exists a constant $C > 0$, depending only on K and f , such that the following statements hold for all $n \geq i \geq 0$, $x, y \in \mathbb{R}^d$,*

$$|\xi_i(x+y) - \xi_i(x)| \leq C \lambda_{i-1}^{-\eta} b_i^{-d-1} \|y\| \quad (21)$$

$$|\xi_i(x)| \leq C \lambda_{i-1}^{-\eta} b_i^{-d} \quad (22)$$

$$\mathbb{E}(\xi_i(x)^2 | \mathcal{F}_{j-1}) \leq C b_i^{-d(1+(1-2\eta)_+)} \lambda_{i-1}^{-(2\eta-1)_+}. \quad (23)$$

Proof. Notice first that the following bound is a consequence of Assumption 1-iii and (4):

$$w_i^\eta \leq (c \lambda_{i-1})^{-\eta}. \quad (24)$$

We establish the first point. Set $\Delta_i(x, y) := K_{b_i}(x+y - X_i) - K_{b_i}(x - X_i)$ and remark that $|\Delta_i(x, y)| \leq L_K \|y\| b_i^{-d-1}$, where L_K the Lipschitz constant of K (cf. Assumption 2-i). Recalling the definition of ξ_i in (19), we obtain:

$$\begin{aligned} |\xi_i(x+y) - \xi_i(x)| &= |w_i^\eta \Delta_i(x, y) - \mathbb{E}(w_i^\eta \Delta_i(x, y) | \mathcal{F}_{i-1})| \\ &\leq w_i^\eta |\Delta_i(x, y)| + \mathbb{E}(w_i^\eta |\Delta_i(x, y)| | \mathcal{F}_{i-1}) \\ &= (w_i^\eta + \mathbb{E}(w_i^\eta | \mathcal{F}_{i-1})) L_K \|y\| b_i^{-d-1}. \end{aligned}$$

We conclude by using (20) and (24). The second point follows from:

$$\begin{aligned} |\xi_i(x)| &= |w_i^\eta K_{b_i}(x - X_i) - \mathbb{E}(w_i^\eta K_{b_i}(x - X_i) | \mathcal{F}_{i-1})| \\ &\leq \|K\|_\infty b_i^{-d} (w_i^\eta + \mathbb{E}(w_i^\eta | \mathcal{F}_{i-1})), \end{aligned}$$

and (20) and (24) again. We establish the third point.

$$\begin{aligned} \mathbb{E}[\xi_{i+1}(x)^2 | \mathcal{F}_i] &\leq \mathbb{E}[w_{i+1}^{2\eta} K_{b_{i+1}}(x - X_{i+1})^2 | \mathcal{F}_i] \\ &= \int f(y)^{2\eta} q_i(y)^{-2\eta} K_{b_{i+1}}(x - y)^2 q_i(y) dy \\ &\leq \|f\|_\infty^{2\eta} \int K_{b_{i+1}}(x - y)^2 q_i(y)^{1-2\eta} dy. \end{aligned}$$

If $\eta < \frac{1}{2}$, we use the Hölder inequality:

$$\begin{aligned} \mathbb{E}[\xi_{n+1}(x)^2 | \mathcal{F}_n] &\leq \|f\|_\infty^{2\eta} \left(\int K_{b_{n+1}}(x - y)^{1/\eta} dy \right)^{2\eta} \left(\int q_i(y) dy \right)^{1-2\eta} \\ &\leq \|f\|_\infty^{2\eta} \left(\int b_{i+1}^{-d/\eta} K(b_{i+1}^{-1}(x - y))^{1/\eta} dy \right)^{2\eta} \\ &= \|f\|_\infty^{2\eta} b_{i+1}^{-2d(1-\eta)} \|K\|_{1/\eta}^2, \end{aligned}$$

which is finite by Assumptions 1- i and 2- i. If $\eta \geq \frac{1}{2}$, we use that $q_i \geq \lambda_i q_0$ (cf. (4)) and Assumption 1-i:

$$\begin{aligned} \mathbb{E}[\xi_{i+1}(x)^2 | \mathcal{F}_i] &\leq \|f\|_\infty^{2\eta} \lambda_i^{1-2\eta} \int K_{b_{n+1}}(x - y)^2 q_0(y)^{1-2\eta} dy \\ &\leq \|f\|_\infty^{2\eta} \|q_0\|_\infty^{1-2\eta} \|K\|_2^2 \lambda_i^{1-2\eta} b_{i+1}^{-d} \end{aligned}$$

In any case, (23) is satisfied. \square

Lemma 3. *There exist a constant C and a rank n_0 depending only on quantities given in the assumptions, such that, for all $A > 0$, $n \geq 1$ and $q \geq 1$:*

$$\begin{aligned} \mathbb{P} \left(\sup_{\|x\| \leq A} |M_n(x)| > Cq\tau_n \sqrt{\log n} \right) &\leq C \left(1 + A \frac{r_n}{\tau_n} \right)^d n^{-q} \quad n \geq n_0 \\ \tau_n &:= \left(n\lambda_n^{-(2\eta-1)+} b_n^{-d(1+(1-2\eta)_+)} \right)^{1/2} \\ r_n &= n\lambda_n^{-\eta} b_n^{-d-1}. \end{aligned}$$

Proof. Let $n \in \mathbb{N}^*$. For any $i \leq n$, and any $x \in \mathbb{R}^d$, we define

$$\zeta_{i,n}(x) := \gamma_i \psi_{i+1}^n \xi_i(x).$$

Note that $\mathbb{E}(\zeta_{i,n}(x) | \mathcal{F}_{i-1}) = 0$ for all $i \leq n$. We will apply Proposition 7. We consider $\varepsilon > 0$ to be chosen later, and, following this proposition, denote

$$\tilde{\zeta}_{i,n}(x) := \sup_{\|y\| \leq \varepsilon} |\zeta_{i,n}(x+y) - \zeta_{i,n}(x)|.$$

By Lemmas 2, 9 and Assumption 3, the following inequalities hold, for all $x \in \mathbb{R}^d$:

$$\begin{aligned} \max_{i=1, \dots, n} |\zeta_{i,n}(x)| &\leq C\gamma_n \lambda_{n-1}^{-\eta} b_n^{-d} =: m \\ \sum_{i=1}^n \mathbb{E}(\zeta_{i,n}(x)^2 | \mathcal{F}_{i-1}) &\leq C\gamma_n^2 \sum_{i=1}^n \left(b_i^{-d-d(1-2\eta)_+} \lambda_{i-1}^{-(2\eta-1)_+} \right) = C\gamma_n^2 \tau_n^2 =: v \\ \sum_{i=1}^n \mathbb{E}(\tilde{\zeta}_{i,n}(x) | \mathcal{F}_{i-1}) &\leq C\varepsilon\gamma_n \sum_{i=1}^n \lambda_{i-1}^{-\eta} b_i^{-d-1} \leq C\varepsilon\gamma_n n\lambda_n^{-\eta} b_n^{-d-1} = C\gamma_n \varepsilon r_n =: u. \end{aligned}$$

