Minimax Number of Strata for Online Stratified Sampling given Noisy Samples

Alexandra Carpentier and Rémi Munos

INRIA Lille Nord-Europe,
40 avenue Halley, 59650 Villeneuve d'Ascq, France
{alexandra.carpentier,remi.munos}@inria.fr
https://sequel.lille.inria.fr/

Abstract. We consider the problem of online stratified sampling for Monte Carlo integration of a function given a finite budget of n noisy evaluations to the function. More precisely we focus on the problem of choosing the number of strata K as a function of the numerical budget n. We provide asymptotic and finite-time results on how an oracle that knows the smoothness of the function would choose the number of strata optimally. In addition we prove a *lower bound* on the learning rate for the problem of stratified Monte-Carlo. As a result, we are able to state, by improving the bound on its performance, that algorithm MC-UCB, defined in [1], is minimax optimal both in terms of the number of samples n and the number of strata K, up to a log factor. This enables to deduce a minimax optimal bound on the difference between the performance of the estimate output by MC-UCB, and the performance of the estimate output by the best oracle static strategy, on the class of Hölder continuous functions, and up to a log factor.

Keywords: Bandit Theory, Online learning, Stratified sampling, Monte Carlo integration, Regret bounds.

Introduction

The objective of this paper is to provide an efficient strategy for Monte-Carlo integration of a function f over a domain $[0,1]^d$. We assume that we can query the function n times. Querying the function at a time t and at a point $x_t \in [0,1]^d$ provides a noisy sample[1]

$$f(x_t) + s(x_t)\epsilon_t, \tag{1}$$

where ϵ_t is an independent noise drawn from ν_{x_t} and $s \geq 0$ is a function on $[0,1]^d$. Here ν_x is a distribution with mean 0, variance 1 and whose shape may depend on x. This model is actually very general (see Section 1).

Stratified sampling is a well-known strategy to reduce the variance of the estimate of the integral of f, when compared to the variance of the estimate

[1] It is the usual model for regression in heterocedastic noise. We emphasize the standard deviation $s(x)$ of the noise at x, in the expression of the noise, since this quantity is very relevant.

provided by crude Monte-Carlo. The principle is to partition the domain in K subsets called *strata* and then to sample in each stratum (see [11][Subsection 5.5] or [6]). If the variances of the samples in the strata are known, there exists an optimal static allocation strategy which allocates the number of samples in each stratum proportionally to the measure of the stratum times the standard deviation in the stratum (see Equation 3 for the variance of the resulting estimate). We refer to this allocation as optimal oracle strategy for a given partition. In the case that the variations of f and the standard deviation of the noise s are unknown, it is not possible to adopt this strategy.

Consider first that the partition of the space is fixed. A way around this problem is to estimate the variations of the function and the amount of noise on the function in the strata *online* (exploration) while allocating the samples according to the estimated optimal oracle strategy (exploitation). This setting is considered in [3, 8, 1]. In the long version [2] of the last paper, the authors describe the MC-UCB algorithm which is based on Upper-Confidence-Bounds (UCB) on the standard deviation. They provide upper bounds for the difference between the mean-squared error(w.r.t. the integral of f) of the estimate provided by MC-UCB and the mean-squared error of the estimate provided by the optimal oracle strategy (optimal oracle variance). The algorithm performs almost as well as the optimal oracle strategy. However, the authors of [2] do not verify nor assess the optimality of their algorithm. As a matter of fact, no lower bound on the rate of convergence (to the oracle optimal strategy) for the problem of stratified Monte-Carlo exists, to the best of our knowledge. Still in the same paper [2], the authors do not discuss how to *stratify* the space. In particular, they do not pose the problem of what an *optimal* partition of the space is, and do not try to answer on whether it is possible or not to achieve this.

The next step is thus to efficiently design the partition. There are some interesting papers on that topic such as [7, 10, 4]. The recent, state of the art, work of [4] describes a strategy that samples *asymptotically* almost as efficiently as the optimal oracle strategy, and at the same time adapts the direction and number of the strata online. This is a very difficult problem. The authors do not provide proofs of convergence of their algorithm. However for static allocation of the samples, they present some properties of the stratified estimate when the number of strata goes to infinity and provide convergence results under the optimal oracle strategy. As a corollary, they prove that the more strata there are, the smaller the optimal oracle variance is.

Contribution: The more strata, the smaller the variance of the estimate computed when following the optimal oracle strategy. However, the more strata there are, the more difficult it is to estimate the variance within each of these strata, and thus the more difficult it is to perform almost as well as the optimal oracle strategy. Choosing the number of strata is thus crucial and this is the problem we address in this paper. This defines a trade-off similar to the one in model selection (such as in e.g. density estimation, regression...): The wider the class of considered models, i.e. the larger the number of strata, the smaller the distance between the true model and the best model of the class, i.e. the

approximation error, but the larger the estimation error.
Paper [4], although proposing no finite time bounds, develops very interesting ideas for bounding the first term, i.e. the approximation error. As pointed out in e.g. [1], it is possible to build algorithms that have a small estimation error. By constructing tight and finite-time bounds for the approximation error, it is thus possible to select a number of strata that minimizes an upper bound on the performance. It is however not clear if this choice is really optimal in some sense. The essential ingredients for choosing efficiently a partition are thus lower bounds *on the estimation error, and on the approximation error.*

The objective of this paper is to propose a method for choosing the minimax-optimal number of strata. Our contributions are the following.

- We first present results on what we call the *quality* $Q_{n,\mathcal{N}}$ of a given partition $\mathcal{N}$ in K strata (i.e., using the previous analogy to model selection, this would represent the approximation error). Using very mild assumptions we compute a lower bound on the variance of the estimate given by the optimal oracle strategy on the optimal oracle partition. Then if the function and the standard deviation of the noise are $\alpha-$Hölder, and if the strata also satisfy some conditions, we prove that $Q_{n,\mathcal{N}} = O(\frac{K^{\alpha/d}}{n})$. This bound is also minimax optimal on the class of $\alpha-$Hölder functions.
- We then present results on the estimation error for the estimate output by algorithm MC-UCB of [1] (pseudo-regret in the terminology of [1]). In this paper, we improve the analysis of the MC-UCB algorithm compared to [1] in terms of the dependence on K. The problem independent bound on the pseudo-regret in [1] is of order[2] $\tilde{O}(Kn^{-4/3})$, and we tighten this bound in this paper so that it is of order $\tilde{O}(K^{1/3}n^{-4/3})$.
- We provide the first *lower bound* (on the pseudo-regret) for the problem of online Stratified Sampling. The bound $\Omega(K^{1/3}n^{-4/3})$ is tight and *matches the upper-bound of MC-UCB both in terms of the number of strata and the number of samples* up to a $\sqrt{\log(nK)}$ factor. We believe that the proof technique for this bound is original.
- Finally, we combine the results on the quality and on the pseudo-regret of MC-UCB to provide a value on the number of strata leading to a minimax-optimal trade-off (up to a $\sqrt{\log(n)}$) on the class of $\alpha-$Hölder functions.

The rest of the paper is organized as follows. In Section 1 we formalize the problem and introduce the notations used throughout the paper. Section 2 states the results on the quality of a partition. Section 3 improves the analysis of the MC-UCB algorithm, and establishes the lower bound on the pseudo-regret. Section 4 reports the best trade-off to choose the number of strata. And in Section 5, we illustrate how important it is to carefully choose the number of strata. We finally conclude the paper and suggest future works.
Due to space constraints, we were not able to incorporate complete proofs of our results in this paper, but they are all available in the Technical Report [12].

[2] Here $\tilde{O}$ is a O up to $\mathbf{poly}(\log(n))$ factor.

1 Setting

We consider the problem of numerical integration of a function $f : [0,1]^d \rightarrow \mathbb{R}$ with respect to the uniform (Lebesgue) measure. We have at our disposal a budget of n queries (samples) to the function, and we can allocate this budget *sequentially*. When querying the function at a time t and at a point x_t, we receive a noisy sample $X(t)$ of the form described in Equation 1.

We now assume that the space is stratified in K Lebesgue measurable strata that form a partition $\mathcal{N}$. We index these strata, called Ω_k, with indexes $k \in \{1, \dots, K\}$, and write w_k their measure, according to the Lebesgue measure. We write $\mu_k = \frac{1}{w_k} \int_{\Omega_k} \mathbb{E}_{\epsilon \sim \nu_x}[f(x) + s(x)\epsilon] dx = \frac{1}{w_k} \int_{\Omega_k} f(x) dx$ their mean and $\sigma_k^2 = \frac{1}{w_k} \int_{\Omega_k} \mathbb{E}_{\epsilon \sim \nu_x}[(f(x) + s(x)\epsilon - \mu_k)^2] dx$ their variance. These mean and variance correspond to the mean and variance of the random variable $X(t)$ when the coordinate x at which the noisy evaluation of f is observed is chosen uniformly at random on the stratum Ω_k.

We denote by $\mathcal{A}$ an algorithm that allocates online the budget by selecting at each time step $1 \leq t \leq n$ the index $k_t \in \{1, \dots, K\}$ of a stratum and then samples uniformly in the corresponding stratum Ω_{k_t}. The objective is to return the best possible estimate $\hat{\mu}_n$ of the integral of the function f. We write $T_{k,n} = \sum_{t \leq n} \mathbb{I}\{k_t = k\}$ the number of samples in stratum Ω_k up to time n. We denote by $\left(X_{k,t}\right)_{1 \leq k \leq K, 1 \leq t \leq T_{k,n}}$ the samples in stratum Ω_k, and we define $\hat{\mu}_{k,n} = \frac{1}{T_{k,n}} \sum_{t=1}^{T_{k,n}} X_{k,t}$ (the empirical means in the stratum). We estimate the integral of f by $\hat{\mu}_n = \sum_{k=1}^{K} w_k \hat{\mu}_{k,n}$.

If we allocate a deterministic number of samples T_k to each stratum Ω_k and if the samples are independent and chosen uniformly on each stratum Ω_k, we have

$$\mathbb{E}(\hat{\mu}_n) = \sum_{k \leq K} w_k \mu_k = \sum_{k \leq K} \int_{\Omega_k} f(u) du = \int_{[0,1]^d} f(u) du = \mu,$$

and also

$$\mathbb{V}(\hat{\mu}_n) = \sum_{k \leq K} \frac{w_k^2 \sigma_k^2}{T_k},$$

where the expectation and the variance are computed according to all the samples that the algorithm collected.

For a given algorithm $\mathcal{A}$ allocating $T_{k,n}$ samples drawn *uniformly* within stratum Ω_k, we call *pseudo-risk* the quantity

$$L_{n,\mathcal{N}}(\mathcal{A}) = \sum_{k \leq K} \frac{w_k^2 \sigma_k^2}{T_{k,n}}. \tag{2}$$

Note that if an algorithm $\mathcal{A}^*$ has access the variances σ_k^2 of the strata, it can choose to allocate the budget in order to minimize the pseudo-risk, i.e., sample each stratum $T_k^* = \frac{w_k \sigma_k}{\sum_{i \leq K} w_i \sigma_i} n$ times (this is the so-called oracle allocation). These optimal numbers of samples can be non-integer values, in which case the proposed optimal allocation is not realizable. But we still use it as a benchmark.

The pseudo-risk for this algorithm (which is also the variance of the estimate here since the sampling strategy is deterministic) is then

$$L_{n,\mathcal{N}}(\mathcal{A}^*) = \frac{\Big(\sum_{k\leq K} w_k\sigma_k\Big)^2}{n} = \frac{\Sigma_{\mathcal{N}}^2}{n}, \tag{3}$$

where $\Sigma_{\mathcal{N}} = \sum_{k\leq K} w_k\sigma_k$. We also refer in the sequel as optimal proportion to $\lambda_k = \frac{w_k\sigma_k}{\sum_{i\leq K} w_i\sigma_i}$, and to optimal oracle strategy to this allocation strategy. Although, as already mentioned, the optimal allocations (and thus the optimal pseudo-risk) might not be realizable, it is still very useful in providing a lower-bound. No static (even oracle) algorithm has a pseudo-risk lower than $L_{n,\mathcal{N}}(\mathcal{A}^*)$ on partition $\mathcal{N}$.

