

Learning to Create Pareto Charts in Python: A Step-by-Step Tutorial

Authored by
Mohammed loot


November 1, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Create Pareto Charts in Python: A Step-by-Step Tutorial*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8090>

The [Pareto chart](#) stands as an indispensable tool in the fields of statistical analysis and process improvement, bridging the gap between descriptive statistics and actionable insights. This specialized [data visualization](#) combines the clarity of a bar chart--displaying categories ordered by frequency--with the interpretative power of a line graph that illustrates the cumulative contribution of these factors. This dual structure is fundamental, allowing analysts in areas like [quality control](#) and business process management to swiftly pinpoint the most critical contributors within any given dataset.

The entire methodology of the Pareto chart is rooted in the powerful concept known as the [Pareto Principle](#), famously dubbed the 80/20 rule. This principle suggests that, across many phenomena, approximately 80% of the total outcomes or effects are attributable to just 20% of the causes. By employing the Pareto visualization technique, practitioners can move past the noise of minor factors and focus their valuable resources and improvement efforts on the "vital few" causes that drive the majority of the results, dramatically enhancing efficiency and impact.



To implement this robust analytical technique, this comprehensive tutorial will guide you through utilizing [Python](#), the world's leading language for data science and statistical computing, alongside its powerful visualization libraries. We will walk through the process step-by-step, from initial data preparation to the final generation of a professional, informative Pareto chart, demonstrating how to leverage Python's capabilities to transform raw data into critical business intelligence.

The Essential Structure and Components of the Pareto Chart

Before initiating the coding process, it is paramount to grasp the distinct structure that separates a Pareto chart from standard bar graphs. The primary difference lies in the mandatory arrangement of the bars: they must always be sorted in strict descending order based on their frequency or count. This careful ordering immediately spotlights the categories that contribute most significantly to the total volume, ensuring that attention is naturally drawn to the factors requiring the most focused intervention.

The second, equally critical component is the cumulative percentage line, which is invariably plotted on a [secondary y-axis](#). This line begins at the top of the first bar and monotonically increases across the chart, culminating precisely at 100%. The cumulative metric is the key to applying the 80/20 rule, as it allows users to rapidly determine the minimum number of categories necessary to reach a predefined threshold--most often 80%--of the total measured effect. This instant interpretation capability makes the Pareto chart invaluable for tasks such as defect analysis and root cause identification in manufacturing and service industries.

Step 1: Structuring and Preparing Data with Pandas

Successful generation of a valid Pareto chart requires meticulous data preparation. Specifically, the data must be formatted to include not only the raw counts for each category but also the calculated cumulative percentages derived from those counts. For this crucial step, we rely on the [pandas](#) library, which serves as the industry gold standard in [Python](#) for efficient data manipulation, cleaning, and statistical calculations. Pandas streamlines the necessary sorting and aggregation tasks, making complex transformations accessible.

Consider a hypothetical scenario where we have surveyed 350 individuals regarding their preferred cereal brand across six options (A through F). Our initial task is to define this raw count data within a pandas DataFrame. Following initialization, we must perform two essential calculations: first, sorting the DataFrame in descending order based on the count column, and second, computing the crucial `cumperc` (cumulative percentage) column. This preparation ensures the data conforms perfectly to the structural requirements of the Pareto chart visualization.

The code snippet below demonstrates how to initialize the DataFrame, apply a descending sort based on the count, and then compute the crucial `cumperc` (cumulative percentage) column.

```
import pandas as pd
```

```
#create DataFrame
```

```
df = pd.DataFrame({'count': })
```

```
df.index =
```

```
#sort DataFrame by count descending
df = df.sort_values(by='count', ascending=False)

#add column to display cumulative percentage
df = df.cumsum()/df.sum()*100

#view DataFrame
df

count cumperc
A 140 40.000000
B 97 67.714286
C 58 84.285714
D 32 93.428571
E 17 98.285714
F 6 100.000000
```

It is important to emphasize that the sorting step is absolutely mandatory for generating a valid [Pareto chart](#). The resulting DataFrame, as shown above, clearly indicates that Brand A holds the highest frequency. Crucially, the newly calculated `cumperc` column is now perfectly prepared for its role as the overlying line graph, providing the cumulative perspective necessary for actionable analysis.

Step 2: Constructing the Dual-Axis Visualization using Matplotlib

Once the data is meticulously prepared, the next phase involves translating this structure into a visual chart using the powerful [matplotlib](#) library, the foundational visualization tool in the Python ecosystem. Since the Pareto chart requires two fundamentally different scales--raw counts on the left and cumulative percentages on the right--implementing a dual-axis structure is essential. This allows both datasets to be accurately represented without distortion.