We apply Proposition 7 with

$$\begin{aligned} t &:= C_0 \sqrt{v \log n} = C_0 \gamma_n \tau_n \sqrt{C \log n} \\ \varepsilon &:= t (C r_n \gamma_n)^{-1}, \end{aligned}$$

for some $C_0 \geq 1$. With this specific choice of ε , we remark that $u = t$. In order to compute the factor of t^2 in the exponent of (43), let us show first that $tm \leq C_0 v$ for n large enough. This is equivalent to $m^2 \log n \leq v$, and is true if

$$\lim_{n \rightarrow \infty} \frac{\lambda_n^{(2\eta-1)+} b_n^{d(1-2\eta)+} \log n}{n \lambda_n^{2\eta} b_n^d} = 0. \quad (25)$$

Since $\gamma_n \geq \frac{C_\gamma}{n}$, it is also true if,

$$\lim_{n \rightarrow \infty} \frac{n \gamma_n^2 \log n}{\lambda_n b_n^d} = 0.$$

The latter is true under Assumption 3-iii. Thus, (25) is true for $n \geq n_0$ with n_0 that depends only on quantities given in assumptions. It leads to (for $n \geq n_0$):

$$\max(v, 2mu) + \frac{2}{3}mt = \max(v, 2mt) + \frac{2}{3}mt \leq 3C_0 v.$$

We obtain for $n \geq n_0$, a bound on the exponent in (43):

$$\frac{t^2}{8(\max(v, 2mu) + \frac{2}{3}mt)} \geq \frac{C_0^2 v \log n}{24C_0 v} = \frac{C_0}{24} \log n.$$

By applying (43) and setting $C_0 = 24q$ for $q \geq 1$, we obtain, recalling $M_n(x) = \sum_{i=1}^n \zeta_{i,n}(x)$,

$$\mathbb{P} \left(\sup_{\|x\| \leq A} |M_n(x)| > 48q \gamma_n \tau_n \sqrt{\log n} \right) \leq C' \left(1 + A \frac{r_n}{\tau_n} \right)^d n^{-q}, \quad n \geq n_0.$$

The conclusion follows, taking $C = \max(48q, C')$. \square

Proof of Proposition 2. By choosing q large enough in Lemma 3, we obtain

$$\begin{aligned} \sum_n \mathbb{P} \left(\sup_{\|x\| \leq n^p} |M_n(x)| > t_n \right) &< \infty \\ t_n &= C \gamma_n q \tau_n \sqrt{\log n}. \end{aligned}$$

For $\omega \in E$ such that $\mathbb{P}(E) = 1$, by Borel-Cantelli's lemma there exists $N(\omega)$ such that

$$|M_n(x)| \leq t_n < \infty \quad \text{if } n > \max(N(\omega), \|x\|^{1/p}).$$

By Assumption 3-iii, $\lim_{n \rightarrow \infty} \frac{t_n^2}{\lambda_n^{2(1-\eta)}} = 0$, which finishes the proof. \square

A.3.1 Proof of Proposition 3

We define $V_n := \bar{Z}_n - 1$. Then, we obtain

$$V_{n+1} = (1 - \gamma_{n+1})V_n + \int \xi_{n+1}$$

Iterating the latter, we obtain

$$V_n = \sum_{i=1}^n \psi_{i+1}^n \gamma_i \int \xi_i$$

We will apply Proposition 7 to the sum $V_n = \sum_{i=1}^n \int \zeta_{i,n}$, where we recall $\int \zeta_{i,n} = \psi_{i+1}^n \gamma_i \int \xi_i = \psi_{i+1}^n \gamma_i w_i^n$. We see that:

$$\mathbb{E} \left[\left(\int \xi_{n+1} \right)^2 \middle| \mathcal{F}_n \right] \leq \mathbb{E} \left[w_{n+1}^{2\eta} \middle| \mathcal{F}_n \right] = \int f^{2\eta} q_n^{1-2\eta}.$$

If $\eta < \frac{1}{2}$, we use Hölder's inequality:

$$\mathbb{E} \left[\left(\int \xi_{n+1} \right)^2 \middle| \mathcal{F}_n \right] \leq \left(\int f \right)^{2\eta} \left(\int q_n \right)^{1-2\eta} = 1.$$

If $\eta \geq \frac{1}{2}$, we use the majoration $q_n \geq \lambda_n q_0 \geq \lambda_n c f$ (cf. (4) and Assumption 1-iii):

$$\mathbb{E} \left[\left(\int \xi_{n+1} \right)^2 \middle| \mathcal{F}_n \right] \leq c^{1-2\eta} \lambda_n^{1-2\eta}.$$

Using Lemma 9, we obtain the bound

$$\sum_{i=1}^n \mathbb{E} \left[\left(\int \zeta_{i,n} \right)^2 \middle| \mathcal{F}_{i-1} \right] \leq \sum_{i=1}^n C \gamma_n^2 \lambda_{i-1}^{-(2\eta-1)_+} \leq C n \gamma_n^2 \lambda_n^{-(2\eta-1)_+}.$$

Moreover, $\max_{i=1\dots n} |\int \zeta_{i,n}| \leq C \gamma_n$. Let $C_0 > 1$, we apply Proposition 7 with $\varepsilon = 1$, $A := 1$, $u := 0$, $v := C n \gamma_n^2 \lambda_n^{-(2\eta-1)_+}$, $m := C \gamma_n$, and $t := C_0 \sqrt{v \log n}$. We claim that $tm \leq v$ for n big enough. This is equivalent to $C_0 m^2 \log n \leq v$. Remaking $\gamma_n \geq \frac{C_0}{n}$, the inequality is true, since by Assumption 3, $n \gamma_n^2 \log n \rightarrow 0$. Consequently, for n big enough

$$\frac{t^2}{8(v + \frac{2}{3}mt)} \geq \frac{C_0^2 v \log n}{16 C_0 v} \geq \frac{C_0}{16} \log n.$$

Consequently we obtain taking $C_0 = 32$

$$\mathbb{P}(|V_n| > 32 \sqrt{v \log n}) \leq C n^{-2}.$$

For $\omega \in E$, such that $\mathbb{P}(E) = 1$, by Borell Cantelli's lemma, there exists $N(\omega)$ such that $|V_n| \leq 32 \sqrt{v \log n}$, for $n \geq N(\omega)$. Since by Assumption 3-iii, $v \log n \rightarrow 0$, Proposition 3 is proven.

