Geometry in Very High Dimensions

Renato Feres
Washington University, St. Louis

September, 2006

1 Balls, Spheres and Cubes

The n-dimensional ball, or *n-ball*, of radius r is the subset of Euclidean n-space defined by:

$$B^n(r) = \{x \in \mathbb{R}^n : \|x\| \leq r\}.$$

Here, $x = (x_1, \dots, x_n)$ is an element of $\mathbb{R}^n$ and $\|x\| = \sqrt{x_1^2 + \cdots + x_n^2}$.

The *n-sphere* is the boundary of the $n+1$-ball. Thus, it is given by the set

$$S^n(r) = \{x \in \mathbb{R}^{n+1} : \|x\| = r\}.$$

The *n-cube* of side $2r$ is the subset of $\mathbb{R}^n$ defined as the n-fold product of intervals $[-r, r]$:

$$K^n(r) = [-r, r]^n = \underbrace{[-r, r] \times \cdots \times [-r, r]}_{n \text{ times}}.$$

We will be interested in the volumes of various kinds of regions in Euclidean space, such as the sets just defined. The volume of a region $\mathcal{R}$ of dimension n will be written as $V(\mathcal{R})$. For example, $V(S^2(r))$ is the surface area (2-dimensional "volume") of the sphere of radius r, so $V(S^2(r)) = 4\pi r^2$, whereas $V(B^3(r)) = \frac{4}{3}\pi r^3$.

If $\mathcal{R}(r)$ is any of the above sets, note that $\mathcal{R}(r) = r\mathcal{R}$. That is, changing r changes the sets by a simple scale transformation having Jacobian determinant r^n. This implies the intuitive fact that

$$V(\mathcal{R}(r)) = r^n V(\mathcal{R}(1)).$$

Note that the vertices of the n-cube lie in the set of all n-tuples of points $(x_1, \dots, x_n)$ such that each x_i belongs to the two-point set $\{-r, r\}$. We can interpret choosing a vertex at random as the outcome of flipping n independent coins. I will have more to say about this point.

2 Testing our geometric intuition

We divide the n cube $K^n(1)$ into 2^n smaller cubes (sectors) as in figure 1. Each sector is a translate of $K^n(1/2)$. Now draw the largest centered ball in each sectors. These are balls of radius $1/2$. Draw the n-ball, centered at the origin, touching the surface of all the other balls. We can do this for each dimension n. Let r_n be the radius of the center ball. What is the limit of the sequence $r_1, r_2, r_3, \ldots$?

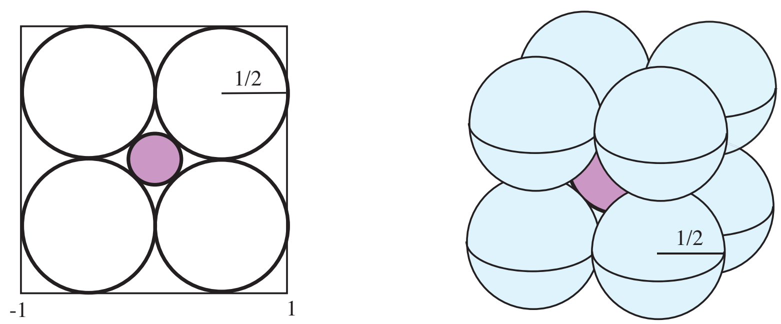

Figure 1: What happens to the radius of the central ball as the dimension n increases to infinity?

You can compute r_n easily by noting that $r_n + 1/2$ is the length of the vector from the origin to the center of one of the sectors. The center of the positive sector is $(1/2, 1/2, \ldots, 1/2)$. Therefore, $r_n + 1/2 = \|(1/2, \ldots, 1/2)\| = \sqrt{n}/2$. this gives the answer

$$r_n = (\sqrt{n} - 1)/2.$$

Does this surprise you? (It did surprise me when I was first told about it.) This is saying that the radius of the center ball goes to infinity at the speed of $\sqrt{n}$ as the dimension grows.

3 Concentration of volumes on shells

There are other surprises in high dimensions related to the distribution of volume in n-dimensional regions. Let us consider volumes of spherical shells. A spherical shell of radius r and thickness a is the region between two concentric balls, the bigger one with radius r and the smaller one with radius $r - a$. A 2-dimensional shell is shown in figure 2.

