

How to Easily Plot a Normal Distribution in Seaborn

Authored by
stats writer

December 2, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Plot a Normal Distribution in Seaborn*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=103750>

The ability to visualize statistical concepts is fundamental to effective data analysis. When dealing with continuous data, the normal distribution, often referred to as the Gaussian distribution or the bell curve, is one of the most critical concepts to understand and plot. It describes how values of a variable are distributed around the mean, providing insight into the underlying probability structure of the dataset. For Python users, the Seaborn library stands out as an exceptional tool for generating informative and aesthetically pleasing statistical graphics with minimal code. This article serves as an expert guide on how to leverage Seaborn, a high-level interface built on top of Matplotlib, to accurately and clearly visualize the normal distribution of your data.

Plotting a distribution involves representing the frequency or probability density of various data points. The goal is not just to draw a curve, but to visually communicate key characteristics such as the central tendency, spread, and symmetry of the dataset. While older Seaborn tutorials might reference the now-deprecated `distplot()` function, modern practice recommends the highly versatile `displot()` function. This function integrates histogram plotting, KDE (Kernel Density Estimation) estimation, and rug plots into a single, cohesive figure-level interface, simplifying the process of statistical data visualization considerably.

To effectively plot a distribution using this powerful data visualization library in Python, we focus on the various configurations of `sns.displot()`. The function accepts the data series as its primary argument and then uses optional parameters like `kind` and `kde` to control the output. Understanding these options allows us to choose the most appropriate visual representation for communicating the characteristics of the normal distribution present in our data. We will explore three distinct, yet complementary, methods for achieving this goal, each offering a unique perspective on the data's shape and density.

The Primary Tool: Utilizing Seaborn's `displot()` Function

The modern approach to visualizing univariate distributions in Seaborn centers entirely on the `displot()` function. Unlike its predecessor, which was an axes-level function, `displot()` is a figure-level function. This distinction is important because it means the function manages the entire figure, which allows for easier creation of multi-panel figures and improved consistency across different types of plots. When plotting a normal distribution, we are primarily interested in showing the density of observations, and `displot()` offers several built-in mechanisms to accomplish this, including histograms, KDE, and Empirical Cumulative Distribution Functions (ECDF).

By default, when you call `sns.displot(x)` where `x` is your data array, Seaborn intelligently defaults to displaying a histogram. A histogram partitions the data space into discrete bins and shows the count of observations falling into each bin, giving a fundamental visual representation of the distribution's shape. For a dataset generated from a normal distribution, the resulting bars will generally form the familiar bell shape, provided enough data points are used and bin selection is

appropriate. This is often the first step in exploratory data analysis.

If we want to represent the distribution using a smooth curve, we must adjust the parameters. The `kind` parameter controls the type of plot drawn. Setting `kind='kde'` tells Seaborn to calculate and plot the KDE, which provides a non-parametric estimate of the probability density function (PDF) of the variable. Alternatively, if we desire both the raw counts (histogram) and the smooth estimate (KDE) overlaid, we can use the boolean parameter `kde=True` within the default histogram call. These three fundamental configurations form the basis of visualizing a normal distribution effectively.

Core Methods for Visualization

We will focus on three practical methods to plot the distribution of a variable, moving from the most basic count-based representation to a sophisticated density curve. Each method serves a slightly different analytical purpose, but all are vital for confirming whether a dataset adheres closely to the assumptions of a normal distribution. These methods utilize the standard `displot()` function, differentiating only by a single parameter setting.

The following methods are available for plotting the distribution using the Seaborn library:

Method 1: Plot Normal Distribution Histogram

This method focuses on visualizing the frequency count of the data partitioned into discrete intervals, using the default behavior of the `displot` function.

`sns.displot(x)`

Method 2: Plot Normal Distribution Curve (KDE)

This approach generates a smooth, estimated probability density function, removing the dependency on bin size inherent in histograms.

