

What is the Roundest Country?

GONZALO CIRUELOS

My home country, Argentina, is long and narrow, measuring 2,268 miles north to south and 889 miles east to west at its widest (our neighbor Chile is even longer and narrower). To our east is Uruguay, which is much more circular. This made me wonder: What country is the most round?

But what do we mean when we say *round*? Roundness can be defined in many ways.

The circle is the shape that maximizes the area for a given perimeter. This fact hints at a definition for roundness: The perimeter of a circle with radius r is $2\pi r$ and its area is πr^2 . Hence, for any circle, the ratio between the square of its perimeter and its area is 4π . Because the circle maximizes the area, any other figure will have a larger ratio—for example, the value of that ratio for a square is 16.

This measure determines the roundness of a figure by how close its ratio is to 4π . The closer it is, the rounder it is. Unfortunately, this definition has a fatal flaw: Countries generally have fractal-like perimeters (also known as borders), so they tend to be much longer than they seem to be. This phenomenon is known as the *coastline paradox*. This would make the ratio too large, effectively rendering this metric useless.

Defining Roundness

Our definition for roundness focuses on how close our shape is to an actual circle. Informally, we look at all the possible circles that intersect the figure and take the one it overlaps with the most. Then we measure how much it overlaps: The more it overlaps, the rounder the figure is.


To compute the roundness value of a country, we represent it as a plane region $C \subset \mathbb{R}^2$ and define

$$\text{roundness}(C) = \frac{\max_x \text{area}(C \cap D_r(x))}{\text{area}(C)}$$

where $D_r(x)$ is the disk with center x and radius r . That is, we look at possible circles ($D_r(x)$), use area ($C \cap D_r(x)$) to measure the overlap with the country, and then normalize the measure so it doesn't depend on the size of the country or circle.

This measure enjoys two nice properties. First, for all C , $0 \leq \text{roundness}(C) \leq 1$. This means that we can estimate how round a country is just by looking at $\text{roundness}(C)$. If it is close to 0, it isn't round at all. If it is close to 1, it is very round. Second, $\text{roundness}(C) = 1$ if and only if C is a circle.

Figure 1. Sierra Leone, the roundest country, has an impressive roundness of 0.934.



This property guarantees that we can achieve maximum roundness only with an actual circle.

Gradient Descent

The question now is how to compute the roundness of a given a country. To describe the algorithm, we will assume that countries are polygons given by a list of the vertices in \mathbb{R}^2 . This is usually how countries are represented in online datasets.

We must find the circle that maximizes the formula. The problem is that we need to try all possible circles—an impossible task. This issue arises often in optimization problems: We need to maximize a function over an infinite domain.

Our solution is to use an algorithm called *gradient descent*. The idea is simple: If we want to find the maximum of a function, we start with a guess and then make small steps in the domain in the direction of maximum increase (the idea for minimizing is the same, but we will concentrate in maximization for simplicity).

Figure 2. Zimbabwe, on the left hand side, is the second roundest country, with a roundness of 0.915. Poland is third, with a coefficient of 0.907.




Figure 3. The continental United States is somewhat round; it has a roundness of 0.773. If we include Alaska, Hawaii, and all the overseas territories, its roundness drops to 0.693. Japan, on the other hand, is not very round; its ratio is 0.416.



More formally, gradient descent is an iterative method for maximizing a function f , in which we start with an initial guess x_0 , and make progressively better guesses x_1, x_2 , and so on. To find a better guess, we define a finite neighborhood for each point, call it $N(x_i)$, and take x_{i+1} to be the best possible in the neighborhood; that is, it is the x -value in $N(x_i)$ that maximizes $f(x)$. We continue until we are satisfied with the last x_i or when we don't improve, which happens when $f(x_{i+1}) = f(x_i)$.

The gradient descent algorithm has a known problem: It gets stuck in local maxima. If it finds a local maximum, then no point in the neighborhood will be better than the current one, which ends the search. The interested reader can look up solutions to this problem, usually called *metaheuristics*.

Another important piece of the puzzle is choosing a smart initial guess, x_0 . This will often depend on the specific function we are trying to optimize.

The Algorithm

Now that we know what gradient descent is, we apply it to find the circle that best overlaps the given country. We first have to define the *neighborhood* of a circle, the circles we will look at to make our guess better. The neighborhood of a circle will be the circles we get by moving the circle a little up, down, left, or right, and by making it a little bigger or smaller.

We used both the incircle and circumcircle of the given country for the initial guesses. The incircle of a polygon P is the largest circle we can fit inside P , and its circumcircle is the smallest circle that contains P . There are efficient algorithms for computing those circles.

We ran gradient descent two times, one for each initial guess, and kept the best result of both runs. The result of applying this algorithm to some countries can be seen in the figures in the article.

For the complete list of countries, more detailed discussion, and references see gciruelos.com/what-is-the-roundest-country.html. In this article we considered countries larger than 50 km² ■

Gonzalo Ciruelos is a computer science and mathematics student from the University of Buenos Aires, in Argentina. He blogs at gciruelos.com.

10.1080/10724117.2018.1524089