

# Learning to Create Ogive Graphs with Python: A Step-by-Step Tutorial

Authored by  
**Mohammed loot**

November 8, 2025

## RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Create Ogive Graphs with Python: A Step-by-Step Tutorial*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=12724>

The [Ogive](#), often referred to as a cumulative frequency graph, stands as an indispensable tool in statistical visualization. Its primary function is to graphically represent the running total of frequencies within a given [dataset](#). This particular visualization is exceptionally useful for rapid percentile estimation, allowing analysts to quickly ascertain how many observations fall above or below a specific data point, thereby providing crucial insight into the overall distribution characteristics. This detailed tutorial guides readers through the precise, step-by-step methodology required to generate a clean, valid [Ogive](#) using the robust programming language, [Python](#), specifically leveraging the computational power of the **NumPy** library and the advanced plotting capabilities of **Matplotlib**.

## Defining the Ogive and Its Statistical Utility

At its foundation, the [Ogive](#) is designed to visualize the **cumulative frequency** distribution. This differs fundamentally from standard frequency plots, such as a [histogram](#), which only display the frequency count contained within discrete class intervals. Instead, the Ogive plots the upper boundary of each class interval against the corresponding total number of observations accumulated up to that point. This inherent structure ensures that the resulting curve is strictly non-decreasing, beginning at zero and culminating precisely at the total number of observations in the sample.

Statistically, the Ogive serves as the graphical proxy for the empirical [cumulative distribution function \(CDF\)](#). By analyzing the slope and overall shape of this curve, researchers can quickly glean essential metrics regarding data dispersion, symmetry, and measures of central tendency. For example, locating the point on the curve that corresponds to 50% of the total frequency immediately provides a highly accurate graphical estimation of the distribution's **median**. This capacity for rapid statistical inference makes the Ogive a powerful diagnostic tool.

Successfully generating this chart within the [Python](#) environment requires a clear, two-stage process: first, calculating the discrete frequencies for defined class intervals, and second, transforming those frequencies into running totals, or [cumulative sums](#). The integrity and accuracy of the final visualization depend critically on meticulous data preparation, utilizing numerical libraries like [NumPy](#) before the final plotting commands are executed via **Matplotlib**.

## Setting Up the Python Environment and Sample Data

Generating an [Ogive](#) in a professional statistical computing environment such as [Python](#) necessitates a structured, multi-step workflow. This process relies heavily on the core functions provided by **NumPy** for efficient numerical operations and **Matplotlib** for high-quality data rendering. Our initial focus must be on establishing a robust, representative [dataset](#) upon which all subsequent frequency analysis will be accurately performed. For the purpose of this

demonstration, we generate a substantial sample of random integers.

We utilize the **NumPy** library's powerful random number generation capabilities to create a large array of observations. Specifically, the `np.random.randint()` function is employed to generate 1,000 random integer observations, constrained to fall within a specified range (0 to 10, exclusive). A crucial best practice in reproducible data science is setting the random seed using `np.random.seed(1)`. This ensures that the exact sequence of random numbers generated is consistent across different executions, which is vital for both verification and collaborative work.

Upon execution, this initialization code provides the raw data array, which serves as the fundamental input for our entire frequency calculation process. We can inspect the initial values to confirm the data structure and composition before proceeding to the complex statistical transformations required to construct the cumulative plot.

```
import numpy as np
```

```
#create array of 1,000 random integers between 0 and 10
```

```
np.random.seed(1)
```

```
data = np.random.randint(0, 10, 1000)
```

```
#view first ten values
```

```
data
```

```
array()
```

## Core Data Transformation: Calculating Cumulative Frequencies

The next essential phase involves transforming the raw data array into the specific metrics required for the cumulative plot. This multi-step transformation relies on two specialized functions provided by the [NumPy](#) library. The initial step is utilizing `np.histogram()` to automatically determine the optimal class boundaries (bins) and calculate the discrete frequency counts within each of those defined classes. While this function is most commonly associated with generating a [histogram](#), here we are interested solely in extracting its calculated frequency and boundary values.

The output of `np.histogram()` consists of two distinct arrays: the `values` array, which contains the calculated frequencies for every bin, and the `base` array, which defines the corresponding bin edges or class boundaries. It is crucial to note the structural relationship between these outputs: the `base` array will always contain one more element than the `values` array, as it must define both the starting point and the ending point for all class intervals.

The second, and arguably most critical, mathematical operation is applying `np.cumsum()` to the

`values` array. This function performs the necessary calculation of the **cumulative sum**. In the resulting array, each element represents the sum of all preceding frequencies combined with the current frequency count. This final cumulative array represents the precise Y-axis data points for our Ogive plot, indicating the total [cumulative distribution function \(CDF\)](#) value reached at the upper limit of each respective class interval.