`sns.displot(x, kind='kde')`

Method 3: Plot Normal Distribution Histogram with Curve Overlay

This provides a powerful hybrid view, showing the raw data counts while simultaneously overlaying the estimated probability density function for context.

`sns.displot(x, kde=True)`

The subsequent sections provide detailed, runnable examples demonstrating the practical application of each of these methods, starting with the necessary step of generating synthetic data using `NumPy`.

Detailed Implementation: Generating Test Data with NumPy

Before plotting any distribution, we require a dataset. For demonstration purposes, it is standard practice to generate synthetic data that we know follows a normal distribution. We use the powerful `NumPy` library for numerical operations in Python, specifically its `random` module, to create this data. The use of `np.random.normal(size=N)` allows us to quickly generate a specified number (N) of samples drawn from a standard normal distribution (mean=0, standard deviation=1).

Crucially, reproducibility is a key aspect of good data science practice. By setting a random seed (`np.random.seed(0)`), we ensure that every time the code is executed, the "random" data generated is identical. This guarantees that the visualizations produced in these examples are consistent and verifiable. We will generate 1000 data points for our variable `x`, which should be sufficient to clearly approximate the bell-shaped curve.

The initialization steps, common across all our examples, involve importing the necessary libraries and generating the reproducible dataset:

```
import numpy as np
import seaborn as sns

# Ensure example reproducibility across sessions
np.random.seed(0)

# Create 1000 data points drawn from a standard normal distribution
x = np.random.normal(size=1000)

# The plotting steps follow this initialization.
```

Example 1: Plotting the Normal Distribution as a Histogram

The histogram is the most intuitive method for displaying the shape of a distribution. By default, `sns.displot(x)` creates a frequency histogram, dividing the range of data into bins and counting how many observations fall into each interval. This visualization is particularly useful for assessing the raw counts and understanding the underlying structure of the distribution without any smoothing or estimation techniques. When applied to our normally distributed data `x`, we expect the counts to peak centrally and taper off symmetrically towards the tails.

The key advantage of the histogram is its fidelity to the raw data; every bar height directly corresponds to the count of observations. However, a potential drawback is the sensitivity to the number of bins chosen. If too few bins are used, crucial details of the distribution's shape might be obscured (oversmoothing); conversely, too many bins can make the plot appear jagged and noisy. Seaborn provides intelligent default bin selections, but these can be customized using parameters like `bins` if fine-tuning is required for publication-quality graphics.


The following code snippet demonstrates the straightforward implementation necessary to generate this count-based representation of the distribution. It requires only the data array `x` as an input to the `displot` function, relying entirely on Seaborn's default configuration for visualization:

```
import numpy as np
import seaborn as sns

#make this example reproducible
np.random.seed(0)

#create data
x = np.random.normal(size=1000)

#create normal distribution histogram
sns.displot(x)
```



Example 2: Visualizing the Distribution using a Kernel Density Estimate (KDE) Curve

While histograms show discrete counts, the Kernel Density Estimate (KDE) provides a smooth, continuous estimate of the probability density function (PDF). This is achieved by placing a small kernel--typically a Gaussian kernel--over each observation and summing up these kernels to form the overall density curve. The result is a smooth representation of the data's distribution, which is often preferable when presenting the theoretical shape of a normal distribution.

To plot the KDE curve exclusively, we utilize the `kind` parameter within `sns.displot()` and set its value to `'kde'`. This tells Seaborn to skip the histogram calculation entirely and focus on generating the smooth density estimate. The resulting plot clearly illustrates the symmetry and peak of the normal distribution without the visual clutter of bins, making it excellent for demonstrating the theoretical properties of the dataset.