A.4 Proof of Lemma 1

The lemma is a direct consequence of the following lemma.

Lemma 4. *Suppose $f = f_u$, let Assumptions 1, 2 and 3 hold true. For all $\alpha > 0$ and $p > 0$, almost surely, there exists $N_{\alpha,p} \in (0, \infty)$ such that,*

$$\min_{f(x) \geq \alpha, \|x\| \leq n^p} \left(\sum_{i=1}^n \psi_{i+1}^n \gamma_i \lambda_{i-1}^{1-\eta} \underline{T} q_0 * K_{b_i} + M_n \right) (x) \geq 0,$$

for every $n \geq N_{\alpha,p}$.

Proof. Let $x \in \mathbb{R}^d$, such that, $f(x) \geq \alpha$. For every $i \geq 1$, $\underline{T} q_0 * K_{b_i}(x) \geq c^{1-\eta} \alpha 2^{-\eta}$. Since $(\lambda_n)_{n \geq 0}$ is nonincreasing,

$$\frac{1}{\lambda_n^{1-\eta}} \sum_{i=1}^n \psi_{i+1}^n \gamma_i \lambda_{i-1}^{1-\eta} \underline{T} q_0 * K_{b_i}(x) \geq c^{1-\eta} \alpha 2^{-\eta} \sum_{i=1}^n \psi_{i+1}^n \gamma_i.$$

By Lemma 10 with $a_i = 1$ and by Proposition 2, almost surely, we obtain

$$\lim_{n \rightarrow \infty} \frac{1}{\lambda_n^{1-\eta}} \min_{f(x) \geq \alpha, \|x\| \leq n^p} \left(\sum_{i=1}^n \psi_{i+1}^n \gamma_i \lambda_{i-1}^{1-\eta} \underline{T} q_0 * K_{b_i} + M_n \right) (x) \geq c^{1-\eta} \alpha 2^{-\eta},$$

which finishes the proof. \square

A.5 Proof of Prop 4

In this subsection, we let Assumptions 1, 2 and 3 hold true. Moreover, we assume $f = f_u$. (v_n) can be viewed recursively as:

$$v_{n+1} = (1 - \gamma_{n+1})v_n + \gamma_{n+1} \frac{(1 - \lambda_n)^{1-\eta}}{\bar{Z}_n^{1-\eta}} \underline{T}g_n * K_{b_{n+1}}, \quad (26)$$

with $v_0 = 0$.

We denote for a real valued function h defines in \mathbb{R}^d :

$$\underline{h}_\varepsilon(x) := \inf_{y \in B(x, \varepsilon)} h(y), \quad (27)$$

for all ε, x . In the case of a sequence of real functions (h_n) , we define

$$\underline{g}_{n, \varepsilon}(x) := \inf_{y \in B(x, \varepsilon)} g_n(y). \quad (28)$$

Let $B \subset \Omega$ such that, Proposition 3 and Lemma 4 hold for every $\omega \in B$ and $\mathbb{P}(B) = 1$. We fix $\varepsilon > 0, x \in \mathbb{R}^d, \omega \in B$ in the whole proof and assume that $\underline{f}_\varepsilon(x) > \alpha$. For any $y \in B(x, \varepsilon)$, we obtain

$$\begin{aligned} \underline{T}g_n(y) &\geq \underline{T}g_n \mathbf{1}_{B(x, \varepsilon')} * K_{b_{n+1}}(y) \\ &\geq 2^{-\eta} \underline{f}_\varepsilon^\eta(x) \underline{g}_{n, \varepsilon'}^{1-\eta}(x) (\mathbf{1}_{B(x, \varepsilon')} * K_{b_{n+1}})(y) \\ &\geq 2^{-\eta} \underline{f}_\varepsilon^\eta(x) \underline{g}_{n, \varepsilon}^{1-\eta}(x) (\mathbf{1}_{B(0, \varepsilon')} * K_{b_{n+1}})(y - x) \\ &\geq 2^{-\eta} \tilde{c}_\varepsilon \underline{f}_\varepsilon^\eta(x) \underline{g}_{n, \varepsilon}^{1-\eta}(x), \end{aligned}$$

where $\tilde{c}_\varepsilon := \inf_{n \geq 0} \inf_{B(0, \varepsilon)} K_{b_{n+1}} * \mathbf{1}_{B(0, \varepsilon)}$ is a non-negative constant by Lemma 8. Going back to (26), for every $y \in B(x, \varepsilon)$, we obtain

$$v_{n+1}(y) \geq (1 - \gamma_{n+1})v_n(y) + \gamma_{n+1} 2^{-\eta} c_{\varepsilon, n} \underline{f}_\varepsilon^\eta(x) \underline{g}_{n, \varepsilon}^{1-\eta}(x),$$

where $c_{\varepsilon, n} := \tilde{c}_\varepsilon \frac{(1 - \lambda_n)^{1-\eta}}{\bar{Z}_n^{1-\eta}}$. By (16), we obtain

$$v_{n+1}(y) \geq (1 - \gamma_{n+1})v_n(y) + \gamma_{n+1} 2^{-\eta} c_{\varepsilon, n} \underline{f}_\varepsilon^\eta(x) \underline{v}_{n, \varepsilon}^{1-\eta}(x),$$

for every $n \geq N_{\alpha, p}(\omega)$. Moreover, there exists $\delta > 0$ and a constant $N_\delta(\omega)$, such that

$$\frac{(1 - \lambda_n)^{1-\eta}}{\bar{Z}_n^{1-\eta}} \geq 1 - \delta$$

for every $n \geq N_\delta(\omega)$ by Prop 3. Consequently, taking the infimum over the ball $B(x, \varepsilon)$, we obtain

$$\underline{v}_{n+1, \varepsilon}(x) \geq (1 - \gamma_{n+1})\underline{v}_{n, \varepsilon}(x) + \gamma_{n+1} (1 - \delta) 2^{-\eta} c_{\varepsilon, n} \underline{f}_\varepsilon^\eta(x) \underline{v}_{n, \varepsilon}^{1-\eta}(x),$$

for every $n \geq N_{\alpha, x, \delta}(\omega) := \max(N_{\alpha, x}(\omega), N_\delta(\omega))$. There exists a constant $c > 0$ small enough satisfying

$$(1 - \delta) 2^{-\eta} c_\varepsilon \alpha c^{1-\eta} \geq c, \quad \text{and} \quad \underline{v}_{N_{\alpha, x, \delta}(\omega), \varepsilon} \geq c.$$

One shows by induction that $\underline{v}_{n, \varepsilon}(x) > c$ for every $n \geq N_{\alpha, x, \delta}(\omega)$. Taking the limit inferior, we obtain Prop 4.

A.6 Proof of Proposition 5

In this subsection, we let Assumptions 1, 2 and 3 hold true. Moreover, we assume $f = f_u$. We also use the notations in (27) and (28).

Lemma 5. *For every $\varepsilon > 0, x \in \mathbb{R}^d$, almost surely, we obtain*

$$\underline{\lim} \underline{g}_{n, \varepsilon}(x) \geq \underline{f}_\varepsilon(x) \quad (29)$$

Proof. Let $B \subset \Omega$ such that, Proposition 3 and 4 holds for every $\omega \in B$, and $\mathbb{P}(B) = 1$. We fix $\varepsilon > 0, x \in \mathbb{R}^d, \omega \in B$ in the whole proof, and assume that $\underline{f}_{2\varepsilon}(x) > 0$, because otherwise (29) is obvious.

Since $Z_n \leq \bar{Z}_n$, we obtain for every $y \in \mathbb{R}^d$,

$$q_n(y) \geq \frac{1 - \lambda_n}{\bar{Z}_n} g_n(y) + \lambda_n q_0(y) \geq \frac{1 - \lambda_n}{\bar{Z}_n} g_n(y).$$

Hence, we obtain $Tq_n(x) \geq c_n Tg_n(y)$, where we defined $c_n := ((1 - \lambda_n)/\bar{Z}_n)^{1-\eta}$. Note that $c_n \rightarrow 1$ a.s., by Proposition 3. From (9), for every $y \in \mathbb{R}^d$, we obtain

$$g_{n+1}(y) \geq (1 - \gamma_{n+1})g_n(y) + \gamma_{n+1}c_n Tg_n * K_{b_{n+1}}(y) + \gamma_{n+1}\xi_{n+1}(y). \quad (30)$$

Let $2\varepsilon \geq \varepsilon' \geq \varepsilon$. Remark that, for $y \in B(x, \varepsilon)$,

$$\begin{aligned} Tg_n * K_{b_{n+1}}(y) &\geq Tg_n \mathbb{1}_{B(x, \varepsilon')} * K_{b_{n+1}}(y) \\ &\geq \underline{f}_{\varepsilon'}^\eta(x) \underline{g}_{n, \varepsilon'}^{1-\eta}(x) (\mathbb{1}_{B(x, \varepsilon')} * K_{b_{n+1}})(y) \\ &\geq \underline{f}_{\varepsilon'}^\eta(x) \underline{g}_{n, \varepsilon'}^{1-\eta}(x) (\mathbb{1}_{B(0, \varepsilon')} * K_{b_{n+1}})(y - x) \\ &\geq \tilde{c}_{\varepsilon', n} \underline{f}_{\varepsilon'}^\eta(x) \underline{g}_{n, \varepsilon'}^{1-\eta}(x), \end{aligned}$$

where $\tilde{c}_{\varepsilon', n} := \inf_{B(0, \varepsilon')} K_{b_{n+1}} * \mathbb{1}_{B(0, \varepsilon)}$. Going back to (30), for every $y \in B(x, \varepsilon)$, we obtain

$$g_{n+1}(y) \geq (1 - \gamma_{n+1})g_n(y) + \gamma_{n+1}c_{\varepsilon', n} \underline{f}_{\varepsilon'}^\eta(x) \underline{g}_{n, \varepsilon'}^{1-\eta}(x) + \gamma_{n+1}\xi_{n+1}(y),$$

where $c_{\varepsilon', n} := c_n \tilde{c}_{\varepsilon', n}$. Iterating the latter equation, for every $y \in B(x, \varepsilon)$, we obtain (since $g_0 = 0$)

$$g_n(y) \geq \sum_{i=2}^n \gamma_i \psi_{i+1}^n c_{\varepsilon', i-1} \underline{f}_{\varepsilon'}^\eta(x) \underline{g}_{i-1, \varepsilon'}^{1-\eta}(x) + M_n(y).$$

Taking the infimum over the ball $B(x, \varepsilon)$, we obtain

$$\underline{g}_{n, \varepsilon}(x) \geq \sum_{i=2}^n \gamma_i \psi_{i+1}^n c_{\varepsilon', i-1} \underline{f}_{\varepsilon'}^\eta(x) \underline{g}_{i-1, \varepsilon'}^{1-\eta}(x) - \sup_{y \in B(x, \varepsilon)} |M_n(y)|$$

Applying Lemma 10, we obtain

$$\underline{\lim} \underline{g}_{n, \varepsilon}(x) \geq \underline{\lim} c_{\varepsilon', n-1} \underline{f}_{\varepsilon'}^\eta(x) \underline{g}_{n-1, \varepsilon'}^{1-\eta}(x) - \overline{\lim} \sup_{y \in B(x, \varepsilon)} |M_n(y)|.$$

Since $c_n \rightarrow 1$, and $\tilde{c}_{\varepsilon', n} \rightarrow 0$ by Lemma 7 and Proposition 3, we obtain $\underline{\lim} c_{\varepsilon', n-1} \underline{g}_{n-1, \varepsilon'}^{1-\eta}(x) = (\underline{\lim} \underline{g}_{n-1, \varepsilon'}(x))^{1-\eta}$. Moreover, by Proposition 2, $\overline{\lim} \sup_{y \in B(x, \varepsilon)} |M_n(y)| = 0$. Hence, we obtain

$$\underline{\lim} \underline{g}_{n, \varepsilon}(x) \geq (\underline{f}_{\varepsilon'}(x))^\eta \underline{\lim} \underline{g}_{n, \varepsilon'}(x)^{1-\eta}.$$

Letting ε decrease to some value $t < t$ and then ε' decrease to t , we get that $\underline{\lim} \underline{g}_{n, t+}(x) = \lim_{s \downarrow t} \underline{\lim} \underline{g}_{n, s}(x)$ satisfies

$$\underline{\lim} \underline{g}_{n, t+}(x) \geq \underline{f}_t^\eta \underline{\lim} \underline{g}_{n, t+}(x)^{1-\eta}.$$

