

# Creating External Legends in R: A Comprehensive Guide

Authored by  
**Mohammed looti**

November 5, 2025

## RECOMMENDED CITATION

Mohammed looti (2025). *Creating External Legends in R: A Comprehensive Guide*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=11123>

## Mastering External Legends in Base R Graphics

The effective visualization of data hinges on the ability to clearly distinguish and identify different data series. In the context of the [Base R Graphics](#) system, placing the plot **legend** outside the primary data area is often a necessary technique. This requirement becomes especially acute when visualizations are complex, containing numerous categories or highly dense data points, where an internal legend would inevitably lead to visual clutter or data occlusion. By positioning the legend externally, we significantly enhance the overall clarity, professional appearance, and interpretability of the data visualization.

Achieving this external placement is not automatic; it demands precise manipulation of the plot's internal structure, specifically concerning margins and drawing boundaries. The most effective and reliable approach involves proactively utilizing the `par()` function to allocate substantial additional space around the graphical device before any plotting commands are executed. This preparatory step is fundamental, as failing to expand the margins will result in the legend being cropped or completely hidden by the default rendering boundaries.

The foundational setup for enabling this extended drawing capability centers around two key parameters: defining the expanded margin space and overriding the default **clipping** behavior. The following code snippet demonstrates the required configuration, specifically allocating eight lines of extra space to the right margin, which is the conventional location for external legends in this setup.

```
par(mar=c(5, 4, 4, 8), xpd=TRUE)
```

The subsequent sections provide a detailed, step-by-step guide on implementing this powerful technique within the [R](#) environment, ensuring that the final graphical output is polished, informative, and adheres to professional standards for data presentation.

### Critical Configuration of Graphical Parameters using `par()`

Prior to drawing any elements onto the graphical device, it is imperative to establish the correct plotting environment using the [par\(\) function](#). This function is responsible for managing global [graphical parameters](#) that control fundamental aspects of how plots are rendered, including dimensions and boundaries. For the specific task of positioning a legend outside the central plotting area, two parameters--`mar` and `xpd`--are absolutely essential and must be configured correctly.

The `mar` parameter dictates the dimensions of the plot margins, specified in lines of text. It requires a vector of four numbers corresponding to the margins in the order: bottom, left, top, and right.

Standard default margins are usually set to `c(5, 4, 4, 2)`. By dramatically increasing the fourth value (the right margin), such as setting it to 8 as shown in our example, we effectively reserve a generous strip of screen real estate dedicated exclusively to housing the external legend. This deliberate expansion is the physical foundation that prevents the legend from overlapping the visualized data.

Equally important is the `xpd` parameter, which governs the graphical behavior known as **clipping**. By default, **R** restricts all graphical output solely to the plot region, meaning anything drawn outside the data boundaries is clipped (`xpd=FALSE`). Setting `xpd=TRUE` overrides this restriction, granting permission to draw elements--including the **legend**--into the expanded margin area (the figure region). Without this critical setting, even if the margin is expanded using `mar`, any attempt to place the legend outside the standard plot boundaries will result in it being truncated or becoming completely invisible.

## Step 1: Preparing Multi-Series Data for Visualization

To effectively demonstrate the mechanics of creating and placing an external legend, we must first generate a sample dataset that requires differentiation. This simulated data will form the basis of a multi-series **scatterplot**, making a clear and unambiguous legend necessary for interpreting the two distinct groups.

We will define two separate data structures, `df1` and `df2`, each containing corresponding `x` and `y` coordinates. These structures simulate two independent datasets that we intend to display simultaneously within the same graphical coordinate system, requiring the legend to act as the key to their identification.

The following R code snippet utilizes the `data.frame()` function to initialize these two essential data structures. This preparation sets the stage for the subsequent steps involving plotting and advanced parameter manipulation.

### #create data frames

```
df1 <- data.frame(x=c(1, 2, 3, 4, 5, 6, 7),  
y=c(2, 7, 19, 26, 24, 29, 31))
```

```
df2 <- data.frame(x=c(1, 2, 3, 4, 5, 6, 7),  
y=c(4, 4, 7, 9, 12, 13, 8))
```

## Step 2: Plot Construction and Initial Legend Placement

With the necessary data prepared and the graphical environment configured--specifically the

increased right margin set via `mar` and clipping disabled via `xpd=TRUE`--we can proceed to render the visualization and call the [legend\(\) function](#). It is critical that the `par()` setup is executed immediately before the plotting commands to ensure the expanded margin settings are active when the plot area is defined.

