

Learning to Test for Normality in Python: A Guide to 4 Methods

Authored by
Mohammed loot

October 29, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Test for Normality in Python: A Guide to 4 Methods*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=5314>

In the rigorous field of [statistics](#), a vast majority of [statistical tests](#), known as parametric tests, rely on a crucial [assumption](#): that the underlying data are sampled from a [normal distribution](#). This concept, often visualized as the bell curve, is fundamental. The validity and reliability of popular analyses--ranging from the simple [t-test](#) to sophisticated techniques like [ANOVA](#) and [linear regression](#) models--are predicated on meeting this requirement. If the data significantly deviates from normality, the resulting conclusions drawn from the analysis may be seriously flawed or misleading.

Fortunately for modern data scientists, the powerful capabilities of [Python](#), coupled with its extensive [data science](#) libraries (such as SciPy and Statsmodels), offer a robust toolkit for assessing whether a given dataset approximates this essential distribution. These methods range from quick, intuitive graphical checks to mathematically rigorous [statistical tests](#), ensuring comprehensive data validation before proceeding with inferential analysis.

This expert guide details four widely adopted and highly effective techniques used to check the assumption of normality in [Python](#), providing practical implementation examples for each:

- 1. (Visual Method) Creating a Histogram:** This technique provides an immediate, graphical assessment of the data's overall shape, symmetry, and central tendency.
- 2. (Visual Method) Generating a Q-Q plot:** A more diagnostic visual tool that rigorously compares the quantiles of your sample data against the quantiles of a theoretical [normal distribution](#).
- 3. (Formal Statistical Test) Performing a Shapiro-Wilk Test:** Recognized as one of the most powerful inferential [statistical tests](#) specifically designed to evaluate normality, particularly useful for smaller samples.
- 4. (Formal Statistical Test) Executing a Kolmogorov-Smirnov Test:** A versatile formal test that determines if a sample's distribution differs significantly from a specified reference distribution, such as the [normal distribution](#).

A robust assessment of normality often involves combining both visual inspection and formal statistical hypothesis testing. The following sections will provide clear demonstrations of how to implement, interpret, and draw reliable conclusions from each of these four essential methods using practical [Python](#) code.

The Crucial Role of the Normal Distribution in Data Analysis

The [normal distribution](#), often synonymously called the Gaussian distribution or the bell curve, is the cornerstone of classical [statistics](#). It is defined by its perfect symmetry, where values cluster tightly around a central point, and the frequency of observations decreases systematically as values move away from the mean. A key characteristic is that the mean, median, and mode are all situated at the exact center of the distribution. This theoretical model is pervasive in nature and social sciences, making it the default expectation for many observed variables.

The assumption of normality is far more than a technical detail; it directly impacts the mathematical foundation of many [statistical tests](#). Parametric tests, such as the [t-tests](#) and [ANOVA](#), derive their formulas based on the known properties of the normal curve to calculate accurate probabilities and confidence intervals. When the input data severely violates this assumption, the test statistics become unreliable, potentially leading to inflated or depressed [p-values](#) and resulting in erroneous conclusions about the relationships or differences within the data.

Therefore, before embarking on any extensive [statistical analysis](#), it is an imperative step in the data preparation workflow to rigorously check this condition. The subsequent methods outlined in this article provide the necessary tools in [Python](#) to conduct these checks effectively, safeguarding the integrity and trustworthiness of your statistical findings.

Initial Assessment: Visual Techniques for Normality

Visual methods serve as the critical first layer of [data analysis](#) when assessing normality. While they cannot provide the definitive quantitative answer of a formal statistical test, they offer invaluable, intuitive insights by allowing the researcher to quickly identify blatant deviations from the characteristic bell shape, such as severe skewness, heavy tails, or the presence of multiple modes (multimodality).

These graphical techniques are fundamental to exploratory data analysis (EDA). A visual inspection can often immediately highlight issues like outliers or structural problems within the data that might necessitate special handling or suggest the use of non-parametric alternatives to standard [statistical tests](#). By plotting the data distribution in specific, standardized formats, we can observe patterns that either confirm or deny alignment with the theoretical normal model.

The two primary graphical tools utilized for this initial assessment are the [histogram](#) and the [Q-Q plot](#). Although both provide a visual summary, they each reveal distinct aspects of the data's distributional properties, making them complementary tools for a thorough visual check.

Method 1: Creating a Histogram for Visual Inspection

A [histogram](#) is the most straightforward graphical representation used to visualize the frequency distribution of continuous numerical data. It organizes observations into a series of contiguous intervals, or "bins," displaying the count or frequency of observations within each bin as vertical bars. When assessing normality, the analyst looks for the classic bell curve profile: the graph should appear symmetrical, smoothly rising to a single peak near the center, and then tapering off equally in both directions (the tails).

