Two cases of counterintuitive high dimensional phenomena

Benne de Weger
Department of Mathematics and Computer Science
Eindhoven University of Technology
b.m.m.d.weger@tue.nl

v0.3, March 24, 2020

1 Introduction

High-dimensional spaces play important roles in several mathematical disciplines, such as combinatorics, complexity of algorithms (the 'combinatorial explosion' phenomenon in algorithms of superpolynomial complexity), coding theory, numerical mathematics, searching, sampling, optimization, lattice based cryptography, etc. Human intuition about dimensions often is rooted in our everyday experience with dimension up to 3. Sometimes this can be very misleading, as is well known among experts. Some authors even speak about "the curse of dimensionality", a term coined by Bellman[1].

This short note describes two simple phenomena in n-dimensional real space $\mathbb{R}^n$, that show surprising behaviour in higher dimensions (like $n = 100$ or even $n \geq 1000$) compared to the more humanly experienced dimensions ($n = 1, 2$ or 3; some people claim to have a good feeling for 4-dimensional space but I tend to distrust that).

This note does not claim any originality. Its main goal is to help students of mathematics improve their intuition about high dimensions if they base it on low dimensional phenomena.

2 Angles of random vectors

Take pairs of random vectors in $\mathbb{R}^2$, and look at the distribution of the angles (in $[0, \pi]$) of the pairs. Without loss of generality one may assume that one of the vectors is the *north pole*, $(0, 1)$, by applying some rotation and scaling, for which angles are invariant. Since the lengths of the vectors are irrelevant, we can sample vectors from a unit circle.

Figure 1 shows the result of an experiment with 1 million random vectors on the 2-dimensional unit circle. Not surprisingly the angles are uniformly distributed over the interval $[0, \pi]$, as is indicated by the graph of a horizontal line at height $1/\pi$.

Next we turn to higher dimensions. When we do a similar experiment on a 3-dimensional

[1]Richard E. Bellman, "Dynamic programming", Princeton University Press, 1957.

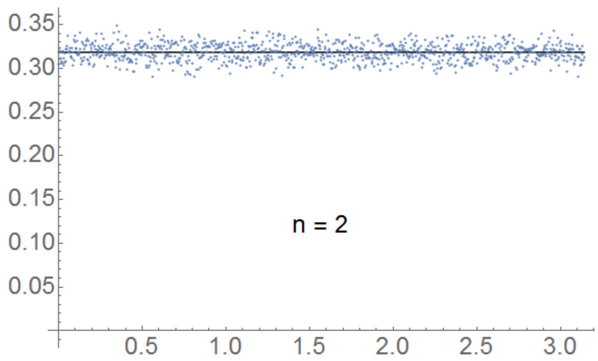

Figure 1: Distribution of 1 million angles between points randomly and uniformly selected on the unit circle; results were tallied in 1000 equally sized bins partitioning the interval $[0, \pi]$.

unit sphere[2], the result may come as a surprise, see Figure 2. The angle distribution is clearly not uniform anymore, and the question is what the exact distribution curve here is.

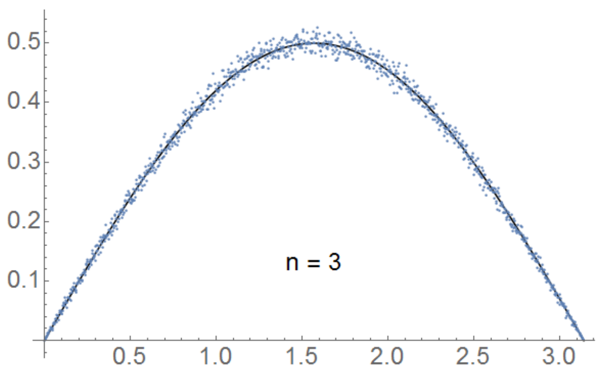

Figure 2: The same experiment on the unit sphere.

