

How to Combine Multiple Excel Sheets into One Pandas DataFrame

Authored by
Mohammed loot

November 7, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *How to Combine Multiple Excel Sheets into One Pandas DataFrame*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=12397>

In contemporary data science and analytical engineering, analysts frequently encounter datasets that are fragmented, often distributed across numerous files or, more commonly, separated into distinct tabs within a single spreadsheet. When leveraging the robust capabilities of the [Pandas](#) library in Python, the fundamental requirement for any subsequent processing or analysis is the successful importation and consolidation of these disparate sources into a unified structure. This consolidated structure is typically referred to as a DataFrame. Successfully executing this aggregation is critical for ensuring data integrity and simplifying complex statistical tasks.

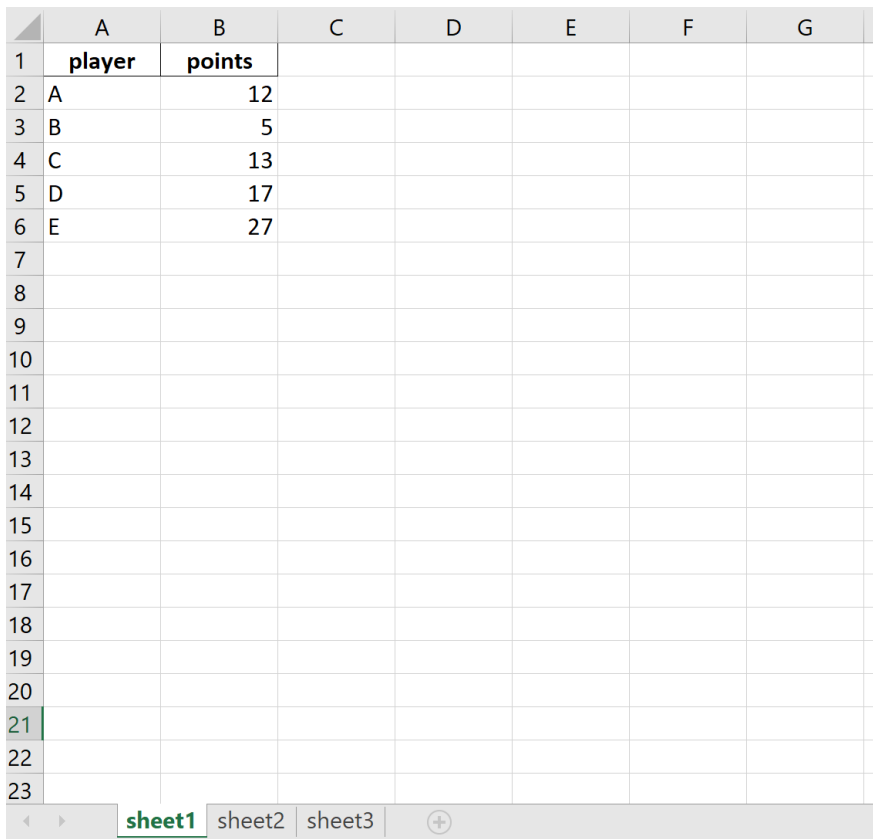
This authoritative tutorial outlines a highly efficient and formal methodology for importing every sheet contained within a single [Excel workbook](#) and merging them into one cohesive object. We will achieve this streamlined workflow by utilizing two core functions provided by Pandas: the [read_excel\(\)](#) function for file ingestion and parsing, and the powerful [concat\(\)](#) function for row-wise aggregation. This approach minimizes boilerplate code, enabling analysts to manage complex data consolidation tasks swiftly and reliably.

Defining the Data Aggregation Challenge

To demonstrate this powerful consolidation technique, let us consider a typical business intelligence scenario: an analyst is tasked with tracking performance metrics across several organizational units. For this specific example, the source data resides within an [Excel workbook](#) named **data.xlsx**. Crucially, this workbook contains three distinct sheets, where each sheet holds detailed, granular data relevant to a specific group of basketball players or teams. The fragmentation of this data across different tabs necessitates a systematic approach to unification before any macro-level analysis can commence.

A prerequisite for the most straightforward implementation of the single-line solution we present is that all three source sheets must possess an identical schema. In our basketball tracking scenario, this means every sheet must contain two columns labeled consistently: one for **player** identification and one for **points** scored. The uniformity of both the column headers and their underlying data types across all sheets is absolutely vital. This consistency allows the [concat\(\)](#) function to perform a clean, simple, and error-free row-wise merge without requiring complex mapping, alignment logic, or manual intervention.

The visual structure below confirms the standard organization we are aiming to resolve, where sheet names might be arbitrarily labeled--such as 'Team A', 'Team B', and 'Team C'--illustrating the distributed nature of the data that we must resolve for centralized reporting. Our core objective is to seamlessly transform this collection of fragmented sheets into a single, comprehensive DataFrame, making it immediately available for advanced statistical analysis and reporting functions.



The image shows a screenshot of an Excel spreadsheet. The spreadsheet has columns labeled A through G and rows numbered 1 through 23. The data is as follows:

	A	B	C	D	E	F	G
1	player	points					
2	A	12					
3	B	5					
4	C	13					
5	D	17					
6	E	27					
7							
8							
9							
10							
11							
12							
13							
14							
15							
16							
17							
18							
19							
20							
21							
22							
23							

The spreadsheet interface shows tabs for 'sheet1', 'sheet2', and 'sheet3' at the bottom, with 'sheet1' selected.

Essential Prerequisites: Installing Dependencies

Before executing the primary Python script designed for mass data ingestion, it is imperative to ensure that your Python environment is equipped with the necessary dependencies. While Pandas handles the majority of the data manipulation, reading binary [Excel files](#), particularly those in the older `.xls` format and reliably handling complex `.xlsx` structures, typically relies on a dedicated external parser. The `xlrd` library fulfills this critical role, acting as the necessary backend engine that the `read_excel()` function calls upon to successfully parse the binary structure of the spreadsheet data.

If you encounter any operational errors related to file reading, such as an `ImportError` or `FileNotFoundError` specifically referencing the engine, the installation or update of the `xlrd` package should be your immediate first step. This prerequisite action guarantees that the Pandas function can robustly interact with and interpret your target Excel file without dependency issues. Without a functional engine, Pandas cannot open and translate the raw spreadsheet data into usable memory objects.

To secure this dependency, execute the following command within your terminal or command prompt environment:

pip install xlrd

Once the successful installation of **xlrd** is confirmed, your development environment is fully prepared to handle the complex requirements of reading multiple sheets simultaneously and translating the raw spreadsheet data into the necessary DataFrame objects required for Python processing.

The Efficient Python Implementation

With all necessary dependencies successfully secured, we can proceed directly to implementing the core logic responsible for the data merger. The entire sheet aggregation process can be remarkably condensed into a single, elegant line of Python code, which powerfully showcases the conciseness and expressiveness of the Pandas library. This advanced approach involves tightly nesting the file reading operation within the concatenation function, thereby establishing a highly readable and streamlined data pipeline.

The script first imports the [Pandas](#) library using the conventional alias `pd`. It then immediately executes the combined read and merge operation, storing the definitive, unified result in a variable named `df`. Notice how the entire workflow is completed without requiring manual loops or explicit sheet name referencing, offering significant time savings for analysts dealing with frequently changing source files.

Load pandas library

```
import pandas as pd
```

```
# Import and combine all sheets into one pandas DataFrame
```

```
df = pd.concat(pd.read_excel('data.xlsx', sheet_name=None), ignore_index=True)
```

```
# View the resulting unified DataFrame
```

```
df
```

```
player points
```

```
0 A 12
```

```
1 B 5
```

```
2 C 13
```

```
3 D 17
```

```
4 E 27
```

```
5 F 24
```

```
6 G 26
```

```
7 H 27
```

```
8 I 27
```

9 J 12
10 K 9
11 L 5
12 M 5
13 N 13
14 O 17

The resulting output, displayed above, is a single DataFrame named `df` that successfully incorporates all 15 individual records, which originated from the three separate source sheets (5 records per sheet). This unified dataset is now ready for streamlined analysis, ensuring that all data points are treated equally within the context of the larger collective dataset. This single-line method effectively bypasses the requirement of manually iterating over sheet names, offering unparalleled conciseness and efficiency for standard aggregation tasks.