The question we ask now is this: what fraction of the total volume of the ball $B^n(r)$ is contained in the shell of radius r and thickness a? This ratio is

$$\begin{aligned}\frac{V(B^n(r)) - V(B^n(r-a))}{V(B^n(r))} &= \frac{V(B^n(1))(r^n - (r-a)^n)}{V(B^n(1))r^n} \\ &= \frac{r^n - (r-a)^n}{r^n} \\ &= 1 - (1 - a/r)^n.\end{aligned}$$

The conclusion is that, as n grows, most of the volume concentrates near the boundary sphere. For example, if $n \geq 500$, more than 99% of the volume of $B^n(r)$ is contained in a shell of thickness corresponding to 1% of the radius!

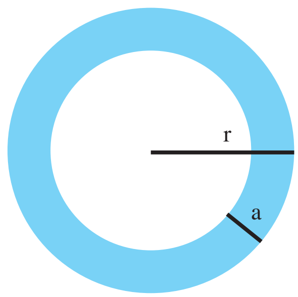

Figure 2: A shell of radius r and thickness a.

4 Concentration of volume on slices

We now show a fact which is remarkable both for its geometric content and because it will give a strong hint of a connection between these geometric properties and probability theory. It has to do with the problem of how volumes in the n-ball are distributed on parallel slices. Figure 3 shows the central slice and a parallel slice at position x.

Before getting to that, though, we need some preliminary results. We haven't used yet the actual value of $V(B^n(r))$. All that was used so far was the fact that $V(B^n(r)) = r^n V(B^n(1))$. This number can be obtained by Calc III methods. I leave it to you as an exercise to obtain the following (by induction):

$$V(B^n(1)) = \frac{\pi^{n/2}}{\Gamma(n/2+1)}.$$

Recall that the gamma function $\Gamma(x)$ is defined by

$$\Gamma(x) = \int_0^\infty e^{-t} t^{x-1} dt$$

and has the following special values: $\Gamma(n+1) = n!$ if n is a positive integer, $\Gamma(1/2) = \pi^{1/2}$, and $\Gamma(x+1) = z\Gamma(x)$.

When x is large, we can use Stirling's formula to approximate $\Gamma(x)$:

$$\Gamma(x) = \sqrt{2\pi} e^x x^{x-1/2}(1 + 1/(12x) + \dots).$$

This provides the approximation of $V(B^n(1))$ for large n:

$$V(B^n(1)) \sim \frac{1}{\sqrt{2\pi}}(2\pi e/n)^{n/2}.$$

Notice the interesting fact that a ball of *volume* equal to 1 must have radius a_n approximately (asymptotically) equal to

$$a_n \sim \sqrt{\frac{n}{2\pi e}}.$$

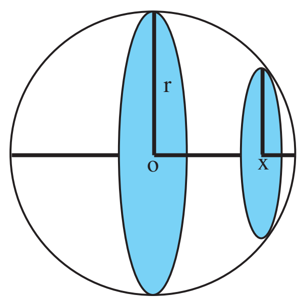

Figure 3: The central slice has the same radius as the ball $B^n(r)$. The parallel slice at position x has radius $\sqrt{r^2 - x^2}$.

We now return to the problem of how volumes are distributed over slices. To simplify the notation, I will write $v_n = V(B^n(1))$. In particular, $a_n = v_n^{-1/n}$. Let $B = B^n(a_n)$ denote the ball of volume 1. The $n-1$-dimensional volume of the central slice of B (which also has radius a_n) can be approximated using Stirling's formula (or the above approximate expression for v_n) as follows (recall the limit $(1 + a/n)^n \to e^a$):

$$
\begin{aligned}
V(B^{n-1}(a_n)) = v_{n-1}a_n^{n-1} &= v_{n-1}v_n^{-(n-1)/n} \\
&\sim \frac{1}{\sqrt{2\pi}}\left(\frac{2\pi e}{n-1}\right)^{(n-1)/2}\left[\frac{1}{\sqrt{2\pi}}\left(\frac{2\pi e}{n}\right)^{n/2}\right]^{-(n-1)/n} \\
&\sim \left(\frac{2\pi e}{n-1}\right)^{(n-1)/2}\left(\frac{2\pi e}{n}\right)^{-(n-1)/2} \\
&\sim \left(\frac{n}{n-1}\right)^{(n-1)/2} = \left(1+\frac{1}{n-1}\right)^{(n-1)/2} \\
&\sim \sqrt{e}.
\end{aligned}
$$

