

Calculate a Rolling Average in R (With Example)

Authored by
Mohammed looti

October 28, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *Calculate a Rolling Average in R (With Example)*.
PSYCHOLOGICAL STATISTICS. Retrieved from
<https://statistics.arabpsychology.com/?p=4654>

Understanding Rolling Averages in Time Series Analysis

Time series analysis involves examining data points collected sequentially over time. A crucial concept within this field is the **rolling average**, often referred to as a **moving average**. This metric represents the average value of a specific, fixed number of previous periods, providing an essential tool for data smoothing and trend identification.

The primary utility of calculating a rolling average is to dampen short-term fluctuations and noise, allowing analysts to discern underlying trends and cycles within the data. For instance, in financial markets, a 50-day moving average helps distinguish genuine shifts in price momentum from daily volatility. Unlike a simple arithmetic mean, which averages all historical data points equally, the rolling average is dynamic; it evolves as new data arrives and older data drops out of the calculation window. This focus on recent history makes it highly relevant for forecasting and monitoring current performance trends.

To effectively manage and compute these calculations within a robust statistical environment, we turn to **R**. R provides specialized packages and optimized functions designed to handle the sequential nature of time series data efficiently. Utilizing these tools ensures that complex calculations, such as determining a rolling average across thousands of data points, can be executed quickly and accurately, forming the backbone of quantitative analysis.

Essential R Packages: zoo and dplyr for Data Manipulation

Calculating a rolling average in R relies predominantly on specialized packages built for handling indexed and ordered data. The most straightforward and widely accepted method utilizes the **rollmean()** function, which is housed within the **zoo** package. The **zoo** package (short for Z's Ordered Observations) provides functionality for creating and manipulating indexed data structures, which are essential for ensuring that calculations respect the temporal order of observations.

Furthermore, while **zoo** handles the mathematical computation, the process is significantly streamlined by integrating it with the **dplyr** package, a core component of the Tidyverse. **dplyr** facilitates powerful data manipulation tasks, particularly the creation of new variables using the **mutate()** function and the use of the pipe operator (`%>%`). This combination allows for clean, readable code where data is sequentially piped through manipulation functions.

The following basic syntax demonstrates how these packages are loaded and utilized to calculate a 3-period rolling average on a generic data column named `values`:

```
library(dplyr)
```

```
library(zoo)
```

```
#calculate 3-day rolling average
df %>%
  mutate(rolling_avg = rollmean(values, k=3, fill=NA, align='right'))
```

This code snippet is designed to calculate a 3-day rolling average for the column `values` within the data frame `df`. The result is stored in a new column called `rolling_avg`. The specific parameters within the `rollmean()` function are critical for defining how the calculation is performed, particularly how the average is aligned and how missing initial values are handled.

Implementing the `rollmean()` Function and Key Parameters

The `rollmean()` function is the engine behind this process, requiring careful configuration through its primary arguments. Understanding these parameters--specifically `k`, `fill`, and `align`--is essential for accurate [rolling average](#) calculation, ensuring the resulting metric accurately reflects the desired time window.

The parameter `k` determines the window size, or the number of periods, used in the averaging calculation. If `k=3`, the function averages the current period and the two preceding periods. If the underlying data represents daily sales, then `k=7` would produce a 7-day average, useful for smoothing out weekly cycles. Choosing the correct value for `k` is often an iterative process that depends entirely on the periodicity of the data and the specific trends the analyst is attempting to isolate. A smaller `k` results in a more responsive average that follows recent price movements closely, while a larger `k` produces a smoother line that better highlights long-term trends.

The `fill` and `align` parameters manage how the rolling average handles the edges of the data set. The `fill=NA` argument specifies that when there are not enough preceding data points to complete the window (which happens at the very beginning of the series), the resulting average should be marked as `NA` (Not Applicable). For a 3-day average, the first two days will contain `NA` values, as a full three-day window cannot yet be established. The `align='right'` argument dictates that the calculation should be right-aligned, meaning the calculated average is attributed to the most recent period in the window. This configuration is standard for time series analysis, as it ensures that the average is backward-looking--the Day 3 average is calculated using Day 1, Day 2, and Day 3 data.

Practical Example: Calculating a 3-Day Rolling Sales Average

To see the `rollmean()` function in action, consider a common scenario in business analytics: tracking product sales over a short period. Suppose we have a data frame in R that records the sales of a specific product across 10 consecutive days. Our objective is to generate a new column that reflects the 3-day rolling average of these sales figures, allowing us to smooth out daily

fluctuations and observe short-term performance trends.

First, we must define our sample data frame. This structure includes two columns: `day` (the index) and `sales` (the metric we wish to analyze). Note that we are using the base R function `data.frame()` to structure our raw data:

```
#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
```

We can now apply the necessary functions to calculate the rolling average. We will use the `%>%` pipe operator to feed our data frame `df` directly into the [mutate\(\)](#) function, which is responsible for creating our new column, `avg_sales3`. Within [mutate\(\)](#), we call [rollmean\(\)](#), setting `k=3` and ensuring right alignment.

```
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

```

The resulting `avg_sales3` column now shows the rolling average value of sales for the preceding 3 periods, including the current day. Notice that the first two entries are `NA`, consistent with our `fill=NA` parameter setting. For demonstration, the first calculated value of **19.66667**, corresponding to Day 3, is mathematically derived from the sales figures of the first three days: $(25 + 20 + 14) / 3 = 19.66667$. Similarly, the value for Day 4 (16.66667) is derived from the sales of Days 2, 3, and 4 $(20 + 14 + 16) / 3$.

Advanced Usage: Calculating Multiple Rolling Windows

One of the strengths of using `dplyr` combined with the `rollmean()` function is the ability to calculate several different rolling averages concurrently within a single, efficient operation. Analysts often require multiple moving averages (e.g., a short-term 3-day average and a medium-term 7-day average) to observe trends across different time scales. Shorter averages are more reactive to immediate changes, while longer averages provide a more stable view of underlying momentum. We can achieve this by simply including multiple `rollmean()` calls, separated by a comma, within the `mutate()` function.

For example, the following code demonstrates how to calculate both the 3-day and the 4-day moving averages of the sales data simultaneously. This comparative approach is valuable for assessing how the selection of the window size (**k**) impacts the smoothing effect observed in the results.

```
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
```

Upon reviewing the output, it is immediately clear that the 4-day average (`avg_sales4`) has three initial `NA` values, as expected, since it requires four periods to complete its window. Furthermore, comparing `avg_sales3` and `avg_sales4` reveals how the 4-day average tends to be slightly smoother, reacting less drastically to the single large jump in sales observed on Day 5 (27 units), illustrating the fundamental trade-off between responsiveness and stability in moving average calculations.

Further Learning and Resources

The calculation of rolling averages is a foundational skill in data science and business intelligence when working with time-series data in R. Mastering the `rollmean()` function and integrating it seamlessly with `dplyr` greatly enhances an analyst's capacity to process and interpret sequential data. The following tutorials offer additional guidance on performing other common data manipulation and visualization tasks within the R environment:

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

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

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