

Learning About the intersect() Function in R: A Tutorial with Examples

Authored by
Mohammed loot

October 29, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning About the intersect() Function in R: A Tutorial with Examples*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=5336>

Introduction to Set Operations and the `intersect()` Function in R

The ability to perform [Set operations](#) is fundamental in data analysis and programming. In the statistical programming environment of [R](#), we frequently need to determine the common elements shared between two distinct objects. This crucial task is efficiently handled by the **`intersect()`** function, which is readily available in base R. Understanding how to utilize this function is essential for tasks ranging from basic list comparisons to complex data filtering and validation.

The concept of an intersection is derived directly from set theory: it identifies the elements that exist simultaneously in both input objects. When dealing with [vectors](#) or lists, the output will be a new object containing only these shared, unique values. It is important to distinguish this tool from related set operations like `union()` (combining all unique elements) or `setdiff()` (finding elements present in one set but not the other).

For simple data structures, such as atomic vectors, the base R version of [intersect\(\)](#) is both powerful and intuitive. However, as we explore more complex scenarios, specifically finding common rows across two [data frames](#), we will see the necessity of leveraging specialized packages, particularly the highly popular **`dplyr`** package, to extend this functionality beyond one-dimensional comparisons.

Understanding the Base R `intersect()` Syntax

The core syntax for the base R implementation of the **`intersect()`** function is extremely concise, requiring only the two objects you wish to compare. This simplicity makes it highly accessible for quick comparisons of any atomic vectors, including numeric, integer, logical, or character types.

The function strictly accepts two arguments, typically defined as `object1` and `object2`. The output is inherently a set, meaning that any duplicate elements present in the original input objects are automatically reduced to a single occurrence in the resulting intersection vector. Furthermore, the order of the resulting elements is generally determined by the order in which they first appear in the first specified object (`object1`).

This function uses the following basic syntax structure in R:

```
intersect(object1, object2)
```

The following examples demonstrate how to utilize the **`intersect()`** function efficiently with both numeric and character vectors, highlighting its behavior regarding data types and duplicate handling.

Example 1: Using `intersect()` with Vectors

To illustrate the practical usage of `intersect()`, we first examine two numeric vectors, `x` and `y`. These vectors contain several shared elements, some of which are duplicated within their respective definitions. The goal is to accurately extract the unique common values, demonstrating the function's utility in numerical comparisons.

The following code shows how to use the `intersect()` function to find the intersection between two vectors in R:

```
# define two numeric vectors, including duplicates
```

```
x <- c(1, 4, 5, 5, 9, 12, 19)
```

```
y <- c(1, 2, 5, 5, 10, 14, 19)
```

```
# find intersection between two vectors
```

```
intersect(x, y)
```

```
1 5 19
```

From the output, we can see that vectors `x` and `y` have three unique values in common: **1**, **5**, and **19**. It is important to notice that while the value 5 appeared multiple times in the source vectors, the result vector includes it only once, confirming the set-based nature of the operation.

Furthermore, the `intersect()` function operates identically with character vectors, allowing for straightforward comparison of textual data, such as identifying shared category labels or identifiers between two lists.

```
# define two character vectors
```

```
x <- c('A', 'B', 'C', 'D', 'E')
```

```
y <- c('C', 'D', 'E', 'F')
```

```
# find intersection between two character vectors
```

```
intersect(x, y)
```

```
"C" "D" "E"
```

The resulting vector reveals that vectors `x` and `y` share the strings **C**, **D**, and **E**. A practical benefit of this function is its independence from vector length; the two input [vectors](#) do not have to be of the same size for the `intersect()` function to yield a valid result.

Example 2: Finding Intersections Across Data Frames Using `dplyr`

While the base R implementation is suitable for one-dimensional vectors, it is insufficient for comparing entire rows across two [data frames](#). To find the exact rows that two data frames have in common--meaning every value in every column must match precisely--we must utilize the **intersect()** function provided by the [dplyr package](#). This function treats each row as a single unit or observation when performing the set comparison.

The **`dplyr::intersect()`** function is essential for tasks like merging datasets or verifying data consistency, as it performs a robust, whole-row comparison. When using this tool, always ensure the **dplyr** package is loaded or explicitly reference the function using the `dplyr::` prefix to avoid confusion with the base R version.

The following code illustrates how to define two sample data frames, `df1` and `df2`, and subsequently use the **dplyr** function to identify their shared rows:

library(dplyr)

```
# define two data frames
```

```
df1 <- data.frame(team=c('A', 'A', 'B', 'B'),  
points=c(12, 20, 25, 19))
```

```
df1
```

```
team points
```

```
1 A 12
```

```
2 A 20
```

```
3 B 25
```

```
4 B 19
```

```
df2 <- data.frame(team=c('A', 'A', 'B', 'C'),
```

```
points=c(12, 22, 25, 32))
```

```
df2
```

```
team points
```

```
1 A 12
```

```
2 A 22
```

```
3 B 25
```

```
4 C 32
```

```
# find intersection between two data frames (rows must match across all columns)
```

```
dplyr::intersect(df1, df2)
```

```
team points
```

```
1 A 12
```

```
2 B 25
```

The output confirms that only two rows are common to both data frames: the row where team is 'A' and points is 12, and the row where team is 'B' and points is 25. Note carefully that this **dplyr::intersect()** function will only return the rows that have the same values in *every single* column between the two [data frames](#).

Calculating the Size of the Intersection

Beyond merely identifying the common elements, analysts often require a quick count of the overlap--the size of the intersection. This metric provides a crucial measure of similarity or redundancy between the two input sets. Fortunately, we can easily determine this count for both vectors and data frames by wrapping the **intersect()** function within the standard **length()** function.

When applied to data frames, calculating the length of the resulting intersection data frame immediately yields the total number of identical observations shared between the two datasets. This method is particularly useful in auditing large datasets for overlap and ensuring that merging operations do not introduce unintended duplicates based on primary keys.

Using the data frames defined previously, we can quickly determine the count of common rows with the following syntax:

```
# find number of rows in common between the two data frames
```

```
length(dplyr::intersect(df1, df2))
```

```
2
```

The output clearly indicates that the two data frames have precisely **2** rows in common. This quantitative approach allows for robust, scalable assessment of data relationships within [R](#).

Additional Resources for Set Operations

The **intersect()** function, whether used in its base R form for vectors or the extended **dplyr** version for data frames, is an indispensable tool for performing fundamental [set operations](#). To further enhance your data manipulation skills in R, it is highly beneficial to explore the related functions that complement **intersect()**, providing a comprehensive toolkit for comparing and manipulating data based on set theory principles.

The following tutorials explain how to use other common set operation functions in R:

union(): Returns all unique elements present in either object, effectively merging the two sets.

setdiff(): Returns elements present in the first object but absent in the second, providing the set difference.

setequal(): A logical function that checks if two sets contain exactly the same elements, regardless of element order.