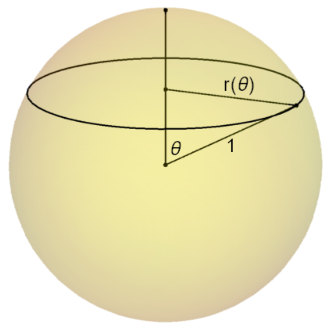

Figure 3: The unit sphere and the circle of points for angle θ.

This is not too hard to analyse. For a given $\theta \in (0, \pi)$, the points on the unit sphere with angle θ from the north pole form a circle, see Figure 3. We denote the radius of this circle by $r(\theta)$, then its circumference is $2\pi r(\theta)$, and the surface area of the entire sphere is $\int_0^\pi 2\pi r(\theta)\, \mathrm{d}\theta = 4\pi$. The distribution function of the angles is the quotient of this circle circumference and sphere surface area, so it is $\frac{1}{2} r(\theta)$. From Figure 3 it is clear that $r(\theta) = \sin\theta$, so that the distribution function of the angles is $\frac{1}{2}\sin\theta$.

How does this continue in higher dimensions? In Figures 4 – 7 below we give experimental results for dimensions $4, 10, 20$ and 50, and we derive the distribution function for any dimension.

To find the exact distribution function, the above argument can be generalized. On the n-

[2] It is well known that a uniform distribution over the surface of a (hyper)sphere can be obtained by sampling normally distributed coordinates, and then scaling the so obtained vectors to norm 1.

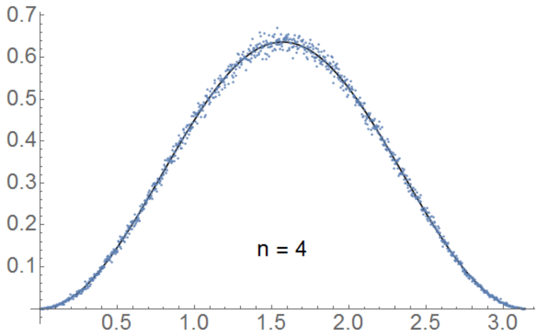

Figure 4: The same for $n = 4$.

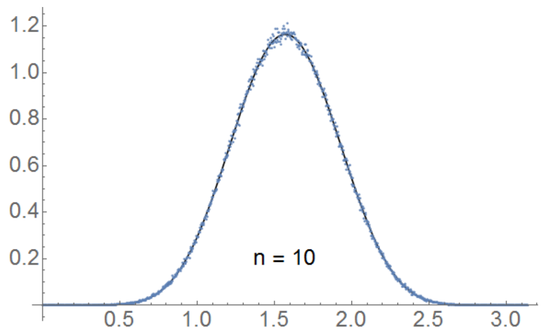

Figure 5: The same for $n = 10$.

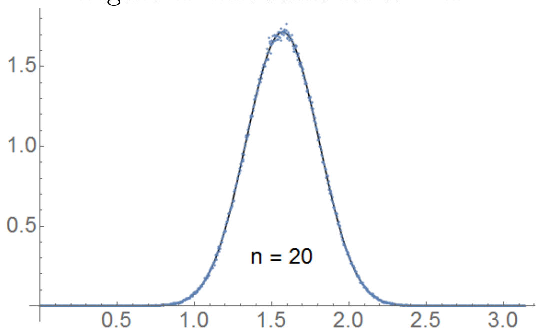

Figure 6: The same for $n = 20$.