The plotting process begins by using the `plot()` command with `df1` to establish the overall coordinate system, axes, and initial data series. Subsequently, the `points()` function is used with `df2` to overlay the second data series onto the existing graph. We differentiate these groups visually by assigning distinct plotting characters (`pch`), such as `pch=1` for the first group and `pch=3` for the second.

The key step is the invocation of the [legend\(\) function](#). Although we specify a starting position (e.g., `"topright"`), the magic of external placement relies entirely on the `inset` argument. By setting `inset` to a vector like `c(-0.2, 0)`, we instruct **R** to horizontally displace the legend by 20% of the plot width in a negative direction (away from the center). This displacement pushes the legend cleanly out of the plot area and into the generous right margin space we previously reserved.

**#add extra space to the right of the plot**

```
par(mar=c(5, 4, 4, 8), xpd=TRUE)
```

```
#plot both data frames
```


```
plot(y ~ x, df1, pch=1, main="Scatterplot with multiple groups")
```

```
points(y ~ x, df2, pch=3)
```

```
#add legend outside of plot
```

```
legend("topright", inset=c(-0.2, 0), legend=c("df1","df2"), pch=c(1,3), title="Data")
```

The immediate outcome of this configuration is a highly legible visualization where the explanatory key is strategically positioned to the right, ensuring graphical elements do not contend for space within the primary data presentation area, as evidenced in the output image below.



### Step 3: Fine-Tuning Horizontal Placement with the `inset` Parameter

The true power and flexibility of external legend placement come from the precise controls offered by the [`legend\(\)` function](#), particularly through the `inset` argument. This argument, a vector of length two `c(x, y)`, specifies the fractional distance the legend should be moved relative to its designated anchor corner. These fractions are calculated based on the total width and height of the plotting region, providing a scalable method for placement.

When aiming for external placement on the right side, we must use a negative value for the x-inset. This negative fraction dictates how far the legend is pushed horizontally away from the plot boundary and into the expanded right margin. For instance, increasing the magnitude of the negative x-inset from `-0.2` to `-0.3` will shift the legend further right. This adjustment is highly valuable when dealing with verbose legend entries or when space needs to be reserved close to the plot boundary for annotations or other graphical elements.

By observing the modification in the code block below, notice that the first element of the `inset` vector has been changed to `-0.3`. This horizontal repositioning highlights the separation between the core visualization and its explanatory key, demonstrating granular control over the layout.

**#add extra space to the right of the plot**

```
par(mar=c(5, 4, 4, 8), xpd=TRUE)
```

```
#plot both data frames
```


```
plot(y ~ x, df1, pch=1, main="Scatterplot with multiple groups")
```

```
points(y ~ x, df2, pch=3)
```

```
#add legend outside of plot
```

```
legend("topright", inset=c(-0.3, 0), legend=c("df1","df2"), pch=c(1,3), title="Data")
```

The resultant visualization clearly shows that the increased negative x-inset successfully shifts the legend further to the right, maximizing the visual separation from the active plot area, as demonstrated in the visual output below.



#### Step 4: Achieving Vertical Alignment via Inset Adjustment

In addition to precise horizontal control, the `inset` argument provides robust capabilities for managing the vertical alignment of the external **legend**. The second element of the `inset` vector, which corresponds to the y-axis adjustment, moves the legend up or down relative to the initial anchor corner specified (e.g., `"topright"`).

Since our current setup anchors the legend at the "topright" position, applying a positive y-inset value will push the legend downwards into the right margin. For instance, setting the y-inset from to 0.5 will vertically displace the legend by half the height of the plotting region. This level of manipulation is often necessary to center the key vertically within the allocated margin space or to adjust its height to accommodate other titles or elements placed in the top or bottom margins of the figure.

It is crucial to remember that the direction of movement is relative to the anchor point. If the legend were instead anchored at the bottom corner (e.g., "bottomright"), a positive y-inset would move it upwards toward the center of the plot, illustrating the importance of understanding the reference system when using the `inset` vector.

The final code block illustrates how to combine a significant horizontal offset (-0.3) with a vertical offset (0.5) to achieve a desired alignment that places the legend lower down within the expanded right margin.