The volume of the slice at x is now given by

$$
\begin{aligned}
V(B^{n-1}(\sqrt{a_n^2-x^2})) &= V(B^{n-1}(a_n))\frac{(a_n^2-x^2)^{(n-1)/2}}{a_n^{n-1}} \\
&\sim \sqrt{e}\left(1-(x/a_n)^2\right)^{(n-1)/2} \\
&\sim \sqrt{e}\left[1-\left(\frac{x}{\sqrt{n/(2\pi e)}}\right)^2\right]^{(n-1)/2} \\
&= \sqrt{e}\left[1-\frac{2\pi e x^2}{n}\right]^{(n-1)/2} \\
&\sim \sqrt{e}e^{-\pi e x^2}.
\end{aligned}
$$

Therefore, the $n-1$-dimensional volume of the slice at x of an n-ball of total volume 1 approaches the value $e^{1/2}e^{-\pi e x^2}$ as n grows to infinity.

A numerical example will help show what is so remarkable about this result. Take the volume of a central slab of $B^n(r)$ (for an arbitrary r) made of all slices corresponding to values of x in the interval $-2.1r/\sqrt{n} \leq x \leq 2.1r/\sqrt{n}$ (the approximate value of $\sqrt{2\pi e}/2$ is 2.1). As n grows, this slab becomes thinner and thinner compared to the radius of the ball, but it contains a volume approximately equal to 96% of the total volume of the ball no matter how big n is. This means that the volume of the ball is concentrating more and more on the central slice! In other words, if we randomly (with uniform probability) choose a point on a ball of radius 1 of very high dimension, we can say with 96% confidence that this point will lie at a distance less than about $2.1/\sqrt{n}$ from the equatorial plane.

5 The central limit of probability theory

The concentration of volume on the central slice of the ball is a special case of a much more general result, which is related to the central limit theorem in probability theory. There was already a hint of the connection with probability in the way we interpreted the concentration phenomenon at the end of the previous section. To make this connection a bit more explicit, we look at the same problem in a somewhat different setting: the distribution of vertices on the n-cube $K^n(1)$.

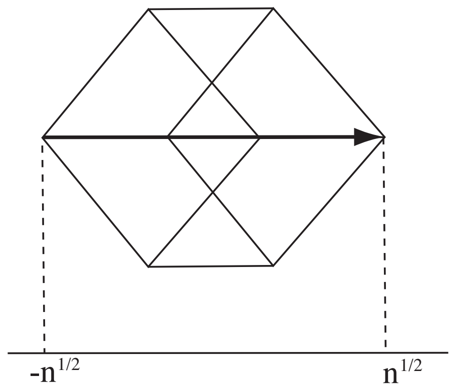

Figure 4: The diagonal of the cube $K^n(1)$ has length $2\sqrt{n}$. The vertices of the cube project orthogonally to the main diagonal, which is in the direction of the unit vector $u = (1, \ldots, 1)/\sqrt{n}$. The orthogonal projection of a vertex $x = (x_1, \ldots, x_n)$ on the direction of the vector u is given by the dot product: $x \cdot u = (x_1 + \cdots + x_n)/\sqrt{n}$.

Figure 4 show the cube $K^n(1)$ and its diagonal axis. The previous problem of calculating the volume of parallel slabs of the ball corresponds here to finding how many vertices of the cube project orthogonally to an interval $[a, b]$ of the diagonal axis. For a fixed n, all those projection points lie on the interval $[-\sqrt{n}, \sqrt{n}]$, since the length of the diagonal is $2\sqrt{n}$. As explained in the legend of the figure, this is the number of points $(x_1, \ldots, x_n)$ such that x_i is either 1 or -1 and

$$a \leq \frac{x_1 + \cdots + x_n}{\sqrt{n}} \leq b.$$

We can interpret this problem probabilistically as follows: if we pick a vertex of the cube randomly (equivalently, flip a fair coin n times independently to obtain a sequence of outcomes $(x_1, \ldots, x_n)$ where -1 represents "heads" and 1 "tails") what is the probability that the orthogonal projection will fall in the interval $[a, b]$? The standard central limit theorem says that this number, for large n, is approximately

$$\frac{1}{\sqrt{2\pi}} \int_a^b e^{-x^2/2} dx.$$

6 A little statistical physics

I'd like to briefly indicate now how these facts about high dimensional geometry are related to the kinetic theory of gases and the notion of temperature. Consider the gas system shown in figure 5. There are two types of gases in the container, which are distinguished by their mass, M_1 and M_2. Let n be the number of molecules of type 1 and m the number of type 2. Their velocity vectors at any particular moment will be written

$$v_1, \ldots, v_n, w_1, \ldots, w_m.$$

We suppose that the container is thermally insulated so that the total energy E is constant (we disregard any potential energy contribution to E). The energy is then given by

$$E = \frac{1}{2}M_1(\|v_1\|^2 + \cdots + \|v_n\|^2) + \frac{1}{2}M_2(\|w_1\|^2 + \cdots + \|w_m\|^2).$$

Note that each square norm is of the form $\|v\|^2 = v_x^2 + v_y^2 + v_z^2$. This equation for the total energy E describes an ellipsoid of dimension $3(n+m)-1$ in Euclidean space of dimension $3(n+m)$ (the *velocity space*).

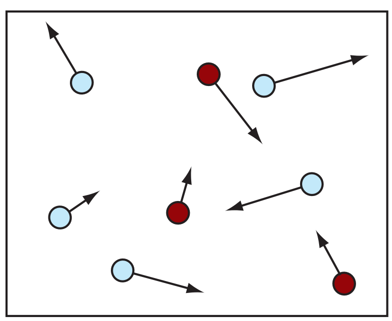

Figure 5: A container has two types of gases distinguished by masses M_1 and M_2. The system is thermally insulated so that the total (kinetic) energy of the system is assumed constant.

In terms of the the energies of the separate gas components, $E = E_1 + E_2$, where

$$\begin{aligned} E_1 &= \frac{1}{2}M_1(\|v_1\|^2 + \cdots + \|v_n\|^2) \\ E_2 &= \frac{1}{2}M_2(\|w_1\|^2 + \cdots + \|w_m\|^2). \end{aligned}$$

The mean energy per molecule of each type is: $T_1 = E_1/n$ and $T_2 = E_2/m$. These are, by definition, the *temperatures* of each gas type. From general physics we know that the system should in due time achieve *thermal equilibrium*, which means $T_1 = T_2$. Our goal is to understand what this means from a geometric viewpoint.

7 Why do the temperatures end up equal?

The *state* of the gas in the container is specified by a point in the ellipsoid of velocities. (We do not need here to keep track of the molecular positions.) The state of the subsystems of molecular types 1 and 2 can be represented by points in spheres of radii $r_1 = \sqrt{2E_1/M_1}$ and $r_2 = \sqrt{2E_2/M_2}$.

The total volume of the ellipsoid can be expressed by integrating a volume density written in terms of the radius r_1 as follows:

$$V(\text{Ellipsoid}) = \int_0^{\sqrt{2E/M_1}} V(S^{3n-1}(r_1))V(S^{3m-1}(r_2))dr_1$$

where r_2 depends on r_1 according to the relation $M_1 r_1^2 + M_2 r_2^2 = 2E$.

Thus the fraction of volume occupied in a slab around the slice of r_1 (which is the Cartesian product of two spheres) per unit thickness is

$$\frac{V(S^{3n-1}(r_1))V(S^{3m-1}(r_2))}{V(\text{Ellipsoid})} = C r_1^{3n-1} r_2^{3m-1}.$$

(The constant C depends on m, n and E.) Expressed as a function of T_1, this quantity takes the form:

$$f(T_1) = T_1^{\alpha}(c - T_1)^{\beta}$$

where $c = E/n$, $\alpha = (3n-1)/2$ and $\beta = (3m-1)/2$.

The maximum value of $f(r_1)$ is easily shown to be attained for

$$T_1 = \frac{\alpha}{\alpha+\beta} c \sim \frac{E}{n+m}$$

if $n + m$ is large. But then $T_1 \sim T_2$, and the approximation improves as the number of molecules increases.

The conclusion is that thermal equilibrium simply corresponds to a thermodynamical state of maximal volume. If we were to pursue the calculations of the previous sections we'd notice that the volume of the ellipsoid is very strongly concentrate around the region corresponding to $T_1 = T_2$ if the number of particles is large.

References

[Ball] Keith M. Ball. *An Elementary Introduction to Modern Convex Geometry*, chapter 1 of *Flavors of Geometry*, ed. Silvio Levy. Cambridge University Press - MSRI Publications, 1997.

[BZ] Ingemar Bengtsson and Karol Zyczkowski. *Geometry of Quantum States*, Cambridge University Press, 2006. (Chapter 1.)

[Conc] Michel Ledoux. *The Concentration of Measure Phenomenon*, Americal Mathematical Society, Mathematical Surveys and Monographs, Volume 89, 2001.