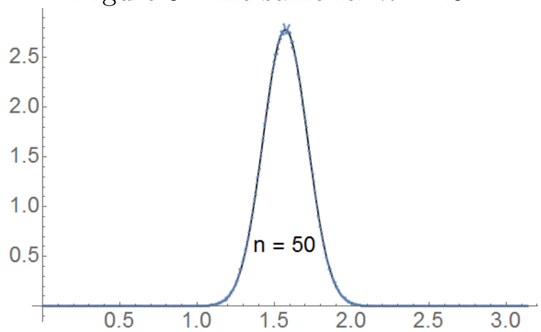

Figure 7: The same for $n = 50$.

sphere with radius 1 the points at an angle θ from the north pole form an $(n-1)$-sphere with radius $r(\theta) = \sin\theta$ (those points $(x_1, \ldots, x_n)$ satisfy $x_1^2 + \ldots + x_n^2 = 1$ and $x_n = \cos\theta$, so $\sqrt{x_1^2 + \ldots + x_{n-1}^2} = \sin\theta$). Its surface area thus is $c_n \sin^{n-2}\theta$ for some constant c_n, which can be computed by some advanced Calculus, but it is not needed. The surface area of the total n-sphere is $\int_0^\pi c_n \sin^{n-2}\theta \, \mathrm{d}\theta$, and so the distribution function for the angles is $b_n \sin^{n-2}\theta$, where $b_n = 1/\int_0^\pi \sin^{n-2}\theta \, \mathrm{d}\theta$. Elementary Calculus shows that

$$b_{2m} = \frac{2m-2}{2m-3}\frac{2m-4}{2m-5}\cdots\frac{4}{3}\frac{2}{1}\frac{1}{\pi}, \quad b_{2m-1} = \frac{2m-3}{2m-4}\frac{2m-5}{2m-6}\cdots\frac{5}{4}\frac{3}{2}\frac{1}{2}.$$

Fortunately the experiments correspond to the theory. For larger n the peak obviously gets higher (approximately $\frac{1}{\sqrt{2\pi}}\sqrt{n}$) and thinner. The takeaway here is that for large dimensions randomly selected vectors have a quite high probability of being almost orthogonal.

3 Hyperspheres and hypercubes

Some mildly surprising phenomena might be that if one samples random points inside an n-sphere, the mass will be concentrated at the sphere's radius. That is simply because the radius of an n-sphere with half the volume of the unit n-sphere is $r = 2^{-1/n}$, and this quickly approaches 1 for growing n. So if one samples random vectors in a unit n-sphere, not only will they tend to be almost orthogonal, but they will also tend to have lengths close to 1.

Here's my favourite example, of a somewhat different nature.

Start with a square in $\mathbb{R}^2$ with corners $(\pm 1, \pm 1)$. Partition it into its four quadrants. Inside each quadrant take the inscribed circle. Position a fifth circle, inscribed in the open space in between these four circles. See Figure 8.

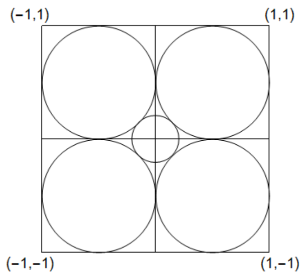

Figure 8: What is the radius of the central circle?

What is the radius r of the central circle? That's easy to compute: the line segment from $(0,0)$ to $(\pm\frac{1}{2}, \pm\frac{1}{2})$ intersects the central circle where it intersects the corresponding quadrant circle, which is centered at $(\pm\frac{1}{2}, \pm\frac{1}{2})$ and has radius $\frac{1}{2}$. So $r+\frac{1}{2}$ equals the distance from $(0,0)$ to $(\pm\frac{1}{2}, \pm\frac{1}{2})$, which according to Pythagoras is $\dfrac{1}{2}\sqrt{2}$. We find $r = \dfrac{1}{2}\left(\sqrt{2}-1\right) = 0.2071\ldots$, which is perfectly according to intuition.

Now we do the same game in dimension 3. Start with a cube in $\mathbb{R}^3$, with corners $(\pm 1, \pm 1, \pm 1)$. Partition it into its eight octants. Inside each octant take the inscribed sphere, and then take as ninth sphere the inscribed sphere in the open space between the eight spheres in the octants. See Figure 9.

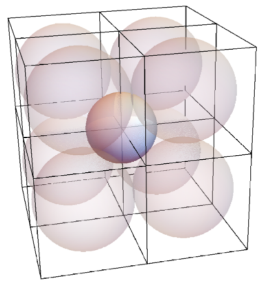

Figure 9: What is the radius of the central sphere?

What is the radius of the central sphere? The same reasoning as in dimension 2 applies: it equals the distance from $(0,0,0)$ to an octant centre, which according to Pythagoras is $\dfrac{1}{2}\sqrt{3}$, minus the radius of that octant sphere, which is again $\dfrac{1}{2}$, so the radius of the central sphere is $\dfrac{1}{2}\left(\sqrt{3}-1\right) = 0.3660\ldots$. This is a bit bigger than in dimension 2, but still quite reasonable according to intuition.

Then we do the same game in dimension n. Start with a hypercube in $\mathbb{R}^n$, with corners $(\pm 1, \ldots, \pm 1)$. Partition it into 2^n hypercubes of size 1. Inside each take the inscribed hypersphere, and then take as central hypersphere the inscribed hypersphere in the open space between the 2^n hyperspheres. The radius of the central hypersphere clearly is $\frac{1}{2}(\sqrt{n}-1)$. No doubt about that, exactly the same argument as in dimensions 2 and 3 applies.

But now things get interesting: at dimension 9 the formula gives $\frac{1}{2}\left(\sqrt{9}-1\right) = 1$, meaning that the central hypersphere inside the big hypercube touches the outer border of the hypercube in the points $(\pm 1, 0, 0, \ldots, 0)$, $(\pm 0, 1, 0, \ldots, 0)$, etc. That's funny. But it gets funnier: when the dimension gets > 9, the central hypersphere turns out to be not that small at all: its radius gets > 1, and it grows outside the big hypercube. When n grows further, the central hypersphere even grows as far out of the hypercube as one wishes. That's weird. Where is the error here? The answer is: not in the math, but in your intuition; the math is rock solid.

It gets even weirder. Compute the volume of the central (not so small anymore) hypersphere relative to the volume of the hypercube (which is 2^n). The volume of an n-sphere of radius r is $\frac{\pi^{n/2}}{\Gamma\left(\frac{1}{2}n+1\right)} r^n$, as can be found in any good Calculus textbook[3]. So the ratio of the volumes of hypersphere and hypercube is

$$\frac{\left(\frac{1}{4}\sqrt{\pi}\,(\sqrt{n}-1)\right)^n}{\Gamma\left(\frac{1}{2}n+1\right)}.$$

Here's a plot of this relative volume for dimensions up to 20.

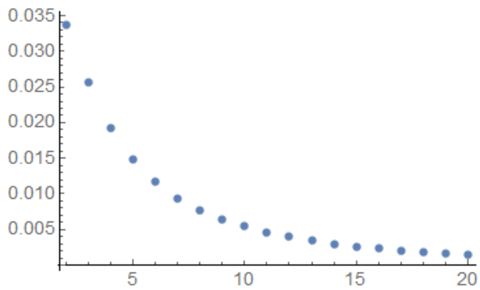

Figure 10: Volume of central hypersphere relative to hypercube up to dimension 20.

That's reassuring: the central hyperspheres get smaller in this sense. Lets extend the experiment a bit further, to dimension 200, say.

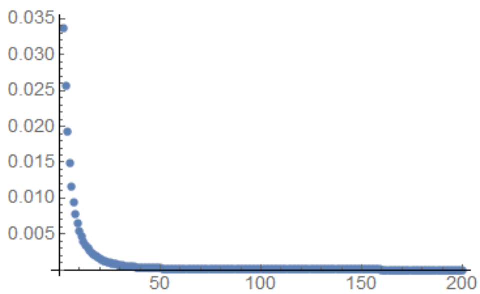

Figure 11: Volume of central hypersphere relative to hypercube up to dimension 200.

[3]For convenience we refer to Wikipedia: `https://en.wikipedia.org/wiki/Volume_of_an_n-ball`.

