Vamshi Jandhyala

# *An Average Shadow*

*July 2026*

A unit cube hangs above a horizontal plane. Light arrives in parallel rays from directly overhead, so the shadow is the vertical projection of the cube. Turn the cube to a random orientation. What is the average area of the shadow?

The answer is 3/2, and it is not really a fact about cubes. It is a case of a theorem of Cauchy: for any convex body, the mean area of the shadow, averaged over orientations, is a quarter of the surface area. The cube has surface area 6.

> *Stan Wagon's Problem of the Week 1417, circulated 29 July 2026, which attributes the result to A.-L. Cauchy, 1850.*

## *Everything depends on one vector*

Fix the plane and turn the cube. Let $R$ be the rotation that carries the cube from its starting position to where it now sits, and let

$$u = R^{\mathsf{T}} k, \qquad k = (0, 0, 1),$$

so that $u$ is the vertical direction expressed in the cube's own frame. If you were sitting inside the cube, $u$ is where you would point to say "up". Its three components are the cosines of the angles between the vertical and the cube's three edge directions.

**Proposition.** The area of the shadow is $|u_x| + |u_y| + |u_z|$.

*Proof.* The argument is a double count. Take a point in the interior of the shadow and send a vertical line up through it. Because the cube is convex, that line meets the boundary in exactly two points, one where it enters and one where it leaves, unless it happens to graze an edge or a vertex. Those exceptional points form a set of area zero, so almost every interior point of the shadow is covered twice as the boundary is projected down, which gives

$$2 \cdot \mathrm{area}(\mathrm{shadow}) \;=\; \sum_{\text{faces } F} \mathrm{area}(F)\, |n_F \cdot k|,$$

where $n_F$ is the unit outward normal of the face $F$. The factor $|n_F \cdot k|$ is the ordinary rule that a flat plate of area $A$ tilted so its normal makes angle $\theta$ with the vertical throws a shadow of area $A|\cos\theta|$.

For the unit cube the six faces each have area 1, and their outward normals are the six vectors $\pm Re_1, \pm Re_2, \pm Re_3$. Now $(Re_j) \cdot k = e_j \cdot R^{\mathsf{T}} k = u_j$, so the six terms pair off and the sum is $2(|u_x| + |u_y| + |u_z|)$. Halve it. $\square$

For the cube the double count has a concrete reading. From each pair of opposite faces exactly one is turned upwards, so three faces are lit; their projections tile the shadow without overlapping, and they contribute $|u_x|$, $|u_y|$ and $|u_z|$. Three orientations are worth carrying forward:

| | | | |
|---|---|---|---|
| $u$ | $(0,0,1)$ | $(1,1,0)/\sqrt{2}$ | $(1,1,1)/\sqrt{3}$ |
| | a face level | an edge up | a diagonal upright |
| area | 1 | $\sqrt{2}$ | $\sqrt{3}$ |

The middle one looks like a throwaway, and it is not: $\sqrt{2}$ turns out to be the threshold at which the distribution changes character.

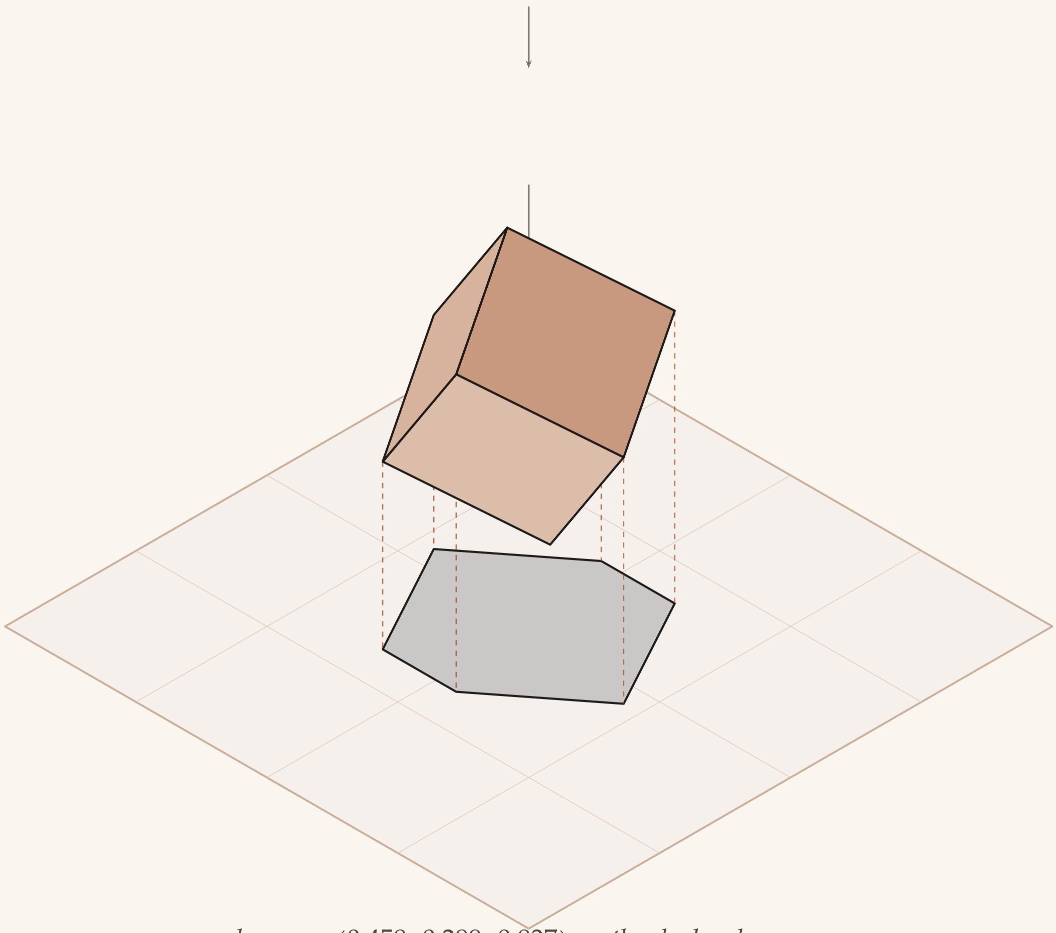


Figure 1: The shadow is the vertical projection, and its area is read off a single vector: the vertical direction as the cube sees it.

The whole problem has now collapsed. There is no geometry left, only a vector on the unit sphere and the function $u \mapsto |u_x| + |u_y| + |u_z|$. Write $A$ for it.

## *What the extremes are*

Both bounds come straight from $u_x^2 + u_y^2 + u_z^2 = 1$. Squaring the sum gives $(|u_x| + |u_y| + |u_z|)^2 = 1 + 2(|u_x u_y| + |u_y u_z| + |u_z u_x|) \geq 1$, and Cauchy-Schwarz in the other direction gives $(|u_x| + |u_y| + |u_z|)^2 \leq 3(u_x^2 + u_y^2 + u_z^2) = 3$. So

$$1 \leq A \leq \sqrt{3},$$

the lower bound holding when $u$ is a coordinate direction and the upper when the three components are equal in absolute value. In the cube those are a face lying level, whose shadow is the unit square, and a space diagonal standing upright, whose shadow is a regular hexagon of area $\sqrt{3} \approx 1.732$. In linear-algebra language $A$ is the $\ell^1$ norm of $u$, and this is the standard comparison with the $\ell^2$ norm.

For a generic orientation the shadow is a centrally symmetric hexagon. When one component of $u$ vanishes the hexagon degenerates to a rectangle, as it does for both extreme cases above and for the edge-up orientation. The area varies continuously with the orientation, though not smoothly: the absolute values put a corner wherever a component crosses zero.

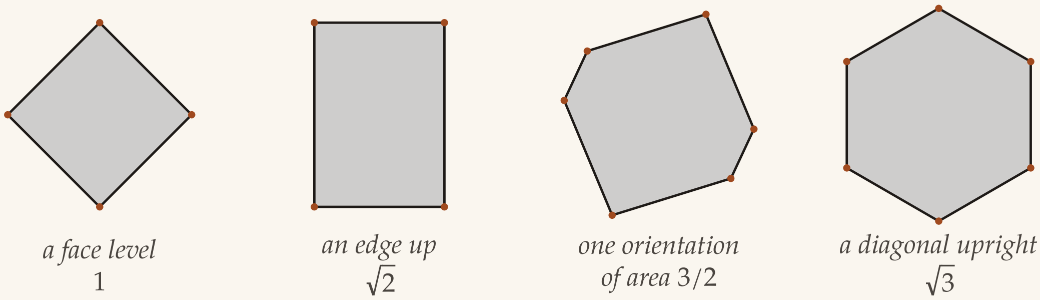


Figure 2: Four shadows seen from directly above. The largest is only about 73 per cent bigger than the smallest: a cube casts a surprisingly steady shadow. Only the third is a hexagon; the other three have a component of $u$ equal to zero, and degenerate to rectangles.

