

How to Calculate a Rolling Average in R Using the rollmean() Function

Authored by
stats writer

November 27, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Calculate a Rolling Average in R Using the rollmean() Function*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=100646>

Calculating a rolling average (also known as a moving average) in R is a fundamental technique used in time series analysis to smooth out short-term fluctuations and highlight longer-term trends or cycles. By averaging data points over a specified window of time, this method significantly mitigates the impact of sudden, isolated outlier values, providing a clearer perspective on the underlying pattern of the data.

The primary tool for this calculation within the R environment is the powerful `rollmean()` function. This function is designed to handle sequential data efficiently, accepting a numeric vector of observed values as its core input. It subsequently generates an output vector of identical length, where each entry represents the calculated rolling mean for that specific data window. Mastery of this function is essential for anyone conducting serious data smoothing or trend visualization in R.

For practical implementation, the `rollmean()` function is housed within the widely utilized zoo package. This package specializes in generalized time series objects and crucial functions for financial and economic data analysis. Integrating `rollmean()` with the data manipulation capabilities of packages like dplyr allows data analysts to integrate moving averages seamlessly into existing data frames, facilitating advanced reporting and visualization workflows.

The Importance of Rolling Averages in Data Analysis

In time series analysis, a **rolling average** represents the average value of a certain number of previous periods. This calculation is vital for uncovering meaningful patterns obscured by daily noise. Unlike a simple arithmetic mean, which averages all historical data equally, the rolling average is dynamic and constantly updates as new data becomes available, making it an invaluable tool for forecasting and trend identification in fields ranging from finance to logistics.

The primary advantage of employing a rolling average is its ability to filter out high-frequency noise. For example, if a company experiences a single day of unusually high or low sales, this outlier event will only briefly influence the trend calculation before it moves out of the defined window. This ensures that strategic decisions are based on sustained performance rather than momentary anomalies.

The easiest way to calculate a rolling average in R is to use the **`rollmean()`** function from the **zoo** package, often combined with **dplyr** for efficient data frame operations. The basic structure for a 3-day calculation is demonstrated below, illustrating how the packages work together within a typical data pipeline:

```
library(dplyr)
```

```
library(zoo)
```

```
#calculate 3-day rolling average
```

```
df %>%  
mutate(rolling_avg = rollmean(values, k=3, fill=NA, align='right'))
```

This particular example calculates a **3-day** rolling average for the column titled **values**, defining the window size (`k=3`) and handling initial incomplete observations by using `fill=NA`.

The Role of the `zoo` Package and `rollmean()` Function

The `rollmean()` function is the cornerstone of moving average calculations in R. It belongs to the **zoo** package, which stands for Z's ordered observations. This package is specifically designed to handle indexed data, making it perfectly suited for processing time series data where observations are ordered sequentially, such as daily sales, stock prices, or temperature readings. Before utilizing `rollmean()`, both the `zoo` package and typically the `dplyr` package (for data manipulation efficiency) must be loaded into the R session using the `library()` function.

The pipeline workflow illustrated above uses the `mutate()` function from the `dplyr` package. The `mutate()` function is invaluable for creating new columns derived from existing ones, allowing us to seamlessly integrate the rolling average calculation into the data frame structure (`df`). The pipe operator (`%>%`) sequentially passes the data frame to the next function, enhancing code readability and maintainability.

Understanding the key parameters within `rollmean()` is crucial for accurate calculation. The primary argument is the data vector itself. The second most important parameter is `k`, which specifies the width of the rolling window--that is, the number of consecutive periods to average. For instance, setting `k=3` ensures that each calculated average represents the mean of the three most recent data points, including the current one. This precision ensures that the smoothing effect is tailored precisely to the desired temporal scale of the analysis.

Understanding the Syntax and Core Parameters

The flexibility of the `rollmean()` function comes from its powerful yet straightforward set of parameters. Analysts must carefully consider three major parameters: `k`, `fill`, and `align`, as they dictate how the rolling average is calculated and how boundary conditions are handled. Incorrect specification of these parameters can lead to misinterpretation of trends, especially at the beginning or end of a time series.

The parameter `k`, the window size, controls the degree of smoothing applied. A larger `k` results in a smoother trend line, emphasizing long-term movements but potentially delaying the detection of recent trend shifts. Conversely, a smaller `k` makes the rolling average more reactive to short-term changes. For comprehensive trend assessment, it is often best practice to experiment with

different `k` values, such as 3, 5, or 7, to determine the optimal level of smoothing for the specific dataset under investigation.

The `fill` parameter addresses how the function handles the initial observations for which a full window size (`k`) is not yet available. If `k=3`, the first two observations cannot have a valid three-period average. By setting `fill=NA` (as shown in the examples), R inserts a `NA` (Not Available) value for these incomplete periods. This preserves the length of the data frame, ensuring that the new column aligns perfectly with the original observations. Users may also utilize other padding methods, but `fill=NA` is the industry standard for maintaining data integrity when dealing with time series.

Finally, the `align` parameter specifies how the rolling window is centered relative to the output observation. When conducting retrospective analysis--calculating the average based on historical data--the essential setting is `align='right'`. This setting ensures that the rolling average calculated for the current time step includes the current value and the `k-1` preceding values. This contrasts with `align='center'`, which requires future data points and is typically reserved for visualization where the average needs to be perfectly centered over the time period it represents.

Practical Example: Setting Up the Sales Data

To demonstrate the practical application of `rollmean()`, consider a scenario involving product sales tracking over a period of ten consecutive days. This small dataset provides a clear, manageable sequence of observations suitable for illustrating how the rolling average smooths the volatility present in daily figures. Our goal is to transform this raw data into actionable insights regarding underlying sales performance.

We begin by constructing a data frame in R. This structure, named `df`, contains two columns: `day` (representing the sequence index from 1 to 10) and `sales` (representing the recorded sales volume for that specific day). The initial raw sales figures show noticeable variation, ranging from a low of 12 to a high of 27, which is precisely the kind of variability that rolling averages are designed to normalize.

The R code used to define and inspect this initial data structure is as follows:

```
#create data frame  
df <- data.frame(day=c(1, 2, 3, 4, 5, 6, 7, 8, 9, 10),  
sales=c(25, 20, 14, 16, 27, 20, 12, 15, 14, 19))  
  
#view data frame  
df
```

```
day sales
1 1 25
2 2 20
3 3 14
4 4 16
5 5 27
6 6 20
7 7 12
8 8 15
9 9 14
10 10 19
```

Upon reviewing the initial output, it is evident that while Day 5 saw a peak sales figure of 27, Day 7 dropped significantly to 12. Such rapid shifts can skew short-term decision-making if viewed in isolation. By applying a rolling average, we aim to calculate a smoother trend line that reflects the sustained performance level across those ten days, rather than focusing solely on the daily volatility.

Calculating a Single Rolling Average (k=3)

To derive the 3-day trend, we apply the `rollmean()` function within the `mutate()` pipeline provided by the `dplyr` package. This streamlined approach ensures that the new rolling average column, which we name `avg_sales3`, is automatically appended to the existing data frame `df`, maintaining the integrity and structure of the sequential data. This method avoids manual looping and is computationally efficient for large datasets.

The core command structure explicitly sets `k=3`, defining the necessary three-day window. Crucially, we use `fill=NA` to handle the boundary condition at the start of the series. Since the rolling average for Day 1 and Day 2 cannot be calculated based on three prior days, these entries are populated with `NA` values. The `align='right'` argument confirms that the calculated average is attributed to the latest day in the window, reflecting the historical trend leading up to that point.

Executing the following code yields the new column, `avg_sales3`:

```
library(dplyr)
```

```
library(zoo)
```

```
#calculate 3-day rolling average of sales
```

```
df %>%
```

```
mutate(avg_sales3 = rollmean(sales, k=3, fill=NA, align='right'))
```

```
day sales avg_sales3
1 1 25 NA
2 2 20 NA
3 3 14 19.66667
4 4 16 16.66667
5 5 27 19.00000
6 6 20 21.00000
7 7 12 19.66667
8 8 15 15.66667
9 9 14 13.66667
10 10 19 16.00000
```

Interpreting the Results and Calculation Breakdown

The resulting data frame clearly shows the effect of the 3-day rolling average calculation. The first two entries in the `avg_sales3` column are `NA`, confirming the use of the `fill=NA` parameter, as there are not enough preceding data points to form a complete 3-day window. This is a crucial aspect of time series smoothing--the initialization phase of the average must be accounted for or ignored, depending on the analytical requirement.

The first valid rolling average appears on Day 3, with a value of **19.66667**. This value is derived by averaging the sales figures from Day 1, Day 2, and Day 3. Specifically, the calculation confirms the retrospective nature of the analysis, where the average is anchored to the latest date in the window:

3-Day Moving Average for Day 3 = (Sales Day 1 + Sales Day 2 + Sales Day 3) / 3

3-Day Moving Average = (25 + 20 + 14) / 3 = **19.66667**

This process is repeated sequentially for every subsequent day. For instance, the value on Day 4 (**16.66667**) is the average of sales on Days 2, 3, and 4. This sequential updating ensures that the average is always reflective of the most recent short-term performance.

By comparing the raw sales figures to the `avg_sales3` column, the smoothing power of the rolling average becomes apparent. For instance, the high sales peak of 27 on Day 5 is tempered by the lower preceding values (20 and 16), resulting in a smoothed average of only 19.00. Similarly, the deep sales trough of 12 on Day 7 does not drag the average down as severely, yielding a 3-day average of 19.66667, suggesting that the long-term trend remains relatively stable despite momentary fluctuations.

Advanced Application: Calculating Multiple Rolling Averages

One of the significant strengths of combining the `dplyr` workflow with the `rollmean()` function is the ease with which multiple rolling windows can be computed simultaneously. Analyzing trends across different time horizons (e.g., short-term 3-day average versus medium-term 4-day average) is crucial for technical analysts and planners, as diverging or converging trends across these metrics can signal important inflection points in the data.

To calculate both the 3-day and the 4-day moving averages in a single step, we simply include a second instance of `rollmean()` within the same `mutate()` call, assigning it to a new column name, `avg_sales4`, and adjusting the `k` parameter accordingly to `k=4`. All other parameters, such as `fill=NA` and `align='right'`, remain consistent to ensure accurate historical calculation for both metrics.

The following code block demonstrates this efficient parallel calculation:

```
library(dplyr)
```

```
library(zoo)
```

```
#calculate 3-day and 4-day rolling average of sales
df %>%
mutate(avg_sales3 = rollmean(sales, k=3, fill=NA, align='right'),
avg_sales4 = rollmean(sales, k=4, fill=NA, align='right'))
```

```
day sales avg_sales3 avg_sales4
```

```
1 1 25 NA NA
```

```
2 2 20 NA NA
```

```
3 3 14 19.66667 NA
```

```
4 4 16 16.66667 18.75
```

```
5 5 27 19.00000 19.25
```

```
6 6 20 21.00000 19.25
```

```
7 7 12 19.66667 18.75
```

```
8 8 15 15.66667 18.50
```

```
9 9 14 13.66667 15.25
```

```
10 10 19 16.00000 15.00
```

Observing the final output, it is important to note the differing starting points for the valid calculations. The 3-day average starts reporting on Day 3, but the 4-day average (`avg_sales4`) only begins reporting on Day 4, as it requires four complete observations (Day 1, 2, 3, and 4) to produce its first mean (18.75). Furthermore, the 4-day average exhibits slightly more inertia and is

generally smoother than the 3-day average, confirming the relationship between the window size (k) and the degree of data smoothing achieved. By implementing these techniques, analysts can gain a comprehensive view of short-term volatility alongside broader structural trends.

Summary of Key Rolling Average Parameters

k: Defines the size of the window used for averaging. A larger k provides greater smoothing.

fill: Determines how incomplete windows at the start of the series are handled. Using `fill=NA` is standard practice to preserve the output vector length.

align: Specifies the orientation of the window relative to the output index. `align='right'` ensures the average uses the current point and past points, making it suitable for retrospective trend tracking.

[How to Plot Multiple Columns in R](#)

[How to Average Across Columns in R](#)

[How to Calculate the Mean by Group in R](#)