Deconstructing the Mechanism: The Read Operation

A deep understanding of this single, powerful command requires us to break down its two primary components, beginning with the inner function call: the specialized use of `pd.read_excel()`. Traditionally, when this function is invoked without explicit sheet parameters, it defaults to reading only the very first sheet of the file, returning a single DataFrame. However, the crucial technique for achieving mass aggregation is unlocked by leveraging the specific argument `sheet_name=None`.

By setting the argument `sheet_name=None`, we fundamentally alter the function's behavior, instructing Pandas to read every single sheet present within the specified [Excel workbook](#) ('data.xlsx'). Critically, instead of returning a single DataFrame, this function call returns an ordered dictionary. In this dictionary structure, the keys correspond exactly to the names of the sheets found within the workbook (e.g., 'Sheet1', 'Team A', etc.), and the values are the corresponding DataFrame objects derived from the data contained within those sheets. This dictionary acts as the perfect intermediate data container, organized and ready for the subsequent merge operation.

```
pd.read_excel('data.xlsx', sheet_name=None)
```

Deconstructing the Mechanism: The Concatenation Process

Once the ordered dictionary containing all individual DataFrames has been generated by the `read_excel()` operation, the outer function, `pd.concat()`, assumes control. The primary responsibility of `concat()` is to stack or join multiple DataFrame or Series objects along a specified axis. In this common use case, where our goal is to append rows vertically, we rely on the function's default axis setting, `axis=0`.

The input provided to `concat()` is the dictionary generated in the previous step. [Pandas](#) exhibits smart behavior here, extracting only the values--which are the individual DataFrame objects--from the dictionary and stacking them sequentially one beneath the other. This process quickly results in the formation of a single, continuous, and integrated dataset.

`pd.concat(DataFrames to concatenate, ignore_index=True)`

The parameter `ignore_index=True` addresses a critical indexing challenge inherent to data concatenation. When DataFrames are simply stacked, they retain their original indices from their source sheets. If every sheet begins its indexing at zero, the final combined DataFrame would contain numerous duplicate index values (e.g., index 0 appearing three times). By setting this parameter to **True**, we explicitly instruct Pandas to discard the original, potentially overlapping indices from the source sheets and assign a completely new, unique, and sequential index (0, 1, 2, 3, ...) to the final, consolidated structure. This indexing normalization is absolutely essential for subsequent data manipulation tasks that rely on unique row identifiers.

Considerations for Heterogeneous Data Structures

It is important to emphasize that the elegant, single-line solution demonstrated throughout this guide is optimally suited for, and highly dependent on, the structural homogeneity of the source [Excel sheets](#). Our basketball example succeeded smoothly because every sheet contained the exact same column names (**player** and **points**) and consistent data types across all tabs.

If the schemas of the sheets were to differ--for example, if Sheet 2 contained a column named 'Score' instead of 'points', or if Sheet 3 introduced an additional column such as 'Assists'--the `concat()` operation would still execute successfully. However, it would align columns based strictly on their name. The resulting DataFrame would contain the union of all column names found across all sheets, filling any missing values with `NaN` (Not a Number) wherever a particular sheet lacked data for a specific column present elsewhere.

For complex scenarios involving such heterogeneous structures, analysts must implement specific pre-processing steps. These necessary steps typically involve working directly with the dictionary returned by `read_excel(sheet_name=None)`. One must iterate through the individual DataFrames within that dictionary and apply critical transformations--such as renaming columns for standardization, dropping irrelevant columns, or enforcing consistent data types--before finally passing the list of standardized DataFrames to the `concat()` function. While this manual preparation is slightly more verbose, it is the only way to ensure complete data integrity and consistency when working with complex, non-uniform datasets.

Additional Resources

For users seeking to expand their proficiency in utilizing Pandas for advanced spreadsheet management and manipulation, the following resources provide valuable technical context and advanced techniques related to file I/O:

[The Ultimate Guide: How to Read Excel Files with Pandas](#)

[How to Write Pandas DataFrames to Multiple Excel Sheets](#)