## *The mean, without evaluating an integral*

The natural move now is to average $|u_x| + |u_y| + |u_z|$ over the sphere. It is worth resisting, because there is a route that never computes an average of $|\cos\theta|$ at all, and it delivers the general theorem rather than the cube.

**Theorem** (Cauchy, 1850)**.** Let $K$ be a convex body with surface area $S$. The mean area of the vertical projection of $K$, over a uniformly random orientation, is $S/4$.

*Proof.* Go back to the double count, which used nothing about cubes. For any convex body a vertical line still meets the boundary twice above almost every point of the shadow, so for a fixed light direction

$$2 \cdot \text{area(shadow)} = \int_{\partial K} |n \cdot k| \, dA,$$

with $n$ the outward unit normal, which exists at almost every boundary point of a convex body. The identity is immediate for polytopes and for smooth bodies, and the general case follows by approximation.

Now average over a uniform orientation. Averaging is linear, so it passes through the integral and lands on $|n \cdot k|$ alone. Here is the step that does the work: $\mathbb{E}|n \cdot k|$ is the same number for every unit vector $n$. A uniform orientation is by definition invariant under rotation, so no direction is distinguishable from any other, and the average cannot depend on which normal it is looking at. Call that common value $m$. Then

$$\mathbb{E}[\,\text{area}(\text{shadow})\,] \;=\; \tfrac{1}{2}m\int_{\partial K} dA \;=\; \tfrac{1}{2}mS,$$

so the average shadow is proportional to the surface area, with one constant good for every convex body, and so far nothing has been integrated.

The constant is then fixed by any single body whose shadows are known without calculation. The sphere of radius 1 is the obvious candidate: turn it however you like and the shadow is a disc of area $\pi$, while its surface area is $4\pi$. So $\frac{1}{2}m \cdot 4\pi = \pi$, giving $m = 1/2$ and a mean shadow of $S/4$. □

The cube has $S = 6$, so its average shadow is $3/2$. The sphere did the work that an integral would otherwise have had to do.

It is worth seeing the cube computed directly as well, since the two routes are independent. Under a uniform orientation $u = R^{\mathsf{T}}k$ is a uniform point on the sphere, and its three coordinates have the same distribution, so the average shadow is $3\,\mathbb{E}|u_z|$. Archimedes settles that: his hat-box theorem says the vertical coordinate of a uniform point on the sphere is itself uniform on $[-1, 1]$, because the sphere and the circumscribed cylinder have equal areas between any two horizontal planes. Hence $\mathbb{E}|u_z| = 1/2$ and the mean is $3/2$ again.

## *The distribution, not just the mean*

The mean is the question asked, but the distribution is where the problem becomes interesting, because it has a closed form on most of its range.

Replacing $u$ by $(|u_x|, |u_y|, |u_z|)$ leaves $A$ alone and lands in the positive octant, and the eight octants have equal area, so it is enough to take $u$ uniform on the positive octant, where $A = u_x + u_y + u_z$. Write $\mathbb{P}(A \geq s)$ for the fraction of orientations whose shadow has area at least $s$.

The set $\{A \geq s\}$ is the part of the octant on the far side of the plane $u_x + u_y + u_z = s$. On the whole sphere that plane cuts off a cap whose pole is the space diagonal $(1, 1, 1)/\sqrt{3}$ and whose angular radius $\alpha$ satisfies $\cos\alpha = s/\sqrt{3}$; what we want is that cap intersected with the octant.

A full spherical cap of angular radius $\alpha$ has area $2\pi(1 - \cos\alpha)$. So as long as the cap sits entirely inside the octant, its area is $2\pi(1 - s/\sqrt{3})$, and dividing by the octant's area $\pi/2$ gives

$$\mathbb{P}(A \geq s) = 4\left(1 - \frac{s}{\sqrt{3}}\right).$$

When does the cap fit? The nearest point of the octant's boundary to the pole lies on one of the three walls $u_x = 0$, $u_y = 0$, $u_z = 0$, at angular distance $\arccos\sqrt{2/3}$ from the diagonal. The cap fits precisely when $\cos\alpha \geq \sqrt{2/3}$, that is when $s \geq \sqrt{2}$, and at $s = \sqrt{2}$ it is inscribed, touching all three walls at once.

The three points where it touches are worth naming: they are $(0, 1, 1)/\sqrt{2}$ and its two permutations, which are exactly the edge-up orientations from the table earlier. That is where the $\sqrt{2}$ in the third column came from. The threshold is not a coincidence of the algebra but the moment the level set reaches the orientations with an edge vertical.

