On the shape of the earth, Part I

John Anderson & Theodore D. Drivas
Notes for MAT 401, Spring 2026

In these notes, we will tell the story of how we have come to know the shape of the earth, to varying degrees of precision, with little bit of data but a lot of ingenuity. It is a story that begins with Aristotle and Eratosthenes, involves Kepler, Newton and continues on today.

1 Aristotle's observation

How do know the shape of the earth? If we could step away from it, and look from any angle, we would know that it's (nearly) a sphere. But if we saw it only from one angle as a circle, the same could be a disk (for instance, this is the premise of Discworld). Unfortunately, we couldn't "step away" from the earth until fairly recently, so how is it that humanity has been so sure of the shape of the earth for so long?

To do this, one needs to use another object, in this case the moon. A lunar eclipse happens when the moon falls into the earth's shadow. Aristotle knew about lunar eclipses, and knew that the earth's shadow on the surface of the moon is always a circular arc.

Figure 1: Compilation of pictures of moon during a lunar eclipse [3]

Theorem 1: Let $S \subset \mathbb{R}^3$ be a closed set. If the "shadow" of S on any plane is a perfect disk, then S is a sphere.

Remark: This result holds for any dimensions $d \geq 3$, with the appropriate replacements. It does not hold in dimension two (the Reuleaux triangle). It may be constructed as follows. Place down the vertices of an equilateral triangle. Using the respective lengths between pairs of vertices as radii of three circles, carve out the curvy triangle defined by the joint intersection. Such a shape has constant width by construction, as can easily be seen in the picture.

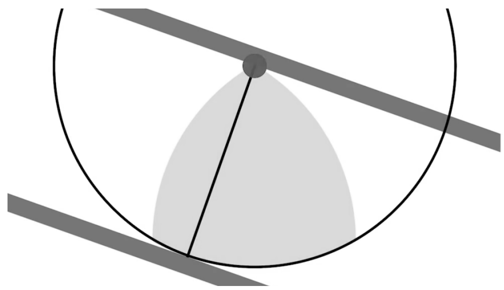

Proof of Theorem 1: First, lets define what a shadow is. Consider a light source infinitely far away from a perfectly light-absorbing solid body (our closed set S), so all rays come in parallel to a given direction. A shadow consists of all points on a plane (not intersecting S, whose normal is that given direction) that are not touched by light. We call a convex body to be ones such that lights parallel to light rays intersect the body either two or zero times, for any possible direction of light.

If a body is not convex, by its *convex closure* one means you add all these line segments to the body to make the aforementioned property for all possible directions. Clearly the shadow of a body and its convex closure are the same.

We first consider the convex closure of the set S. In this case, it follows that, given any point p on the boundary of the body, there is a plane through p such that the body lies entirely on one side of that plane.

Denote by d the maximum diameter achieved by the circular shadows. Then there are points a and b of the boundary of the body, distance d apart. Denote by c the center of the segment ab. Consider any straight line from c, and let it exit the body at a point p. Now, the distance cp can't be greater than $d/2$, for that would violate circularity of the shadow to the plane parallel to the plane abp. Indeed, if the distance was greater than $d/2$ and the shadow was still circular, it would contradict the fact that d is the largest diameter of all possible shadows.

But if distance cp were less than $d/2$, then

i) draw the plane, P, through p such that the entire body lies on one side,

ii) draw a plane P' containing the line ab, but tipped slightly – so that it passes above or below p, depending on what P is,

iii) take the shadow of the body on a plane P' parallel to P.

The resulting shadow will not be circular: It will be indented slightly near p. This contradicts our assumption. Therefor the distance cp must equal $d/2$ for any point p. Thus S must be a sphere and c must be its center.

In order to complete the proof, we now argue that any closed set S whose convex closure is a sphere must contain the sphere.

Lemma 1: Let S be a closed set whose convex closure contains the unit sphere. Then, S contains the unit sphere.