By Proposition 4, $\underline{\lim} \underline{g}_{n, t+}(x) \geq \underline{\lim} \underline{g}_{n, \varepsilon}(x) > 0$. Consequently, $\underline{\lim} \underline{g}_{n, t+}(x) \geq \underline{f}_t$. The latter implies that for any $\varepsilon > s > t$, close enough to t , $\underline{\lim} \underline{g}_{n, t}(x) \geq \underline{\lim} \underline{g}_{n, s+}(x) \geq \underline{f}_s$, hence $\underline{\lim} \underline{g}_{n, t}(x) \geq \underline{f}_t$ by continuity of $t \mapsto \underline{f}_t$. \square

Proof of Proposition 5. Suppose that the proposition doesn't hold. Then, there exists a bounded sequence (x_n) and $\delta > 0$ such that:

$$(g_n(x_n) - f(x_n)) \leq -\delta.$$

(x_n) admits a converging subsequence. So, we can suppose without loss of generality that $x_n \rightarrow x$. By continuity of f there exists $\varepsilon > 0, n_0 \geq 0$ such that:

$$f(x_n) \leq \underline{f}_\varepsilon(x) + \frac{\delta}{2}, \quad n \geq n_0.$$

And for n_1 big enough:

$$g_n(x_n) \geq \underline{g}_{n,\varepsilon}(x), \quad n \geq n_1.$$

Then, for all $n \geq \max(n_0, n_1)$, we have:

$$-\delta \geq \left(\underline{g}_{n,\varepsilon}(x) - \underline{f}_\varepsilon(x) \right) - \frac{\delta}{2}.$$

Using Lemma 5, we obtain $-\delta \geq -\delta/2$, which is a contradiction. \square

A.6.1 Proof of Proposition 6

In this subsection, we let Assumptions 1, 2 and 3 hold true. Moreover, we assume $f = f_u$.

We recall (12)

$$g_n = \sum_{i=1}^n \psi_{i+1}^n \gamma_i T q_{i-1} * K_{b_i} + M_n.$$

We define:

$$u_n := \sum_{i=1}^n \psi_{i+1}^n \gamma_i T q_{i-1} * K_{b_i}. \quad (31)$$

Recursively, we have:

$$u_{n+1} = (1 - \gamma_{n+1})u_n + \gamma_{n+1} T q_n * K_{b_{n+1}}, \quad (32)$$

with $u_0 = 0$.

Lemma 6. *With probability one, for any $\varepsilon > 0$ and any $x \in \mathbb{R}^d$*

$$\overline{\lim} \bar{g}_{n,\varepsilon}(x) \leq \bar{f}_\varepsilon(x). \quad (33)$$

Proof. Let $B \subset \Omega$ such that $\mathbb{P}(B) = 1$, and 18 and Proposition 2 hold for every $\omega \in B$. In the whole proof, we fix $\omega \in B$. From (31), we get for any $x \in \mathbb{R}^d$

$$u_{n+1}(x) \leq \sum_{i=1}^n \psi_{i+1}^n \gamma_i \|K_{b_i} * T q_{i-1}\|_\infty. \quad (34)$$

Set

$$\theta_n = \max_{\|x\| \leq n^p} q_n(x)$$

and notice that

$$\begin{aligned} \|K_{b_{n+1}} * T q_n\|_\infty &= \sup_{x \in \mathbb{R}^d} \int K_{b_{n+1}}(x-y) f(y)^\eta q_n(y)^{1-\eta} dy \\ &\leq \|f\|_\infty^\eta \theta_n^{1-\eta} + \sup_{x \in \mathbb{R}^d} \int_{\|y\| > n^p} K_{b_n}(x-y) f(y)^\eta q_{n-1}(y)^{1-\eta} dy. \end{aligned}$$

For any $x \in \mathbb{R}^d$, using Assumption 1-ii:

$$\begin{aligned} \int_{\|y\| > n^p} K_{b_n}(x-y) f(y)^\eta q_{n-1}(y)^{1-\eta} dy &\leq b_n^{-d} \|K\|_\infty \int_{\|y\| > n^p} f(y)^\eta q_{n-1}(y)^{1-\eta} dy \\ &\leq b_n^{-d} \|K\|_\infty \left(\int_{\|y\| > n^p} f(y) dy \right)^\eta \quad (\text{H\"older}) \\ &\leq \|K\|_\infty b_n^{-d} C_f^\eta n^{-p\eta r}. \end{aligned}$$

Finally, (34) yields

$$u_{n+1}(x) \leq C \sum_{i=1}^n \psi_{i+1}^n \gamma_i \theta_i^{1-\eta} + C \sum_{i=1}^n \psi_{i+1}^n \gamma_i b_i^{-d} i^{-p\eta r},$$

and for $\|x\| \leq n^p$, using 18, there exists a constant $C_0 > 0$ such that:

$$\theta_n(x) \leq C_0 \sup_{\|y\| \leq n^p} g_n(y) \leq C_0 \sup_{\|y\| \leq n^p} u_n(y) + C_0 \sup_{\|y\| \leq n^p} |g_n(y) - u_n(y)|.$$

Then, with $C' = CC_0$ and $g_n - u_n = M_n$, we obtain:

$$\theta_{n+1} \leq C' \sum_{i=1}^n \psi_{i+1}^n \gamma_i \theta_i^{1-\eta} + C' \sum_{i=1}^n \psi_{i+1}^n \gamma_i b_i^{-d} i^{-p\eta r} + C_0 \sup_{\|x\| \leq (n+1)^p} |M_{n+1}(x)|.$$

By Proposition 2, the last term tends to zero. Lemma 10 implies that the second term tends to zero for p large enough. Denote by r_n the sum of these terms, and the first sum by S_n , that is: $S_n := \sum_{i=1}^n \psi_{i+1}^n \gamma_i \theta_i^{1-\eta}$. In order to prove that θ_n is bounded, it suffices to prove that S_n is bounded. We obtain:

$$\begin{aligned} S_{n+1} &= (1 - \gamma_{n+1})S_n + \gamma_{n+1}C'\theta_n^{1-\eta} \\ &\leq (1 - \gamma_{n+1})S_n + \gamma_{n+1}C'(C'S_n + r_n)^{1-\eta} \\ &\leq (1 - \gamma_{n+1})S_n + \gamma_{n+1}C'^{2-\eta}S_n^{1-\eta} + \gamma_{n+1}C'r_n^{1-\eta} \end{aligned}$$

Consider a constant C_1 such that

$$\begin{aligned} C'^{2-\eta}C_1^{1-\eta} + C'r_n^{1-\eta} &\leq C_1, \quad n \geq 1 \\ S_1 &\leq C_1. \end{aligned}$$