The plotting process begins by creating the primary bar plot for the frequency counts. We then leverage the crucial `twinx()` function from matplotlib to establish a mirrored axis, conventionally labeled `ax2`. This secondary axis shares the exact same x-axis categories but provides an independent y-axis scale specifically calibrated for the cumulative percentage line. To ensure maximum clarity and professionalism, we utilize the `PercentFormatter()` function, imported from `matplotlib.ticker`, which automatically formats the right-hand axis ticks to display values correctly as percentages, significantly enhancing the chart's readability for stakeholders.

```
import matplotlib.pyplot as plt
from matplotlib.ticker import PercentFormatter
```

```
#define aesthetics for plot
color1 = 'steelblue'
color2 = 'red'
line_size = 4

#create basic bar plot
fig, ax = plt.subplots()
ax.bar(df.index, df, color=color1)


#add cumulative percentage line to plot
ax2 = ax.twinx()
ax2.plot(df.index, df, color=color2, marker="D", ms=line_size)
ax2.yaxis.set_major_formatter(PercentFormatter())

#specify axis colors
ax.tick_params(axis='y', colors=color1)
ax2.tick_params(axis='y', colors=color2)

#display Pareto chart
plt.show()
```

Step 3: Interpreting the Pareto Chart and Applying the 80/20 Rule

The resulting visualization offers immediate, actionable insight into the distribution of preferences among the surveyed population. Visually, the x-axis organizes the cereal brands (A through F) in the necessary descending order of popularity. The blue bars correspond to the left y-axis, quantifying the absolute frequency or count for each brand, while the overlaid red line corresponds to the right y-axis, providing the crucial cumulative percentage data. It is this cumulative line that unlocks the true analytical power of the chart, enabling rapid identification of the 'vital few' factors.



By carefully tracking where the cumulative line intersects specific thresholds, we can derive powerful conclusions that directly inform strategic decision-making:

Brand A alone dominates the market, contributing to exactly 40% of the total survey responses, establishing it as the single most important factor.

The combined popularity of Brands A and B accounts for approximately 68% of all recorded survey responses, demonstrating a clear concentration of preference early in the distribution.

Most significantly, Brands A, B, and C collectively account for roughly 84% of the total responses. This critical finding indicates that over 80% of the consumer preference distribution is concentrated among only three of the six available brands--meaning 50% of the categories drive 84% of the effect.

This type of insight is invaluable for business planning, particularly regarding inventory management, product development, and resource allocation. The analysis suggests that maximum impact will be achieved by heavily concentrating resources, such as marketing budgets or manufacturing focus, on these top three brands. Conversely, the remaining brands (D, E, and F) represent the "trivial many," where investment yields diminishing returns, confirming the application of the [Pareto Principle](#) in practice.

Step 4: Enhancing Visual Communication through Custom Aesthetics

While the basic Pareto chart constructed in the previous step is statistically functional, refining its

visual aesthetics is often necessary to adhere to corporate branding standards, improve accessibility, or simply increase its persuasive power during presentations. The [Matplotlib](#) library offers exceptional flexibility, providing analysts with extensive options to adjust visual properties such as colors, line styles, and marker appearances.

Customization is straightforward: by simply redefining the variables `color1` (for the bars), `color2` (for the cumulative line), and `line_size` (for line thickness) prior to the plotting commands, we can instantly transform the chart's appearance. For example, changing the bar color to pink and the line color to purple dramatically alters the visual identity. This level of control ensures that analysts can produce charts that are not only statistically rigorous but also highly effective and polished when communicating complex findings to a diverse audience.

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

```
from matplotlib.ticker import PercentFormatter
```

```
#define aesthetics for plot
```

```
color1 = 'pink'
```

```
color2 = 'purple'
```

```
line_size = 6
```

```
#create basic bar plot
```

```
fig, ax = plt.subplots()
```

```
ax.bar(df.index, df, color=color1)
```

```
#add cumulative percentage line to plot
```

```
ax2 = ax.twinx()
```

```
ax2.plot(df.index, df, color=color2, marker="D", ms=line_size)
```

```
ax2.yaxis.set_major_formatter(PercentFormatter())
```


```
#specify axis colors
```

```
ax.tick_params(axis='y', colors=color1)
```

```
ax2.tick_params(axis='y', colors=color2)
```

```
#display Pareto chart
```

```
plt.show()
```



Conclusion: Leveraging Pareto Charts for Strategic Decision-Making

The entire workflow for creating a [Pareto chart](#) in [Python](#) is highly streamlined, fundamentally relying on two core technological components: meticulous data structuring and calculation facilitated by [pandas](#), and the robust visualization capabilities provided by [matplotlib](#), particularly its effective implementation of dual-axis plotting. Mastering this visualization technique empowers analysts to transform raw frequency data into critical intelligence for [process optimization](#), effective resource allocation, and advanced quality control. By invariably drawing attention to the high-impact factors, the Pareto chart ensures that analytical efforts are always focused where they matter most.

Additional Resources for Data Visualization in Python

To further develop your skills in statistical visualization and data storytelling, the following tutorials explain how to create other common charts in [Python](#), building directly upon the foundational programming and plotting techniques utilized throughout this guide: