

# How to Create a Nested For Loop in R (Including Examples)

Authored by  
**stats writer**

December 19, 2025

## RECOMMENDED CITATION

stats writer (2025). *How to Create a Nested For Loop in R (Including Examples)*.  
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107998>

The concept of a **nested for loop** is a fundamental structure in programming, particularly valuable when working with multi-dimensional data structures in statistical environments like R. Essentially, a nested loop involves placing one iteration structure, known as the inner loop, entirely within the body of another iteration structure, the outer loop. This design allows programmers to systematically iterate through two or more dimensions simultaneously, enabling complex operations across rows and columns of a dataset, or within hierarchical lists.

## Understanding Nested Loops in R: Fundamentals and Mechanics

The operational mechanics of a **nested for loop** are crucial to understand for efficient coding. The execution begins with the outer loop initiating its first iteration. Crucially, before the outer loop can proceed to its second iteration, the inner loop must execute completely from start to finish. This nested execution flow allows you to loop through elements in multiple vectors or multiple dimensions of a matrix and perform operations based on both indices.

If the outer loop is set to run  $N$  times and the inner loop is set to run  $M$  times, the total number of operations performed within the inner loop's body will be  $N$  times  $M$ . This multiplicative nature explains why nested loops, while powerful, must be used judiciously, especially when dealing with very large datasets where performance optimization is paramount and highly optimized R functions are available.

In the context of R, nested loops are frequently employed when native vectorized solutions are not immediately apparent or when dealing with highly specific indexing requirements. Common use cases include iterating over multiple columns of an R data frame to apply transformations, traversing a list of lists, or initializing and manipulating values within a two-dimensional structure. Mastering nested loops provides the necessary explicit control for sequential processes that require interaction between indices from different dimensions.

## The Anatomy of a Basic R For Loop

Before diving into the complexities of nesting, it is helpful to reaffirm the basic structure and purpose of a standard for loop in R. A for loop is designed to execute a block of code repeatedly for a fixed number of times, typically by iterating over a sequence, such as a vector of numbers or elements in a list. The structure requires three primary components: the initialization of the iterator variable (conventionally denoted as  $i$ ), the sequence over which the iteration occurs, and the block of code contained within curly braces `{ }` that specifies the task to be performed during each cycle.

In R, the syntax utilizes the `for` keyword followed by the iterator variable and the `in` keyword, which defines the sequence. This sequence is often created using the colon operator (e.g., `1:4`) to generate a sequence of integers. Understanding this fundamental structure is the prerequisite for

implementing any form of iteration. The loop progresses sequentially, ensuring that the operation defined inside the body is completed for the current value of the iterator before moving to the next element in the sequence.

Consider the following simple illustration where the loop iterates four times, printing the current value of the index  $i$  during each cycle. This serves as the foundational building block upon which the nested structure relies. Notice how the output clearly shows the linear progression of the loop through the defined sequence.

```
for(i in 1:4) {  
  print (i)  
}
```

```
1  
2  
3  
4
```

## Constructing the Nested For Loop Structure

When transitioning from a single loop to a nested setup, we introduce a second iterator, typically  $j$ , which is managed by the inner loop. The inner loop is placed entirely within the execution scope of the outer loop. This hierarchical arrangement creates a matrix-like iteration pattern, where the outer loop controls one dimension (e.g., rows) and the inner loop controls the second dimension (e.g., columns). It is vital that the scope of the inner loop is clearly defined and that the variable names for the iterators (e.g.,  $i$  and  $j$ ) are distinct to prevent naming conflicts and maintain code clarity.

The control flow dictates that the outer loop runs once, setting the value for its iterator ( $i$ ). While  $i$  remains constant for that iteration, the program enters the inner loop. The inner loop then executes all of its cycles completely, changing its iterator ( $j$ ) through its entire defined sequence. Only once the inner loop is exhausted does the program return control to the outer loop, which then increments  $i$  and restarts the entire process of the inner loop. This deep, sequential execution is the defining characteristic of a **nested for loop** and makes it ideal for operations requiring pairwise combination or cell-by-cell manipulation of structured data.

The code below demonstrates this structure. The outer loop iterates four times ( $i$  from 1 to 4), and for each of those iterations, the inner loop iterates twice ( $j$  from 1 to 2). The resulting output shows 4 groups of 2 operations, totaling 8 executions, where the calculation  $i*j$  is performed. This clear example solidifies the concept of the inner loop completing its full cycle before the outer loop

advances.

```
for(i in 1:4) {  
  for(j in 1:2) {  
    print (i*j)  
  }  
}
```

```
1  
2  
2  
4  
3  
6  
4  
8
```

## Practical Application 1: Initializing and Populating an R Matrix

One of the most intuitive applications of a **nested for loop** in R is the initialization and population of values within a two-dimensional matrix. A matrix is inherently suited for this structure, as its elements are addressed using two indices: row (typically controlled by the outer loop,  $i$ ) and column (controlled by the inner loop,  $j$ ). This method offers granular control over every single cell, allowing complex calculations or assignments based directly on the position of the cell itself, something that might be less straightforward using certain vectorized approaches.

In the following demonstration, we first establish a 4x4 matrix initialized with missing values (`NA`). The goal is to replace these placeholders with calculated values, specifically the product of the row index and the column index (i.e.,  $i * j$ ). The outer loop iterates from 1 to 4, representing the rows, and the inner loop also iterates from 1 to 4, representing the columns. During each of the 16 total iterations, the cell corresponding to the current coordinates is accessed and updated with the result of the multiplication.