It is straightforward to see that the more refined the partition $\mathcal{N}$ the smaller $L_{n,\mathcal{N}}(\mathcal{A}^*)$ (see e.g. [7]). We thus define the *quality of a partition* $Q_{n,\mathcal{N}}$ as the difference between the variance $L_{n,\mathcal{N}}(\mathcal{A}^*)$ of the estimate provided by the optimal oracle strategy on partition $\mathcal{N}$, and the infimum of the variance of the optimal oracle strategy on *any* partition (optimal oracle partition) (with an arbitrary number of strata):

$$Q_{n,\mathcal{N}} = L_{n,\mathcal{N}}(\mathcal{A}^*) - \inf_{\mathcal{N}' measurable} L_{n,\mathcal{N}'}(\mathcal{A}^*). \tag{4}$$

We also define the *pseudo-regret* of an algorithm $\mathcal{A}$ on a given partition $\mathcal{N}$, as the difference between its pseudo-risk and the variance of the optimal oracle strategy:

$$R_{n,\mathcal{N}}(\mathcal{A}) = L_{n,\mathcal{N}}(\mathcal{A}) - L_{n,\mathcal{N}}(\mathcal{A}^*). \tag{5}$$

We will assess the performance of an algorithm $\mathcal{A}$ by comparing its pseudo risk to the minimum possible variance of an optimal oracle strategy on the optimal oracle partition:

$$L_{n,\mathcal{N}}(\mathcal{A}) - \inf_{\mathcal{N}' measurable} L_{n,\mathcal{N}'}(\mathcal{A}^*) = R_{n,\mathcal{N}}(\mathcal{A}) + Q_{n,\mathcal{N}}. \tag{6}$$

Using the analogy of model selection mentioned in the Introduction, the quality $Q_{n,\mathcal{N}}$ is similar to the approximation error and the pseudo-risk $R_{n,\mathcal{N}}(\mathcal{A})$ to the estimation error.

Motivation for the model $f(x) + s(x)\epsilon_t$. Assume that a learner can, at each time t, choose a point x and collect an observation $F(x, W_t)$, where W_t is an independent noise, that can however depend on x. It is the general model for representing evaluations of a noisy function. There are many settings where one needs to integrate accurately a noisy function without wasting too much budget, like for instance pollution survey. Set $f(x) = \mathbb{E}_{W_t}[F(x, W_t)]$, and $s(x)\epsilon_t = F(x, W_t) - f(x)$. Since by definition ϵ_t is of mean 0 and variance 1, we have in fact $s(x) = \sqrt{\mathbb{E}_{W_t}[(F(x, W_t) - f(x))^2]}$ and $\epsilon_t = \frac{F(x,W_t)-f(x)}{s(x)}$. Observing $F(x, W_t)$ is thus equivalent to observing $f(x) + s(x)\epsilon_t$, and this implies that the model that we choose is also very general.

There is also an important setting where this model is relevant, and this is for the integration of a function F in high dimension d^*. Stratifying in dimension d^*

seems hopeless, since the budget n has to be exponential with d^* if one wants to stratify in every direction of the domain: this is the curse of dimensionality. It is necessary to reduce the dimension by choosing *a small number* of directions $(1, \ldots, d)$ that are particularly relevant, and control/stratify only in these d directions[3]. Then the control/stratification is only on those d coordinates, so when sampling at a time t, one chooses $x = (x_1, \ldots, x_d)$, and the other $d^* - d$ coordinates $U(t) = (U_{d+1}(t), \ldots, U_{d^*}(t))$ are uniform random variables on $[0,1]^{d^*-d}$ (without any control). When sampling in x at a time t, we observe $F(x, U(t))$. By writing $f(x) = \mathbb{E}_{U(t) \sim \mathcal{U}([0,1]^{d^*-d})}[F(x, U(t))]$, and $s(x)\epsilon_t = F(x, U(t)) - f(x)$, we obtain that the model we propose is also valid in this case.

2 The quality of a partition: Analysis of the term $Q_{n,\mathcal{N}}$.

In this Section, we focus on the *quality* of a partition defined in Section 1.

Convergence under very mild assumptions As mentioned in Section 1, the more refined the partition $\mathcal{N}$ of the space, the smaller $L_{n,\mathcal{N}}(\mathcal{A}^*)$, and thus $\Sigma_{\mathcal{N}}$. Through this monotony property, we know that $\inf_{\mathcal{N}} \Sigma_{\mathcal{N}}$ is also the limit of the $(\Sigma_{\mathcal{N}_p})_p$ of a sequence of partitions $(\mathcal{N}_p)_p$ such that the diameter of each stratum goes to 0. We state in the following Proposition that for *any* such sequence, $\lim_{p \to +\infty} \Sigma_{\mathcal{N}_p} = \int_{[0,1]^d} s(x)dx$. Consequently $\inf_{\mathcal{N}} \Sigma_{\mathcal{N}} = \int_{[0,1]^d} s(x)dx$.

Proposition 1. *Let $(\mathcal{N}_p)_p = (\Omega_{k,p})_{k \in \{1,\ldots,K_p\}, p \in \{1,\ldots,+\infty\}}$ be a sequence of measurable partitions (where K_p is the number of strata of partition $\mathcal{N}_p$) such that*

- AS1: $0 < w_{k,p} \leq \upsilon_p$, *for some sequence $(\upsilon_p)_p$, where $\upsilon_p \to 0$ for $p \to +\infty$.*
- AS2: *The diameters according to the $||.||_2$ norm on $\mathbb{R}^d$ of the strata are such that* $\mathbf{Diam}(\Omega_{k,p}) \leq D(w_{k,p})$*, for some real valued function $D(\cdot)$, such that $D(w) \to 0$ for $w \to 0$.*

If the functions m and s are in $\mathbb{L}_2([0,1]^d)$, then

$$\lim_{p \to +\infty} \Sigma_{\mathcal{N}_p} = \inf_{\mathcal{N} measurable} \Sigma_{\mathcal{N}} = \int_{[0,1]^d} s(x)dx,$$

which implies that $n \times Q_{n,\mathcal{N}_p} \to 0$ for $p \to +\infty$.