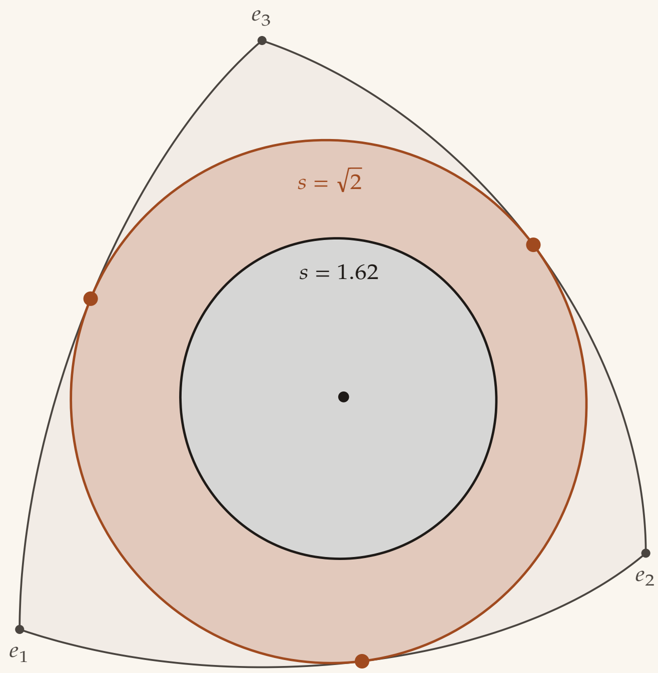


Figure 3: The superlevel sets $\{A \geq s\}$ are spherical caps about the space diagonal, marked at the centre, intersected with the octant. At $s = \sqrt{2}$ the cap is inscribed, touching each wall once at an edge-up orientation (the three marked points); above $\sqrt{2}$ it lies wholly inside, so its area is linear in $s$.

So on $[\sqrt{2}, \sqrt{3}]$ the tail is a straight line and the density is the constant $4/\sqrt{3}$. That stretch is no thin sliver of the range:

$$\mathbb{P}(A \geq \sqrt{2}) = 4 - \frac{4\sqrt{6}}{3} \approx 0.7340.$$

So $[\sqrt{2}, \sqrt{3}]$ carries about 73.4 per cent of the probability with a constant density: conditional on $A \geq \sqrt{2}$, the area is uniformly distributed on that interval. Two more values fall out of the same formula. Putting $s = 3/2$ gives $\mathbb{P}(A \geq 3/2) = 4 - 2\sqrt{3} \approx 0.5359$, so the mean sits below the median, and solving $4(1 - s/\sqrt{3}) = 1/2$ puts the median at

$$\operatorname{med}(A) = \frac{7\sqrt{3}}{8} \approx 1.5155.$$

Below $\sqrt{2}$ the octant walls clip the cap and the linear formula fails. What survives is a closed form for the density. Differentiating the tail gives

$$f_A(s) = \frac{4}{\sqrt{3}}\left[1 - \frac{3}{\pi}\arccos\left(\frac{s}{\sqrt{2(3-s^2)}}\right)\right], \qquad 1 \le s \le \sqrt{2},$$

and $f_A(s) = 4/\sqrt{3}$ on $[\sqrt{2}, \sqrt{3}]$. The arccosine is exactly the angular correction for the part of the cap the walls cut away. Integrating this density gives the tail values

| | | | | | |
|---|---|---|---|---|---|
| $s$ | 1.05 | 1.1 | 1.2 | 1.3 | $\sqrt{2}$ |
| $\mathbb{P}(A \ge s)$ | 0.9975 | 0.9897 | 0.9552 | 0.8870 | 0.7340 |

and integrating $sf_A(s)$ across the whole range returns $3/2$, as it must.

