GAUSS' CIRCLE PROBLEM

1. Gauss' circle problem

We begin with a very classical problem: how many lattice points lie on or inside the circle centered at the origin and with radius r? (In keeping with the classical terminology, we will in this section speak of lattice points "in the circle" rather than the more careful modern terminology "in the closed disk of radius r".) Let $L(r)$ be the number of such points.

If one gathers a bit of data, it becomes apparent that $L(r)$ grows quadratically with r, which leads to consideration of the function $L(r)/r^2$. Here a little bit of data makes the conclusion inescapable:

$$R(10)/10^2 = 3.17.$$

$$R(100)/100^2 = 3.1417.$$

$$R(1000)/1000^2 = 3.141549.$$

$$R(10^4)/10^8 = 3.14159053.$$

Thus the result we want is surely the following:

Theorem 1. *As $r \to \infty$, we have $L(r) \sim \pi r^2$.*

In fact, if we did not guess this result before, we may feel a little sheepish: πr^2 is of course the area of the region bounded by the circle $x^2 + y^2 = r^2$, and it stands to reason that the area of a circle is a good approximation to the number of lattice points it contains. (Why? A picture would surely help. The key is to associate to each lattice point the unit square with that lattice point as the southwestern corner.)

But how to prove it? If we are in a "calculus mood," our first thought is that the area of the circle of radius r is obtained as a Riemann integral, i.e., by integrating 1 over the interior of the circle. This involves integration over a region with curved sides. An equivalent but technically easier approach is to integrate the characteristic function $\chi(r)$ of the disk of radius r over the entire plane, or say, over the square $[-r, r]^2 = [-r, r] \times [-r, r]$. Recall that this function χ is by definition 1 at a point (x, y) inside the disk – so with $x^2 + y^2 \leq r^2$ – and 0 otherwise.

Another nice trick is to take seriously the rescaling $L(r) \mapsto L(r)/r^2$. Namely, counting the number of standard lattice points inside the circle $x^2 + y^2 = r^2$ is equivalent to counting the number of points (x, y) inside the unit disk $x^2 + y^2 = 1$ such that $rx, ry \in \mathbb{Z}$ (i.e., rational numbers whose denominators divide r). Namely, we would like to see that the number of these "$\frac{1}{r}$-lattice points" divided by r^2 approaches π as $r \to \infty$. But now consider that we can divide the square $[-1, 1]$ into $(2r)^2 = 4r^2$ subsquares of width $\frac{1}{r}$. Consider the sum $\Sigma(r)$ of the form

$$\sum_{i,j} \epsilon(i, j) \frac{1}{r^2},$$

where $-r \leq i, j < r$ so the sum ranges over the subsquares, and we will put $\epsilon(i, j) = 1$ if the southwestern corner of the subsquare is on or inside the unit circle and zero otherwise. On the one hand we have $\Sigma(r) = \frac{L(r)-2}{r^2}$, since each $(\frac{1}{r})$-lattice point inside the circle is the southwestern corner of a unique subsquare of $[-1,1]^2$ except the right-most point $(0,1)$ and the top-most point $(1,0)$. This little discrepancy hardly matters, because of course

$$\lim_{r\to\infty} \Sigma(r) = \lim_{r\to\infty} \frac{L(r)}{r^2} - \lim_{r\to\infty} \frac{2}{r^2} = \lim L(r)r^2 - 0 = \lim_{r\to\infty} \frac{L(r)}{r^2};$$

where the first limit exists iff the last limit does. But indeed $\lim_{r\to\infty} \Sigma(r)$ does exist, since $\Sigma(r)$ can be interpreted as a sequence of Riemann sums for the integral $\int_{-1}^{1}\int_{-1}^{1} \chi(1)dxdy$ corresponding to the partition of $[-1,1]^2$ into r^2 subsquares. Since the mesh of this sequence of partitions approaches zero, the sums approach the integral of the characteristic function $\chi(1)$ of the unit circle, namely π.

This completes the proof of the theorem, except perhaps that we ought to reflect a little (now, if not in the multivariable calculus class when the characteristic function device was first introduced to us) on the Riemann integrability of the characteristic function $\chi(1)$ of the unit disk. It is a bounded function, but it is not everywhere continuous: indeed it is discontinuous precisely along its boundary, the unit circle $x^2 + y^2 = 1$. Now it is a famous (and not so easy) result from advanced calculus that a bounded function is Riemann-integrable iff the set S of its discontinuities has **measure zero** in the sense that for every $\epsilon > 0$ there is a (possibly infinite) union of open disks B_{r_i} covering S such that the sum of the areas of all the disks is at most ϵ. It is also well-known (and much easier to check) that any piecewise smooth curve has measure zero.

Remark: We can take further advantage of the flexibility of choice of a sample point in the Riemann sum to see that if we counted either (i) the number of unit squares lying entirely inside $x^2 + y^2 = r^2$ or (ii) the number of unit squares any portion of which lies inside $x^2 + y^2 = r^2$, then these quantities are also asymptotic to πr^2. Just use lower sums in the first case and upper sums in the second case.

In fact the argument proves something much more general: in place of the unit circle, we could take any subset S of $[-1,1]$ whose boundary ∂S has measure zero in the above sense – a **Jordan set**. For instance, take S to be a region bounded by a piecewise smooth simple closed curve.[1] Then the argument shows that as $r \to \infty$ the number $L_s(r)$ of $(\frac{1}{r})$-lattice points inside S divided by r^2 approaches the area $A(S)$ of S: here $A(S)$ is defined to be equal to the integral of the characteristic function of S.[2]

To get a result analogous to Theorem 1, we scale back: for any subset S of the plane and α a positive real number, let $\alpha S = \{(\alpha x, \alpha y) \mid (x,y) \in S\}$ be the

[1]In fact we could take a much more rugged "fractal" set – say, the interior of the Koch snowflake – and the result would still hold!

[2]There is one small change: rather than there being 2 $\frac{1}{r}$-squares whose SW corners lie in S even though the squares do not lie in $[-1,1]^2$, we may now have as many such subsquares as fill the entire northern and eastern corners of $[-1,1]^2$, namely $2(2r+1) - 1 = 4r + 1$. This still goes to zero when divides by r^2, so no matter.

"α-dilate" of S. Moreover define $L_S(r)$ to be the number of standard lattice points inside rS, the r-dilate of S. Then we have shown:

Theorem 2. *Suppose that $S \subset [-1,1]$ is a Jordan subset (i.e., with Riemann-integrable characteristic function). Then as $r \to \infty$,*

$$L_S(r) \sim \text{Area}(S) \cdot r^2.$$

Let us now come back to the circle problem. Can we use Theorem 1 to compute π? Well, yes and no: taking larger and larger r we will get a sequence of values, $L(r)/r^2$, which approach π, but in order to compute π to a certain number of decimal places we need to have an estimate on the error term $E(r)$, where

$$E(r) = |N(r) - \pi r^2|.$$

Theorem 3. *For all $r \geq 8$, $E(r) \leq 10r$.*

Proof: Let $P = (x, y)$ be a lattice point with $x^2 + y^2 \leq r^2$. As above, consider the unit square $S(P) = [x, x+1] \times [y, y+1]$ uniquely associated to P. Since the diameter of the unit square is $\sqrt{2}$, although $S(P)$ may or may not lie entirely inside the circle of radius r, it necessarily lies entirely on or inside the circle of radius $r + \sqrt{2}$. Thus we get

$$N(r) \leq \pi(r + \sqrt{2})^2 = \pi r^2 + 2\pi\sqrt{2}r + 2\pi.$$

Similarly, suppose (x, y) is any point in the plane with $\sqrt{x^2 + y^2} \leq r - \sqrt{2}$. Then the entire unit square $([x], [x]+1) \times ([y], [y]+1)$ lies on or inside the circle of radius r, which gives the estimate

$$N(r) \geq \pi(r - \sqrt{2}) = \pi r^2 - 2\pi\sqrt{2}r + 2\pi.$$

Thus

$$|N(r) - \pi r^2| \leq 2\pi + 2\sqrt{2}\pi r \leq 7 + 9r < 10r,$$

the last inequality holding for all $r \geq 8$.

Corollary 4. *For any $r \geq 8$,*

$$|\pi - \frac{L(r)}{r^2}| \leq \frac{10}{r}.$$

Thus one can really compute π by counting lattice points in circles! However, the bound of the preceding corollary makes for an excruciatingly slow computation: if we take $r = 1000$ then the Corollary guarantees (roughly) 2 decimal places of accuracy. In fact, looking at the above value of $L(r)/r^2$, we see that we in fact have 4-place accuracy, and in general the accuracy seems noticeably better than what our bounds guarantee.

2. Connections to average values

Recall that were interested in the arithmetical function $r_2(n)$, which counts the number of pairs of integers (x, y) such that $n = x^2 + y^2$. We did not (at least not officially) give an exact formula for $r_2(n)$, but we will see now that it is much easier to compute the function "on average". Indeed, this is essentially what we've already done.

For f an arithmetical function, we define

$$f_{\text{ave}}(n) = \frac{1}{n}(f(1) + f(2) + \ldots + f(n)),$$

i.e., f_{ave} at n is simply the average ("arithmetic mean") of the first n values of f. The idea here is that f_{ave} is roughly of the same size of f but is, at least in many cases, "more regular."[3] In other words, many classical arithmetical functions f have distinct lower and upper orders, but there is often a simple asymptotic expression for f_{ave}.

A trivial but useful remark is that in practice one often passes between f_{ave} and simply the function which is the sum of the first n values of f, $n \mapsto \sum_{i=1}^{n} f(i)$. It is tempting to denote this summatory function by F, in analogy to the antiderivative, but of course we have already used this notation for the sum over (only) the divisors of n. So let us call it instead $S(f)$, a rather awkward notation that we hope not to have to use too frequently. In particular, if we know $f_{\text{ave}} \sim g$ then $S(f)(n) \sim n \cdot g$, and conversely.

Theorem 5. *If $f(n) = r_2(n)$, then $f_{\text{ave}}(n) \sim \pi$.*

Proof: Note first that $r_2(1) + \ldots + r_2(n)$ counts the number of solutions to the inequality $x^2 + y^2 \leq n$. So for $n = r^2$ this is precisely the number $L(r)$, and we get that

$$S(f)(r^2) \sim \pi r^2.$$

We would like to have this with r instead of r^2. For this, we leave it to the reader to go back and verify that in the work of last section, we didn't need to consider merely integer dilates of the unit circle (or of any set S). Defining, as we did, the α dilate αS for any $\alpha \in \mathbb{R}^{>0}$, one can modify the proofs to show that

$$\lim_{\alpha \to \infty} \frac{L_S(\alpha)}{\alpha^2} = \text{Area}(S).$$

(We remark that the second proof goes through with no change whatsoever and gives the same error bound. If α is not an integer, it is no longer literally possible to divide the square $[-1, 1]^2$ into α^2 equally sized subsquares, but a bit of thought shows that this is close enough to being true to make the argument go through.) Thus we have the desired

$$S(f)(n) \sim \pi n,$$

and dividing through by n, we get that

$$f_{\text{ave}}(n) = \frac{1}{n}(r_2(1) + \ldots + r_2(n)) \sim \pi.$$

Note that since for every nonsquare n (and hence "most n"), if $r_2(n) > 0$ then $r_2(n) \geq 4$, since if $x^2 + y^2 = n$ is one representation, $(-x)^2 + y^2$, $x^2 + (-y)^2$ and $(-x)^2 + (-y)^2$ are three more representations. From this it follows that there must be a lot of values of n for which $r_2(n) = 0$. Ironically, since π is slightly larger than 3, we cannot quite conclude from this that the proportion of n for which $r_2(n)$ is at least $\frac{1}{4}$, although we know this to be the case (since we know that $r_2(4k + 3) = 0$). However, looking a bit more carefully we see that if n is also not of the form $2m^2$

[3]This must be taken *cum grano salis*; in the Exercises you will look at $f(n) = n$.

(which is still "zero percent of all numbers") then also by switching the roles of x and y we see that $r_2(n) > 0$ implies $r_2(n) \geq 8$, and this implies for instance that at $r_2(n) = 0$ at least 60% of the time – the average value of a non-negative quantity which is at least 8 at least 40% of the time is at least $.4 * 8 = 3.2 > \pi$.

In fact this only hints at the truth: truthfully, $r_2(n)$ is equal to 0 most of the time. In other words, if we pick a very large number N and choose at random an element $1 \leq n \leq N$, then the probability that n is a sum of 2 squares goes to 0 as $N \to \infty$. This exposes one of the weaknesses of the arithmetic mean: without further assumptions it is unwarranted to assume that the average value is a "typical" value of the function in any reasonable sense. To better capture this typicality one can import further statistical methods and study the **normal order** of an arithmetical function. With regret, we shall have to pass this concept over entirely as being too delicate for our first course.

3. The prospect of improved error bounds

Suppose we are given a nice region S in the plane (say connected, compact, convex, equal to the closure of its interior and with boundary a piecewise smooth curve), and consider the function $L(r)$ which counts the number of lattice points in the rth dilate rS of S, as above. By the Riemann sum argument, we know that $L(r) \sim A(S)r^2$, and we saw that we could do somewhat better in the case of a unit circle: the error

$$E(r) = |L(r) - A(S)r^2|$$

was bounded by a constant times r. Let us look for further improvements.

First, since the boundary of S is a piecewise smooth curve, it is "one-dimensional", and this makes it geometrically clear that the number of unit squares intersecting the boundary of rS is less than a constant times r. Thus it should be true that $E(r) \leq Cr$ in all cases. This observation was first made by Gauss. It is indeed true, but we will not attempt a rigorous proof here.

The next observation to make is that in the above level of generality one cannot do any better. Namely, suppose S is the square $[-1, 1] \times [-1, 1]$, so that $rS = [-r, r] \times [-r, r]$. In this case we can compute both $L(r)$ and $A(rS)$ exactly, getting

$$L(r) = (2r + 1)^2 = 4r^2 + 4r + 1$$

and

$$A(r) = (2r)^2 = 4r^2,$$

so

$$E(r) = 4r + 1.$$

In other words, the number of lattice points always exceeds the area by a linear function. Further experimentation with convex lattice polygons – i.e., convex polygons all of whose vertices are lattice points – reveals that this is a general phenomenon. However, in this case there is a beautiful exact formula. The key observation is that the excess comes from the lattice points on the boundary:

Theorem 6. *(Pick) Let S be a convex lattice polygon together with its interior. Let $\partial(S)$ be the number of lattice points on the perimeter of S. Then*

$$L(r) = \text{Area}(S)r^2 + \frac{\partial(S)}{2}r + 1.$$

An equivalent – and more common – formulation is that the area of a lattice polygon is equal to the number of interior lattice points plus half of the boundary lattice points minus 1. We leave the proof of Pick's Theorem as an exercise; the internet is replete with sketched proofs.

Note that the function $L(r)$ turns out just to be a quadratic polynomial with rational coefficients!

There is a beautiful generalization of this to **lattice polytopes** in higher-dimensional Euclidean spaces, due to Ehrhart, which unfortunately lies outside the scope of our course.

Let us now return to the case of the unit circle. Recall that it seemed from our (admittedly meager) data that one could do better than $E(r) \leq Cr$. In fact improving the error is (or was) an industry unto itself. The bound $E(r) \leq Cr^{\frac{2}{3}}$ was attained by Sierpinski in 1906. Various incremental improvements on the exponent have been made over the years. For a while the best known exponent has been $\frac{46}{73} \approx .63$, due to Huxley in 1990. On the other hand, it is a result of Hardy and Landau that there is *not* a positive constant C such that $E(r) \leq Cr^{\frac{1}{2}}$, so for a long time the standard conjecture has been the following:

Conjecture 7. *For every $\epsilon > 0$, there exists a number C (depending on ϵ such that*

$$E(r) \leq Cr^{\frac{1}{2}+\epsilon}.$$

We mention that a recent (February 2007) preprint due to Cappell and Shaneson claims a full proof of this conjecture. As of this writing, their paper has not yet been sufficiently checked by the experts (I am not one).

There are yet further results on regions of other shapes: for instance, for any set other than the circle, one can rotate S and get a different set, and if by some unlucky accident we have "too many" lattice points on the boundary, it is plausible that a small rotation will improve the situation. Moreover, one can pursue the problem in higher-dimensions. For instance, the analogous argument involving lattice points on or inside the sphere of radius r in $\mathbb{R}^3$ gives:

Theorem 8. *The number of integer solutions (x, y, z) to $x^2 + y^2 + z^2 \leq r^2$ is asymptotic to $\frac{4}{3}\pi r^3$, with the error being bounded by (e.g.!) Cr^2.*

Corollary 9. *The average value of the function $r_3(n)$ which counts representations by sums of three squares is asymptotic to $\frac{4}{3}\pi\sqrt{n}$.*

We can similarly compute nice asymptotic expressions for the average value of $r_k(n)$ for $k \geq 4$ provided only we know a formula for the volume of the unit ball in $\mathbb{R}^k$. Of course such formulas are known (internet exercise: look them up!); alternately it might be fun to guess them by actually counting $(\frac{1}{r})$-lattice points.