The full proof of this Proposition (omitted due to space constraints) is available in the Technical Report [12].

In Proposition 1, even though the optimal oracle allocation might not be realizable (in particular if the number of strata is larger than the budget), we can still compute the quality of a partition, as defined in Equation 4. It does not correspond to any reachable pseudo-risk, but rather to a lower bound on any (even oracle) static allocation.

When f and s are in $\mathbb{L}_2([0,1]^d)$, for any appropriate sequence of partitions $(\mathcal{N}_p)_p$, $\Sigma_{\mathcal{N}_p}$ (which is the principal ingredient of the variance of the optimal oracle allocation) converges to the smallest possible $\Sigma_{\mathcal{N}}$ for given f and s. Note however that this condition is not sufficient to obtain a *rate* of convergence.

[3] This is actually a very common technique for computing the price of options, see [6].

Finite-Time analysis under Hölder assumption: We make the following assumption on the functions f and s.

Assumption 1 *The functions f and s are $(M, \alpha)-$Hölder continuous, i.e., for $g \in \{f, s\}$, for any x and $y \in [0,1]^d$, $|g(x) - g(y)| \leq M||x - y||_2^{\alpha}$.*

The Hölder assumption enables us to consider arbitrarily non-smooth functions (for small α, the function can vary arbitrarily fast), and is thus a fairly general assumption.

We also consider the following partitions in K hyper-cubes.

Definition 1. *We write $\mathcal{N}_K$ the partition of $[0,1]^d$ in K hyper-cubic strata of measure $w_k = w = \frac{1}{K}$ and side length $(\frac{1}{K})^{1/d}$: we assume for simplicity that there exists an integer l such that $K = l^d$.*

The following Proposition holds.

Proposition 2. *Under Assumption 1 we have for any partition $\mathcal{N}_K$ as defined in Definition 1 that*

$$\Sigma_{\mathcal{N}_K} - \int_{[0,1]^d} s(x)dx \leq \sqrt{2d}MK^{-\alpha/d}, \tag{7}$$

which implies

$$Q_{n,\mathcal{N}_K} \leq \frac{2\sqrt{2d}M\Sigma_{\mathcal{N}_1}}{n}K^{-\alpha/d},$$

where $\mathcal{N}_1$ stands for the "partition" with one stratum.

The full proof of this Proposition (omitted due to space constraints) is available in the Technical Report [12].

2.1 General comments

The impact of α and d: The quantity $Q_{n,\mathcal{N}_K}$ increases with the dimension d, because the Hölder assumption becomes less constraining when d increases. This can easily be seen since a squared strata of measure w has a diameter of order $w^{1/d}$. $Q_{n,\mathcal{N}_K}$ decreases with the smoothness α of the function, which is a consequence of the Hölder assumption. Note also that when defining the partitions $\mathcal{N}_K$ in Definition 1, we made the crucial assumption that $K^{1/d}$ is an integer. This is of little importance in small dimension, but matters in high dimensions, as we will highlight in the last remark of Section 4.

Minimax optimality of this rate: The rate $n^{-1}K^{-\alpha/d}$ is minimax optimal on the class of $\alpha-$Hölder functions since for any n and K one can easily build a function with Hölder exponent α such that the corresponding $\Sigma_{\mathcal{N}_K}$ is at least $\int_{[0,1]^d} s(x)dx + cK^{-\alpha/d}$ for some constant c.

Discussion of the shape of the strata: Whatever the shape of the strata, as long as their diameter goes to zero[4], $\Sigma_{\mathcal{N}_K}$ converges to $\int_{[0,1]^d} s(x)dx$. The shape of the

[4] And note that in this *noisy* setting, if the diameter of the strata does not go to 0 on non homogeneous part of m and s, then the standard deviation corresponding to the allocation is larger than $\int_{[0,1]^d} s(u)du$.

strata have an influence only on the negligible term, i.e. the speed of convergence to this quantity. This result was already made explicit, in a different setting and under different assumptions, in [4]. Choosing small strata of same shape and size is also minimax optimal on the class of Hölder functions. Working on the shape of the strata could, however, improve the speed of convergence in some specific cases, e.g. when the noise is very localized. It could also be interesting to consider strata of varying size, and have this size depend on the specific problem.

The decomposition of the variance: The variance σ_k^2 within each stratum Ω_k comes from two sources. First, σ_k^2 comes from the noise, that contributes to it by $\frac{1}{w_k}\int_{\Omega_k} s(x)^2 dx$. Second, the mean f is not a constant function, thus its contribution to σ_k^2 is $\frac{1}{w_k}\int_{\Omega_k}\big(f(x)-\frac{1}{w_k}\int_{\Omega_k} f(u)du\big)^2 dx$. Note that when the size of Ω_k goes to 0, this later contribution vanishes, and the optimal allocation is thus proportional to $\sqrt{w_k\int_{\Omega_k} s(x)^2 dx + o(1)} = \int_{\Omega_k} s(x)dx + o(1)$. This means that for small strata, the variation in the mean are negligible when compared to the variation due to the noise.

3 Algorithm MC-UCB and a matching lower bound

3.1 Algorithm $MC-UCB$

In this Subsection, we describe a slight modification of the algorithm $MC-UCB$ introduced in [1]. The only difference is that we change the form of the high-probability upper confidence bound on the standard deviations, in order to improve the elegance of the proofs, and we refine their analysis. The algorithm takes as input two parameters b and $f_{\max}$ which are linked to the distribution in the strata, δ which is a (small) probability, and the partition $\mathcal{N}_K$. We remind in Figure 1 the algorithm $MC-UCB$.