Such a C_1 exists, and one shows by induction that $S_n \leq C_1$ for all n . So, θ_n is bounded as well for all p . Consider $0 < \varepsilon$, taking the sup on $B(x, \varepsilon)$ in (32), we get

$$\bar{u}_{n+1, \varepsilon}(x) \leq (1 - \gamma_{n+1})\bar{u}_{n, \varepsilon}(x) + \gamma_{n+1} \sup_{y \in B(x, \varepsilon)} K_{b_{n+1}} * Tq_n(y).$$

But, for any $\varepsilon' > \varepsilon$

$$\begin{aligned} K_{b_{n+1}} * Tq_n(y) &= \int K_{b_n}(y-z) f(z)^\eta q_{n-1}(z)^{1-\eta} \mathbf{1}_{\|z-x\| \leq \varepsilon'} dz \\ &\quad + \int K_{b_n}(y-z) f(z)^\eta q_{n-1}(z)^{1-\eta} \mathbf{1}_{\|z-x\| > \varepsilon'} dz \\ &\leq \bar{f}_{\varepsilon'}(x)^\eta \bar{q}_{n, \varepsilon'}(x)^{1-\eta} (K_{b_{n+1}} * \mathbf{1}_{\|\cdot\| \leq \varepsilon'})(y-x) \\ &\quad + \|f\|_\infty \left(\int K_{b_n}(y-z)^{1/\eta} \mathbf{1}_{\|z-x\| > \varepsilon'} dz \right)^\eta. \end{aligned}$$

Using Lemma 7, the convolution in term first term converges to one uniformly on $y \in B(x, \varepsilon)$. As far as the second term is concerned, one has, for $y \in B(x, \varepsilon)$, using Assumption 2.(ii)

$$\begin{aligned} \int K_{b_n}(y-z)^{1/\eta} \mathbf{1}_{\|z-x\| > \varepsilon'} dz &\leq \int_{\|t\| \geq \varepsilon' - \varepsilon} K_{b_n}(t)^{1/\eta} dt \\ &\leq \int_{\|t\| \geq \varepsilon' - \varepsilon} \left(b_n^{-d} (\|t\|/b_n)^{-(r+d)} \right)^{1/\eta} dt \\ &\leq C(\varepsilon, \varepsilon') b_n^{r/\eta} \end{aligned}$$

which tends to zero, since $\int_{\mathbb{R}^d} \frac{1}{1+\|t\|^{d+\delta}} dt$ is bounded for every $\delta > 0$. Finally we obtain (we omit the argument x)

$$\bar{u}_{n+1, \varepsilon} \leq (1 - \gamma_{n+1})\bar{u}_{n, \varepsilon} + \gamma_{n+1} \kappa_{n+1} \left(\bar{f}_{\varepsilon'}^\eta \bar{q}_{n, \varepsilon'}^{1-\eta} + r_n \right)$$

for some sequence $r_n \rightarrow 0$ depending on $\varepsilon, \varepsilon'$, and $\kappa_n \rightarrow 1$. The boundness of θ_n guarantees that the l.h.s. is finite. By Lemma 10 we have, omitting the argument x ,

$$\overline{\lim} \bar{u}_{n,\varepsilon} \leq \bar{f}_{\varepsilon'}^\eta (\overline{\lim} \bar{q}_{n,\varepsilon'})^{1-\eta}.$$

Since $\overline{\lim} \bar{u}_{n,\varepsilon} = \overline{\lim} \bar{q}_{n,\varepsilon} = \overline{\lim} \bar{g}_{n,\varepsilon}$ (cf. 18 and Proposition 2), the bounded numbers $u_\varepsilon = \overline{\lim} \bar{g}_{n,\varepsilon}$ satisfy

$$u_\varepsilon \leq \bar{f}_{\varepsilon'}^\eta u_{\varepsilon'}^{1-\eta}$$

for any $\varepsilon' > \varepsilon$. The function $\varepsilon \mapsto u_\varepsilon$ is increasing. Letting ε decrease to some value t and then ε' decrease to t , we get that $u_{t+} := \lim_{s \downarrow t} u_s$ satisfies

$$u_{t+} \leq \bar{f}_t^\eta u_{t+}^{1-\eta}$$

Hence $u_{t+} \leq \bar{f}_t$ for all t , which implies that $u_t \leq \bar{f}_t$. \square

Proof of Proposition 6. Suppose that the proposition doesn't hold. Then, there exists a bounded sequence (x_n) and $\delta > 0$ such that:

$$(g_n(x_n) - f(x_n)) \geq \delta.$$

(x_n) admits a converging subsequence. Hence, we can suppose without loss of generality that $x_n \rightarrow x$. By continuity of f there exists $\varepsilon > 0, n_0 \geq 0$ such that:

$$f(x_n) \geq \bar{f}_\varepsilon(x) - \frac{\delta}{2}, \quad n \geq n_0.$$

And for n_1 big enough:

$$g_n(x_n) \leq \bar{g}_{n,\varepsilon}(x), \quad n \geq n_1.$$

Then, for all $n \geq \max(n_0, n_1)$, we have:

$$\delta \leq (\bar{g}_{n,\varepsilon}(x) - \bar{f}_\varepsilon(x)) + \frac{\delta}{2}.$$

Now, using Lemma 6, we obtain $\delta \leq \delta/2$, which finishes the proof. \square

A.7 Proof of Theorem 1

When $f = f_u$, for any compact set $A \subset \mathbb{R}^d$, we obtain

$$\begin{aligned} \lim_{n \rightarrow \infty} \sup_{x \in A} |g_n(x) - f(x)| &\leq \overline{\lim} \max \left(\sup_{x \in A} (g_n(x) - f(x)), - \inf_{x \in A} (g_n(x) - f(x)) \right) \\ &\leq \max \left(\overline{\lim} \sup_{x \in A} (g_n(x) - f(x)), \overline{\lim} - \inf_{x \in A} (g_n(x) - f(x)) \right) \\ &\leq \max \left(\overline{\lim} \sup_{x \in A} (g_n(x) - f(x)), - \underline{\lim} \inf_{x \in A} (g_n(x) - f(x)) \right). \end{aligned}$$

By Propositions 5 and 6, Theorem 1 is proven.

When $f \neq f_u$, according to the Step 0, in Sec. A.2 the result still holds.