## Visualizing the Ogive Curve with Matplotlib

Once the class boundaries (X-axis data) and the [cumulative frequency](#) array (Y-axis data) are accurately prepared, the final technical step is to render the results graphically using [Matplotlib](#). The Ogive is fundamentally a specialized line plot where points are connected, showing the progression of the running total. These connection points mark where the cumulative frequency intersects the upper boundary of the corresponding class interval. It is essential that the arrays used for the X and Y axes are meticulously aligned and of matching length before the plotting command is executed.

The core plotting instruction is `plt.plot(base, cumulative, 'ro-')`. A critical detail here is the slicing notation `base`. Since the `base` array contains N+1 bin edges (boundaries) and the `cumulative` array contains N frequencies (sums), we must intentionally exclude the very last element of the `base` array. This slicing action ensures perfect alignment, pairing the N cumulative counts with the N corresponding upper boundaries of the class intervals on the X-axis, thereby forming the correct Ogive curve.

The string argument `'ro-'` passed within the plot function serves to define the specific aesthetic style of the rendered output. This concise notation dictates the color, the marker type, and the line style, ensuring a visually clear and professional representation of the statistical distribution that is easy for viewers to interpret.

```
import numpy as np
```

```
import matplotlib.pyplot as plt
```

```
#obtain histogram values with 10 bins
```


```
values, base = np.histogram(data, bins=10)
```

```
#find the cumulative sums
```

```
cumulative = np.cumsum(values)
```

```
# plot the ogive
```

```
plt.plot(base, cumulative, 'ro-')
```



## The Critical Role of Bin Selection in Ogive Interpretation

The resulting visual appearance and the ultimate interpretability of the Ogive chart are profoundly sensitive to the number of bins specified in the `numpy.histogram` function. The chosen bin count directly governs the granularity of the class intervals, which in turn determines the smoothness or jaggedness of the final cumulative curve. Selecting an appropriate number of bins is not merely an aesthetic choice; it is a critical step in effective data visualization and accurate statistical reporting.


When an analyst selects a small number of bins (such as the default 10 used in the initial example), the resulting Ogive curve tends to appear stepped or highly jagged. This occurs because the cumulative frequency only registers significant increases across relatively large intervals. While this stepped representation clearly illustrates the discrete nature of the frequency accumulation, it possesses a notable limitation: it may inadvertently obscure finer, subtle details regarding the underlying data distribution, making interpolation less precise.

Conversely, for the purpose of generating a cleaner, seemingly more continuous curve--one that more closely approximates the shape of the theoretical [cumulative distribution function \(CDF\)](#)--it is generally necessary to significantly increase the number of bins. For example, setting the bin count to 30 or higher would make the class intervals much narrower. This forces the Ogive to register smaller, more frequent increases in cumulative frequency across the range of the [dataset](#), resulting in a substantially smoother visual output. Although the code block below retains the initial value for continuity, conceptually, increasing the bin count drastically improves the detail and resolution of how the cumulative counts accumulate.

```
#obtain histogram values with 30 bins  
values, base = np.histogram(data, bins=10)
```

```
#find the cumulative sums
cumulative = np.cumsum(values)

# plot the ogive
plt.plot(base, cumulative, 'ro-')
```



## Customizing Aesthetics for Professional Statistical Plots

[Matplotlib](#) is renowned for providing flexible and concise methods for controlling the visual aesthetics of any generated plot, including the Ogive. The style string argument, such as `'ro-'`, which we employed in the `plt.plot` function, is a powerful shorthand convention that efficiently dictates the visual characteristics of the resulting cumulative curve. Understanding this shorthand allows for rapid customization suitable for various presentation needs.

This succinct plotting string is composed of three distinct and independent elements, each governing a specific aspect of the visualization:

The first character, `'r'`, functions as the color code, specifying that **red** should be used for both the connecting line and the markers drawn on the plot.

The second character, `'o'`, designates the marker type, indicating that a circular **marker** must be drawn precisely at every data point. These markers highlight the exact location of the cumulative frequency value corresponding to the upper boundary of each class interval.

The final character, `'-'`, specifies the line style, dictating the use of a solid **line** to smoothly connect all the defined markers, thereby ensuring the continuous, non-decreasing nature of the cumulative graph is clearly visible.

Professional analysts are encouraged to experiment extensively with these style options to tailor

the chart's aesthetics to their specific reporting requirements. For instance, substituting the style string with `'b--'` would render a blue dashed line connecting the points, whereas using `'g^'` would produce green triangular markers without any connecting lines. This level of customization is paramount for ensuring that the statistical findings presented via the Ogive are not only accurate but also visually impactful and formatted to the highest professional standard.