

# Create a Nested DataFrame in Pandas (With Example)

Authored by  
**Mohammed loot**

November 16, 2025

## RECOMMENDED CITATION

Mohammed loot (2025). *Create a Nested DataFrame in Pandas (With Example)*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=2622>

## Introduction to the Concept of Nested DataFrames

In the expansive ecosystem of [Python](#) programming, especially when focused on advanced [data analysis](#), the [Pandas](#) library stands out as the fundamental tool. It is primarily utilized for its highly versatile and robust [DataFrame](#) object, which traditionally excels at managing two-dimensional tabular data, meticulously organized into distinct rows and columns. While this standard structure suffices for many common tasks, contemporary data science workflows frequently encounter complex, multi-layered, and interconnected datasets that demand a more sophisticated organizational paradigm than simple, flat tables. This evolution in data complexity necessitates exploring advanced data structures, specifically leading us to the powerful concept of a **nested DataFrame**.

A nested [DataFrame](#) is a structure where entire DataFrame objects are stored as individual elements within a column of a larger, encompassing parent DataFrame. This structure provides an elegant and highly efficient mechanism for handling complex, [hierarchical data](#) relationships that are often challenging to manage using conventional methods. Instead of fragmenting your project space with dozens of separate DataFrame variables (e.g., `regional_sales_q1`, `regional_sales_q2`, etc.), the nested approach allows you to consolidate all these related entities into a single, cohesive, high-level structure. This consolidation is not merely for cosmetic tidiness; it profoundly improves operational efficiency, significantly elevates code readability by logically grouping related data elements, and streamlines programmatic access that relies on a unified identifier or index.

The fundamental principle driving this structure is the construction of a new parent [DataFrame](#) where one of its designated columns is intentionally populated solely by other Pandas DataFrame objects. By treating a collection of related datasets as a single, manageable entity, developers gain the ability to iterate through them systematically, apply custom or complex functions across all nested elements simultaneously, or quickly retrieve specific sub-DataFrames based on contextual identifiers. This methodology proves invaluable when dealing with dynamically generated or contextually linked datasets that benefit immensely from being contained within a singular, easily accessible structure. As we proceed, we will meticulously detail the precise syntax and walk through a practical application of this technique, ensuring you can seamlessly integrate nested DataFrames into your personal data manipulation toolkit.

## The Fundamental Syntax for DataFrame Nesting

The technique for embedding multiple distinct Pandas [DataFrames](#) into a new container DataFrame is remarkably simple and relies directly on the standard DataFrame constructor provided by the [Pandas](#) library. This crucial process is accomplished by passing a Python dictionary to the constructor, where the dictionary keys define the column headers of the new

parent DataFrame, and the associated values are lists representing the column data. Critically, in the context of nesting, one of these value lists will be composed entirely of the actual DataFrame objects you intend to embed.

To provide a clear conceptual framework, consider the following fundamental syntax used to instantiate a nested DataFrame object:

```
df_all = pd.DataFrame({'idx':, 'dfs':})
```

In this specific construction, `df_all` is established as the principal, overarching container [DataFrame](#). It is defined with two primary columns: `idx` and `dfs`. The `idx` column typically functions as a metadata field, providing a simple numerical index or identifier (represented here by ) used for tracking or referencing the nested datasets. The true architectural innovation resides in the `dfs` column, which is assigned a list containing pre-existing DataFrame objects--specifically, . This design effectively yields a columnar structure where each individual cell within the `dfs` column holds an entire DataFrame object, transforming `df_all` into an efficient and highly structured repository for these sub-DataFrames.

This implementation beautifully showcases an application of [object-oriented programming](#) (OOP) principles within the [Pandas](#) framework, allowing complex, specialized objects, such as DataFrames, to be treated as standard values within other primary [data structure](#) elements. The specified setup successfully nests three distinct DataFrames (labeled **df1**, **df2**, and **df3**) inside the single parent DataFrame, **df\_all**. The inherent simplicity and intuitive nature of this syntax provide a clean and powerful mechanism for organizing and managing conceptually related tabular data, which sets the stage perfectly for conducting more advanced and modular analyses.

## Practical Example Setup: Preparing Individual DataFrames

To fully appreciate the practical utility and organizational benefits of DataFrame nesting, we will now proceed through a concrete, illustrative example. Consider a common business scenario where a data analyst must manage sales figures that are segregated either by different product regions, specific quarterly periods, or distinct operational units. Our objective is to begin by creating three distinct and fully independent Pandas DataFrames, which we designate as **df1**, **df2**, and **df3**. Although independent, each of these DataFrames will share a common internal structure, comprising the columns `item` and `sales`, but will contain unique data entries relevant to their specific context--like sales records from three different months.

The initial requirement for this exercise is importing the necessary [Pandas](#) library using its common alias and then defining the content and structure for these three individual DataFrames. This foundational setup is absolutely essential, as it clearly demonstrates that we are starting with

separate, independent data entities that will subsequently be unified into a single, highly organized nested structure. The subsequent code snippet explicitly illustrates the construction of these three DataFrames, followed immediately by their respective outputs. This allows you to clearly visualize the contents and verify the independence of each DataFrame before the consolidation procedure takes place.

### **import pandas as pd**

```
#create first DataFrame (e.g., January Sales)
```

```
df1 = pd.DataFrame({'item': ,  
'sales': })
```

```
print(df1)
```

```
item sales
```

```
0 A 18
```

```
1 B 22
```

```
2 C 19
```

```
3 D 14
```

```
4 E 30
```

```
#create second DataFrame (e.g., February Sales)
```

```
df2 = pd.DataFrame({'item': ,  
'sales': })
```

```
print(df2)
```

```
item sales
```

```
0 F 10
```

```
1 G 12
```

```
2 H 13
```

```
3 I 13
```

```
4 J 19
```

```
#create third DataFrame (e.g., March Sales)
```

```
df3 = pd.DataFrame({'item': ,  
'sales': })
```

```
print(df3)
```

```
item sales
```

```
0 K 41
```

1 L 22  
2 M 28  
3 N 25  
4 O 18

At this specific juncture, `df1`, `df2`, and `df3` are all standard, fully independent Pandas [DataFrames](#), each maintaining its own isolated internal index and column definitions. Our ultimate goal is now to demonstrate a robust methodology where these independent DataFrames--which are conceptually related (all sales data)--can be logically grouped, efficiently accessed, and centrally managed through a single, overarching structure. This transition simplifies subsequent processing operations and promotes a highly scalable approach to [data modeling](#) when dealing with complex, multi-layered information systems.

## Constructing and Verifying the Unified Nested DataFrame

With our three separate DataFrames (`df1`, `df2`, and `df3`) prepared and verified, the subsequent and most critical phase involves consolidating them into a single, cohesive nested DataFrame. This step requires creating a new parent DataFrame that is specifically engineered to act as a unified container, dedicating one of its columns exclusively to holding the existing DataFrame objects. This centralized organization technique is incredibly valuable when multiple related datasets are conceptually or contextually linked, yet still require the ability to be individually maintained and accessed within a broader, centralized organizational framework.

To successfully execute this consolidation, we utilize the standard `pd.DataFrame()` constructor. We must supply this constructor with a dictionary that defines the structure of the new parent object. This dictionary contains two essential key-value pairs: the first key is for an index or identifier column (named `idx` in our example) and the crucial second key (named `dfs`) designates the column that will hold our embedded DataFrames. The value associated with the `dfs` key is simply the Python list `.` By passing this list, we efficiently embed these entire DataFrame objects directly into the corresponding cells of the new parent DataFrame.

The necessary syntax required to create our nested DataFrame, which we have named `df_all`, is identical to the fundamental example previously discussed, demonstrating the simplicity of the approach:

```
df_all = pd.DataFrame({'idx':, 'dfs':})
```

Immediately after this command is executed, `df_all` is instantiated, containing two columns: `idx`, which provides a simple numerical identifier for each nested record, and `dfs`, which contains the actual DataFrame objects themselves. This structural transformation is foundational to the nested

approach: instead of managing three distinct variables, we now efficiently manage a single `df_all` DataFrame that encapsulates all of them. This unified structure drastically simplifies data management, greatly enhances programmatic access, and is especially advantageous when automating operations across a dynamic or expanding set of conceptually related datasets, such as time-series data or regional reports.

## Precision Retrieval: Accessing Individual Nested DataFrames

Once the `df_all` nested [DataFrame](#) has been successfully established, the practical focus shifts to understanding how to efficiently retrieve and interact with the individual DataFrames housed within its structure. The [Pandas](#) library provides robust [indexing](#) and selection mechanisms that make this process straightforward and reliable. The fundamental key to retrieving a specific nested DataFrame involves two sequential steps: first, selecting the specific column containing the DataFrame objects (the `dfs` column in our example), and second, employing positional [indexing](#) to pinpoint the desired sub-DataFrame element within that column.

The primary and most effective tool for this precise extraction is the [iLoc](#) accessor, which is specifically designed for integer-location based [indexing](#) in Pandas. Since the nested DataFrames reside as standard Python objects within a Pandas Series (the selected column), we first isolate this Series using standard column selection (`df_all`). We then immediately apply the [iLoc](#) accessor to retrieve the content--which is the entire DataFrame object--at a specific numerical position. This reliance on positional indexing ensures highly reliable and fast access, particularly when the sequential order of the nested DataFrames (like chronological reports) is significant to the analysis.

For instance, to retrieve and display the first nested DataFrame (which corresponds to our original `df1`, representing the data at index position 0), the following precise and concise syntax is utilized:

```
#display first nested DataFrame (df1)  
print(df_all.iLoc)
```

```
item sales  
0 A 18  
1 B 22  
2 C 19  
3 D 14  
4 E 30
```

This command chain first selects the `dfs` column from `df_all`, yielding a Pandas Series object. The subsequent application of [iLoc](#) successfully extracts the element located at index position 0,

which is the first nested DataFrame. Similarly, accessing the second nested DataFrame (our original `df2`) is accomplished simply by adjusting the index within the `iloc` accessor to `1`. This demonstrates the high degree of flexibility and systematic access provided by this method, where each successful retrieval effectively "unpacks" the DataFrame object, allowing it to be treated as an independent DataFrame for any subsequent analysis, filtering, or modification required.

### #display second nested DataFrame (df2)

```
print(df_all.iloc
```

```
item sales
```

```
0 F 10
```

```
1 G 12
```

```
2 H 13
```

```
3 I 13
```

```
4 J 19
```

## Expanding Your Expertise: Further Pandas Resources

Mastering the creation, efficient management, and precise manipulation of nested DataFrames in [Pandas](#) represents a substantial and crucial step forward in your overall data manipulation capabilities. This advanced technique is indispensable for effectively handling intricate data relationships, significantly improving the modularity and organization of your data processing pipelines, and enhancing the overall clarity and maintainability of your analytical code. The ability to consolidate related datasets into a single structure dramatically increases productivity when dealing with complex projects.

To further solidify your understanding of complex data structures and explore complementary functionalities within the powerful Pandas library, we strongly recommend focusing your continued learning on the following related topics and essential tutorials. Expanding your knowledge base in these areas will comprehensively prepare you to tackle a wider and more challenging spectrum of real-world data science problems with greater efficiency and confidence:

**Exploring MultiIndex DataFrames:** Understanding this concept serves as a common and powerful alternative mechanism for managing complex [hierarchical data](#) within a single DataFrame. Grasping the principles of MultiIndex allows you to accurately select the most appropriate data structure based on the specific analytical requirements of your project.

**Applying Functions to DataFrame Columns:** Become proficient in the effective use of methods such as `apply()`, `map()`, and `applymap()`. These functions are fundamental for performing vectorized and custom operations on data, including the necessary techniques for systematically iterating through and processing data within your nested DataFrames.

Merging and Joining DataFrames: Develop essential skills for reliably combining data sourced from disparate origins based on common keys or indices. This is a frequent and critical requirement in large-scale data integration projects where information must be synthesized from multiple tables.

Reshaping DataFrames with `pivot_table` and `melt`: Acquire the expert knowledge necessary to transform the layout of your data, optimizing it specifically for improved analytical performance, streamlined reporting, and advanced visualization tasks.

Handling Missing Data: Cultivate robust and systematic techniques for identifying, managing, and imputing null or missing values (NaNs) to ensure the statistical integrity, accuracy, and overall reliability of your DataFrames before analysis.

By dedicating focused time and effort to mastering these advanced areas, you will substantially enhance your overall proficiency in data manipulation using [Python](#) and Pandas, enabling you to address the most complex data analysis challenges with superior effectiveness, precision, and professional confidence.