A.8 Proof of Corollary 2

Let $w_{i,n}$ be defined by

$$w_{i,n} = \frac{1}{\sqrt{n}} \left(w_i h(X_i) - \int f h \right),$$

and check that $\mathbb{E}[w_{i,n} | \mathcal{F}_{i-1}] = 0$ a.s. Following from [HH14, Corollary 3.1], it suffices to show

$$\lim_{n \rightarrow \infty} \sum_{i=1}^n \mathbb{E}[w_{i,n}^2 | \mathcal{F}_{i-1}] = \int f h^2 - \left(\int f h \right)^2, \text{ in probability} \quad (35)$$

$$\lim_{n \rightarrow \infty} \sum_{i=1}^n \mathbb{E}[w_{i,n}^2 \mathbb{1}_{|w_{i,n}| > \varepsilon} | \mathcal{F}_{i-1}] = 0, \text{ in probability.} \quad (36)$$

Note that

$$\mathbb{E}[w_{i+1,n}^2 | \mathcal{F}_i] = \frac{1}{n} \left(\int \frac{f^2}{q_i} h^2 - \left(\int fh \right)^2 \right).$$

Hence, using the Cesaro theorem, (35) is a consequence of

$$\lim_{i \rightarrow \infty} \int \frac{f^2}{q_i} h^2 = \int fh^2, \text{ almost surely.} \quad (37)$$

Since $\inf_{x \in A} f(x) > 0$, by Theorem 1, there exists $i_0 \geq 0$, such that $\inf_{i \geq i_0, x \in A} q_i(x) > 0$, for every $\omega \in B$ such that $\mathbb{P}(B) = 1$. Moreover, for $i \leq i_0$, $q_i \geq \inf_{i \leq i_0} \lambda_i > 0$. Hence, we obtain, for all $\omega \in B$, $\inf_{i \geq 1, x \in A} q_i(x) > 0$ and therefore $M := \sup_{i \geq 1, x \in A} (f(x)/q_i(x)) < \infty$. Hence, for all $\omega \in B$, the function $f^2/q_i h$ is bounded by Mfh^2 which is integrable. In addition, by Theorem 1, for all $\omega \in B$ and all $x \in A$, $q_i(x) \rightarrow f(x)$. The Lebesgue dominated convergence theorem implies (37).

For the rest of the proof, we fix an arbitrary $\omega \in B$. Remarking that $\int |h|f$ is bounded, when n is large enough ($\varepsilon\sqrt{n} \geq \int |h|f$) we obtain that

$$\{x \in \mathbb{R}^d : \left| \frac{1}{\sqrt{n}} \left(\frac{fh}{q_i}(x) - \int fh \right) \right| \geq \varepsilon\} = \{x \in \mathbb{R}^d : \frac{fh}{q_i}(x) \geq \sqrt{n}\varepsilon + \int fh\}$$

Taking $|\int fh| \leq \varepsilon\sqrt{n}/2$ yields

$$\begin{aligned} \{x \in \mathbb{R}^d : \left| \frac{1}{\sqrt{n}} \left(\frac{fh}{q_i}(x) - \int fh \right) \right| \geq \varepsilon\} &\subset \{x \in A : fh(x) \geq \frac{\varepsilon}{2}\sqrt{n}q_i(x)\} \\ &\subset \{x \in A : Mh(x) \geq \frac{\varepsilon}{2}\sqrt{n}\} =: A_n. \end{aligned}$$

Hence, using Young inequality, we obtain

$$\begin{aligned} \mathbb{E}[w_{i+1,n}^2 \mathbf{1}_{|w_{i+1,n}| > \varepsilon} | \mathcal{F}_i] &\leq \frac{2}{n} \int \frac{f^2(x)h^2(x)}{q_i(x)} \mathbf{1}_{A_n}(x) dx + \frac{2}{n} \left(\int fh \right)^2 \int \mathbf{1}_{A_n}(x) dx \\ &\leq \frac{2}{n} M \int f(x)h^2(x) \mathbf{1}_{A_n}(x) dx + \frac{2(\int fh)^2}{n} \int_A \mathbf{1}_{A_n}(x) dx. \end{aligned}$$

Finally, we obtain

$$\sum_{i=0}^{n-1} \mathbb{E}[w_{i+1,n}^2 \mathbf{1}_{|w_{i+1,n}| > \varepsilon} | \mathcal{F}_i] \leq 2(M \vee (\int fh)^2) \left(\int f(x)h^2(x) \mathbf{1}_{A_n}(x) dx + \int_{A_n} 1 \right).$$

By the Lebesgue dominated convergence theorem, the r.h.s. of the above inequality converges to 0 as $n \rightarrow \infty$. Since ω is arbitrary fixed in B , (36) holds. Consequently, the proof is finished.

B Technical results

Lemma 7. Let $K : \mathbb{R}^d \rightarrow [0, +\infty)$ be a continuous function such that $\int K = 1$ and $K(0) > 0$. Define $K_b(x) = b^{-d}K(x/b)$ for every $b > 0$. For every $\varepsilon' > \varepsilon > 0$,

$$\liminf_{b \downarrow 0} \inf_{B(0, \varepsilon')} K_b * \mathbf{1}_{B(0, \varepsilon')} = 1.$$

Proof. Choose $0 < r < \varepsilon' - \varepsilon$. For every $\|x\| < \varepsilon$ and every $\|y\| < \frac{r}{b}$, $\|x - by\| < \varepsilon'$, by the triangular inequality. Therefore,

$$\inf_{B(0, \varepsilon)} K_b * \mathbf{1}_{B(0, \varepsilon')} = \inf_{\|x\| < \varepsilon} \int \mathbf{1}_{\|x-by\| < \varepsilon'} K(y) dy \geq \int_{\|y\| < \frac{r}{b}} K(y) dy.$$

Letting b converge to zero, the monotone convergence theorem implies that the righthand side of the above inequality converges to one. \square

Lemma 8. Let $K : \mathbb{R}^d \rightarrow [0, +\infty)$ be a continuous function such that $\int_{\mathbb{R}^d} K(x) dx = 1$ and $K(0) > 0$. Define $K_b(x) = b^{-d}K(x/b)$. There exists $c > 0$ such that for every $\varepsilon > 0$,

$$\liminf_{b \downarrow 0} \inf_{B(0, \varepsilon)} K_b * \mathbf{1}_{B(0, \varepsilon)} \geq c.$$

