The Metric Spaces Around Us

Aryan Raj Dhawan

April 17 2025

Introduction

Imagine you're trying to describe how far apart two things are—whether it's cities on a map, points on a graph, or even abstract concepts like how different two files are on your computer. In mathematics, there's a structure that helps formalize this idea of "distance," and it's called a **metric space**.

This paper offers an intuitive introduction to metric spaces: and will explore what they are, the rules they follow, and why they're useful. We'll also explore a few examples to show how these structures help us understand space, similarity, and more.

What Is a Metric Space?

At its core, a **metric space** is just a way of talking about distance between different objects. More precisely, it's a set of objects (let's call this S), along with a rule—called a *metric*—that tells us how far apart any two objects in the set are. A **metric** is simply a function d that takes in two objects from the set S and returns a number, which represents this notion of "distance" between the two objects.

To qualify as a true metric, this function d_s has to follow three important criteria. Note that these criteria must hold true for any elements A and B picked from the set S.

1. **Non-negativity:** The distance between two objects can't be negative. Also, the only way two things can be distance zero apart is if they are actually the same object. For all possible objects A, B in S:

$$d(A, B) \geq 0$$

$$d(A, B) = 0 \text{ if and only if } A = B$$

2. **Symmetry:** For any objects A, B in S, the distance from A to B should be the same as the distance from B to A.

$$d(A, B) = d(B, A)$$

3. **Triangle inequality:** Formally, given any 3 objects A, B, C in S, the direct distance from A to B should be no greater than the sum of the distances from A to C and then from C to B.

$$d(A, B) \leq d(A, C) + d(C, B)$$

   Intuitively, this means that if you're going from point A to point B, taking a detour through point C shouldn't be shorter than going directly. In other words, going straight is always the shortest route.

If you have a set and a rule/function that follows all three of these guidelines, then you have a metric space.

The Euclidean Plane

One of the most familiar examples of a metric space is the flat, two-dimensional world of high school geometry—the coordinate plane, denoted $\mathbb{R}^2$. This is merely the set of all pairs of real numbers!

Imagine two points in R^2: $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$. You might recall that the distance between them is given by the Pythagorean Theorem:

$$d(P_1, P_2) = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}.$$

This formula defines what's called the **Euclidean distance**. And it turns out, this distance rule follows all three of our metric guidelines:

- *Non-negativity:* Squared terms are always zero or positive, so the square root must be zero or positive. The only way to get a zero distance is if $x_1 = x_2$ and $y_1 = y_2$—that is, the two points are the same.
- *Symmetry:* Swapping the points doesn't change the distance, because $(x_2 - x_1)^2 = (x_1 - x_2)^2$ and $(y_2 - y_1)^2 = (y_1 - y_2)^2$
- *Triangle inequality:* This one is trickier to prove, but it turns out the Euclidean distance always satisfies it. This is why "the shortest distance between two points is a straight line."

So the pair $(\mathbb{R}^2, d)$—all 2D points, with this distance formula d—is a metric space.

A Real-World Example

Let's move beyond the physical and into something more abstract.

Consider a music app trying to recommend songs. Suppose we represent each song by a list of features—lyrics, instrumentals, rhythm, etc. Each song is now a point in a multi-dimensional space, where each coordinate represents one of these features. Let's say we have a function that takes in a song and converts it to a vector (basically a structure with many elements), which represents each of these features.

Take, for instance, the song "All Too Well" by Taylor Swift. Suppose it rates 0.95 out of 1 for lyrics, 0.7 out of 1 for instrumentals, and 0.65 out of 1 on rhythm. We can define a function that takes in a song name as an input and converts it to a vector, so that we can carry out computations with it.

$$f(\text{All Too Well}) = \begin{bmatrix} 0.95 \\ 0.7 \\ 0.65 \end{bmatrix}$$

To measure how "similar" two songs are, this app could use a version of the Euclidean distance—just like we did for points on the plane. So, the set we're considering is S, a list of songs (each converted to a vector using this function f). The distance function d, which gives the "distance" between any 2 songs (both of which are in vector format) is given by:

$$d(S_1, S_2) = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2 - z_1)^2}$$

where x, y, z refer to the vector's 1st, 2nd, and 3rd components respectively.

The closer two songs are in this space, the more alike they are. This turns the app's database into a metric space, where songs are "closer" if they sound more similar. For illustration, suppose we have another Taylor Swift song "The Great War", whose vector representation is given by the following.

$$f(\text{The Great War}) = \begin{bmatrix} 0.82 \\ 0.65 \\ 0.70 \end{bmatrix}$$

Then,

$$d(\text{All Too Well , The Great War}) = \sqrt{(0.82 - 0.95)^2 + (0.65 - 0.7)^2 + (0.7 - 0.65)^2}$$
$$= 0.148$$

This is a powerful idea, as we can use it as the basis for recommending songs to users that are "similar" to the songs they're listening to! For instance, if a user listens to "All Too Well" frequently, the app could recommend songs whose distance from $f(\text{All Too Well})$ is less than 0.5. Thus, metric spaces are not just about geometry—they also underpin recommendation systems, search engines, and even spellcheck algorithms (similarity between words instead of songs)!

Other Kinds of Metrics

Not all metric spaces use the Euclidean distance. Different situations call for different ways to measure distance. For instance:

- **Discrete Metric:** This is particularly used in identity checks, specifically when the only important piece of information we need is whether or not 2 objects or the same.

  It is one of the simplest distance functions. For any objects A, B in S, where S is any set of objects:

  $$d(A,B) = \begin{cases} 0 & \text{if } A = B \\ 1 & \text{if } A \neq B \end{cases}$$

  However, we must still verify that the 3 axioms of a metric space are being satisfied by such a distance function.

  1. *Non-negativity*: Clearly, $d(A, B)$ is either 0 or 1 (by definition), so it's always non-negative. Further, $d(A, B) = 0$ only if $A = B$, by the very definition of the metric.
  2. *Symmetric*:
     (a) *Case 1:* If $A = B$, then $B = A$. Thus, we have $d(A, B) = 0 = d(B, A)$, in such a case.
     (b) *Case 2:* If $A \neq B$, then $B \neq A$. Thus, we have $d(A, B) = 1 = d(B, A)$.

     In either scenario, $d(A, B) = d(B, A)$
  3. *Triangle Inequality*: Again, a case by case approach seems suitable. Suppose we have A, B, C as arbitrary objects in set S,
     (a) *Case 1:* If $A = B$, then $d(A, B) = 0 = 0+0 \leq d(A, C)+d(C, B)$, regardless of the object C we choose, since the discrete metric is non-negative.
     (b) *Case 2:* If $A \neq B$, then $d(A, B) = 1$. For any object C, if we consider $d(A, C) + d(C, B)$, we know both terms are either 0 or 1 (by definition of the metric). Thus, the only the triangle inequality would not hold in such a case is if both terms are 0, i.e $d(A, C) = 0$ and $d(C, B) = 0$ (in all other cases the sum of the two terms would be at least 1, which would be greater than equal to $d(A, B) = 1$). However, this would mean $A = C$ and $C = B$ (since both distances are 0), which effectively means $A = B$ using the chain of equalities. However, this is a contradiction as $A \neq B$ in this particular case. Hence, the triangle inequality holds in case 2 as well.

     In either case, $d(A, B) \leq d(A, C) + d(C, B)$

  Therefore, this is indeed a metric space. For instance, S could be the set of all stored user logins (in a database) for a particular app, and the discrete metric could be used to check whether the user login matches an existing object in the database - essentially an identity check!.

- **Manhattan distance:** Instead of measuring the straight-line distance, you sum up the horizontal and vertical distances. Precisely, for 2 points $P_1(x_1, y_1)$ and $P_2(x_2, y_2)$ in the 2D plane (R^2),

  $$d(P_1, P_2) = |x_2 - x_1| + |y_2 - y_1|$$

  This is often used in grid-like city layouts, such as in Manhattan! For example, the Manhattan distance from cell 1 to 9 in the grid below is $d(1, 9) = 2 + 2 = 4$.

1	2	3
4	5	6
7	8	9

  We can verify that this is indeed a valid metric.

  1. *Non-negativity*: The metric is defined as the sum of two absolute value terms. Both of these terms are non-negative and hence so is their sum. Further, their sum (the distance) is 0 only if both terms are 0, i.e $x_1 = x_2$ and $y_1 = y_2$, meaning both points are the same.
  2. *Symmetric*: Again, swapping the points doesn't change the distance under this metric, because $|x_2 - x_1| = |x_1 - x_2|$ and $|y_2 - y_1| = |y_1 - y_2|$
  3. *Triangle inequality*: The easiest way to view this is algebraically. We use the fact that for any real numbers x, y, $|x+y| \leq |x|+|y|$, an important fact called the **triangle inequality (for absolute values)**. We now consider arbitrary points $A = (x_1, y_1), B = (x_1, y_1), C = (x_3, y_3)$ in R^2.

     $$\begin{aligned} d(A, B) &= |x_2 - x_1| + |y_2 - y_1| \\ &= |(x_2 - x_3) + (x_3 - x_1)| + |(y_2 - y_3) + (y_3 - y_1)| \\ &\leq |(x_2 - x_3)| + |(x_3 - x_1)| + |(y_2 - y_3)| + |(y_3 - y_1)| \\ &= (|(x_3 - x_1)| + |(y_3 - y_1)| + (|(x_2 - x_3)| + |(y_2 - y_3)|)) \\ &= d(A, C) + d(A, B) \end{aligned}$$

     Note that in the steps above we're simply adding $0 = x_3 - x_3$ to both the absolute value terms (which means the equation still holds true), after which we apply the triangle inequality and rearrange the resulting terms.

  One might ponder the usefulness of such a metric - which is used extensively in AI applications! Suppose an AI agent starts at cell 1 and must determine the sequence of actions needed to reach cell 9 (the goal). At each step, it could use the manhattan distance as a **heuristic**, which measures the estimated distance from the goal (Russell and Norvig). For instance, $d(1, 9) = 2 + 2 = 4$, whereas $d(2, 9) = 1 + 2 = 3 < 4$. Thus, the agent will move to cell 2 (a more favorable position), as it's closer to the goal (using the aforementioned metric).

- **Edit distance:** When comparing words, like "cat" and "cut," you can define distance d as the number of letter changes needed to turn one into the other. This metric is widely used in spellcheckers and DNA sequence analysis. Here, the set S is simply the set of all words in our dictionary.

  For instance, $d(\text{cat}, \text{cut}) = 1$. If a user incorrectly writes the word "cur", the spellchecker could recommend all words which have a distance of 1 from "cur" under this metric - such as "cut", "car", etc.

  We can also verify that this satisfies the metric space axioms.

  1. *Non-negativity*: Clearly one can't have a negative number of letter changes to change one word to another, thus the metric must be non-negative. Further, distance d between 2 words w_1, w_2 is 0 only if there are 0 letter changes between them, meaning they're the same word.
  2. *Symmetric*: Edits can be reversed, and thus the number of changes from word w_1 to w_2 is the same as that from w_2 to w_1.
  3. *Triangle inequality*: Let words w_1, w_2, w_3 be in S. If the edit distance from w_1 to w_2 is A, and the distance from w_2 to w_3 is B, then one possible way to transform w_1 into w_3 is by first converting w_1 to w_2, and then w_2 to w_3. This combined sequence involves $A + B$ edits. While this may not always be the shortest possible path from w_1 to w_3, it provides an upper bound on the true edit distance (Gressman). Hence,
  $$d(w_1, w_3) \leq A + B = d(w_1, w_2) + d(w_2, w_3).$$

Hence, each of these examples defines a valid metric space—we just need to ensure that the 3 rules of a valid metric are being followed.

Conclusion

A metric space is more than just a mathematical curiosity. It's a framework for comparing things—whether they're locations, sounds, words, or data points. As long as we can represent an object mathematically, we can measure distances between them within a metric space!

Hence, metric spaces can help us navigate both real and abstract worlds. We can notices them all around us—from the maps we follow, to the spellcheckers we use, all the way to the playlists we jam to.

Works Cited

Russell, Stuart, and Peter Norvig. Artificial Intelligence: A Modern Approach, Global Edition. 2016. pp. 114-115

Gressman, Philip T. Metric Spaces I. Department of Mathematics, University of Pennsylvania, www2.math.upenn.edu/ gressman/analysis/08a-metricspace1.html. Accessed 12 Apr. 2025