

Learning Manhattan Distance: A Comprehensive Guide with R Examples

Authored by
Mohammed loot

November 6, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning Manhattan Distance: A Comprehensive Guide with R Examples*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=11623>

Introduction: Understanding Manhattan Distance (L1 Norm)

The calculation of dissimilarity between data points is fundamental to almost every discipline within [data science](#) and statistical analysis. While most practitioners are familiar with the standard [Euclidean distance](#), which determines the shortest straight line between two points, a powerful alternative exists: the **Manhattan distance**. Also known as Taxicab Geometry or the [L1 norm](#), this metric measures the path taken exclusively along the axes at right angles, simulating movement on a strict grid structure.

The intuition behind the Manhattan distance is derived directly from its namesake--navigating the street grid of Manhattan, where one must travel north-south and east-west, rather than cutting diagonally through buildings. This constraint makes the metric highly valuable in specific contexts, such as evaluating dissimilarity in image processing (where pixels form a grid) or when analyzing feature spaces where the cost of movement is proportional to the distance traveled along each dimension independently. Its application is widespread, proving particularly effective in scenarios where traditional Euclidean measurement might overemphasize the impact of outliers.

In this comprehensive guide, we will explore the mathematical definition of this crucial metric and demonstrate two robust methods for calculating it efficiently using the [R programming language](#). We will cover both the creation of a flexible custom function for singular comparisons and the utilization of R's powerful built-in functions for large-scale matrix calculations, equipping you with the necessary tools for integrating **Taxicab Geometry** into complex analytical workflows.

Mathematical Foundation of Taxicab Geometry

To fully leverage the Manhattan distance, it is essential to understand its formal mathematical definition. It provides a measure of vector dissimilarity by focusing on the absolute differences across all corresponding dimensions. This approach ensures that the metric is robust and interpretable, regardless of the number of features or the scale of the input data. This conceptual clarity is one reason the L1 norm is frequently preferred over other distance metrics in high-dimensional space.

Formally, the [Manhattan distance](#), D , between two points or [vectors](#), A and B , each possessing n dimensions, is defined by the summation of the absolute differences of their [Cartesian coordinates](#) across every dimension. This calculation yields a scalar value representing the total minimal distance required to move from point A to point B while adhering strictly to axial movements.

The formula is explicitly defined as:

$$\sum |a_i - b_i|$$

Here, the index i iterates from 1 to n , representing the i th element in each corresponding vector. The use of the absolute difference, denoted by the vertical bars, is crucial, as distance must always be a non-negative value. Because this metric does not involve squaring the differences (unlike Euclidean distance), large deviations in a single dimension do not disproportionately influence the final result, making it a valuable tool for algorithms where resistance to outlying data points is a priority, such as certain types of **clustering** and [classification algorithms](#).

Practical Implementation in R: Creating a Custom Function

When the task requires calculating the Manhattan distance between just two specific [vectors](#), defining a customized function within R offers unparalleled transparency and control. This method allows the user to explicitly define the steps required by the L1 norm formula, ensuring that the logic is perfectly aligned with the mathematical definition. Furthermore, creating a custom function enhances code modularity, allowing it to be easily reused across various scripts and projects within the [R environment](#).

The following R code block outlines the structure of a simple, yet robust, function named `manhattan_dist`. This function accepts two input vectors, `a` and `b`, calculates the vector of absolute differences between their corresponding elements, and then sums these differences to produce the final scalar distance. This is the most straightforward way to implement the definition of the **L1 norm** when dealing with individual observations.

Define a function to calculate Manhattan distance (L1 norm)

```
manhattan_dist <- function(a, b){  
  dist <- abs(a-b)  
  dist <- sum(dist)  
  return(dist)  
}  
  
# Define two example vectors for calculation  
a <- c(2, 4, 4, 6)  
  
b <- c(5, 5, 7, 8)  
  
# Execute the custom function  
manhattan_dist(a, b)
```

9

Upon execution, the custom function processes the four-dimensional input vectors a and b and returns a single value of **9**. This result concisely quantifies the total axial separation between these

two points in the four-dimensional feature space. This method is ideal for quick checks or when building algorithms where the distance calculation is encapsulated within a larger iterative process, offering high performance and readability.

Validating Results and Understanding the L1 Metric

In statistical programming, especially when defining custom functions that handle core mathematical operations, rigorous verification of the output is a non-negotiable step. Ensuring the accuracy of the `manhattan_dist` function confirms that our implementation correctly reflects the principles of **Taxicab Geometry**. This process involves manually applying the summation of absolute differences to the example vectors used in the previous step, providing a crucial self-check mechanism.

We confirm the calculation for the input vectors $a = (2, 4, 4, 6)$ and $b = (5, 5, 7, 8)$ as follows, adhering precisely to the formula $\sum |a_i - b_i|$:

$$\sum |a_i - b_i| = |2-5| + |4-5| + |4-7| + |6-8|$$

This expands to the sum of the absolute differences across each dimension: $3 + 1 + 3 + 2$. The final total is **9**.

This detailed, step-by-step manual confirmation validates that our R function, `manhattan_dist`, accurately calculates the [Manhattan distance](#). Understanding the L1 metric also extends to appreciating its mathematical properties. Unlike the [Euclidean distance](#) (L2 norm), which minimizes the straight-line distance, the L1 norm minimizes the number of axial steps, making it particularly useful in scenarios where the dimensions represent distinct, independent features, and movement between them must be assessed linearly.

Calculating Pairwise Distances Using R's Built-in `dist()` Function

While a custom function is excellent for calculating the distance between two specific [vectors](#), analytical tasks often require the computation of pairwise distances among hundreds or even thousands of observations simultaneously. For these large-scale operations, R provides the highly optimized built-in function, `dist()`, which is part of the base statistical package. Utilizing `dist()` ensures computational efficiency and leverages R's internal optimizations for handling matrix algebra.

To use the `dist()` function effectively, the individual vectors (observations) must first be organized into a single rectangular data structure, typically a [matrix](#) or data frame, where each row represents an observation (a point in space) and each column represents a dimension (a feature). We use the `rbind()` function in R to row-bind our example vectors into a unified data structure. The critical

step for obtaining the L1 norm specifically is to set the argument `method` equal to `"manhattan"` within the function call, overriding the default Euclidean method.

Define four example vectors

```
a <- c(2, 4, 4, 6)
```

```
b <- c(5, 5, 7, 8)
```

```
c <- c(9, 9, 9, 8)
```

```
d <- c(1, 2, 3, 3)
```

```
# Bind vectors into one matrix (each row represents an observation)
```

```
mat <- rbind(a, b, c, d)
```

```
# Calculate Manhattan distance between each vector in the matrix
```

```
# We utilize the built-in dist() function
```

```
dist(mat, method = "manhattan")
```

```
a b c
```

```
b 9
```

```
c 19 10
```

```
d 7 16 26
```

The output generated by the [dist\(\)](#) function is a distance object, which R displays as a lower triangular matrix. This format is standard because the distance matrix is inherently symmetric (distance from A to B is the same as B to A), and the distance of any point to itself is zero. This matrix contains all possible pairwise Manhattan distances for the four vectors (a, b, c, and d), making it the definitive tool for preparing data for clustering algorithms, such as k-nearest neighbors or hierarchical clustering, which rely heavily on comprehensive dissimilarity metrics.

Interpreting the Distance Matrix Output

Interpreting the distance matrix generated by the **dist()** function is crucial for drawing meaningful conclusions from the data. Each cell in the output represents the L1 norm between the corresponding row and column vectors. For example, the value listed under column `a` and next to row `b` represents the distance between vector `a` and vector `b`. Since the structure is lower triangular, the distances are read by matching the row name (the second vector) to the column name (the first vector).

The resulting output clearly shows the dissimilarities in the four-dimensional space, providing a comprehensive basis for subsequent analysis. For instance, in a clustering context, vectors with

smaller distances are considered more similar and would likely be grouped together. Conversely, vectors with larger distances indicate greater dissimilarity, suggesting they belong to different clusters or categories.

A detailed breakdown of the pairwise distances calculated in the example matrix is as follows:

The Manhattan distance between vector *a* and *b* is confirmed to be **9**.

The Manhattan distance between vector *a* and *c* is **19**, indicating a moderate separation.

The Manhattan distance between vector *a* and *d* is **7**, suggesting *a* and *d* are the closest pair to each other.

The Manhattan distance between vector *b* and *c* is **10**.

The Manhattan distance between vector *b* and *d* is **16**.

The Manhattan distance between vector *c* and *d* is **26**, confirming they are the most dissimilar pair in this set.

This matrix is the standard input format for many statistical functions in R, highlighting the utility and necessary efficiency of the built-in **dist()** function when managing complex [matrix](#) operations within the R ecosystem.

Conclusion: Mastering Dissimilarity Metrics in R

The **Manhattan distance**, or L1 norm, stands as an indispensable metric in the toolkit of any analyst working with multi-dimensional data, particularly in fields like feature engineering and pattern recognition. It provides a robust and geometrically intuitive measure of dissimilarity, offering a valuable alternative to the commonly used Euclidean distance, especially when dealing with high-dimensional data or constrained movement environments.

Through this tutorial, we have established two highly effective methods for calculating this metric within the [R programming language](#). The first method, defining a custom function, offers precision and clarity for single comparisons, ensuring the mathematical definition is perfectly implemented. The second method, utilizing R's powerful and highly optimized [dist\(\)](#) function with the `method = "manhattan"` argument, provides the speed and scalability necessary for calculating comprehensive pairwise distance matrices across large datasets.

Mastering both the custom implementation and the efficient built-in utility ensures that you can confidently select the appropriate computational method based on the scale and complexity of your analytical needs. Integrating the L1 norm effectively allows for more accurate and robust modeling, particularly in applications sensitive to the geometry of the feature space.

Additional Resources

[How to Calculate Euclidean Distance in R](#)

[How to Calculate Mahalanobis Distance in R](#)

[How to Calculate Minkowski Distance in R](#)