

Learning R: A Practical Guide to Random Number Generation with `rnorm()` and `runif()`

Authored by
Mohammed loot

November 15, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning R: A Practical Guide to Random Number Generation with `rnorm()` and `runif()`*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=2419>

In the expansive field of statistical computing and [data analysis](#), the R programming language is recognized globally as an essential environment for generating and managing [random numbers](#). At the heart of this capability lie two fundamental functions: `rnorm()` and `runif()`. These tools are critical for a wide array of computational tasks, including running complex [simulations](#), performing advanced statistical inference, and rigorously testing scientific hypotheses. Although both are designed to produce random values, they operate under fundamentally different mathematical structures, drawing samples from distinct underlying [probability distributions](#). Understanding these differences is paramount for accurate application.

This detailed article serves as a guide to clarifying the key statistical differences between `rnorm()` and `runif()`. We will systematically examine the foundational statistical principles that define each function's output, provide practical details on their required syntax, and illustrate their implementation through concrete examples in R. Developing a profound comprehension of the unique characteristics of these functions empowers users to select the optimal tool for their specific [statistical modeling](#) or data generation needs. All concepts will be reinforced using clear, step-by-step, and fully reproducible R code.

Foundational Principles: Pseudo-Randomness and Reproducibility

A crucial prerequisite for working with R's distribution functions is understanding the nature of computational [random number generation](#). Unlike true randomness observed in physical systems, digital computers are inherently deterministic machines. They cannot produce genuinely random sequences. Instead, they employ sophisticated mathematical processes to produce [pseudo-random numbers](#). These sequences are extremely complex, designed to satisfy rigorous statistical criteria for randomness, yet they originate from a deterministic algorithm. This algorithm requires an initial numerical value, commonly referred to as the "seed," to start its computation.

In the context of statistical research and scientific publication, the concept of reproducibility holds supreme importance. Reproducibility ensures that if the same analysis code is executed repeatedly, whether by the original researcher or a third party verifying the results, the exact sequence of "random" numbers used will be regenerated, guaranteeing consistent outcomes. In R, this essential consistency is guaranteed through the strategic use of the `set.seed()` function.

When `set.seed()` is called with a specific integer argument (the seed value), the pseudo-random number generator is initialized to a predefined state. Consequently, all subsequent calls to random sampling functions--including `rnorm()` and `runif()`--will yield an identical, verifiable sequence of values every single time the script is run. Adopting the practice of setting a seed is non-negotiable for establishing transparent and trustworthy scientific workflows.

Generating Data with Central Tendency: The `rnorm()` Function

The `rnorm()` function is specifically engineered within R to generate random variates that accurately follow a [normal distribution](#). This distribution, also widely known as the Gaussian distribution, is statistically ubiquitous and is often visualized as the [bell curve](#). Its immense importance in statistics and the natural sciences stems from its ability to model continuous phenomena where observations naturally gravitate toward a central value. Examples include biological measurements, random measurement errors, and various psychometric test scores.

The entire structure and location of any given [normal distribution](#) are precisely dictated by just two fundamental parameters. Firstly, the [mean](#) (represented by the symbol μ) defines the exact center of the distribution, establishing the expected average value of the dataset. Secondly, the [standard deviation](#) (σ) measures the dispersion, or spread, of the data points relative to that central [mean](#). Critically, a smaller [standard deviation](#) results in data that is tightly clustered near the [mean](#), producing a tall, peaked bell shape, while a larger value indicates greater variability.

The required syntax for utilizing the `rnorm()` function is highly intuitive, aligning directly with these statistical definitions: `rnorm(n, mean, sd)`. Here, `n` specifies the desired number of random values to be generated, `mean` sets the population average for the simulation, and `sd` defines the measure of data dispersion. If the `mean` and `sd` arguments are deliberately omitted during the function call, R automatically defaults to generating values from the standard [normal distribution](#), which is standardized with a [mean](#) of 0 and a [standard deviation](#) of 1.