We shall study this using the interpretation of supporting hyperplanes in terms of linear functionals. Let p be a point on the unit sphere which is not contained in S (we shall derive a contradiction). Now, let P be a supporting hyperplane of the sphere at p, and let ℓ be a corresponding affine functional satisfying $\ell(p) = 0$ and $\ell(x) \geq 0$ for all $x \in B$, where B is the unit ball. Because S is a closed set not containing p, there is a $\delta > 0$ such that $\ell(x) \geq \delta$ for all $x \in S$. Now, by one of the previous exercises, any point in the convex closure of S can be realized as lying on a line segment joining two points of x. Let q be a point lying on the line segment between some such $a \in S$ and $b \in S$. Because q lies on the line segment between a and b, we can write that $q = at + b(1-t)$ for some $0 \leq t \leq 1$. Now, we can evaluate That

$$\ell(q) = \ell(at + b(1-t)) = t\ell(a) + (1-t)\ell(b) \geq t\delta + (1-t)\delta = \delta.$$

This shows that p cannot be a limit point of such points q, which contradicts that p is in the convex closure of S.

Let us press this a little further. I claim that, essentially, if you give me the shadows of a compact, convex object, I can reconstruct the whole object.

Theorem 2: The collection of shadows of a compact surface $S \subset \mathbb{R}^3$ with convex closure onto any plane determines the convex hull of S.

Argument for Theorem 2: Here is a sketch of the most direct proof which, although definitive, is perhaps not so illuminating about what the shape looks like. After giving this simple proof, we will spend some time thinking about how to understand a bit better what the shadows are telling us and how we can actually use the shadows to get precise information about the shape of the object.

We look at the shadow of the surface S on an arbitrary plane P. Because S has convex closure, this will be a convex subset of P. We can then form the cylinder over this shadow. This is constructed by considering all translates of the shadow in directions orthogonal to the plane P. We can also think of this as taking all points in the shadow and adding in lines orthogonal to P through that point. We note that the convex closure of S must be contained in this cylinder.

We can now consider all such cylinders. Because the convex closure of S is in every such cylinder, it must be in the intersection. We can in fact see that the convex closure of S must be identically equal to this intersection. Thus, the shadows uniquely determine the convex closure of S.

Exercise Fill in the details of this proof. In particular, show that the intersection of convex sets is convex, and show that every point which is not in the convex closure of S must lie outside of one of the cylinders described above. This allows one to rigorously complete the proof of the result.

Now, let's try to make this a bit more precise. We would like a way to actually calculate what the body looks like. This involves, for example, calculating the coordinates of points on the surface. Let's think about this as follows. We will assume that the set S is itself convex, and that there are no "flat faces" (that is, every plane is the supporting hyperplane of only one point). We now imagine approaching S by translating a given plane P from infinity. The plane will eventually intersect S, and it must be a supporting hyperplane at the first point where it does so. Take two orthogonal unit vectors e_1 and e_2 on the supporting hyperplane (which we shall denote also by P), and let e_3 denote the unit normal of the plane P. The vectors (e_1, e_2, e_3) generate a basis of $\mathbb{R}^3$, and they are naturally associated to a coordinate system. We denote by x^i the linear function whose gradient is given by e_i for $1 \leq i \leq 3$. Then, we note that all translates of the plane P are simply level sets of x^3.

Now, if p denotes the point on S having P as its supporting hyperplane, it will have some coordinates (x_0^1, x_0^2, x_0^3) relative to the (x^1, x^2, x^3) coordinate system described above. We can now consider an orthogonal projection onto the x^1x^3 and x^2x^3 planes. We are assuming that we know these shadows. Let p_1 be the point in the shadow on the x^1x^3 plane which has largest x^3 coordinate, and let p_2 be the point in the shadow on the x^2x^3 coordinate which has the largest x^3 coordinate. Because we know the shadows, we can assume that we know the coordinates of p_1 and p_2 within their respective planes. They must be the image of p under the projection, for they are the unique points in their respective projections with the same x^3 coordinate as p. Thus, direct computation tells us that in fact $p_1 = (x_0^1, x_0^3)$ and $p_2 = (x_0^2, x_0^3)$. Thus, we know from the coordinates on these two shadows what (x_0^1, x_0^2, x_0^3) is, meaning we know the coordinates of the point p, as desired.

Because we can recover the coordinates of points on the surface in this way, we have an algorithm for finding points on the boundary of the surface whose supporting hyperplanes are orthogonal to any given direction. It is not so clear from this how to compute properties of what the boundary looks like, but we can figure this out with a bit more thought. This tells us more than just the coordinates of different points. Indeed, it tells us what the tangent planes look like locally, so it tells us a bit about the shape of the surface.