This process highlights the precise control offered by nested iteration. While simple matrix initialization might be handled by functions like `outer()`, using a loop allows for insertion of custom, complex logic that might depend on previous calculations or external conditions specific to the row and column indices. Observing the output confirms that the matrix is successfully populated, demonstrating complete traversal and targeted assignment across the entire structure.

```
#create matrix
```

```
empty_mat <- matrix(nrow=4, ncol=4)
```

```
#view empty matrix
```

```
empty_mat
```

```
NA NA NA NA
```

```
NA NA NA NA
```

```
NA NA NA NA
```

```
NA NA NA NA
```

```
#use nested for loop to fill in values of matrix
```

```
for(i in 1:4) {
```

```
  for(j in 1:4) {
```

```
    empty_mat = (i*j)
```

```
  }
```

```
}
```

```
#view matrix
```

```
empty_mat
```

```
1 2 3 4
```

```
2 4 6 8
```

```
3 6 9 12
```

```
4 8 12 16
```

## Practical Application 2: Transforming Elements within an R Data Frame

Beyond matrices, **nested for loops** are highly useful for iterating over and modifying elements within an R data frame. Although data frames are collections of vectors, treating them generically as a two-dimensional structure allows for systematic, cell-by-cell manipulation. In this scenario, the outer loop often indexes the rows (using `nrow(df)` to determine the row count), and the inner loop indexes the columns (using `ncol(df)` to determine the column count). This setup ensures every single entry in the data frame is visited precisely once.

The subsequent example illustrates how to apply a transformation--specifically squaring the numerical value--to every element in a predefined data frame named `df`. We initialize the data frame with two variables (columns) and three observations (rows). The nested loop structure is then applied. The outer loop iterates from 1 to 3 (rows), and the inner loop iterates from 1 to 2 (columns). Inside the innermost block, the existing value at `df` is retrieved, squared using the exponent operator `^2`, and then immediately reassigned back to the same cell, effectively modifying the data frame in place.

While this method effectively achieves the desired transformation, it is important to acknowledge that R is optimized for vectorization. For simple operations like squaring all elements, a vectorized approach would be significantly faster and more idiomatic in R. However, if the operation requires referencing the row index  $i$  or column index  $j$  explicitly within the calculation, or if the calculation depends on sequential states, the explicit control of a nested loop becomes a necessary mechanism.

#### **#create data frame**

```
df <- data.frame(var1=c(1, 7, 4),  
var2=c(9, 13, 15))
```

```
#view data frame
```

```
df
```

```
var1 var2
```

```
1 1 9
```

```
2 7 13
```

```
3 4 15
```

```
#use nested for loop to square each value in the data frame
```

```
for(i in 1:nrow(df)) {
```

```
for(j in 1:ncol(df)) {
```

```
df = df^2
```

```
}
```

```
}
```

```
#view new data frame
```

```
df
```

```
var1 var2
```

```
1 1 81
```

```
2 49 169
```

```
3 16 225
```

## **Performance Considerations: Vectorization vs. Iteration in R**

A critical consideration for any R programmer is performance, especially when scaling up operations to handle big data. While nested for loops provide flexibility and explicit control, they are inherently less efficient than R's built-in vectorized operations. R is designed to perform operations on entire vectors or arrays at once, primarily because the underlying calculations for these vectorized functions are executed in optimized C or Fortran code, avoiding the overhead

associated with interpreting R code repeatedly during a loop iteration.

The speed discrepancy becomes markedly apparent when dealing with datasets that contain thousands or millions of elements. A nested loop, which requires executing the inner code block  $N$  times  $M$  times, incurs significant cost due to constant function calls, variable assignments, and index lookups within the R environment for every single step. For small arrays, this overhead is negligible, but as the dimensions grow, the quadratic complexity of nested loops quickly leads to unacceptable execution times compared to vectorized equivalents.

Therefore, a core tenet of effective R programming is to always prioritize vectorization over explicit iteration using **for loops** whenever a functional equivalent exists. Functions like `rowSums()`, `colMeans()`, or simply applying an arithmetic operator directly to a data structure leverage R's optimization capabilities. Nested loops should generally be reserved for situations where the iteration logic is too complex or conditional to be easily expressed through built-in vectorized functions, or when the dataset size is guaranteed to remain minimal.

## Alternatives to Nested Loops: The Power of Apply and Data.table

Given the performance limitations of standard iterative structures on large datasets, R provides several powerful alternatives engineered for speed and clarity. Foremost among these are the family of `apply` functions (including `lapply`, `sapply`, `vapply`, and `tapply`). These functions essentially abstract the looping process, pushing the iteration logic down into optimized C code. They allow the user to apply a specified function across the rows, columns, or elements of a list or array, resulting in substantially faster processing times than manually constructed loops.

For operations on rows or columns of a matrix or data frame, the `apply()` function is particularly relevant, allowing the user to specify whether the function should operate across the margins (rows or columns). Similarly, `lapply()` is optimized for iterating over lists. Embracing these high-level functions is essential for writing efficient and idiomatic R code.

For big data, the family of apply functions tend to be much quicker and the data.table package has many built-in functions that perform efficiently on larger datasets. Utilizing these advanced tools ensures that analyses are both robust and scalable, moving away from the performance bottlenecks inherent in manual nested iteration.

## Further Reading on R Iteration and Data Manipulation

[How to Loop Through Column Names in R](#)

[How to Append Rows to a Data Frame in R](#)