

Understanding the Kolmogorov-Smirnov Test: A Practical Guide with R Examples

Authored by
Mohammed Iooti

November 7, 2025

RECOMMENDED CITATION

Mohammed Iooti (2025). *Understanding the Kolmogorov-Smirnov Test: A Practical Guide with R Examples*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=11975>

The **Kolmogorov-Smirnov test** (often referenced as the [KS test](#)) is recognized as a highly versatile [non-parametric statistical tool](#) essential for assessing foundational distributional assumptions in data analysis. Its primary function is twofold: first, to determine if a given sample plausibly originates from a specific theoretical [statistical distribution](#) (the one-sample case, or goodness-of-fit), and second, to evaluate whether two independent datasets were drawn from the same underlying distribution (the two-sample case). This test operates by quantifying the maximum absolute difference between the empirical cumulative distribution function (ECDF) of the sample(s) and the theoretical or comparative cumulative distribution function (CDF).

In the widely used [R statistical environment](#), executing both variants of the Kolmogorov-Smirnov test is handled efficiently by the core function `ks.test()`, which resides within the standard `stats` package. This function is designed for straightforward application, requiring minimal parameters while delivering precise results necessary for formal hypothesis testing.

This tutorial serves as a comprehensive guide, detailing the methodological steps required to implement the `ks.test()` function in R, and providing practical, reproducible examples for both the one-sample and two-sample scenarios, ensuring a robust understanding of this powerful statistical technique.

Understanding the Hypotheses and Mechanics of the KS Test

Before initiating any statistical test, a clear understanding of the null (H_0) and alternative (H_1) hypotheses is fundamental, as these define the interpretation of the resulting [p-value](#). The KS test is predicated on the principle of no difference (the null hypothesis), and deviations from this assumption lead to the rejection of H_0 , indicating a statistically significant distributional difference.

For the **One-Sample KS Test**, designed to check the goodness-of-fit against a specified distribution (e.g., testing if the data is truly [normally distributed](#)), the hypotheses are structured as follows:

Null Hypothesis (H₀): The observed sample data originates from the specified theoretical distribution.

Alternative Hypothesis (H₁): The observed sample data does not originate from the specified theoretical distribution.

Conversely, when applying the **Two-Sample KS Test**, the focus shifts to comparing two independent datasets to see if their underlying structures are identical:

Null Hypothesis (H₀): Both samples are drawn from the exact same underlying distribution.

Alternative Hypothesis (H₁): The two samples are drawn from different underlying distributions.

A decision rule is then applied: if the calculated p-value falls below the chosen significance level (conventionally $\alpha = 0.05$), the null hypothesis is rejected, signifying a statistically significant difference in the distributions being compared. The test statistic, D , measures the maximum distance between the relevant cumulative distribution functions (CDFs).

Implementing the One-Sample Kolmogorov-Smirnov Test

The one-sample KS test is utilized here to determine if a specific set of sample data adheres to a hypothesized theoretical distribution. For demonstration purposes, we will deliberately generate data known to follow a [Poisson distribution](#) (a discrete distribution) and then test it against the assumption of a continuous [normal distribution](#). This deliberate mismatch ensures we expect a clear rejection of the null hypothesis, illustrating the test's capability to detect distributional discrepancies.

We initiate the process by generating our sample dataset in R. The `rpois()` function is employed to create 20 observations consistent with a Poisson distribution, using a mean parameter (λ) of 5. By setting the random seed, we guarantee that the specific data generated is reproducible across different sessions, which is a standard practice in statistical programming.

```
#make this example reproducible  
seed(0)
```

```
#generate dataset of 100 values that follow a Poisson distribution with mean=5  
data <- rpois(n=20, lambda=5)
```

To execute the one-sample **Kolmogorov-Smirnov test** against the assumption of normality, we must specify the dataset and the cumulative distribution function (CDF) of the theoretical distribution. In R, the string `"pnorm"` represents the CDF of the standard normal distribution. We pass this string as the second argument to `ks.test()` to perform the goodness-of-fit assessment.

```
#perform Kolmogorov-Smirnov test  
ks.test(data, "pnorm")
```

One-sample Kolmogorov-Smirnov test

```
data: data  
D = 0.97725, p-value < 2.2e-16  
alternative hypothesis: two-sided
```

Analyzing the Results of the One-Sample Test

The output generated by R provides two essential metrics: the test statistic, **D**, and the corresponding [p-value](#). The D statistic, recorded here as **0.97725**, is the core measure of the KS test; it quantifies the maximum vertical distance observed between the sample's Empirical Cumulative Distribution Function (ECDF) and the theoretical Cumulative Distribution Function (CDF). A significantly large D value suggests a pronounced deviation from the assumed distribution model.

Crucially, the calculated p-value is reported as extremely small (less than **2.2e-16**). Given that this p-value is magnitudes smaller than the standard significance level ($\alpha = 0.05$), we must decisively reject the null hypothesis (H_0). The statistical conclusion is clear: there is overwhelming evidence to assert that the sample data does **not** conform to a [normal distribution](#).

This outcome successfully validates our experimental design. Because the sample data was intentionally synthesized using the `rpois()` function, which models a discrete [Poisson distribution](#), it is fundamentally incompatible with the continuous properties of the normal distribution. The KS test accurately identified this significant difference in distributional form.

Executing the Two-Sample Kolmogorov-Smirnov Test

The two-sample **Kolmogorov-Smirnov test** serves a different but equally vital purpose: determining if two independently drawn datasets originate from the same probability distribution, without requiring the specification of what that distribution is. This test focuses exclusively on assessing the identity versus difference between the two empirical distributions.

To illustrate a clear rejection case, we will generate two datasets from fundamentally different statistical processes: `data1` will be Poisson-distributed, and `data2` will be normally distributed. We hypothesize that the KS test will effectively detect the significant lack of homogeneity between their underlying distributions.

```
#make this example reproducible  
seed(0)
```

```
#generate two datasets  
data1 <- rpois(n=20, lambda=5)  
data2 <- rnorm(100)
```

The implementation of the two-sample test using the `ks.test()` function in the [R statistical environment](#) is straightforward. When two data vectors, `data1` and `data2`, are supplied as arguments, R automatically recognizes the request as a two-sample comparison. No theoretical

CDF needs to be specified, as the test compares the two empirical distributions directly.

#perform Kolmogorov-Smirnov test

```
ks.test(data1, data2)
```

Two-sample Kolmogorov-Smirnov test

data: data1 and data2

D = 0.99, p-value = 1.299e-14

alternative hypothesis: two-sided

Interpreting the Two-Sample Comparison Results

Upon reviewing the output, we observe the test statistic **D** is **0.99**. In the context of comparing two samples, **D** represents the greatest absolute difference observed between the Empirical Cumulative Distribution Functions (ECDFs) of the two datasets. A **D** value approaching 1 (the maximum possible value) is a strong indicator of a massive disparity between the two underlying distributions.

The accompanying p-value is calculated as **1.299e-14**. Since this value is significantly lower than the standard significance threshold of 0.05, we confidently reject the null hypothesis. This provides strong evidence to conclude that the two sample datasets, one derived from a Poisson process and the other from a normal process, do **not** originate from the same underlying [statistical distribution](#). The high **D** value visually confirms the separation between the discrete Poisson data and the continuous normal data.

Critical Limitations of the Kolmogorov-Smirnov Test

While the **Kolmogorov-Smirnov test** offers a robust, non-parametric method for comparisons, practitioners must be aware of specific limitations, particularly concerning the one-sample goodness-of-fit application. Recognizing these nuances ensures appropriate use and interpretation of the results.

A major concern arises when the parameters of the hypothesized theoretical distribution (such as the mean or standard deviation for a normal distribution) are estimated directly from the sample data being tested. The standard KS test calculations assume that the theoretical Cumulative Distribution Function (CDF) and its parameters are completely known beforehand. When parameters are estimated from the sample, the test becomes "conservative"--meaning the true p-value is often smaller than the one reported by the standard `ks.test()` function. This conservatism results in reduced power, making the test less sensitive to genuine distributional deviations.

Furthermore, the KS test is known to exhibit higher sensitivity to discrepancies occurring around the center (median) of the distribution compared to deviations found in the extreme tails. For scenarios demanding high sensitivity in the tails, alternative goodness-of-fit tests, such as the Shapiro-Wilk test (specifically for [normality](#)) or the Anderson-Darling test, often provide greater statistical power. Therefore, a complete assessment of distributional fit should ideally combine the formal results of the KS test with visual diagnostics, such as Q-Q plots.

Complementary Resources for Distributional Analysis in R

For analysts seeking alternative or complementary methods to assess distributional assumptions within the [R statistical environment](#), exploring other statistical tests can provide a more comprehensive view of data fitness. These resources offer guidance on powerful, specialized tests often used alongside or instead of the KS test:

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

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

[How to Perform Multivariate Normality Tests in R](#)