

Perform a Shapiro-Wilk Test in R (With Examples)

Authored by
Mohammed looti

November 7, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *Perform a Shapiro-Wilk Test in R (With Examples)*.
PSYCHOLOGICAL STATISTICS. Retrieved from
<https://statistics.arabpsychology.com/?p=12041>

The [Shapiro-Wilk test](#) stands as one of the most powerful and frequently utilized statistical procedures for assessing [normality](#). Its core function is to rigorously determine whether an observed set of data points plausibly originates from a population that adheres to a [normal distribution](#), often referred to as a Gaussian distribution. This test is crucial for ensuring the foundational assumptions required by many advanced analytical techniques are met before proceeding with inference.

This assessment is fundamentally critical because the assumption of normally distributed data underpins numerous common [parametric statistical tests](#). Procedures such as t-tests, Analysis of Variance (ANOVA), and linear regression models rely heavily on this assumption to produce valid, reliable, and unbiased results. If the data significantly deviates from normality, the calculated p-values and confidence intervals from these parametric methods may be inaccurate, potentially leading to erroneous conclusions and flawed research outcomes. Therefore, verifying normality using the Shapiro-Wilk test is an essential prerequisite in any robust statistical workflow.

Fortunately, conducting the Shapiro-Wilk test within the [R](#) environment is exceptionally straightforward, thanks to a built-in function that simplifies the entire process. R's statistical capabilities make it an ideal platform for implementing and interpreting this test efficiently.

Implementing the Shapiro-Wilk Test in R

The specific function dedicated to this task in R's base package is named `shapiro.test()`. This function is designed to handle the complex calculations of the test statistic (W) quickly and return the corresponding probability value necessary for hypothesis testing. Researchers and analysts can integrate this function seamlessly into their data analysis pipelines with minimal coding effort.

The syntax for calling the function is elegantly simple, requiring just a single mandatory argument:

shapiro.test(x)

This function requires only one primary parameter to execute the analysis:

x: This input must be a numeric vector, containing all the data values from the sample that are to be tested for distributional normality.

When executed, the function returns a formatted output containing two crucial pieces of information: the test statistic W , which measures how closely the sample distribution resembles a normal distribution, and its corresponding [p-value](#). The interpretation of the results relies entirely on comparing the p-value to a pre-established significance level, conventionally denoted as α (typically set at 0.05).

Interpreting Results and Sample Constraints

The fundamental principle of interpreting the Shapiro-Wilk test lies in the context of the [null hypothesis](#) (H_0). For this specific test, the null hypothesis asserts that the data sample was drawn from a normally distributed population. The alternative hypothesis (H_A) posits that the data is not normally distributed.

The decision rule is straightforward: If the calculated p-value is less than or equal to the chosen significance level ($\alpha=0.05$), we reject the null hypothesis. Rejecting H_0 means there is sufficient statistical evidence to conclude that the sample data does not follow a normal distribution. Conversely, if the p-value is greater than α , we fail to reject the null hypothesis, suggesting the data is consistent with having been drawn from a normally distributed population. It is critical to note that failing to reject H_0 does not definitively prove normality, but rather indicates a lack of evidence of non-normality.

Important Procedural Note: Users of the standard `shapiro.test()` implementation in R must be aware of an inherent constraint regarding sample size. This function is designed to operate only on samples containing between 3 and 5,000 observations. For datasets falling outside this range--either extremely small or very large--alternative normality tests (such as the Kolmogorov-Smirnov test or Anderson-Darling test) must be employed, or specialized packages should be utilized to handle the computation.

The following practical examples demonstrate how to apply and interpret the `shapiro.test()` function effectively within the R environment, first using data intentionally generated from a normal population, and subsequently using data generated from a known non-normal population.

Example 1: Testing Data Known to Be Normal

In this first illustrative scenario, our objective is to simulate a dataset that is explicitly known to be normally distributed and then apply the Shapiro-Wilk test to confirm the expected outcome. We will generate a moderate sample size of $n=100$ observations to ensure we meet the operational requirements of the test, providing a clear baseline for interpretation.

The code snippet provided below utilizes the powerful R function `rnorm()`, which is specifically designed to create random values drawn from a standard [normal distribution](#) (with default parameters of mean = 0 and standard deviation = 1). We also set a seed to ensure that our results are perfectly reproducible across different computational sessions, a standard practice in transparent data analysis.

```
# Ensure reproducibility of the random sample  
set.seed(0)
```

```
# Create dataset of 100 random values generated from a normal distribution
```

```
data <- rnorm(100)
```

```
# Perform Shapiro-Wilk test for normality
```

```
shapiro.test(data)
```

```
Shapiro-Wilk normality test
```

```
data: data
```

```
W = 0.98957, p-value = 0.6303
```


Upon executing the test, the resulting output yields a p-value of **0.6303**. This value is significantly higher than our standard significance threshold ($\alpha = 0.05$). Consequently, based on the decision rule, we do not have sufficient evidence to reject the null hypothesis of normality. The statistical conclusion is that the sample data provides no meaningful evidence to suggest that it deviates from a normally distributed population. This result strongly confirms the anticipated outcome, as the data was deliberately constructed to follow a Gaussian distribution.

To provide a robust confirmation of this statistical finding, it is always recommended to visually inspect the data distribution. A histogram is the most common graphical tool for this purpose, allowing us to examine the shape, symmetry, and peakedness of the sample.

Related: [A Guide to dnorm, pnorm, qnorm, and rnorm in R](#)

We generate the histogram using the following R command:

```
hist(data, col='steelblue')
```



The resulting visual distribution clearly displays the characteristic "bell-shape" associated with normal data, exhibiting a distinct peak near the center and symmetrical tails on both sides. This visual evidence perfectly corroborates the high p-value obtained from the Shapiro-Wilk test, confirming that the data is consistent with the assumption of normality.

Example 2: Testing Data Known to Be Non-Normal

In sharp contrast to the previous demonstration, we now investigate a scenario where the sample data is derived from a known non-normal source. We maintain the same sample size of $n=100$, but this time the values are randomly generated from a [Poisson distribution](#). The Poisson distribution is typically used for modeling discrete count data and often exhibits significant skewness, particularly when the rate parameter (λ) is small, thus violating the assumptions of continuous normality.

The R code below employs the `rpois()` function to generate this count-based data, setting the rate parameter λ to 3. Since the Poisson distribution deals with counts (non-negative integers) and has inherent skewness, we strongly anticipate that the Shapiro-Wilk test will lead to the rejection of the null hypothesis.

```
# Ensure reproducibility of the random sample  
set.seed(0)
```

```
# Create dataset of 100 random values generated from a Poisson distribution
```

```
data <- rpois(n=100, lambda=3)
```

```
# Perform Shapiro-Wilk test for normality
```

```
shapiro.test(data)
```

Shapiro-Wilk normality test

```
data: data
```

```
W = 0.94397, p-value = 0.0003393
```


The output of this execution reveals an exceptionally low p-value of **0.0003393**. Given that this value is far below our significance threshold ($\alpha = 0.05$), we are compelled to reject the null hypothesis of normality. This result provides conclusive statistical evidence that the sample data does *not* originate from a population that is normally distributed, confirming the anticipated non-normal nature of the Poisson-generated counts.

This low p-value is entirely consistent with our data generation method. The use of `rpois()` produces data that is inherently discrete and often skewed, thereby violating the critical assumptions required for the Shapiro-Wilk test to suggest normality. The statistical result confirms the necessity of treating this data using non-parametric methods or applying appropriate transformations before using parametric tests.

Related: [A Guide to dpois, ppois, qpois, and rpois in R](#)

As before, we proceed with a visual check using a histogram to ensure that the graphical representation aligns with the statistical test conclusion:

```
hist(data, col='coral2')
```



The histogram visually confirms the non-normality of the sample. It clearly displays a right-skewed distribution, where the bulk of the data is concentrated on the left, with a long tail extending toward the right. This shape fundamentally lacks the symmetry and centralized peak characteristic of a normal distribution. Both the visual evidence from the histogram and the low p-value from the Shapiro-Wilk test lead to the unified conclusion that the sample is non-normal.

Strategies for Addressing Non-Normal Data

When a dataset fails the Shapiro-Wilk test, indicating a significant departure from normality, statisticians must choose a remedial approach. One of the most common and powerful remedies is the application of mathematical [data transformations](#) to the response variable. The goal of these transformations is to change the scale of the variable in a way that reduces skewness and stabilizes variance, thereby making the distribution closer to normal and satisfying the necessary assumptions for subsequent parametric tests.

If transformations do not succeed, the analyst may resort to utilizing non-parametric alternatives, which do not require the assumption of normality (e.g., Mann-Whitney U test instead of a t-test). However, transformation is often the first line of defense. The following are three common mathematical transformations used to mitigate skewness and improve the distributional properties of data:

Log Transformation: This technique involves transforming the response variable y to $\log(y)$. It is highly effective in dealing with data that is strongly right-skewed, particularly when the data consists of positive values and the variation increases with the mean. The logarithmic scale compresses large values more than small values.

Square Root Transformation: This method involves transforming the response variable y to \sqrt{y} . It is a milder transformation than the log transformation and is frequently recommended for count data (such as Poisson data) where the variance is often proportional to the mean, helping to stabilize that relationship while reducing skewness.

Cube Root Transformation: The transformation is achieved by converting the response variable y to $y^{1/3}$. This method is even less aggressive than the square root transformation and offers the advantage of being applicable to data that includes negative values, whereas log and square root transformations are typically restricted to non-negative data.

By meticulously implementing and testing these transformations, the distribution of the response variable can often be adjusted to become more symmetrical and approximately normally distributed, thereby validating the use of powerful parametric statistical methods that were previously unsuitable.

For a practical, step-by-step guide detailing how to implement these specific data transformations within the R environment, readers are encouraged to consult [this detailed tutorial](#).

Additional Resources for Distributional Testing

While the Shapiro-Wilk test is a gold standard for assessing normality, it is only one of several crucial methods available to statisticians. Understanding other tests and their implementations across various platforms is vital for a comprehensive grasp of distributional analysis.

For those interested in exploring alternative statistical procedures for assessing whether a dataset adheres to a specific theoretical distribution, the following resources offer practical guides and comparisons:

[How to Conduct an Anderson-Darling Test in R](#)

[How to Perform a Shapiro-Wilk Test in Python](#)