Any significant departure from this symmetric, bell-shaped form--such as a long tail extending heavily to the right (positive skew), a pronounced flat top (platykurtosis), or multiple distinct peaks--


serves as strong visual evidence of non-normality. The following [Python](#) implementation demonstrates how to generate a [histogram](#). For instructional clarity, we simulate data drawn specifically from a [log-normal distribution](#), which is inherently positively skewed, ensuring a clear visual example of non-normality. We leverage [NumPy](#) for numerical array handling, [SciPy](#) for distribution functions, and [Matplotlib](#) for rendering the plot.

```
import math
import numpy as np
from scipy.stats import lognorm
import matplotlib.pyplot as plt

#make this example reproducible
np.random.seed(1)

#generate dataset that contains 1000 log-normal distributed values
lognorm_dataset = lognorm.rvs(s=.5, scale=math.exp(1), size=1000)

#create histogram to visualize values in dataset
plt.hist(lognorm_dataset, edgecolor='black', bins=20)
```



The resulting [histogram](#) confirms that the simulated data is not normally distributed. Instead of the expected symmetric bell shape, the distribution is acutely skewed to the right (positively skewed),

characterized by a sharp rise and a long, drawn-out tail towards the higher values. This visual asymmetry definitively indicates that the data violates the assumption of normality, consistent with its origin as a [log-normal distribution](#).

Method 2: Interpreting Quantile-Quantile (Q-Q) Plots

The [Q-Q plot](#), or Quantile-Quantile plot, is a sophisticated graphical method that provides a more precise diagnostic assessment than a simple histogram. It works by plotting the observed data quantiles against the theoretical quantiles expected if the data were drawn from a specific distribution, typically the [normal distribution](#). If the data perfectly adheres to the normal model, the plotted points will form a single, straight diagonal line.

Interpreting the [Q-Q plot](#) involves examining deviations from this reference line. S-shaped curves often suggest heavy or light tails (kurtosis issues), while curved tails or points that peel away from the line at the extremes indicate significant skewness or the presence of outliers. The more closely the points hug the straight line, the stronger the evidence for normality.

The [Python](#) code below utilizes the specialized `qqplot` function available in the excellent [Statsmodels](#) library. We apply this function to the same positively skewed [log-normal distribution](#) dataset used previously. The inclusion of a 45-degree reference line is crucial for visual comparison.


```
import math
import numpy as np
from scipy.stats import lognorm
import statsmodels.api as sm
import matplotlib.pyplot as plt

#make this example reproducible
np.random.seed(1)

#generate dataset that contains 1000 log-normal distributed values
lognorm_dataset = lognorm.rvs(s=.5, scale=math.exp(1), size=1000)

#create Q-Q plot with 45-degree line added to plot
fig = sm.qqplot(lognorm_dataset, line='45')

plt.show()
```



As clearly demonstrated by the plot, the data points deviate significantly from the straight red reference line. Instead of lying along the 45-degree axis, the points form a distinct upward curve, particularly noticeable in the upper quantiles. This severe and systematic departure confirms the conclusion drawn from the [histogram](#): the dataset is not normally distributed. The [Q-Q plot](#) provides definitive visual proof that the sample quantiles do not match the expected theoretical quantiles of a [normal distribution](#).

Formal Statistical Hypothesis Testing for Normality

While visual methods are excellent for preliminary assessments, formal [statistical tests](#) are required to provide an objective, quantitative measure of non-normality. These tests operate under the framework of hypothesis testing. Specifically, they establish a [null hypothesis](#) (H_0) which always states that the data sample is drawn from a [normal distribution](#).

The test then calculates a test statistic and a corresponding [p-value](#). The decision rule is standardized: if the calculated [p-value](#) falls below a predefined significance level (alpha, α , typically 0.05), we reject H_0 , concluding there is sufficient statistical evidence that the data is not normal. Conversely, if the [p-value](#) is greater than α , we fail to reject H_0 , meaning there is insufficient evidence to claim non-normality. It is important to note that failing to reject the null hypothesis should not be interpreted as proving normality, but rather as concluding that the test did not find a significant difference from the normal distribution.

The two most widely used and robust formal statistical checks for normality are the [Shapiro-Wilk](#)

[Test](#) and the [Kolmogorov-Smirnov Test](#) (K-S test). The choice between them often depends on the size of the dataset, as each test exhibits different sensitivities to sample size and distribution deviations.

Method 3: Executing the Shapiro-Wilk Test

The [Shapiro-Wilk Test](#) is widely regarded as the most effective and powerful [statistical test](#) for evaluating the normality of a dataset, especially when dealing with small to moderately sized samples (often up to $N=5000$). The test calculates a W statistic that evaluates how closely the sample data correlates with the expected values of a normal distribution. The [null hypothesis](#) (H_0) posits that the sample comes from a [normal distribution](#).