Input: b, $f_{\max}$, δ, $\mathcal{N}_K$. Set $A = 2\sqrt{(1+3b+4f_{\max}^2)\log(2nK/\delta)}$
Initialize: Sample 2 states in each strata.
for $t = 2K+1,\dots,n$ **do**
  Compute $B_{k,t} = \frac{w_k}{T_{k,t-1}}\Big(\hat{\sigma}_{k,t-1} + A\sqrt{\frac{1}{T_{k,t-1}}}\Big)$ for each stratum $k \le K$
  Sample a point in stratum $k_t \in \arg\max_{1\le k\le K} B_{k,t}$
end for
Output: $\hat{\mu}_n = \sum_{k=1}^K w_k\hat{\mu}_{k,n}$

Fig. 1. The pseudo-code of the MC-UCB algorithm. The empirical standard deviations and means $\hat{\sigma}_{k,t}^2$ and $\hat{\mu}_{k,t}$ are computed using Equation 8.

The estimates of $\hat{\sigma}_{k,t-1}^2$ and $\hat{\mu}_{k,t-1}$ are computed according to

$$\hat{\sigma}_{k,t-1}^2 = \frac{1}{T_{k,t-1}}\sum_{i=1}^{T_{k,t-1}}(X_{k,i}-\hat{\mu}_{k,t-1})^2, \ \ and \ \ \hat{\mu}_{k,t-1} = \frac{1}{T_{k,t-1}}\sum_{i=1}^{T_{k,t-1}} X_{k,i}\ . \quad (8)$$

3.2 Upper bound on the pseudo-regret of algorithm MC-UCB.

We first state the following Assumption on the noise ϵ_t:

Assumption 2 *There exist $b > 0$ such that $\forall x \in [0,1]^d$, $\forall t$, and $\forall \lambda < \frac{1}{b}$,*

$$\mathbb{E}_{\nu_x}\Big[\exp(\lambda \epsilon_t)\Big] \leq \exp\Big(\frac{\lambda^2}{2(1-\lambda b)}\Big), \text{ and } \mathbb{E}_{\nu_x}\Big[\exp(\lambda \epsilon_t^2 - \lambda)\Big] \leq \exp\Big(\frac{\lambda^2}{2(1-\lambda b)}\Big).$$

This is a type of sub-Gaussian assumption, satisfied for e.g., Gaussian as well as bounded distributions. We also state an assumption on f and s.

Assumption 3 *The functions f and s are bounded by $f_{\max}$.*

Note that since the functions f and s are defined on $[0,1]^d$, if Assumption 1 is satisfied, then Assumption 3 holds with $f_{\max} = \max(f(0), s(0)) + \sqrt{2d}M$. We now prove the following bound on the pseudo-regret. Note that we state it on partitions $\mathcal{N}_K$, but that it in fact holds for any partition in K strata.

Proposition 3. *Under Assumptions 2 and 3, on partition $\mathcal{N}_K$, when $n \geq 4K$, we have*

$$\mathbb{E}[R_{n,\mathcal{N}_K}(\mathcal{A}_{MC-UCB})] \leq C\frac{K^{1/3}}{n^{4/3}}\sqrt{\log(nK)} + \frac{14K\Sigma_{\mathcal{N}_K}^2}{n^2},$$

where $C = 24\sqrt{2}\Sigma_{\mathcal{N}_K}\sqrt{(1+3b+4f_{\max}^2)}\Big(\frac{f_{\max}+4}{4}\Big)^{1/3}$.

The proof of this Proposition is close to the one of MC-UCB in [1]. But an improved analysis leads to a better dependency in terms of number of strata K. Recall that in [1], the bound is of order $\tilde{O}(Kn^{-4/3})$. This improvement is crucial here since the larger K is, the closer $\Sigma_{\mathcal{N}_K}$ is to $\int_{[0,1]^d} s(x)dx$. The full proof of this Proposition is available in the Technical Report [12]. The next Subsection states that the rate $K^{1/3}\tilde{O}(n^{-4/3})$ of MC-UCB is optimal both in terms of K and n.

3.3 Lower Bound

We now study the minimax rate for the pseudo-regret of any algorithm on a given partition $\mathcal{N}_K$.

Theorem 1. *Let $K \in \mathbb{N}$. Let* inf *be the infimum taken over all online stratified sampling algorithms on $\mathcal{N}_K$ and* sup *represent the supremum taken over all environments, then:*

$$\inf\sup \mathbb{E}[R_{n,\mathcal{N}_K}] \geq C\frac{K^{1/3}}{n^{4/3}},$$

where C is a numerical constant.

Proof (Proof sketch (the full proof of this Theorem is available in the Technical Report [12])). We consider a partition with $2K$ strata. On the K first strata, the samples are drawn from Bernoulli distributions of parameter μ_k where $\mu_k \in \{\frac{\mu}{2}, \mu, 3\frac{\mu}{2}\}$, and on the K last strata, the samples are drawn from a Bernoulli of parameter 1/2. We write $\sigma = \sqrt{\mu(1-\mu)}$ the standard deviation of a Bernoulli of parameter μ. We index by υ a set of 2^K possible environments, where $\upsilon = (\upsilon_1, \ldots, \upsilon_K) \in \{-1, +1\}^K$, and the K first strata are defined by $\mu_k = \mu + \upsilon_k \frac{\mu}{2}$. Write $\mathbb{P}_\upsilon$ the probability under such an environment, also consider $\mathbb{P}_\sigma$ the probability under which all the K first strata are Bernoulli with mean μ.

We define Ω_υ the event on which there are less than $\frac{K}{3}$ strata not pulled correctly for environment υ (i.e. for which $T_{k,n}$ is larger than the optimal allocation corresponding to μ when actually $\mu_k = \frac{\mu}{2}$, or smaller than the optimal allocation corresponding to μ when $\mu_k = 3\frac{\mu}{2}$). See the Appendix D in [12] for a precise definition of these events. Then, the idea is that there are so many such environments that any algorithm will be such that for at least one of them we have $\mathbb{P}_\sigma(\Omega_\upsilon) \leq \exp(-K/72)$. Then we derive by a variant of Pinsker's inequality applied to an event of small probability that $\mathbb{P}_\upsilon(\Omega_\upsilon) \leq \frac{KL(\mathbb{P}_\sigma, \mathbb{P}_\upsilon)}{K} = O(\frac{\sigma^{3/2} n}{K})$. Finally, by choosing σ of order $(\frac{K}{n})^{1/3}$, we have that $\mathbb{P}_\upsilon(\Omega_\upsilon^c)$ is bigger than a constant, and on Ω_υ^c we know that there are more than $\frac{K}{3}$ strata not pulled correctly. This leads to an expected pseudo-regret in environment υ of order $\Omega(\frac{K^{1/3}}{n^{4/3}})$.

This is the first lower-bound for the problem of online stratified sampling for Monte-Carlo. Note that this bound is of same order as the upper bound for the pseudo-regret of algorithm MC-UCB. It means that this algorithm is, up to a factor $\sqrt{\log(nK)}$, minimax optimal, both in terms of the number of samples and in terms of the number of strata. It however holds only on the partitions $\mathcal{N}_K$ (we conjecture that a similar result holds for *any* measurable partition $\mathcal{N}$, but with a bound of order $\Omega\Big(\sum_{x\in\mathcal{N}} \frac{w_x^{2/3}}{n^{4/3}}\Big)$).

4 Minimax-optimal trade-off between $Q_{n,\mathcal{N}_K}$ and $R_{n,\mathcal{N}_K}(\mathcal{A}_{MC-UCB})$

4.1 Minimax-optimal trade-off

We consider in this Section the hyper-cubic partitions $\mathcal{N}_K$ as defined in Definition 1, and we want to find the minimax-optimal number of strata K_n as a function of n. Using the results in Section 2 and Subsection 3.1, it is possible to deduce an optimal number of strata K to give as parameter to the algorithm MC-UCB. Note that since the performance of the algorithm is defined as the sum of the quality of partition $\mathcal{N}_K$, i.e. $Q_{n,\mathcal{N}_K}$ and of the pseudo-regret of the algorithm MC-UCB, namely $R_{n,\mathcal{N}_K}(\mathcal{A}_{MC-UCB})$, one wants to (i) on the one hand take many strata so that $Q_{n,\mathcal{N}_K}$ is small but (ii) on the other hand, pay attention to the impact this number of strata has on the pseudo-regret $R_{n,\mathcal{N}_K}(\mathcal{A}_{MC-UCB})$. A good way to do that is to choose K_n in function of n such that $Q_{n,\mathcal{N}_{K_n}}$ and $R_{n,\mathcal{N}_{K_n}}(\mathcal{A}_{MC-UCB})$ are of the same order.

Theorem 2. *Under Assumptions 1 and 2 (since Assumption 1 implies Assumption 3, by setting* $f_{\max} = X(1) + \sqrt{2d}M$*), with* $K_n = \Big(\lfloor (n^{\frac{d}{d+3\alpha}})^{1/d} \rfloor\Big)^d (\leq n^{\frac{d}{d+3\alpha}} \leq$

n), we have

$$\mathbb{E}[L_n(\mathcal{A}_{MC-UCB})] - \frac{1}{n}\Big(\int_{[0,1]^d} s(x)dx\Big)^2 \leq c d^{\frac{2\alpha}{3d}+\frac{1}{2}} \sqrt{\log(n)} n^{-\frac{d+4\alpha}{d+3\alpha}} (1 + d^{\alpha} n^{-\frac{\alpha}{d+3\alpha}}),$$

where $c = 70(1+M)\Sigma_{\mathcal{N}_K}\sqrt{(1+3b+4(f(0)+s(0)+M)^2)}\Big(\frac{(f(0)+s(0)+M)+4}{4}\Big)^{1/3}$.
If $d \ll n$, *then* $\mathbb{E}[L_n(\mathcal{A}_{MC-UCB})] - \frac{1}{n}\Big(\int_{[0,1]^d} s(x)dx\Big)^2 = \tilde{O}(n^{-\frac{d+4\alpha}{d+3\alpha}})$.

We can also prove a matching (up to $\sqrt{\log(n)}$) minimax lower bound using the results in Theorem 1.

Theorem 3. *Let* sup *represent the supremum taken over all* $\alpha-$*Hölder functions and* inf *be the infimum taken over all algorithms that partition the space in convex strata of same shape, then the following holds true:*

$$\inf \sup \mathbb{E}L_n(\mathcal{A}) - \frac{1}{n}\Big(\int_{[0,1]^d} s(x)dx\Big)^2 = \Omega(n^{-\frac{d+4\alpha}{d+3\alpha}}).$$

4.2 Discussion

Optimal pseudo-risk. The dominant term in the pseudo-risk of MC-UCB with the proper number of strata is $\frac{(\inf_{\mathcal{N}}\Sigma_{\mathcal{N}})^2}{n} = \frac{1}{n}\big(\int_{[0,1]^d} s(x)dx\big)^2$ (the other term is negligible). This means that algorithm MC-UCB is almost as efficient as the optimal oracle strategy on the optimal oracle partition. In comparison, the variance of the estimate given by crude Monte-Carlo is $\int_{[0,1]^d}\big(f(x) - \int_{[0,1]^d} f(u)du\big)^2 dx + \int_{[0,1]^d} s(x)^2 dx$. Thus MC-UCB enables to have the term coming from the variations in the mean vanish, and the noise term decreases (since by Cauchy-Schwarz, $\big(\int_{[0,1]^d} s(x)dx\big)^2 \leq \int_{[0,1]^d} s(x)^2 dx$).

Minimax-optimal trade-off for algorithm MC-UCB. The optimal trade-off on the number of strata K_n of order $n^{\frac{d}{d+3\alpha}}$ depends on the dimension and the smoothness of the function. The higher the dimension, the more strata are needed in order to have a decent speed of convergence for $\Sigma_{\mathcal{N}_K}$. The smoother the function, the fewer strata are needed.
It is yet important to remark that this trade-off is not exact. We provide an almost minimax-optimal order of magnitude for K_n, in terms of n, so that the rate of convergence of the algorithm is minimax-optimal up to a $\sqrt{\log(n)}$ factor.

Link between risk and pseudo-risk. It is important to compare the pseudo-risk $L_n(\mathcal{A}) = \sum_{k=1}^{K} \frac{w_k^2\sigma_k^2}{T_{k,n}}$ and the true risk $\mathbb{E}[(\hat{\mu}_n - \mu)^2]$. Note that these quantities are in general not equal for an algorithm $\mathcal{A}$ that allocates the samples in a dynamic way: indeed, the quantities $T_{k,n}$ are in that case stopping times and the variance of estimate $\hat{\mu}_n$ is not equal to the pseudo-risk. However, in the paper [2], the authors highlighted for $MC-UCB$ some links between the risk and the pseudo-risk. More precisely, they established links between $L_n(\mathcal{A})$ and $\sum_{k=1}^K w_k^2 \mathbb{E}[(\hat{\mu}_{k,n}-\mu_k)^2]$. This step is possible since $\mathbb{E}[(\hat{\mu}_{k,n}-\mu_k)^2] \leq \frac{w_k^2\sigma_k^2}{\underline{T}_{k,n}^2}\mathbb{E}[T_{k,n}]$, where $\underline{T}_{k,n}$ is a lower-bound on the number of pulls $T_{k,n}$ on a high probability event. Then they bounded the cross products $\mathbb{E}[(\hat{\mu}_{k,n} - \mu_k)(\hat{\mu}_{p,n} - \mu_p)]$ and provided some upper bounds on these terms. A tight analysis of these terms as a function of the number of strata K remains to be investigated.

Knowledge of the Hölder exponent. In order to be able to choose properly the number of strata to achieve the rate in Theorem 2, it is needed to possess a proper lower bound on the Hölder exponent of the function: indeed, the rougher the function is, the more strata are required. On the other hand, such a knowledge on the function is not always available and an interesting question is whether it is possible to estimate this exponent fast enough. There are interesting papers on that subject like [9] where the authors tackle the problem of regression and prove that it is possible to adapt to the unknown smoothness of the function. The authors in [5] add to that (in the case of density estimation) and prove that it is even possible under the assumption that the function attain its Hölder exponent to have a proper estimation of this exponent and thus adaptive confidence bands. An idea would be to try to adapt these results in the case of finite sample.

MC-UCB On a noiseless function. Consider the case where $s = 0$ almost surely, i.e. the collected samples are noiseless. Proposition 1 ensures that $\inf_{\mathcal{N}} \Sigma_{\mathcal{N}} = 0$: it is thus possible in this case to achieve a pseudo-risk that has a faster rate than $O(\frac{1}{n})$. If the function m is smooth, e.g. Hölder with an exponent α which is not too small, it is efficient to use low discrepancy methods to integrate the functions. An idea is to stratify the domain in n hyper-rectangular strata of minimal diameter, and to pick at random one sample per stratum. The variance of the resulting estimate is of order $O(\frac{1}{n^{1+2\alpha/d}})$. Algorithm MC-UCB is not as efficient as a low discrepancy scheme: it needs a number of strata $K < n$ in order to be able to estimate the variance within each stratum. Its pseudo-risk is then of order $O(\frac{1}{nK^{2\alpha/d}})$.
However, this only holds when the samples are noiseless. Otherwise, the variance of the estimate is of order $1/n$, no matter what strategy the learner chooses.

In high dimension. The first bound in Theorem 2 expresses precisely how the performance of the estimate output by MC-UCB depends on d. The first bound states that the quantity $L_n(\mathcal{A}) - \frac{1}{n}\Big(\int_{[0,1]^d} s(x)dx\Big)^2$ is negligible when compared to $1/n$ when n is exponential in d. This is not surprising since our technique aims at stratifying equally in every direction. It is not possible to stratify in every directions of the domain if the function lies in a very high dimensional domain.
This is however *not* a reason for not using our algorithm in high dimension. Indeed, stratifying even in a small number of strata already reduces the variance, and in high dimension, any variance reduction techniques are welcome. As mentioned at the end of Section 1, the model that we propose for the function is suitable for modeling d^* dimensional functions that we only stratify in $d < d^*$ directions (and $d \ll n$). A reasonable trade-off for d can also be inferred from

the bound, but we believe that a good choice of d depends heavily on the problem. We then believe that it is a good idea to select the number of strata in the minimax way according to our results. Again, having a very high dimensional function that one stratifies in only a few directions is a very common technique in financial mathematics, for pricing options (practitioners stratify an infinite dimensional process in only 1 to 5 carefully chosen dimensions). We illustrate this in the next Section.

5 Numerical experiment: influence of the number of strata in the Pricing of an Asian option

We consider the pricing problem of an Asian option introduced in [7] and later considered in [10, 3]. This uses a Black-Scholes model with strike C and maturity T. Let $(W(t))_{0\leq t\leq T}$ be a Brownian motion. The discounted payoff of the Asian option is defined as a function of W, by:

$$F((W)_{0\leq t\leq T}) = \exp(-rT)\max\Big[\int_0^T S_0\exp\Big((r-\tfrac{1}{2}s_0^2)t+s_0W_t\Big)dt - C, 0\Big],$$

where S_0, r, and s_0 are constants.

We want to estimate the price $p=\mathbb{E}_W[F(W)]$ by Monte-Carlo simulations (by sampling on W). In order to reduce the variance of the estimated price, we can stratify the space of W. [7] suggest to stratify according to a one dimensional projection of W, i.e., by choosing a time t and stratifying according to the quantiles of W_t (and simulating the rest of the Brownian according to a Brownian Bridge, see [10]). They further argue that the best direction for stratification is to choose $t = T$, i.e., to stratify according to the last time of T. This choice of stratification is also intuitive since W_T has the highest variance, the largest exponent in the payoff and thus the highest volatility. We stratify according to the quantiles of W_T, that is to say the quantiles of a normal distribution $\mathcal{N}(0,T)$. When stratifying in K strata, we stratify according to the $1/K$-th quantiles (so that the strata are hyper-cubes of same measure).

We choose the same numerical values as [10]: $S_0 = 100$, $r = 0.05$, $s_0 = 0.30$, $T = 1$ and $d = 16$. We discretize also, as in [10], the Brownian motion in 16 equidistant times, so that we are able to simulate it. We choose $C = 120$.

In this paper, we only do experiments for MC-UCB, and exhibit the influence of the number of strata. For a comparison between MC-UCB and other algorithms, see [1]. By studying the range of the $F(W)$, we set the parameter of the algorithm MC-UCB to $A = 150\log(n)$.

For $n = 200$ and $n = 2000$, we observe the influence of the number of strata in Figure 2 (the number of strata varying from 2 to 100). We plot results for MC-UCB, uniform stratified Monte-Carlo (that allocates a number of samples in each stratum proportional to the measure of the stratum), and also for crude, unstratified, Monte-Carlo. We observe the trade-off that we mentioned between pseudo-regret and quality, in the sense that the mean squared error of the estimate output by MC-UCB (when compared to the true integral of f) first decreases with K and then increases. Note that, without surprise, for a large n

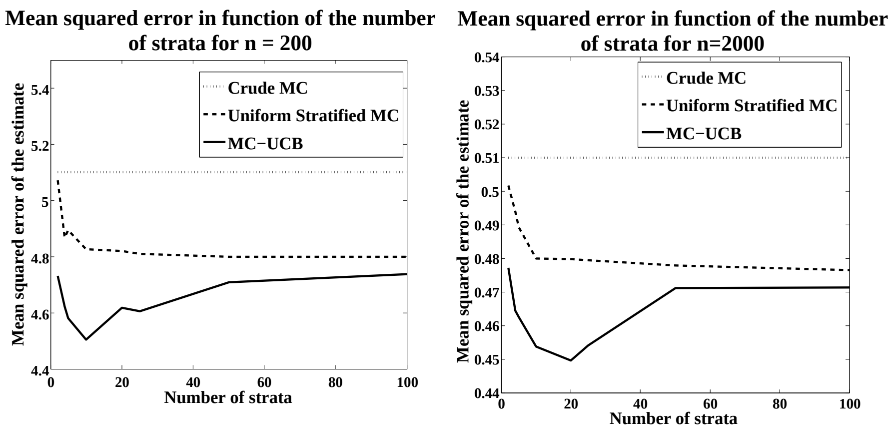

Fig. 2. Mean squared error for crude Monte-Carlo, uniform stratified sampling and MC-UCB, for different numbers of strata, for (Left:) n=200 and (Right:) n=2000.

the minimum of mean squared error is reached with more strata. Finally, note that our technique is never outperformed by uniform stratified Monte-Carlo.

Conclusion

In this paper we studied the problem of online stratified sampling for the numerical integration of a function given noisy evaluations, and more precisely we discussed the problem of choosing the *minimax-optimal number* of strata.

We explained why, to our minds, this is a crucial problem when one wants to design an efficient algorithm. We highlighted the trade-off between having many strata (and a good approximation error, i.e. quality of a partition), and not too many, in order to perform almost as well as the optimal oracle allocation on a given partition (small estimation error, i.e. pseudo-regret).

When the function is noisy, the noise is the dominant quantity in the optimal oracle variance on the optimal oracle partition. Indeed, decreasing the size of the strata does not diminish the (local) variance of the noise. In this case, the pseudo-risk of algorithm MC-UCB is equal, up to negligible terms, to the mean squared error of the estimate output by the optimal oracle strategy on the best (oracle) partition, at a rate of $O(n^{-\frac{d+4\alpha}{d+3\alpha}})$ where α is the Hölder exponent of s and m. This rate is minimax optimal on the class of α-Hölder functions: it is not possible, to do better on simultaneously all α-Hölder functions.

There are (at least) three very interesting remaining open questions:

- The first one is to investigate whether it is possible to estimate online the Hölder exponent *fast enough*. Indeed, one needs it in order to compute the proper number of strata for MC-UCB, and the lower bound on the Hölder exponent appears in the bound. It is thus a crucial parameter.

- The second direction is to build a more efficient algorithm in the noiseless case. We remarked that MC-UCB is not as efficient in this case as a simple non-adaptive method. The problem comes from the fact that in the case of a noiseless function, it is important to sample the space in a way that ensures that the points are as spread as possible.
- Another question is the relevance of fixing the strata in advance. Although it is minimax-optimal on the class of $\alpha-$Hölder functions to have hyper-cubic strata of same measure, it might in some cases be more interesting to focus and stratify more finely at places where the function is rough.

Acknowledgements

This research was partially supported by Region Nord-Pas-de-Calais Regional Council, French ANR EXPLO-RA (ANR-08-COSI-004), the European Communitys Seventh Framework Programme (FP7/2007-2013) under grant agreement 270327 (project CompLACS), and by Pascal-2.

References

[1] A. Carpentier and R. Munos. Finite-time analysis of stratified sampling for monte carlo. In *In Neural Information Processing Systems (NIPS)*, 2011a.

[2] A. Carpentier and R. Munos. Finite-time analysis of stratified sampling for monte carlo. Technical report, INRIA-00636924, 2011b.

[3] P. Etoré and B. Jourdain. Adaptive optimal allocation in stratified sampling methods. *Methodol. Comput. Appl. Probab.*, 12(3):335–360, September 2010.

[4] P. Etoré, G. Fort, B. Jourdain, and É. Moulines. On adaptive stratification. *Ann. Oper. Res.*, 2011. to appear.

[5] E. Giné and R. Nickl. Confidence bands in density estimation. *The Annals of Statistics*, 38(2):1122–1170, 2010.

[6] P. Glasserman. *Monte Carlo methods in financial engineering.* Springer Verlag, 2004. ISBN 0387004513.

[7] P. Glasserman, P. Heidelberger, and P. Shahabuddin. Asymptotically optimal importance sampling and stratification for pricing path-dependent options. *Mathematical Finance*, 9(2):117–152, 1999.

[8] V. Grover. Active learning and its application to heteroscedastic problems. *Department of Computing Science, Univ. of Alberta, MSc thesis*, 2009.

[9] M. Hoffmann and O. Lepski. Random rates in anisotropic regression. *Annals of statistics*, pages 325–358, 2002.

[10] R. Kawai. Asymptotically optimal allocation of stratified sampling with adaptive variance reduction by strata. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 20(2):1–17, 2010. ISSN 1049-3301.

[11] R.Y. Rubinstein and D.P. Kroese. *Simulation and the Monte Carlo method.* Wiley-interscience, 2008. ISBN 0470177942.

[12] A. Carpentier and R. Munos. Minimax Number of Strata for Online Stratified Sampling given Noisy Samples. Technical report, INRIA-00698517, 2012.