There are a few ways to get this information, and we will start with one which is based on infinitesimal rotations. This also gives an understanding of what nearby projections are telling us. For simplicity, we consider a convex body which is cut through by the z axis. We will examine what rotations around the z axis can tell us about what the shape looks like locally. We can parameterize the boundary of the body in cylindrical coordinates adapted to the z axis, meaning we can write $r = \overline{r}(z, \theta)$. The Cartesian coordinates of the surface are then given by $(\overline{r}(z, \theta)\cos(\theta), \overline{r}(z, \theta)\sin(\theta), z)$. Now, take a point on this surface whose tangent plane is parallel to the xz plane. We can find the coordinates of this point using the projection method described above, that is, we can find the values of r, θ, and z at this point, which we denote by $r_0 = \overline{r}(\theta_0, z_0)$, θ_0, and z_0. Now, imagine rotating the figure a little bit by angle ϑ and looking at the projections onto the yz and xy planes. If we otherwise hold the z coordinate fixed (i.e., $z = z_0$), these projections let us recover $\overline{r}(\theta_0 + \vartheta, z_0)$. Because we can do this for ϑ arbitrarily small, we can figure out what $(\partial_\theta \overline{r})(\theta_0, z_0)$ is equal to. In fact, it gives us information about what all θ derivatives of $\overline{r}$ look like at this point. We can also determine what higher order z derivatives look like at this point by holding θ fixed at θ_0 (i.e., don't rotate at all) and by instead varying z. Varying both can give us mixed derivatives. This can tell us things like the curvature of the body. Of course, all of this is under the assumption that the function $\overline{r}$ parameterizing the boundary is sufficiently smooth for the derivatives to make sense.

Exercise: Let $j \geq 0$ and $k \geq 0$ be integers such that $j + k \leq 2$. Find algorithms for approximating $(\partial_z^j \partial_\theta^k \overline{r})(\theta_0, z_0)$ arbitrarily well with finitely many measurements near $\theta = \theta_0$ and $z = z_0$. As a hint, think about Taylor expansions! If you are particularly ambitious, you can find algorithms for general j and k.

Exercise: The above arguments rely on the assumption that every supporting hyperplane of the convex set intersect the convex set in only one point. Figure out where this is used, and what changes if we drop this assumption.

2 Idealizing the non ideal nature of the problem

Up until now, all of our considerations have involved an extremely idealized situation. Aristotle used the fact that the shadow of the Earth on the moon is always a portion of a disk to conclude that the Earth must be a ball. Our mathematically idealized situation asked what we can recover about a body in terms of its projections (its shadows) on all possible planes. There are all kinds of problems with this idealization in practice. Let's list a few:

1. We could never measure the shadow of the Earth on infinitely many (let alone uncountably many!) directions. More generally, we can only ever take finitely many measurements!

2. The Earth is not an exact ball, so we shouldn't expect that it's shadow is an exact disk. There should be some slight deviation. Moreover, we can't measure things to infinite precision, so even if the shadow was a perfect disk, we should only be able to confidently say that the shadow is a disk to within some error. So we have to allow for variations in the shape of the object (the Earth) and in the shape of the shadow, and we have to understand how sensitive our predictions are to such variations.

3. The moon is not an infinite plane, but rather a roughly spherical ball.

By introducing the idealization, we have thus introduced additional problems. It is somewhat amusing that these practical problems can themselves be idealized and lead to other interesting (at least to us) mathematical problems.

In this Section, we will describe some of the questions that appear, and describe some of the things that can be done. Some proofs will be provided, while others will be relegated to another document [2]. Specifically, we will deal with the following questions and their variants:

1. Can we estimate the diameter of a body in terms of its shadows?
2. Can we estimate the volume of a body in terms of its shadows?
3. Can we get *quantitative* versions of the above points? That is, if we have two bodies whose shadows are "close" in some sense, can we say that our estimates on geometric properties of the body are also close?

These questions deal with various aspects of the non ideal nature of the problem in different ways. First, these questions have interpretations that allow for only finitely many measurements. Indeed, variants of these questions could be asked assuming we have finitely many measurements on the shadows. Second, they are much more precise, asking for some kind of understanding of an explicit geometric property of the body (like the diameter or the volume). These precise geometric questions give us a specific quantity to look at. This allows us to focus on what various geometric properties of the shadows tell us about specific geometric properties of the body. Third, the desire to get *quantitative* versions of things is a way of measuring how stable predictions are to imperfections. Examples of imperfections include things like errors in our measurements. It guarantees that we can still say something even if we don't have perfect precision. Indeed, a prediction which requires a perfect measurement is useless from a practical perspective!

Turning to these kinds of questions, we start with something which is even more basic.

Question: Let S be a subset of $\mathbb{R}^3$. Suppose we know the shadows of S on finitely many planes. Is it possible to recover S?

With some thought, we can come up with some examples showing that this is not possible. In order to see this, we assume we are given finitely many projections which match the case of the sphere (so the projections are all disks).

We can take the ball which has these projections as its shadows and modify it a bit to produce another convex set without changing the shadows. This example will in fact show that the question cannot be saved by making all of the objects involved convex.

In order to construct this example, we start with the ball and we look at the finite collection of shadows at hand. Now, each shadow is bounded by a circle. We can consider the set of points on the original ball which project down to this circle. Each such circle is a great circle on the sphere (like the equator on the Earth). Now, because there are finitely many shadows, there are finitely many such circles on the sphere. If we pick any point not on one of these circles, we can make the ball "bulge" out a bit near this point to form another convex body. We can directly see that changing the ball in this way does not alter any of the finite collection of shadows we are given.

Although we cannot know exactly the shape of the body in terms of finitely many shadows, it is at least possible to measure certain properties of the bodies using the shadows. Here is one much more basic question.

Question: Let S be a closed subset of $\mathbb{R}^3$. Is it possible to determine that S is finite using only the shadows it produces on finitely many planes? If so, how many planes are needed, and must these planes satisfy any special condition?

Thinking of the cylinders generated by the shadows makes the picture much more clear. In fact, we see that the following is true.

Proposition 1: Let S be a closed subset of $\mathbb{R}^3$, and suppose that there are two non parallel planes P_1 and P_2 such that the shadow of S on P_1 and P_2 both have finite diameter. Then, the set S itself has finite diameter.

Proof. We look at the cylinder generated by the shadow on P_1. This cylinder is foliated by a bunch of lines orthogonal to P_1. Because the shadow on P_1 has finite diameter, call it d_1, we can pick a line L_1 in this cylinder, and we note that every point in S has to be of distance at most d_1 from L_1. If d_2 denotes the diameter of the shadow on P_2, we can similarly find a line L_2 such that every point of S is at distance at most d_2 from L_2. Now, because P_1 and P_2 are not parallel, the lines L_1 and L_2 (being parallel to the normal vectors) are themselves not parallel. Because of this, there exists a sufficiently large ball in $\mathbb{R}^3$ centered at the point where L_1 and L_2 are closest to each other outside of which the distance between L_1 and L_2 is larger than $2(d_1 + d_2)$. Thus, outside of this ball, it is certainly not possible for there to be any points in S. □

Note the importance of the fact that P_1 and P_2 are not parallel in the above argument. If they were parallel, we would basically just have the information of the projection of S in one direction, and we would just have one cylinder generated by a shadow. For all we know, the set S could be this entire cylinder! This means that, in the limit of the two planes becoming parallel to each other, things must degenerate. We can actually figure out how this happens by making the above argument a bit more precise and tracking how things depend on the angle between the planes P_1 and P_2. Before proceeding, we introduce a bit of language which is used in math somewhat differently than in every day life.

Language: When we say we want an *estimate* for something, we really mean an upper bound.

The first example of this will be the following proposition. We can think that this is giving us an estimate on the diameter of a body in terms of the diameters of the projections on two planes and the angle between the two planes.

Proposition 2: Let $\delta = \arccos(|n_1 \cdot n_2|) > 0$ where n_1 and n_2 denote the normal vectors of P_1 and P_2, respectively. Let d_1 and d_2 denote the diameters of the shadows on P_1 and P_2, respectively. There exists a point $p \in \mathbb{R}^n$ such that S is contained in the ball of radius δ^{-1} centered at p, denoted by $B_p(\delta^{-1})$. Thus, the diameter d of S satisfies $d \leq C\frac{\delta_1+\delta_2}{\sin(\delta)}$ for some universal C.

Before describing the proof, let's note a few things about the statement. First, δ is exactly the angle between the planes P_1 and P_2, as you can check for yourself. The fact that it is nonzero tells us that we are not in the degenerate case where the planes match. It is not surprising, given the discussion above, that the estimate on the diameter degenerates as we approach the degenerate case, that is, as $\delta \to 0$ and the two planes become parallel.

Proof. The proof will only be sketched, and it will be an exercise to fill in the details. We consider the two cylinders generated by the shadows on P_1 and P_2. The cylinders must have nonempty intersection (otherwise, they cannot arise as the cylinders generated by the shadows of an object!). We pick a point p in this intersection, and we look at the lines L_1 and L_2 for both cylinders passing through p. Using Theorem 7 from [2], we can estimate how long it takes for these two cylinders to separate from each other. This will indeed give us the desired estimate. □

Exercise: Fill in the details for this proof, and find the sharp value of C (that is, find the smallest value of C for which the inequality is true).

We can use this opportunity to introduce a bit more notation which is common in math.

Notation: Fill in the details for this proof, and find the sharp value of C (that is, find the smallest value of C for which the inequality is true).

Thus, our estimate on the diameter d could be written as $d \lesssim \frac{d_1+d_1}{\sin(\delta)}$. Another related piece of language is *sharp*.

Language: When we say we say that something is *sharp* we mean that it cannot be improved upon with the given information.

As a simple example, we can estimate the volume V of a set given its diameter d by saying that $V \leq 100d^3$. This is certainly true, but it is far from sharp. Indeed, we can actually say that $V \leq \frac{\pi}{6}d^3$, and the case of a ball shows that this inequality cannot be improved upon. Thus, the constant $\frac{\pi}{6}$ is *sharp* for this estimate.

Show that the volume V of a set of diameter d satisfies $V \leq \frac{\pi}{6}d^3$, and show that this is sharp (i.e., it cannot be improved upon).

Now, returning to the setting of the proposition estimating the diameter of a set in terms of the diameters of its projections, we can also get a lower bound.

Exercise: It is possible to get a straightforward lower bound for the diameter in terms of d_1 and d_2. Formulate and prove the sharp lower bound.

We just talked a bit about volumes compared with diameters when describing what it means for something to be sharp. Let's think about how we can estimate the volume of a body in terms of geometric properties of its shadows.

Exercise: Using the above discussion and the above exercises, show that, in the setting of Proposition 1, the volume V of the body in question satisfies $V \lesssim \frac{(d_1+d_2)^3}{\sin^3(\delta)}$.

This shows us that we can estimate the volume of the body in terms of the diameters of its shadows. Let's start to think about how good our estimate is, that is, how sharp it is. This could mean finding the optimal constant in the $\lesssim$ above, like we described before, but we could also mean "sharp" in terms of the dependence on the parameters δ, d_1, and d_2. There are then several limiting cases to consider. Now, by symmetry, we can assume that $0 < \delta \leq \pi/2$, meaning that the only interesting limiting value for δ is 0. This corresponds to the degenerate case of the two planes being parallel. Other interesting cases are when d_1/d_2 becomes very large or very small (without loss of generality, we can assume that $d_1 \leq d_2$, so we only have to consider the case of it being very small). Being a bit more precise, given our body, we can rescale space by d_1 in order to normalize $d_1 = 1$. This rescaling will distort volumes by a predictable amount (it will change them by d_1^3!), and it will not distort angles. Thus, we are really interested in considering the limiting behavior as $\delta \to 0$ and as $d_2 \to \infty$ when $d_1 = 1$ and $d_2 \geq 1$. Indeed, if we consider the example of the unit ball and where the two planes are at right angles, we see that the estimate is sharp in terms of these parameters whenever δ is large and $d_1 \approx d_2$.

This leaves situations where δ is very small and where d_2 is very large. It turns out that the above estimate is far from sharp in both of these cases!

Proposition 3: In the same setting as Proposition 2, we can bound the volume V of the body in question by $V \lesssim \delta^{-1} d_1 d_2 \min(d_1, d_2)$.

Proof. Let's begin by comparing this estimate with the one above. First, note that they agree when $d_1 \approx d_2$ and δ is not close to 0. However, when δ is close to 0, the first estimate is like δ^{-3} (as can be seen by Taylor expanding $\sin(\delta)$) while the second is like δ^{-1}. Of course δ^{-1} is much better in this case! In addition, note that $d_1 d_2 \min(d_1, d_2)$ is much smaller than $(d_1 + d_2)^3$ when d_2 is much larger than d_1. Thus, we can see that this improves on both fronts.

The proof of this ends up being rather simple. We begin by noting that the shadow of the body on P_1 and P_2 has area bounded by Cd_1^2 and Cd_2^2, respectively. Without loss of generality, we can assume that $d_1 \leq d_2$, so we now consider the cylinder generated by the shadow on P_1. By the diameter bound $d \lesssim \delta^{-1}(d_1 + d_2)$ coming from Proposition 2, we note that the entire body must in fact be contained in a cylinder of height d! This means that the total volume of the body is bounded by the height of the cylinder d multiplied by the area of the cross section, which is bounded by Cd_1^2. Thus, we altogether have that

$$V \lesssim \delta^{-1} d_1^2 (d_1 + d_2) \lesssim \delta^{-1} d_1^2 d_2,$$

where we have used the fact that we are assuming that $d_1 \leq d_2$. If instead $d_2 \leq d_1$, we get the same inequality with d_1 and d_2 interchanged. This establishes the desired result. □

Let's turn to the question about understanding the volume of the body in terms of the areas of its shadows.

Exercise: Let A_1 and A_2 be positive numbers, and let P_1 and P_2 be two coordinate planes. Let $M > 0$ be an arbitrary positive number. Give an example of a body whose shadow on P_i has area A_i, but whose volume is at least as large as M.

This exercise shows that we cannot hope to bound the volume in terms of the areas of the shadows which are cast on two planes. Indeed, in the above exercise, we can take $A_1 = A_2 = 1$ and make M as large as we want. This shows the existence of bodies whose volume is arbitrarily large, but whose shadows on two planes are always not too big in terms of area.

However, knowing the projected area on three planes ends up being enough. This is related to the Loomis–Whitney inequality, which we now state and whose proof can be found in [2].

Theorem 4: Let S be a measurable subset of $\mathbb{R}^3$, and let $\Pi_x(S)$, $\Pi_y(S)$, and $\Pi_z(S)$ denote the projections of S onto the coordinate hyperplanes orthogonal to x, y, and z, respectively. Then, we have that

$$|S| \leq |\Pi_x(S)|^{1/2} |\Pi_y(S)|^{1/2} |\Pi_z(S)|^{1/2}.$$

We can check that the Loomis–Whitney inequality is saturated by rectangular prisms.

Exercise: Let $R \subset \mathbb{R}^3$ be a rectangular prism, that is, let R be the Cartesian product of intervals. Show that R saturates the Loomis–Whitney inequality in the sense that the inequality becomes equality.

We expect the earth to be roughly round, so it is useful to think about the extent to which the Loomis–Whitney inequality is sharp for balls.

Exercise: Let $B_r(x) \subset \mathbb{R}^3$ be the ball of radius r centered at x. Show that the bound from the Loomis–Whitney inequality is not sharp for a ball, and compute how far off it is.

If we think about this, the Loomis–Whitney inequality is adapted to projections on three planes. This is well adapted to things like rectangular prisms because they are characterized by three flat faces, meaning there are three special planes. However, this is not well adapted to something like a ball, which treats every plane the same. Fortunately, there is an "averaged" version of Loomis–Whitney. Finding this is a two step process. We first write the surface area of the body in terms of the projected area averaged over all possible directions. We hope that this is well adapted to the ball because it treats all directions as being the same. This expression is called the Cauchy-Crofton formula (the proof can also be found in [2]).

Theorem 5: Let $B \subset \mathbb{R}^d$ be a convex body bounded by surface S. Then

$$\frac{\mathrm{Area}_{d-1}(S)}{\sigma_{d-1}} = \frac{1}{\omega_{d-1}} \fint_{\mathbb{S}^{d-1}} \mathrm{Area}_{d-1}(B|u^{\perp})\mathrm{d}u \tag{1}$$

where $\omega_d = \mathrm{Vol}_d(B_1(0))$ and $\sigma_d := \mathrm{Area}_d(\mathbb{S}^d)$, and $\mathrm{Area}_{d-1}(B|u^{\perp})$ is the area of the projection of B onto the plane orthogonal to $u \in \mathbb{S}^{d-1}$.

After expressing the surface area in terms of this average projected area, we can then use the celebrated isoperimetric inequality to bound the volume of the body in terms of its surface area (and, thus, in terms of its average projected area). The isoperimetric inequality is stated below, and proved in [2].

Theorem 6: Let $B \subset \mathbb{R}^d$ be a simply connected body bounded by a continuously differentiable surface S. Then

$$\frac{\text{Area}_{d-1}(S)}{\sigma_{d-1}} \geq \left(\frac{\text{Vol}_d(B)}{\omega_d}\right)^{\frac{d-1}{d}} \tag{2}$$

where $\omega_d = \text{Vol}_d(B_1(0))$ is the volume of the unit ball. Moreover, for a fixed surface area, then the unique body that maximizes volume is the sphere.

Combining Theorem 5 and Theorem 6 directly gives us an estimate on the volume of a body in terms of its averaged projected area. We will call this an averaged Loomis–Whitney inequality.

Theorem 7: Let $B \subset \mathbb{R}^d$ be a convex body. We have that

$$\frac{\text{Vol}_d(B)}{\omega_d} \leq \left(\frac{1}{\omega_{d-1}} \fint_{\mathbb{S}^{d-1}} \text{Area}_{d-1}(B|u^\perp)\mathrm{d}u\right)^{\frac{d}{d-1}}. \tag{3}$$

To arrive at the above form using that $\sigma_{d-1} = d\omega_d$. We can check that this inequality is saturated by a ball.

Exercise: Show that the bound from the averaged Loomis–Whitney inequality is saturated by the ball.

It is not, however, saturated by a rectangular prism.

Exercise: Let Q be the unit cube in $\mathbb{R}^3$. Show that the bound from the averaged Loomis–Whitney inequality is not saturated by the cube, and compute how far off it is.

This second exercise is a bit harder, and we recommend looking at [2]. This is also the topic of [1, Chapter 33] in Arnold's book!

3 Eratosthenes's calculation

"Eratosthenes peered into a well here (Alexandria) at noon and came up with the diameter and circumference of our planet! The summer solstice sun and a trip to Syene was all it took..........."

– papyrus from 230 B.C.

Around 240–230 BCE Eratosthenes measured the circumference of the earth as **250,000 stadia**. This translates to 46,250 km. That's 13% higher than the true distance, which is about 40,008 km.

Here is how he did it. Eratosthenes, a Hellenic mathematician living in Alexandria Egypt, knew that in a place called Syene there was one day of the year (the summer solstice) when you could look down a well at midday and see the full sun reflected in the water. In Alexandria, on the summer solstice, Eratosthenes peered down a well and did not see the sun! Instead of disbelieving the news from Syene, Eratosthenes reasoned his way to a measurement of the radius of the earth. He know from Aristotle that the earth was round. Therefore, he realized that, on that day, the sun was not actually vertically above Alexandria.[1]. Note that, if you draw a line between the center of the earth and the sun, there will always be one location of the earth experiencing the "Syene" phenomenon, where the sun is directly overhead. The special thing

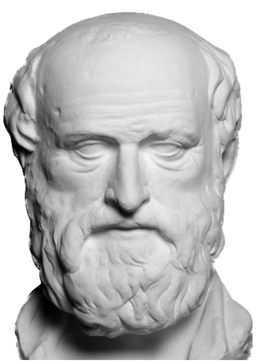

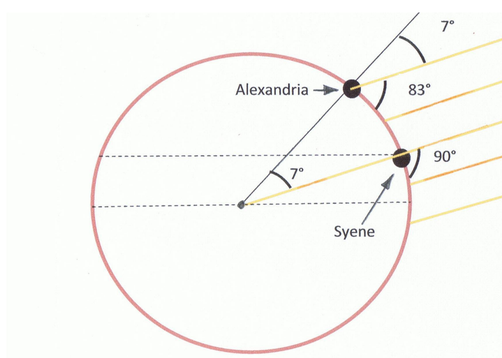

about the solstice in particular is that the axis of earths rotation tilts directly towards the sun. As such, there is an entire line of latitude which experiences the "Syene" phenomenon at noon their time (this line is called the Tropic of Cancer, Syene is on it). Alexandria, on the other hand, is slightly north of the Tropic of Cancer so on the solstice, so the direction he perceives as "straight

[1]We will think of the sun as being so far from the earth all of its rays of light are effectively parallel.

up" is not pointed towards the sun, rather making some angle to its rays. Eratosthenes had a gnomon (or portable sundial), and using this measure that the sun was 7° off from the vertical (maybe not, see [5]). He also knew that, at this exact moment in Syene, the sun was exactly overhead. Thus, he deduced that the arc along earth between Alexandria and Syene was about 7°. Then [7]

$$\frac{\text{dist(Alexandria,Syene)}}{\text{Circumference of the earth}} = \frac{7^\circ}{360^\circ}. \tag{4}$$

Eratosthenes knew that dist(Alexandria, Syene) as about 5000 *stadia* [6] (a *stadion* is a unit of length based on the size of a stadium, about 185m [4]). Thus, he predicted the circumference of earth to be $250,000$ stadia (about 46,250 km). The radius is computed as $39789 \approx 250,000/2\pi$ stadia, or about 7361 km (current measurements estimate $6,373$ km).

4 Distance between the earth and the moon

The longest duration of a lunar eclipse is about 3.5 hours. From this we can get the distance of the moon to the earth through

$$\frac{2\pi\text{dist(moon,earth)}}{2\text{(radius of the earth)}} = \frac{\text{length of a lunar month}}{\text{length of a lunar eclipse}} \approx \frac{28\text{days}}{3.5\text{hours}} \approx 61.$$

This computation was done first by Aristarchus of Samos. The true orbit of the moon (which is elliptical) varies between 58 and 62 earth radii, so the prediction is quite excellent!

How about the size of the moon? The Ancient Greeks devised the following method to estimate it. Given the distance between the earth and moon, they deduced it by:

$$\frac{2\text{(radius of the moon)}}{2\pi\text{dist(moon,earth)}} = \frac{\text{length of rise of full moon}}{\text{apparent time for moon to cycle the earth}} \approx \frac{2\text{minutes}}{24\text{hours}} \approx 61.$$

This works because when you see the moon rise, it is really your line of sight which is sweeping the moon (a distance of twice its radius) as the earth rotates. Crucially, for the apparent time for the moon to cycle around the earth, we note that this is caused by earths rotation, not the moon actually orbiting the earth (which takes roughly 28 days).

References

[1] Arnol'd, V I., Mathematical understanding of nature: essays on amazing physical phenomena and their understanding by mathematicians. American Math. Soc., 2014.

[2] John Anderson and Theodore D. Drivas, Geometric inequalities, unpublished notes (2026)
https://www.math.stonybrook.edu/~tdrivas/Courses/MAT401/geometric_inequalities.pdf

[3] Terence Tao on the cosmic distance ladder, 3Blue1Brown, https://www.youtube.com/watch?v=YdOXS_9_P4U&t=1209s

[4] Peter Gainsford, How Eratosthenes measured the earth, Part 2 https://kiwihellenist.blogspot.com/2023/06/eratosthenes-2a.html

[5] Peter Gainsford, How Eratosthenes measured the earth, Part 4 https://kiwihellenist.blogspot.com/2023/07/eratosthenes-4.html
https://kiwihellenist.blogspot.com/2023/07/eratosthenes-4.html

[6] Eratosthenes, Geography fr. 34 Roller (Strabo 2.5.7)

[7] Eratosthenes, Measurement M6 Roller (Kleomedes 1.7, 100,19–23 Ziegler)

[8] Chandrasekhar, S. 1969, Ellipsoidal Figures of Equilibrium, (New Haven: Yale University Press; reprinted edition 1987, New York: Dover Publication

John Anderson
Department of Mathematics
Stony Brook University, Stony Brook NY 11790
john.r.anderson@stonybrook.edu

Theodore D. Drivas
Department of Mathematics
Stony Brook University, Stony Brook NY 11790
tdrivas@math.stonybrook.edu