This method is particularly useful when comparing multiple distributions on the same axis or when the visual precision of a smooth curve is desired over the granularity of a histogram. However, it is important to remember that the KDE is an estimate, and its smoothness depends heavily on the bandwidth parameter (which Seaborn handles automatically by default), potentially introducing slight biases if not tuned properly. The code below illustrates how to generate this smooth density estimate:

```
import numpy as np
import seaborn as sns

#make this example reproducible
np.random.seed(0)

#create data
x = np.random.normal(size=1000)

#create normal distribution curve
sns.displot(x, kind='kde')
```



Example 3: Combining Histogram and KDE Curve for Comprehensive Analysis

Often, the most informative visualization is one that combines the benefits of both the raw data representation and the statistical estimate. By overlaying the KDE curve onto the histogram, we achieve a comprehensive view: the histogram shows the discrete frequency counts, anchoring the visualization to the actual dataset, while the smooth curve provides the estimated underlying probability density function, helping to visually confirm the fit to a normal distribution.

To achieve this combined plot using `sns.displot()`, we simply use the default histogram plot (by omitting the `kind` parameter) and explicitly set the boolean argument `kde` to `True`. Seaborn then automatically handles the scaling and axis configuration necessary to display both elements clearly on the same figure. This combined approach is highly recommended for exploratory data analysis, as it allows analysts to simultaneously assess noise (from the histogram) and trend (from the KDE).


This visualization is particularly effective for large datasets where the histogram bins naturally align well with the smooth theoretical curve. If the underlying data is truly normally distributed, the KDE curve should trace the outline formed by the top of the histogram bars closely. Any significant deviations between the curve and the histogram indicate skewness, kurtosis, or multimodality that suggest the data is not perfectly normal. The required code for this dual plot is straightforward:

```
import numpy as np
import seaborn as sns

#make this example reproducible
np.random.seed(0)

#create data
x = np.random.normal(size=1000)

#create normal distribution curve
sns.displot(x, kde=True)
```



Beyond Basic Plotting: Customization and Refinement

While the basic implementations of `sns.displot()` are powerful, Seaborn allows extensive customization to enhance the visual appeal and analytical clarity of the plot. For example, users can modify the color of the bars or the curve using the `color` parameter, adjust the number of bins in the histogram using the `bins` parameter, or control the appearance of the KDE using parameters specific to the kernel estimation process, such as `bw_adjust`.

Furthermore, since `displot()` returns a `FacetGrid` object, advanced users can access the underlying Matplotlib axes objects to perform even deeper modifications. This includes setting custom titles, adjusting axis limits, and adding annotations that highlight specific features of the

normal distribution, such as the mean or standard deviation lines. Mastering these customization options is essential for creating publication-ready figures that effectively communicate complex statistical findings.

Effective statistical data visualization relies not just on generating the plot but on ensuring it is labeled clearly and accurately. When plotting distributions, always ensure that the axes are correctly labeled to indicate what is being measured (the variable on the X-axis) and what the height represents (count or density on the Y-axis). When using the KDE curve, the Y-axis represents probability density, not raw counts, a distinction critical for accurate interpretation.

Conclusion and Next Steps in Seaborn Analysis

Visualizing the normal distribution using Seaborn's `displot()` function is a fundamental skill in statistical computing. Whether you choose the raw count representation of the histogram, the smooth theoretical estimate of the KDE curve, or a hybrid combination of the two, Seaborn provides a powerful and intuitive interface. By mastering the distinction between these three primary plotting methods, analysts can gain deeper insight into the characteristics of their datasets and confirm assumptions vital for downstream statistical modeling.

The principles demonstrated here--importing libraries, generating reproducible data with NumPy, and utilizing the flexible `displot()` function--are transferrable to plotting many other types of univariate and bivariate distributions. For those seeking to expand their knowledge, exploring other Seaborn functions for analyzing relationships between variables, such as scatter plots (`scatterplot`) or correlation heatmaps (`heatmap`), would be a logical next step.

For further learning on advanced techniques within this library, the following tutorials explain how to perform other common operations in Seaborn: