Learning to Create Stacked Bar Plots with Seaborn

Authored by

Mohammed looti

November 1, 2025

The ability to craft compelling visualizations is a fundamental requirement in modern data visualization and comprehensive analytical reporting. When tackling categorical data that needs to be broken down into constituent parts, the **stacked bar plot** emerges as an exceptionally effective tool. This chart type is expertly designed to display two critical pieces of information simultaneously: the overall magnitude of a primary category and the proportional or absolute distribution of its underlying segments.

Fundamentally, a stacked bar plot functions as an extension of the standard bar chart, where each bar is systematically segmented into smaller units representing distinct groups. These segments are placed sequentially, or "stacked," upon one another. This arrangement allows analysts and stakeholders to quickly ascertain the total value associated with the main category while efficiently observing the contribution and relative size of each sub-category that comprises that total.

This rigorous, step-by-step tutorial provides a comprehensive methodology for constructing a publication-ready stacked bar plot using the industry-standard Python ecosystem. We will harness the robust data management capabilities of the Pandas DataFrame library for preparing and manipulating data, along with the versatile plotting power of Seaborn and Matplotlib for sophisticated graphic generation. The objective is to produce a clear, informative visualization, similar to the final example shown below, which effectively contrasts daily customer traffic across two distinct time periods.

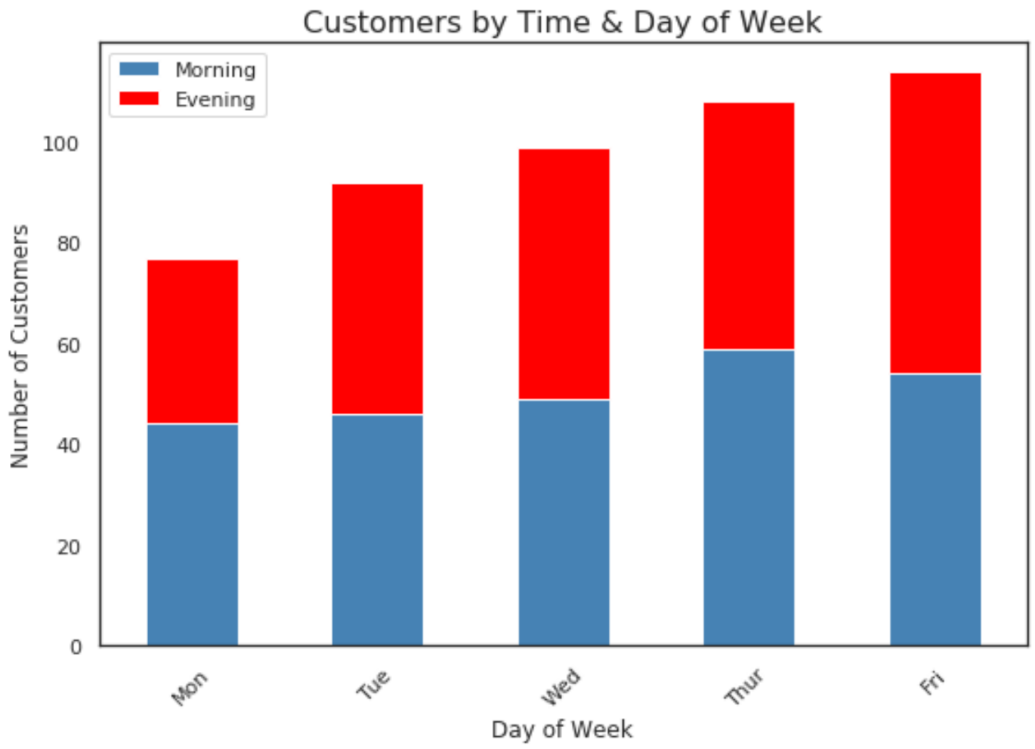

The Analytical Power of the Stacked Bar Plot

Effective data visualization begins with selecting the most appropriate chart type tailored to the specific data structure and the underlying analytical objective. The stacked bar chart is particularly strong in scenarios requiring the comparison of totals across several categories while simultaneously demanding an examination of the internal composition that makes up those totals. For instance, consider a business case where we visualize monthly sales data across various geographic regions (the primary categories) and require a breakdown of those sales figures by specific product lines (the sub-categories). The stacked format provides this layered insight instantly.

In the Python environment, the foundation for any complex visualization relies on meticulous data preparation, a task overwhelmingly handled by the **Pandas DataFrame** structure. The actual rendering of the visualization is primarily managed by a synergistic pair of robust libraries: **Matplotlib**, which serves as the foundational layer providing granular control over underlying plotting elements, and Seaborn, which is built on top of Matplotlib and offers high-level, statistically focused interfaces designed to generate aesthetically pleasing graphics with minimal code. Although we initially utilize Pandas' plotting function in this guide--which itself is a convenient wrapper--it is essential to recognize that Pandas depends entirely on Matplotlib internally. Therefore, we will employ explicit Matplotlib calls for detailed customization and refinement.

The specific example we will walk through involves tracking customer visits to a restaurant, categorized first by the day of the week and then further segmented by the time of day (Morning vs. Evening). This complex, two-dimensional structure necessitates the stacked bar format. This choice allows us to illustrate which days are the busiest overall (total bar height) and, simultaneously, reveal the relative popularity and contribution of the morning versus evening shifts on any given day.

Laying the Foundation: Python Libraries and Environment Setup

Before diving into the coding steps, it is imperative to ensure that all necessary Python libraries are properly installed within your environment. If you do not yet have them, you can perform a quick installation using the standard terminal command: `pip install pandas matplotlib seaborn`. Once the installation is complete, the crucial first step in any data analysis workflow is to import these libraries, conventionally assigning standard aliases (e.g., `pd` for Pandas, `plt` for Matplotlib's pyplot module, and `sns` for Seaborn) for streamlined coding.

We rely heavily on Pandas DataFrame for efficiently structuring the raw data into a usable tabular format that plotting functions can easily interpret. **Matplotlib** is indispensable for the detailed manipulation of plot elements, such as adjusting titles, setting axis limits, and adding text

annotations. Finally, Seaborn is utilized to apply consistent, professional plotting aesthetics, ensuring the visualization is clean, readable, and visually appealing. Importing these libraries correctly ensures that every required function is available throughout the subsequent data preparation and plotting stages.

A critical preparatory step, especially when integrating the high-level Seaborn library with the foundational Matplotlib, is setting the global plotting style. While not technically mandatory for a plot to render, employing a Seaborn style setting--such as `sns.set(style='white')`--standardizes the appearance of the plot by defining the background, managing grid lines, and selecting default font styles. This preparatory measure often results in cleaner, more professional visualizations by default, reducing the amount of manual aesthetic tuning required later.

Step 1: Structuring Data for Visualization using Pandas

The success of our visualization rests entirely on the structure of the foundational data. For the purpose of plotting a stacked bar chart where the segments (Morning and Evening) are layered vertically, the data must be organized in a specific **wide format**. This format requires that each segment intended for stacking must be represented by its own dedicated column, and the independent categorical variable (in our case, the Day of the Week) must be designated as the index or a clear categorical column.

We begin by creating a Pandas DataFrame that meticulously details the total number of customers received across our five primary business days (Monday through Friday), broken down into the two required time periods: Morning counts and Evening counts. This structure is perfectly suited for direct plotting, as the plotting function will inherently recognize 'Morning' and 'Evening' as the numerical variables whose values are intended to be aggregated and stacked within each categorical bar.

The code snippet below initializes this essential DataFrame. Notice the immediate clarity provided by the column names ('Day', 'Morning', 'Evening'), which unambiguously define the categorical axis and the two numerical data series that will form the stacked segments of the resulting bars.

import pandas as pd

#create DataFrame
df = pd.DataFrame({'Day': ,
'Morning': ,
'Evening': })

#view DataFrame
df

Day Morning Evening
0 Mon 44 33
1 Tue 46 46
2 Wed 49 50
3 Thur 59 49
4 Fri 54 60

Once the DataFrame is successfully instantiated and verified, we can inspect the output to confirm that the customer counts are correctly loaded and mapped to their respective columns. This clean, tabular structure provided by the Pandas DataFrame is the immediate prerequisite for generating the visualization in the subsequent steps.

Step 2: Generating the Initial Stacked Plot

With the data correctly prepared and structured, we move directly to generating the foundational visualization. For both simplicity and efficiency, we leverage the convenient, built-in plotting functionality provided by the Pandas library. This function acts as a straightforward wrapper around Matplotlib, enabling us to quickly construct the basic plot structure by chaining the plotting command directly onto the DataFrame object.

Two key parameters are crucial within the plotting command: `kind='bar'`, which explicitly defines the chart type as a bar chart, and `stacked=True`. The latter is the essential instruction that tells the underlying plotting engine to stack the numerical series--in this case, the 'Morning' and 'Evening' columns--vertically upon one another, rather than displaying them in a typical side-by-side grouped bar format.

Crucially, before calling the plot function, we execute `df.set_index('Day')`. This mandatory step ensures that the 'Day' column is designated as the index of the DataFrame. By setting it as the index, the plotting function automatically interprets 'Day' as the categorical axis (the x-axis) for the plot, while all remaining numerical columns ('Morning' and 'Evening') are automatically interpreted as the data series intended to be plotted and stacked.

import matplotlib.pyplot as plt
import seaborn as sns

#set seaborn plotting aesthetics
sns.set(style='white')

#create stacked bar chart
df.set_index('Day').plot(kind='bar', stacked=True, color=)

Executing the code above generates a preliminary visualization. While functionally correct--it displays the stacked data--this initial plot frequently lacks the crucial context, informative labels, and general polish required for professional communication, as evidenced in the example shown immediately below.

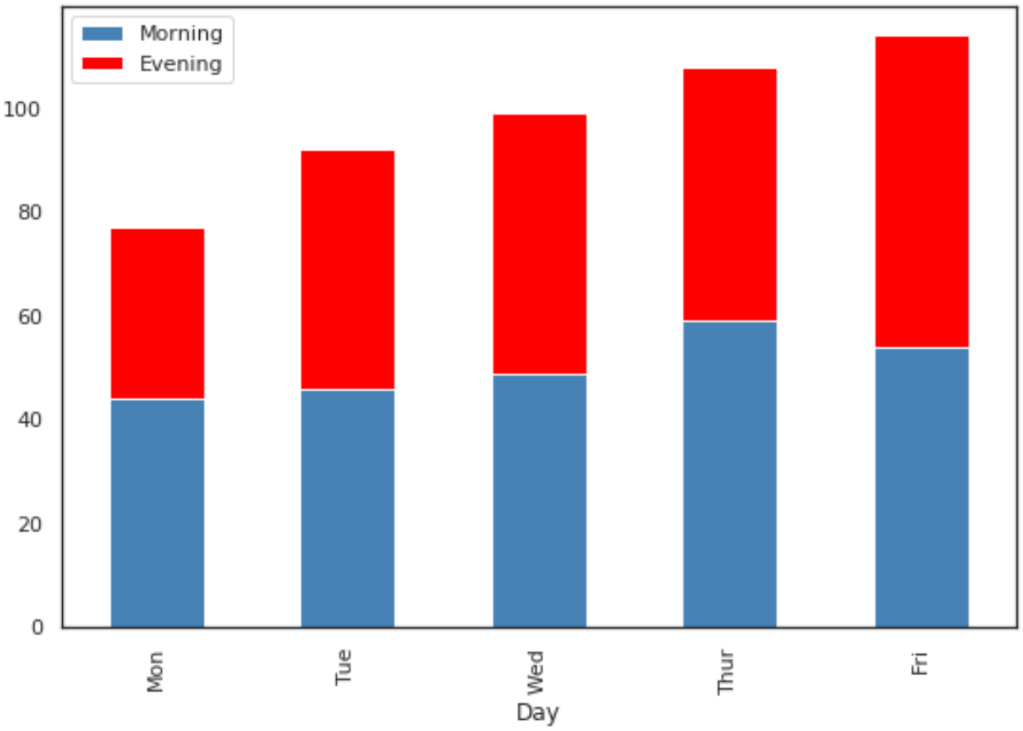

This resulting plot successfully conveys the total number of customers for each day of the week, with the internal division clearly separating the morning and evening segments. The x-axis correctly displays the categorical data (Day of the Week), and the vertical bars accurately represent the combined volume of customers. However, without descriptive titles and labels, its meaning remains ambiguous to an external audience.

Step 3: Achieving Publication Quality with Matplotlib Customization

A truly successful visualization must transcend mere data display; it must communicate its inherent insights clearly, immediately, and without ambiguity. The basic plot generated in Step 2, while structurally sound, must be enhanced through customization, primarily by incorporating informative titles, precisely labeling the axes, and adjusting the orientation of elements for optimized readability. These finishing touches are exclusively handled using the `pyplot` module from the Matplotlib library, which we have aliased as `plt`.

We utilize `plt.title()` to establish the overall context and subject matter of the entire chart. Next, `plt.xlabel()` and `plt.ylabel()` are essential for defining the categorical and quantitative units

of measurement on the x and y axes, respectively. Furthermore, a common practical concern in bar charts is the potential for labels on the categorical axis (Mon, Tue, Wed, etc.) to overlap, especially if the dataset were larger or the labels longer. To proactively prevent this visual collision and significantly improve the readability of the visualization, we introduce `plt.xticks(rotation=45)`. This command rotates the categorical labels by 45 degrees, ensuring they are distinctly separated.

Integrating these vital customization steps directly into the plotting script guarantees that the final output is professional, self-explanatory, and easily interpretable by any audience, regardless of their familiarity with the raw data. This iterative process--generating the functional base plot and then meticulously refining the aesthetics--represents the standard and best practice in professional data visualization.

import matplotlib.pyplot as plt
import seaborn as sns

#set seaborn plotting aesthetics
sns.set(style='white')

#create stacked bar chart
df.set_index('Day').plot(kind='bar', stacked=True, color=)

#add overall title
plt.title('Customers by Time & Day of Week', fontsize=16)

#add axis titles
plt.xlabel('Day of Week')
plt.ylabel('Number of Customers')

#rotate x-axis labels
plt.xticks(rotation=45)

The successful execution of this refined code yields the final, customized stacked bar plot, complete with appropriate labels, contextual titles, and necessary orientation adjustments, providing a clear and professional analysis of the customer traffic data.

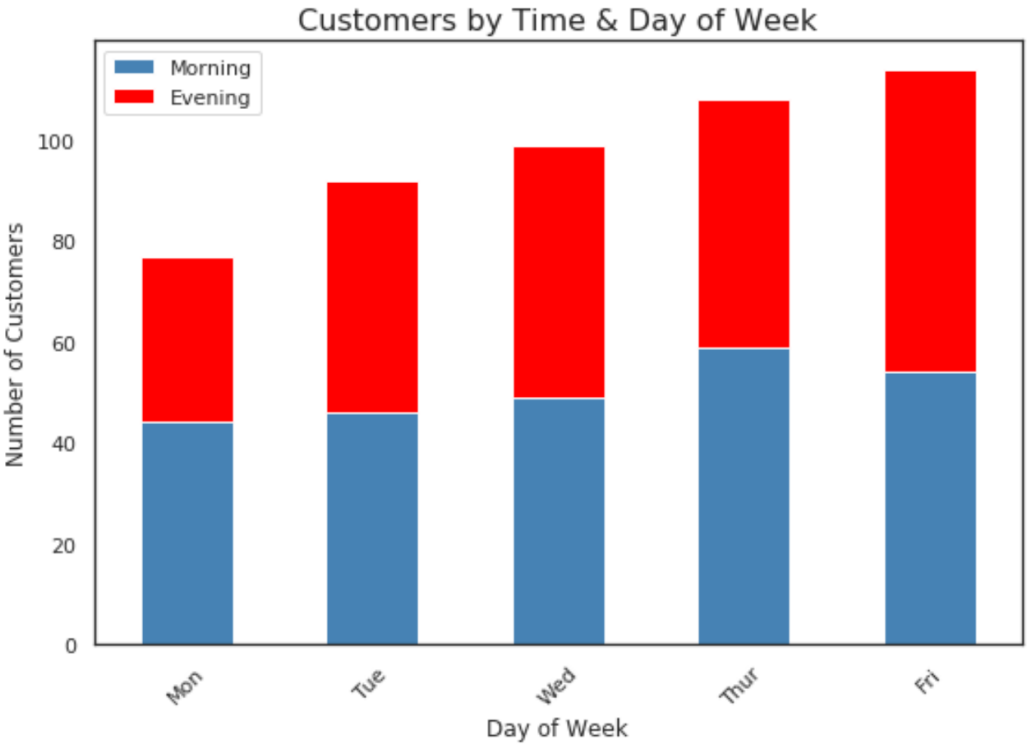

Interpreting and Expanding the Visualization

The final stacked bar plot is a powerful analytical tool because it condenses two dimensions of comparison into one view. By first observing the total height of each bar, we gain an immediate and comprehensive understanding of the busiest and slowest days of the week in terms of overall customer volume. For instance, a quick visual inspection comparing the bars suggests that Thursday and Friday are the undeniable peak days for total visitors, whereas Monday appears to be the quietest.

Beyond the totals, the internal stacking allows for a crucial comparison of the segments (Morning vs. Evening). This feature reveals how the composition of customers shifts across the week. For example, on Thursday, the morning segment (represented by the steelblue color) is noticeably larger than the evening segment (red), indicating a strong daytime rush. Conversely, on Friday, the evening segment dominates, suggesting that the primary volume of customers arrives after the typical morning hours. This high level of comparative detail is critically important for operational planning, such as optimizing staff schedules, adjusting kitchen inventory, or targeting marketing efforts.

The choice of color, while simple in this example (steelblue and red), is strategically managed through the `color` parameter passed within the Pandas plot function. Utilizing distinct colors is paramount for clearly differentiating the stacked categories and ensuring that the legend is easily and accurately mapped to the corresponding data series. Furthermore, the `style='white'` setting

inherited from Seaborn provides a clean, neutral background that effectively minimizes visual noise and distractions, focusing the viewer's attention entirely on the quantitative data represented by the bars themselves.

Further Exploration and Resources

While the stacked bar plot based on raw counts is highly effective for visualizing absolute totals, analysts often need to move beyond raw figures to visualize proportions. This necessity leads to the creation of a **percentage stacked bar chart**, a variation where all bars are normalized to reach 100%. This chart type excels at illustrating the relative contribution of each segment rather than the absolute count. The primary technical step toward creating this variation involves converting the raw counts stored in the Pandas DataFrame into calculated percentages relative to the daily total.

For users interested in achieving deeper and more intricate customization, the official Matplotlib documentation provides exhaustive resources for manipulating virtually every element of the plot. This includes fine-tuning font styles, precisely positioning the legend, and adding complex annotations directly onto the bars. Similarly, the Seaborn library offers a wealth of advanced plotting aesthetics and sophisticated statistical plotting functions that can be used in powerful conjunction with the plotting method outlined in this tutorial.

It is important to remember that effective visualization is fundamentally an iterative process. Experimenting with alternative color palettes, making subtle adjustments to bar width, and modifying axis scales or ranges are all standard methods used by data professionals to refine a chart until it perfectly and compellingly communicates the intended analytical narrative.

Official Seaborn Documentation: Explore advanced styling options, statistical visualizations, and high-level plotting function guides.

Matplotlib Gallery: Find inspiring examples of complex customizations for axes, titles, legends, and adding detailed annotations.

Pandas User Guide: Learn more about efficient data manipulation, advanced indexing techniques, and data reshaping required for complex plotting scenarios.