Ok, that's pretty good evidence, right? It *must* go to 0. We have made these pictures with Mathematica. But Mathematica is computer algebra software that can do symbolic calculations, and certainly will know how to exactly compute the limit for $n \to \infty$. So let's do a "formal" check that the limit indeed equals 0.

$$\texttt{Limit}\left[\frac{\left(\frac{1}{4}\sqrt{\pi}\left(\sqrt{n}-1\right)\right)^{n}}{\texttt{Gamma}\left[\frac{n}{2}+1\right]},\ n \to \infty\right]$$

$$\infty$$

Wow. Not 0 at all, even ∞. Again: what is wrong here? Again: not the math, not even Mathematica, but your intuition!

To see what is happening, let's first extend the picture further, from dimension 200 to 300:

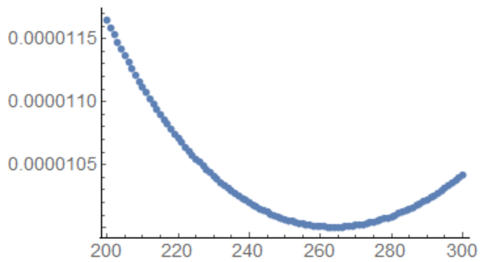

Figure 12: Volume of central hypersphere relative to hypercube in dimensions 200 to 300.

Interesting. It starts going up again, there must be a minimum somewhere. Indeed, a minimum occurs at dimension 264, with as value $\dfrac{(2\sqrt{66}-1)^{264}\pi^{132}}{2^{528}132!} = 0.0000100042800279\ldots$. If one allows non-integer n, the "real" minimum is at $n = 263.655538323\ldots$, with as value $0.0000100042412856\ldots$. Are these special numbers? Plouffe's inverter[4] gave no results.

Let's go even further with drawing these nice pictures:

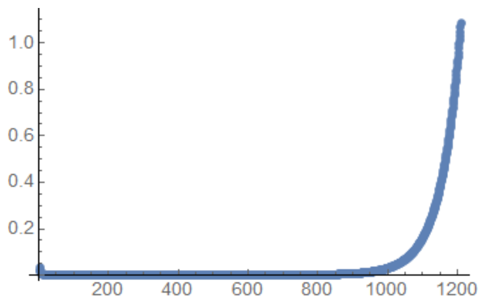

Figure 13: Volume of central hypersphere relative to hypercube up to dimension 1210.

At dimension 1206 the ratio of the two volumes even passes 1, meaning that the volume of the central hypersphere gets even larger than the volume of the hypercube. I bet you didn't see that one coming.

[4] `http://wayback.cecm.sfu.ca/projects/ISC/ISCmain.html`

Here's an "explanation": using Stirling's series for the Gamma function we get

$$\frac{\left(\frac{1}{4}\sqrt{\pi}\left(\sqrt{n}-1\right)\right)^n}{\Gamma\left(\frac{1}{2}n+1\right)} = \frac{\left(\sqrt{\frac{1}{8}\pi e}\right)^n}{e^{\sqrt{n}}\sqrt{\pi e n}}\left(1-\frac{1}{3\sqrt{n}}+O\left(\frac{1}{n}\right)\right).$$

The asymptotically dominating term is the exponential term in the numerator, because its base $\sqrt{\frac{1}{8}\pi e} = 1.033\ldots$ is above 1. But because this base is so close to 1, for small n the subexponential term $e^{\sqrt{n}}$ still kicks in strongly. And this explains the strange behaviour, with the bizarrely delayed asymptotics.

The takeaway: high dimensions maybe start to show their "natural behaviour" only well above dimension 1000.

Final funny remarks: the radius of the $n-1$-dimensional hypersphere which is the intersection of our central hypersphere and an $n-1$-dimensional face of the hypercube equals $\frac{1}{2}\sqrt{n-2\sqrt{n}-3}$. For $n \geq 15$ this gets > 1, so then this intersection hypersphere also grows out of the hypercube's face.

But still: the 2^n hyperspheres in the "corners" of the hypercube still all are inside the hypercube but outside the central hypersphere...