

Perform Multivariate Normality Tests in R

Authored by
Mohammed loot

November 7, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Perform Multivariate Normality Tests in R*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=12113>

The Necessity of Multivariate Normality Testing

In the pursuit of reliable quantitative research, the assumption of normality is foundational. When conducting rigorous [statistical hypothesis testing](#), researchers must first ascertain whether their data aligns with a normal distribution. For datasets involving only a single dependent variable, this process is straightforward, relying on standard [normality tests](#). Diagnostic tools frequently employed in this context include visualizing the data via a [Q-Q plot](#) or implementing formal procedures such as the Anderson-Darling Test or the Jarque-Bera Test.

However, statistical methodology becomes significantly more complex when multiple dependent variables are analyzed simultaneously. Advanced techniques like [Multivariate Analysis of Variance \(MANOVA\)](#), canonical correlation analysis, or structural equation modeling require a much stronger prerequisite: the data must adhere to a **multivariate normal distribution**. This means that the variables are not merely individually normal, but that their joint distribution conforms to the multivariate bell curve.

The failure to satisfy the assumption of [multivariate normality](#) can severely compromise the reliability and validity of inferential statistics derived from these complex models. Because the interpretation of p-values and confidence intervals hinges on this assumption, verifying multivariate normality is a critical, non-negotiable step in the data preparation phase for multivariate statistical analysis.

This guide is dedicated to mastering robust multivariate normality assessment within the powerful **R statistical programming environment**. We will provide detailed implementations and interpretations for two highly respected statistical tools designed for this purpose: **Mardia's Test** and the **Energy Test**. We will also introduce the associated concepts of multivariate [Skewness](#) and [Kurtosis](#), which form the basis of moment-based testing.

Mardia's Test in R: Evaluating Moments and Distribution Shape

Mardia's Test stands as a classical and widely accepted statistical procedure for assessing whether a set of variables collectively follows a multivariate normal distribution. Its strength lies in its ability to evaluate the multivariate extensions of two crucial distributional moments: skewness (a measure of symmetry) and kurtosis (a measure of tailedness). By analyzing these two aspects, the test provides distinct statistics that must be considered together to draw a comprehensive conclusion regarding the overall distribution shape.

The theoretical framework for Mardia's Test utilizes the standard logic of statistical hypothesis testing. The objective is to determine whether there is sufficient evidence to reject the baseline assumption of normality:

H0 (Null Hypothesis): The observed variables follow a **multivariate normal distribution**.

Ha (Alternative Hypothesis): The observed variables *do not* follow a multivariate normal distribution.

To execute Mardia's test effectively in R, we rely on the specialized functions provided by the **QuantPsyc** package. The following code snippet demonstrates the process, starting with the generation of a synthetic dataset. For illustrative purposes, we intentionally generate three variables (x1, x2, x3) from standard normal distributions, expecting this dataset to pass the normality test, thus simulating an ideal scenario.

library(QuantPsyc)

```
# Create a synthetic dataset with three variables (x1, x2, x3)
set.seed(0)
```

```
data <- data.frame(x1 = rnorm(50),
  x2 = rnorm(50),
  x3 = rnorm(50))
```

```
# Perform the Multivariate Normality Test using the mult.norm() function
mult.norm(data)$mult.test
```

Beta-hat kappa p-val

Skewness 1.630474 13.5872843 0.1926626

Kurtosis 13.895364 -0.7130395 0.4758213

Interpreting the Mardia's Test Output

The `mult.norm()` function, sourced from the `QuantPsyc` package, provides a concise table summarizing the results of Mardia's test, detailing statistics related to both multivariate skewness and multivariate kurtosis. Interpreting this output requires close attention to the 'p-val' column for each component, using a predefined significance threshold, typically $\alpha = 0.05$.

The decision rule is simple: if the p-value associated with either the Skewness statistic or the Kurtosis statistic is less than 0.05, we reject the null hypothesis (H_0). Rejecting H_0 implies that the dataset significantly deviates from multivariate normality, either due to asymmetry (skewness) or excessively light/heavy tails (kurtosis). If both p-values are above 0.05, we conclude that there is insufficient evidence to claim a significant deviation from normality.

Analyzing the results from our synthetic example, we observe the following: the p-value for the Skewness component is approximately 0.193, and the p-value for the Kurtosis component is

approximately 0.476. Since both of these values are considerably larger than the conventional 0.05 significance level, we confidently fail to reject the null hypothesis for both components. This outcome suggests that the three variables in our dataset do not exhibit any statistically significant deviation from a multivariate normal distribution based on the moment statistics provided by Mardia's Test.

The Energy Test in R: A Robust Non-Parametric Alternative

While Mardia's Test is a cornerstone of multivariate diagnostics, it relies on the calculation of moments, which can be highly sensitive to the presence of outliers or subtle deviations in the tails of the distribution. For researchers seeking a more globally robust and flexible assessment, the **Energy Test** presents an excellent, non-parametric alternative. This powerful test determines whether a collection of variables conforms to a multivariate normal distribution by utilizing concepts based on characteristic functions and [Euclidean distances](#).

The Energy Test, implemented in R through the specialized [energy package](#), is particularly valuable in situations involving small sample sizes or when the underlying distribution is suspected to be heavy-tailed or otherwise non-standard. Despite its different mathematical foundation, the hypothesis structure remains identical to that of Mardia's Test:

H₀ (Null Hypothesis): The variables follow a multivariate normal distribution.

H_a (Alternative Hypothesis): The variables *do not* follow a multivariate normal distribution.

The following demonstration illustrates how to perform the Energy Test using the `mvnorm.etest()` function, applying it to the exact same synthetic dataset used in the previous section for direct comparison:

```
library(energy)
```

```
# create dataset (using the same seed for reproducibility)
set.seed(0)
```

```
data <- data.frame(x1 = rnorm(50),
  x2 = rnorm(50),
  x3 = rnorm(50))
```

```
# perform Multivariate normality test, specifying R for bootstrap replicates
mvnorm.etest(data, R=100)
```

Energy test of multivariate normality: estimated parameters

data: x, sample size 50, dimension 3, replicates 100

E-statistic = 0.90923, p-value = 0.31

Interpreting the Energy Test Results and Critical Considerations

The output generated by the `mvnrm.etest()` function provides two primary metrics: the E-statistic (Energy statistic) and its corresponding p-value. The **E-statistic** serves as a quantitative measure of the discrepancy between the empirical characteristic function derived from the sample data and the theoretical characteristic function of a perfect multivariate normal distribution. Logically, a higher E-statistic indicates a greater observed deviation from the assumed normality.

In our example, the calculated p-value is **0.31**. Applying the standard decision rule, since 0.31 is substantially greater than the 0.05 significance level, we once again fail to reject the null hypothesis of multivariate normality. This result is consistent with the findings from Mardia's Test, confirming that we lack sufficient statistical evidence to conclude that our three variables are not jointly multivariate normally distributed.

A vital aspect of this implementation is the parameter `R=100`. This argument dictates the number of **bootstrapped replicates** utilized to estimate the distribution of the test statistic and subsequently calculate the p-value. For robust results, particularly when working with smaller sample sizes or data of questionable quality, it is strongly recommended that researchers increase the number of replicates (e.g., setting `R=500` or `R=1000`). A greater number of replicates yields a more stable and reliable p-value, which is essential for ensuring the accuracy of the hypothesis test, especially given the non-parametric nature of the Energy Test.

It is crucial to adopt a holistic perspective when assessing normality. A statistical test merely confirms whether there is *enough evidence* to reject the null hypothesis; failing to reject H_0 does not constitute proof of perfect normality. Therefore, the most comprehensive assessment involves combining the results of formal statistical tests (like Mardia's or the Energy Test) with visual diagnostics, such as examining marginal [Q-Q plots](#) for each individual variable and examining scatterplot matrices for bivariate relationships.

Beyond Normality Testing: Outlier Detection and Diagnostics

While tests for multivariate normality are essential for validating model assumptions, they often need to be supplemented by related techniques designed to diagnose specific issues within complex multivariate datasets, most notably the identification of multivariate outliers. Outliers can severely distort the results of moment-based tests like Mardia's and impact the parameter estimates in subsequent multivariate models.

To effectively identify extreme observations in a multi-dimensional context, researchers should

utilize the **Mahalanobis distance**. This metric measures the distance of an observation from the mean center of the distribution while simultaneously accounting for the covariance structure and intercorrelations among all variables. Unlike simple Euclidean distance, the Mahalanobis distance adjusts for the shape of the data cloud, providing a standardized, robust measure of unusualness in multi-dimensional space. Identifying and treating these outliers is often a prerequisite step before performing any definitive multivariate normality test.

For those seeking to expand their knowledge of univariate testing and related diagnostic tools in R, the following resources offer excellent guidance for achieving statistical expertise:

[How to Create & Interpret a Q-Q Plot in R](#)

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

[How to Conduct a Jarque-Bera Test in R](#)

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