Illustrating `rnorm()`: Code and Visualization

To demonstrate the practical application of `rnorm()`, we will execute a simple simulation to generate a sample containing 100 random values. For this specific scenario, we instruct R to draw values from a [normal distribution](#) defined by a central [mean](#) of 10 and a specified [standard deviation](#) of 2. This configuration is ideal for creating synthetic data that reliably models a real-world population where the average value and known variability are constants.

```
#make this example reproducible
```

```
set.seed(0)
```

```
#create vector of 100 random values from normal distribution
```

```
random_values <- rnorm(n=100, mean=10, sd=2)
```

```
#view first six values
```

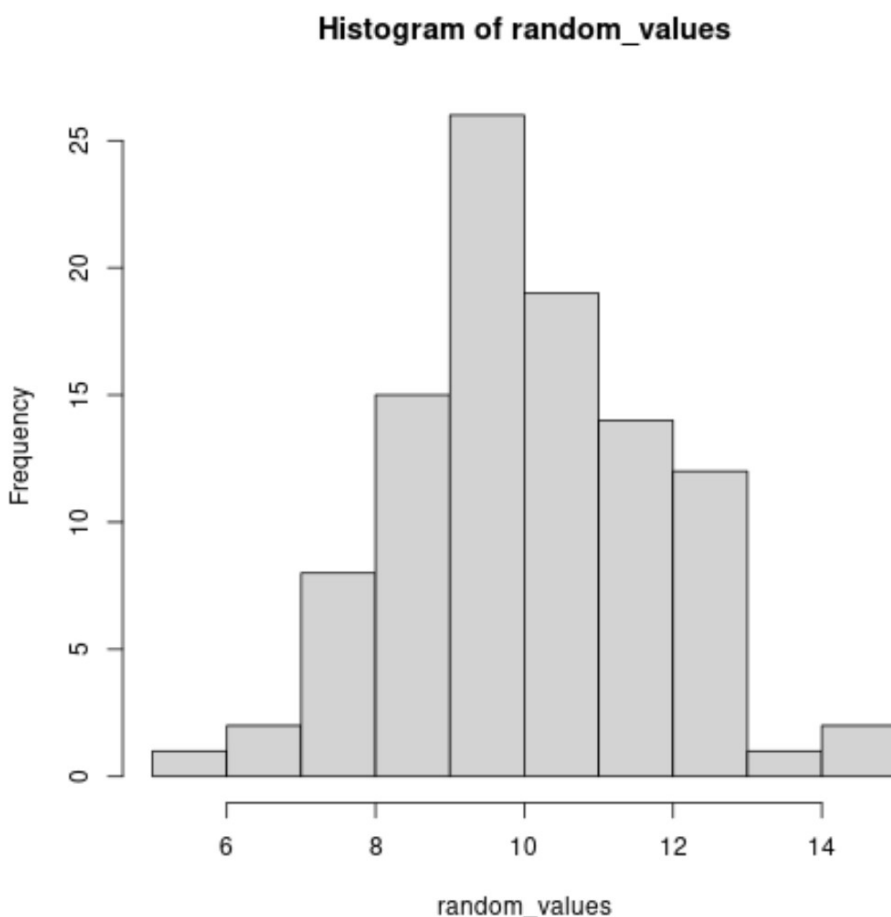
```
head(random_values)
```

```
12.525909 9.347533 12.659599 12.544859 10.829283 6.920100
```

Executing `set.seed(0)` ensures that the resulting sequence of random numbers is exactly the same every time the code is run, guaranteeing consistency. The subsequent `rnorm()` command successfully generates 100 data points that are theoretically centered at 10, with a calculated spread (variability) of 2. We then utilize the [head\(\) function](#) to quickly examine the leading elements of the new numerical vector, verifying that the output is continuous and varies around the expected center.

To visually confirm that the generated data accurately reflects the theoretical distribution shape, visual inspection is indispensable. We use the base R [hist\(\) function](#) to construct a frequency plot, or [histogram](#). This graphical representation allows for an immediate comparison between our simulated data and the expectations of a [normal distribution](#).

```
#create histogram to visualize distribution of values  
hist(random_values)
```



The resulting [histogram](#) clearly exhibits the classical characteristics of a [bell shape](#). The plotted data demonstrates clear symmetry around a distinct central peak, with the highest frequency of values clustering precisely near 10, which confirms the [mean](#) parameter provided to `rnorm()`. The

graceful decline of the bars toward the tails of the distribution graphically represents the specified [standard deviation](#), validating the function's adherence to the Normal distribution model.

Generating Data with Equal Probability: The `runif()` Function

Standing in sharp opposition to `rnorm()`, the `runif()` function is utilized to generate random values specifically drawn from a [uniform distribution](#). The critical defining characteristic of a [uniform distribution](#) is the principle of equiprobability: every single possible outcome within a clearly defined, finite interval possesses an exactly equal chance of occurring. Graphically, the probability density function of a [uniform distribution](#) forms a perfect rectangle, maintaining constant height across the designated range and immediately dropping to zero outside those limits.

Unlike the Normal distribution which relies on central tendency and variability, the [uniform distribution](#) is defined exclusively by its boundary values: a minimum (`min`) and a maximum (`max`). Crucially, the Uniform distribution lacks any inherent bias toward clustering; values are not weighted toward a central point. This neutrality makes `runif()` an essential tool for [simulations](#) requiring genuine randomness across a range, such as implementing Monte Carlo methods, generating unbiased samples, or mapping random coordinates in a defined geometric space where all points must be equally likely.

The syntax structure for `runif()` is simple and focuses purely on the range constraints: `runif(n, min, max)`. The argument `n` determines the total quantity of random values to be produced, `min` establishes the lower bound of the distribution interval, and `max` sets the upper bound. If the user chooses not to specify the `min` and `max` arguments, R automatically defaults to the standard uniform distribution, generating values that are spread uniformly across the interval between 0 and 1.

Illustrating `runif()`: Code and Visualization

To effectively demonstrate the functionality of the `runif()` function, we will follow the previous structure and generate 100 random numbers. However, this time, we mandate that these numbers must be sampled from a [uniform distribution](#) constrained by a minimum value of 5 and a maximum value of 25. This setup simulates any scenario where outcomes must be drawn randomly and without bias from a specific, limited range.

```
#make this example reproducible
```

```
set.seed(0)
```

```
#create vector of 100 random values from uniform distribution
```

```
random_values <- runif(n=100, min=5, max=25)
```


```
#view first six values  
head(random_values)
```

```
22.933944 10.310173 12.442478 16.457067 23.164156 9.033639
```

As established for reliable research, the `set.seed(0)` command guarantees the exact reproducibility of the generated sequence. The `runif()` function then produces 100 observations, all rigorously confined within the closed interval `.` A brief inspection using the [head\(\) function](#) confirms that the initial values generated strictly adhere to the defined minimum and maximum boundaries.

To visually verify the expected flat, rectangular shape--a hallmark of the absence of central tendency--we generate a [histogram](#) of our newly simulated data using the [hist\(\) function](#).

```
#create histogram to visualize distribution of values  
hist(random_values)
```



The resulting [histogram](#) compellingly confirms the characteristic rectangular profile of a [uniform distribution](#). The frequency bars across the interval from 5 to 25 maintain roughly equal heights,

providing visual evidence that all values within this range occurred with approximately the same frequency. Furthermore, the complete absence of data points outside the specified `min` and `max` arguments validates the function's precise boundary constraints.

Synthesizing the Differences: `rnorm()` vs. `runif()`

Although both `rnorm()` and `runif()` are cornerstones of [random number generation](#) in R, their practical utility is entirely dependent on correctly matching the function to the required [probability distribution](#) of the phenomenon being studied. A clear understanding of these fundamental statistical distinctions is mandatory for executing accurate [statistical modeling](#) and ensuring the validity of complex [simulations](#).

Distribution Shape and Central Tendency: `rnorm()` generates data that results in the classic symmetrical [bell shape](#), where values are highly concentrated around the mean (central tendency is inherent). In contrast, `runif()` produces a rectangular or flat distribution, ensuring that every data point is spread equally across the specified interval (no central tendency).

Controlling Parameters: The output parameters for `rnorm()` govern location and variability: `mean` and `sd` (standard deviation). The output parameters for `runif()` define strict boundaries: `min` and `max`.

Primary Use Cases: Use `rnorm()` to model natural, continuous phenomena that are expected to fluctuate around an average, such as simulating sample data for hypothesis testing in [statistical inference](#), or generating realistic noise components in data. Choose `runif()` when modeling scenarios where every outcome in a finite range must be equally probable, such as geometric random sampling, certain Monte Carlo [simulations](#), or tasks requiring non-biased random selection.

Conclusion and Further Exploration

The `rnorm()` and `runif()` functions are indispensable tools within R's expansive framework for [random number generation](#). Each function serves a precise purpose dictated by the mathematical properties of its underlying [probability distribution](#). Successful statistical computing requires not only the technical skill to use these functions but also the conceptual understanding of whether the data should exhibit central clustering (Normal distribution via `rnorm()`) or uniform randomness across a range (Uniform distribution via `runif()`).

By diligently applying the principles of reproducibility--specifically by utilizing `set.seed()`--and accurately parameterizing the distributions, data scientists can ensure their data simulations are robust, transparent, and capable of supporting highly reliable statistical analyses within the R environment. Mastering these core concepts is the first step toward advanced statistical modeling.

To further expand your proficiency in R programming and statistical analysis techniques, we encourage you to explore additional resources that cover related common data manipulation and modeling tasks: