

# Beating Monte Carlo

If you work long enough in any mathematical science, at some point you will need to estimate an integral that does not have a simple closed form. Maybe your function is really complicated. Maybe it's really highly dimensional. Often you cannot even write it down: it could be a quantity associated with a complex system, that you can only "query" at certain points by running an experiment. But you still need your integral, and then you turn to the trustworthy old *Monte Carlo method*. (Check this article by Nicholas Metropolis for the history of the method and what it has to do with Stanislaw Ulam's uncle's gambling habit.) My goal in this post is to tell you a little bit about how you can do better than Monte Carlo using discrepancy theory.

## 1 The Problem and the Monte Carlo Method

Let us fix some notation and look at the simplest possible setting. We have a function  $f$ , that maps the real interval  $[0, 1]$  to the reals, and we want to know

$$\int_0^1 f(x)dx.$$

We will estimate this integral with the average  $\frac{1}{n} \sum_{i=1}^n f(y_i)$ , where  $y := y_1, y_2, \dots$  is a sequence of numbers in  $[0, 1]$ . The error of this estimate is

$$\text{err}(f, y, n) := \left| \int_0^1 f(x)dx - \frac{1}{n} \sum_{i=1}^n f(y_i) \right|.$$

And here is the main problem I will talk about in this post: *How do we choose a sequence  $y$  of points in  $[0, 1]$  so that the average  $\frac{1}{n} \sum_{i=1}^n f(y_i)$  approximates the integral  $\int_0^1 f(x)dx$  as closely as possible?* Intuitively, for larger  $n$  the approximation will be better, but what is the best rate we can achieve? Notice that we want a single sequence, so that if we want a more accurate approximation, we just take a few more terms and re-normalize, rather than start everything from scratch.

The insight of the Monte Carlo method is to take each  $y_i$  to be a fresh random sample from  $[0, 1]$ . Then for any  $n$ , the expectation of  $\frac{1}{n} \sum_{i=1}^n f(y_i)$  is exactly  $\int f(x)dx$  (from now on I will skip the limits of my integrals: they will

all be from 0 to 1). The standard deviation is

$$\frac{1}{\sqrt{n}} \left( \int f(x)^2 dx \right)^{1/2} = \frac{\|f\|_{L_2}}{\sqrt{n}},$$

and standard concentration inequalities tell us that, with high probability,  $\text{err}(f, y, n)$  will not be much larger than the latter quantity.

## 2 Quasi-Monte Carlo and Discrepancy

For a fixed function of bounded  $L_2$  norm, the Monte Carlo method achieves  $\text{err}(f, y, n)$  roughly on the order of  $n^{-1/2}$ . In other words, if we want to approximate our integral to within  $\varepsilon$ , we need to take about  $\varepsilon^{-2}$  random samples. It's clear that we in general, even for smooth functions, we cannot hope to average over fewer than  $\varepsilon^{-1}$  points, but is  $\varepsilon^{-2}$  really the best we can do? It turns out that for reasonably nice  $f$  we can do a lot better using discrepancy.

We define the *star discrepancy function* of a sequence  $y := y_1, y_2, \dots$  as

$$\delta^*(y, n) := \max_{0 \leq t \leq 1} \left| t - \frac{|\{i : i \leq n, y_i < t\}|}{n} \right|.$$

Notice that this is really just a special case of  $\text{err}(f, y, n)$  where  $f$  is the indicator function of the interval  $[0, t)$ . A beautiful inequality due to Koksma shows that in a sense these are the only functions we need to care about:

$$\text{err}(f, y, n) \leq V(f) \delta^*(y, n).$$

$V(f)$  is the *total variation* of  $f$ , a measure of smoothness, and for continuously differentiable functions it is equal to  $\int |f'(x)| dx$ . The important part is that we have bounded the integration error by the product of a term that quantifies how nice  $f$  is, and a term that quantifies how “random” the sequence  $y$  is. With this, our task is reduced to finding a sequence  $y$  with small star discrepancy, hopefully smaller than  $n^{-1/2}$ . This is the essence of the quasi-Monte Carlo method.

## 3 The van der Corput Sequence

A simple sequence with low star discrepancy was discovered by van der Corput in the beginning of the 20th century. Let us write the integer  $i$  in binary as  $i = i_k \dots i_0$ , i.e.  $i = i_0 2^0 + i_1 2^1 + i_2 2^2 + \dots + i_k 2^k$ . Then we define  $r(i)$  to be number we get by flipping the binary digits of  $i$  around the radix point:  $r(i) := i_0 2^{-1} + i_1 2^{-2} + \dots + i_k 2^{-k-1}$ , or, in binary,  $r(i) = 0.i_1 i_2 \dots i_k$ . The *van der Corput sequence* is  $r(0), r(1), r(2), \dots$

I have plotted the first few terms of the van der Corput sequence on the left in Figure 1: the index  $i$  is on the  $x$ -axis and the value  $r(i - 1)$  on the  $y$ -axis. The terms alternate between  $[0, 1/2)$  and  $[1/2, 1)$ ; they also visit each of  $[0, 1/4)$ ,  $[1/4, 1/2)$ ,  $[1/2, 3/4)$ ,  $[3/4, 1)$  exactly once before they return, and




Figure 1: The first eight terms of the van der Corput sequence.

so on. For example, each shaded rectangle on the right in Figure 1 contains exactly one point (the rectangles are open on the top). The key property of the van der Corput sequence then is that  $r(i) \in [j2^{-k}, (j+1)2^{-k})$  if and only if  $i \equiv j \pmod{2^k}$ . So for any such *dyadic interval*, the discrepancy is at most  $1/n$ : the number of integers  $i$  less than  $n$  such that  $i \equiv j \pmod{2^k}$  is either  $\lfloor n2^{-k} \rfloor$  or  $\lceil n2^{-k} \rceil$ . We can greedily decompose an interval  $[0, t)$  into  $O(\log n)$  dyadic intervals plus a leftover interval of length  $o(1/n)$ ; therefore the *star discrepancy of the van der Corput sequence  $y$*  is  $O((\log n)/n)$ .

Remember that, together with Koksma's inequality, this means that we can estimate the integral of any function  $f$  with bounded variation with error  $\text{err}(f, y, n) \ll (V(f) \log n)/n$ , which, for sufficiently smooth  $f$ , is almost quadratically better than the Monte Carlo method! And with a deterministic algorithm, to boot.

This was not that hard, so maybe we can achieve the ultimate discrepancy bound  $O(1/n)$ ? This is the question (asked by van der Corput) which essentially started discrepancy theory. The first proof that  $O(1/n)$  is not achievable was given by van Aardenne-Ehrenfest. Klaus Roth simplified her bound and strengthened it to  $\Omega(\sqrt{\log n}/n)$  using a brilliant proof based on Haar wavelets. Schmidt later proved that van der Corput's  $O((\log n)/n)$  bound is asymptotically the best possible.

## 4 Higher Dimension, Other Measures, and Combinatorial Discrepancy

Quasi-Monte Carlo methods are used in real world applications, for example in quantitative finance, because of the better convergence rates they offer. But there are many complications that arise in practice. One issue is that we usually need to estimate integrals of *high-dimensional functions*, i.e. functions of a large number of variables. The Koksma inequality generalizes to higher dimensions

(the generalization is known as the Koksma-Hlawka inequality), but we need to redefine both discrepancy and total variation for that purpose. The star discrepancy  $\delta^*(y, n)$  of a sequence  $y$  of  $d$ -dimensional points measures the worst-case absolute difference between the  $d$ -dimensional volume (Lebesgue measure) of any anchored box  $[0, t_1] \times \dots \times [0, t_d]$  and the fraction of points in the sequence  $y_1, \dots, y_n$  that fall in the box. The generalization of total variation is the Hardy-Krause total variation. Naturally, the best achievable discrepancy increases with dimension, while the class of functions of bounded total variation becomes more restrictive. However, we do not even know what is the best achievable star discrepancy for 2 or higher dimensional sequences! We know that no  $d$ -dimensional sequence has discrepancy better than  $\Omega((\log^{d/2+\eta_d} n)/n)$ , where  $\eta_d > 0$  is some constant that goes to 0 with  $d$ . The van der Corput construction generalizes to higher dimensions and gives sequences with discrepancy  $O((\log^d n)/n)$  (the implied constants here and in the lower bounds depend on  $d$ ). The discrepancy theory community refers to closing this significant gap as “The Great Open Problem”.

Another issue in high dimension has to do with asymptotics: for very high dimension the star discrepancy bounds above are inferior to what we get from a random sequence unless we take very large  $n$ . For this reason, finding sequences with discrepancy that behaves like  $d^c n^{1/2-\varepsilon}$ , for constants  $c$  and  $\varepsilon > 0$  independent of  $d$ , might have an even larger impact than solving the great open problem. On the positive side, I should point out that in practice Quasi Monte Carlo methods are used for large dimensional estimation: it appears that the worst case error bound given by the Koksma-Hlawka inequality can be too pessimistic.

Everything so far was about integration with respect to the Lebesgue measure, but in practice we are often interested in a different measure space. We could absorb the measure into the function to be integrated, but this can affect the total variation badly. We could do a change of variables, but, unless we have a nice product measure, this will result in a notion of discrepancy in which the test sets are not boxes anymore. Maybe the most natural solution is to redefine star discrepancy with respect to the measure we care about. But how do we find a low-discrepancy sequence with the new definition? It turns out that *combinatorial discrepancy* is very helpful here. A classical problem in combinatorial discrepancy, Tuszady’s problem, asks for is the smallest function  $\Delta_d(n)$  such that any set of  $n$  points in  $\mathbb{R}^d$  can be colored with red and blue so that in any axis-aligned box  $[0, t_1] \times \dots \times [0, t_d]$  the absolute difference between the number of red and the number of blue points is at most  $\Delta_d(n)$ . A general transference theorem in discrepancy theory shows that *for any probability measure in  $\mathbb{R}^d$*  there exists a sequence  $y$  with star discrepancy at most  $O(\Delta_{d+1}(n)/n)$ . The best bound for  $\Delta_d(n)$  is  $O(\log^{d+0.5} n)$ , which gives star discrepancy  $O((\log^{d+1.5} n)/n)$  for any probability measure, only slightly worse than the best star discrepancy for Lebesgue measure. This transference result has long been seen as purely existential, because most non-trivial results in combinatorial discrepancy were not constructive, but recently we have seen amazing progress in algorithms for minimizing combinatorial discrepancy. While even

with these advances we don't get sequences that are nearly as explicit as the van der Corput sequence, there certainly is hope we will get there.

## 5 Conclusion

I have barely scratched the surface of Quasi Monte Carlo methods and geometric discrepancy. Koksma-Hlawka type inequalities, discrepancy with respect to various test sets and measures, combinatorial discrepancy are each a big topic in itself. The sheer breadth of mathematical tools that bear on discrepancy questions is impressive: diophantine approximation to construct low discrepancy sequences, reproducing kernels in Hilbert spaces to prove Koksma-Hlawka inequalities, harmonic analysis to prove discrepancy lower bounds, convex geometry for upper and lower bounds in combinatorial discrepancy. Luckily, there are some really nice references available. Matoušek has a very accessible book on geometric discrepancy. Chazelle focuses on computer science applications. A new collection of surveys edited by Chen, Srivastav, and Travaglini has many of the latest developments.