Proof. By continuity of K , there exists $\delta > 0$ such that $K(x) \geq K(0)/2$ for every $x \in B(0, \delta)$. Denoting by Leb_d the Lebesgue measure on \mathbb{R}^d ,

$$\begin{aligned} \kappa_b(\varepsilon, \varepsilon) &= \inf_{y \in B(0, \varepsilon)} \int_{\mathbb{R}^d} \mathbf{1}_{B(0, \varepsilon)}(y - bx) K(x) dx \\ &\geq \inf_{y \in B(0, \varepsilon)} \frac{K(0)}{2} \text{Leb}_d(B(y/b, \varepsilon/b) \cap B(0, \delta)) \\ &= \frac{K(0)}{2} \text{Leb}_d(B(\varepsilon \bar{v}/b, \varepsilon/b) \cap B(0, \delta)) \end{aligned}$$

where \bar{v} denotes any unit norm vector of \mathbb{R}^d . The sequence of sets $(\varepsilon/b)B(\bar{v}, 1)$ is increasing and converges to the half space $E := \{x : \langle \bar{v}, x \rangle \geq 0\}$ as $b \downarrow 0$. Passing to the limit, the result follows by setting $c := (K(0)/2) \text{Leb}_d(E \cap B(0, \delta))$. \square

Lemma 9. *Let $(\psi_k^n)_{n \geq k \geq 1}$ be defined by (11) and $\psi_n^{n-1} = 1$. Let $(\gamma_n)_{n \geq 1}$ be a sequence satisfying Assumption 3-ii. Then, there exists a constant $C > 0$ satisfying*

$$\psi_{n+1}^n \gamma_i \leq C \gamma_n,$$

for every $1 \leq i \leq n$.

Proof. Let n_0 be the rank given by Assumption 3-ii, such that $\gamma_i - \gamma_{i+1} \leq \gamma_i \gamma_{i+1}$ for every $i \geq n_0$. Let $n \geq n_0$. The result holds when $i = n$. Now, suppose that for a given $n_0 + 1 \leq i_0 \leq n$,

$$\psi_{i_0+1}^n \gamma_{i_0} \leq \gamma_n.$$

We obtain

$$\begin{aligned} \psi_{i_0}^n \gamma_{i_0-1} &= (1 - \gamma_{i_0})(\gamma_{i_0} + (\gamma_{i_0-1} - \gamma_{i_0})) \psi_{i_0+1}^n \\ &\leq \gamma_n (1 - \gamma_{i_0}) + \psi_{i_0+1}^n (\gamma_{i_0-1} - \gamma_{i_0})(1 - \gamma_{i_0}) \\ &\leq \gamma_n (1 - \gamma_{i_0})(1 + \gamma_{i_0-1}) \\ &\leq \gamma_n (1 - \gamma_{i_0} \gamma_{i_0-1} + (\gamma_{i_0-1} - \gamma_{i_0})) \\ &\leq \gamma_n. \end{aligned}$$

Iterating the latter result, $\psi_{i+1}^n \gamma_i \leq \gamma_n$ for every, $n \geq i \geq n_0$. For a given $i < n_0$,

$$\psi_{i+1}^n \gamma_i = \psi_{n_0+1}^n \gamma_{n_0} \frac{\psi_{i+1}^n \gamma_i}{\psi_{n_0+1}^n \gamma_{n_0}} \leq \gamma_n \frac{\gamma_1}{\gamma_{n_0}}.$$

Hence, Lemma 9 is proven. \square

Lemma 10. *Consider a real sequence $(a_n)_{n \geq 1}$, and let $(\gamma_n)_{n \geq 1}$ be a positive sequence converging to zero, and such that $\sum_n \gamma_n = +\infty$. Let $(\psi_k^n)_{n \geq k \geq 1}$ be defined by (11) and $\psi_n^{n-1} = 1$. For any $s_0 \in \mathbb{R}$, the sequence $(s_n)_{n \geq 0}$ given by:*

$$s_n = (1 - \gamma_n) s_{n-1} + \gamma_n a_n,$$

satisfies for every $n \geq 1$:

$$s_n = \psi_1^n s_0 + \sum_{i=1}^n \psi_{i+1}^n \gamma_i a_i. \quad (38)$$

In addition, $\underline{\lim} a_n \leq \underline{\lim} s_n \leq \overline{\lim} s_n \leq \overline{\lim} a_n$.

Proof. An elementary induction shows that for any $k > 0$

$$s_n = \psi_{k+1}^n s_k + \sum_{i=k+1}^n \psi_{i+1}^n \gamma_i a_i \quad n > k.$$

In particular $\sum_{i=k}^n \psi_{i+1}^n \gamma_i = 1 - \psi_k^n$, and we get

$$\overline{\lim}_{n > k} s_n \leq 0 + \sup_{i > k} a_i.$$

This proves that $\overline{\lim} s_n \leq \overline{\lim} a_n$. The inequality $\underline{\lim} s_n \geq \underline{\lim} a_n$ is proved similarly. \square

We state here Theorem 19 of [DP21]:

Proposition 7. *Let $(\Omega, \mathcal{F}, (\mathcal{F}_j)_{j \geq 1}, \mathbb{P})$ be a filtered space. Let $(\xi_j)_{j \geq 1}$ be a sequence of real valued stochastic processes defined on \mathbb{R}^d , adapted to $(\mathcal{F}_j)_{j \geq 1}$, such that for any $x \in \mathbb{R}^d$,*

$$\mathbb{E}[\xi_j(x) | \mathcal{F}_{j-1}] = 0, \quad \text{for all } j \geq 1.$$

Consider $\varepsilon > 0$ and let $(\tilde{\xi}_j)_{j \geq 1}$ be another $(\mathcal{F}_j)_{j \geq 1}$ -adapted sequence of non-negative stochastic processes defined on \mathbb{R}^d such that for all $j \geq 1$ and $x \in \mathbb{R}^d$

$$\sup_{\|y\| \leq \varepsilon} |\xi_j(x+y) - \xi_j(x)| \leq \tilde{\xi}_j(x). \quad (39)$$

Let $n \geq 1$ and assume that for some $A \geq 0$, there exist $m, v, u \in \mathbb{R}^+$ such that for all $\omega \in \Omega$ and $\|x\| \leq A$,

$$\max_{j=1, \dots, n} |\xi_j(x)| \leq m \quad (40)$$

$$\sum_{j=1}^n \mathbb{E}[\xi_j(x)^2 | \mathcal{F}_{j-1}] \leq v \quad (41)$$

$$\sum_{j=1}^n \mathbb{E}[\tilde{\xi}_j(x) | \mathcal{F}_{j-1}] \leq u. \quad (42)$$

Then, for all $t \geq 0$,

$$\mathbb{P}\left(\sup_{\|x\| \leq A} \left| \sum_{j=1}^n \xi_j(x) \right| > t + u\right) \leq 4(1 + 2A\varepsilon^{-1})^d \exp\left(-\frac{t^2}{8(\tilde{v} + \frac{2}{3}mt)}\right), \quad (43)$$

with $\tilde{v} = \max(v, 2mu)$.