**#add extra space to the right of the plot**

```
par(mar=c(5, 4, 4, 8), xpd=TRUE)
```

```
#plot both data frames
```

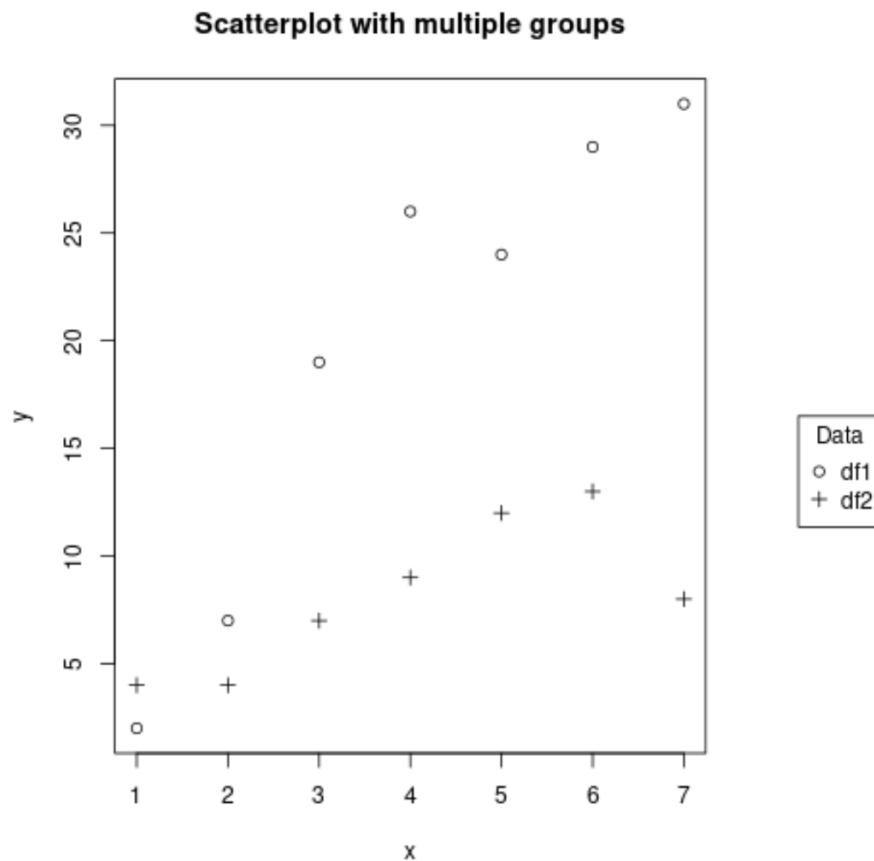
```
plot(y ~ x, df1, pch=1, main="Scatterplot with multiple groups")
```

```
points(y ~ x, df2, pch=3)
```

```
#add legend outside of plot
```

```
legend("topright", inset=c(-0.3, .5), legend=c("df1","df2"), pch=c(1,3), title="Data")
```

As depicted in the image below, this adjustment successfully repositions the legend significantly lower down the right margin. Such precise control ensures that the final visualization adheres to complex layout requirements often demanded by professional reports or academic publications.



## Conclusion and Advanced Considerations for Base R Layouts

The successful placement of a legend outside the main plotting area within **Base R Graphics** relies upon the harmonious execution of a three-part strategy. First, expanding the necessary margin (in this case, the right margin) using the `mar` parameter within the [par\(\) function](#) reserves the physical space. Second, setting the `xpd=TRUE` parameter overrides the default clipping mechanism, permitting drawing into this newly expanded area. Finally, the strategic use of the `inset` argument within the [legend\(\) function](#) provides the necessary fractional offset to push the legend completely into the reserved margin space.

A crucial professional consideration when dealing with `par()` settings is persistence. Graphical parameters modified by the [par\(\) function](#) often remain active across subsequent plots within the same session unless explicitly reset. To ensure consistency and prevent unintended side effects in later visualizations, best practice dictates saving the original `par()` settings before any modifications are made. This is typically achieved using `op <- par(no.readonly = TRUE)`, allowing the user to restore the default environment later in the script using `par(op)`.

By mastering the relationship between margin expansion (`mar`), drawing permissions (`xpd`), and relative positioning (`inset`), **R** users can create sophisticated, clean, and highly professional

visualizations. This methodology ensures that explanatory elements enhance, rather than interfere with, the core message of the data presentation. Further tutorials and advanced R techniques are available to deepen your data visualization expertise.