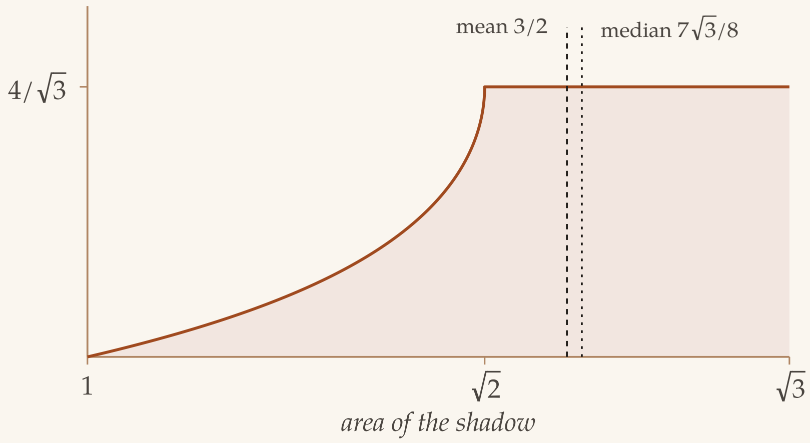


Figure 4: The exact density. It climbs from zero at 1, meets the plateau $4/\sqrt{3}$ continuously at $\sqrt{2}$ with a vertical tangent, and is constant from there to $\sqrt{3}$. The mean and the median differ by about one per cent.

The two pieces agree at $\sqrt{2}$: the arccosine vanishes there, so the rising branch arrives at exactly $4/\sqrt{3}$ and the density is continuous. What it is not is smooth. The approach has a vertical tangent, and the curve turns a corner onto the plateau, which is why sampling the density a little short of $\sqrt{2}$ is so misleading: at $s = \sqrt{2} - 0.002$ it still reads 2.107 against a limit of 2.309, and it only closes the gap in the last thousandth.

## *A trap in the sampling*

The problem statement offers a recipe for simulating a random orientation. Plainly stated, it is: choose a rotation axis uniformly on the unit sphere, then choose the rotation angle uniformly from 0 to $2\pi$. It is an inviting prescription, and it does not sample orientations uniformly. A simulation that follows it converges to 1.4583, not 1.5.

The discrepancy is easy to see. By Rodrigues' formula, rotating $k$ by an angle $\varphi$ about a unit axis $a$ gives a vector whose vertical component is

$$u_z = \cos\varphi + m^2(1 - \cos\varphi), \qquad m = a \cdot k.$$

Under the recipe, $m$ is uniform on $[-1,1]$ and $\varphi$ is independent of it, so $\mathbb{E}[\cos\varphi] = 0$ and

$$\mathbb{E}[u_z] = 0 + \mathbb{E}[m^2] \cdot 1 = \tfrac{1}{3}.$$

Under a genuinely uniform orientation $u$ is uniform on the sphere, so upward and downward tilts are equally likely and $\mathbb{E}[u_z] = 0$. A value of $1/3$ is therefore proof of bias, with no further calculation needed. The recipe leaves the cube leaning towards where it started, and since shadows are smallest near the starting position, the average shadow comes out too small.

The cause is the angle, not the axis. Uniform measure on $SO(3)$, written in axis-angle coordinates, is a uniform axis together with an angle drawn from the density $(1-\cos\varphi)/\pi$ on $[0,\pi]$, not a uniform one. Small rotations are rarer than a flat angle supposes. Keep the uniform axis, replace the flat angle by that density, and the simulation returns $3/2$ to within its standard error. Drawing a unit quaternion uniformly on the 3-sphere does the same job in fewer lines.

How wrong is the recipe? Its construction never refers to any direction but the vertical, so its law for $u$ is invariant under rotations about the vertical; the azimuth is therefore uniform and independent, and the mean shadow is $\mathbb{E}|u_z| + \frac{4}{\pi}\mathbb{E}\sqrt{1-u_z^2}$, with $u_z$ distributed as above. Both expectations are integrals of an explicit function of $m$ and $\varphi$ over $[-1,1]\times[0,2\pi)$. The second is exactly $2/3$; the first has no closed form I could find, and numerical integration gives 0.6094757. The mean shadow is therefore

$$0.6094757 + \frac{8}{3\pi} = 1.4583020, \qquad \text{against } 1.5,$$

a shortfall of 2.78 per cent.

That is not a subtle error. The standard deviation of the shadow area under the recipe is 0.1806, so the shortfall of 0.0417 is about 7.3 standard errors at a thousand samples, and only about 76 samples are needed before it reaches two. The bias could pass for noise in a few dozen draws and should be conspicuous well before a thousand.

The lesson generalises past this problem. "Uniformly at random" is a statement about a measure, and on a group the invariant measure is rarely the one that a natural-looking construction produces. Building the rotation out of two independently uniform ingredients feels symmetric, and the feeling is misleading. Cauchy's answer is exact and holds for every convex body; whether a simulation finds it depends entirely on which measure the sampler was quietly using.

*An interactive version of the cube and its shadow, where the two samplers can be compared directly, is at vamshij.com/mathematics/average_shadow.*