To implement the [Shapiro-Wilk Test](#) in [Python](#), we utilize the dedicated `shapiro` function available within the `scipy.stats` module. The following code applies this powerful test to the same 1000 observations generated from the skewed [log-normal distribution](#).

```
import math
import numpy as np
from scipy.stats import shapiro
from scipy.stats import lognorm

#make this example reproducible
np.random.seed(1)

#generate dataset that contains 1000 log-normal distributed values
lognorm_dataset = lognorm.rvs(s=.5, scale=math.exp(1), size=1000)

#perform Shapiro-Wilk test for normality
shapiro(lognorm_dataset)

ShapiroResult(statistic=0.8573324680328369, pvalue=3.880663073872444e-29)
```

The test output reveals a W statistic of **0.857** and a corresponding [p-value](#) of **3.88e-29**. This [p-value](#) is astronomically small, effectively zero. When compared against the standard significance level ($\alpha = 0.05$), we find that $3.88e-29$ is significantly less than 0.05. This finding compels us to **reject the null hypothesis**. The [Shapiro-Wilk test](#) therefore provides strong statistical evidence that the sample data is definitively not drawn from a [normal distribution](#), reinforcing the conclusions drawn from our visual assessments.

Method 4: Applying the Kolmogorov-Smirnov Test

The [Kolmogorov-Smirnov Test](#) (K-S test) is a non-parametric [statistical test](#) used to evaluate whether a sample distribution is significantly different from a theoretical reference distribution (like the normal distribution). The K-S test operates by calculating the maximum distance between the Empirical Cumulative Distribution Function (ECDF) of the sample data and the Theoretical Cumulative Distribution Function (CDF) of the reference distribution. The [null hypothesis](#) (H0) assumes that the sample data aligns with the specified theoretical distribution. The K-S test is generally considered a good all-purpose test, often preferred for very large sample sizes where the [Shapiro-Wilk test](#) may become overly sensitive.

In [Python](#), the K-S test is performed using the `kstest` function from the `scipy.stats` module. When applying this test for normality, we must specify the theoretical distribution as 'norm' (normal). We again apply this procedure to our simulated, non-normal [log-normal distribution](#) dataset to complete the battery of tests.

```
import math
import numpy as np
from scipy.stats import kstest
from scipy.stats import lognorm

#make this example reproducible
np.random.seed(1)

#generate dataset that contains 1000 log-normal distributed values
lognorm_dataset = lognorm.rvs(s=.5, scale=math.exp(1), size=1000)

#perform Kolmogorov-Smirnov test for normality
kstest(lognorm_dataset, 'norm')

KstestResult(statistic=0.84125708308077, pvalue=0.0)
```

The resulting output from the [Kolmogorov-Smirnov Test](#) shows a test statistic of **0.841** and a [p-value](#) of **0.0**. Since the [p-value](#) (0.0) is drastically lower than our conventional significance threshold ($\alpha = 0.05$), we must firmly reject the [null hypothesis](#). This rejection confirms that the sample data's distribution is significantly different from a theoretical [normal distribution](#). The findings across all four methods--both visual and formal--are perfectly consistent, collectively demonstrating the non-normal nature of the simulated data.

Strategies for Addressing Non-Normal Data

Identifying non-normality in a dataset is a frequent outcome in real-world [data analysis](#), often stemming from intrinsic data properties (e.g., highly skewed distributions like reaction times or economic data), measurement limitations, or the presence of significant outliers. When the normality assumption is violated, simply proceeding with [parametric tests](#) risks invalidating the statistical inference.

One of the most practical and widespread solutions for non-normal data is [data transformation](#). This involves applying a specific mathematical function to every data point with the goal of altering the distribution's shape, making it more symmetric and bell-shaped, thereby approximating a [normal distribution](#). Common transformations are selected based on the nature and direction of the skew:

- 1. [Log Transformation](#):** This strong transformation converts values x to $\log(x)$ (using natural or base-10 logarithms). It is highly effective for reducing pronounced positive skewness, such as that observed in the [log-normal distribution](#) examples used in this guide.
- 2. [Square Root Transformation](#):** A milder transformation where values x are converted to \sqrt{x} . It is useful for data that is moderately positively skewed and can also help stabilize variance.
- 3. [Cube Root Transformation](#):** Values are transformed from x to $x^{1/3}$. Being stronger than the square root, this transformation is useful for more severe skewness and has the advantage of being applicable to both positive and negative values.

Alternatively, if transformation fails or is inappropriate, analysts should consider using [non-parametric statistical tests](#) (e.g., Mann-Whitney U test, Kruskal-Wallis H test), which do not require the assumption of normality. Regardless of the chosen strategy, it is essential to re-check the normality assumption after any data transformation using the visual and formal methods detailed above. For detailed code examples on implementing these mathematical transformations in [Python](#), consult [this comprehensive tutorial](#).