

How to add new columns to existing DataFrame in Pandas?

Authored by
stats writer

November 29, 2025

RECOMMENDED CITATION

stats writer (2025). *How to add new columns to existing DataFrame in Pandas?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=101570>

Working with data often requires dynamic manipulation, and one of the most fundamental operations in Pandas is adding new fields or restructuring datasets. While the title focuses on adding new columns, this process often goes hand-in-hand with creating specialized, smaller datasets--or subsets--from a larger, existing DataFrame. To successfully integrate new columns, developers typically employ direct assignment, using the powerful DataFrame.assign() method, or leveraging the flexibility of DataFrame.insert() to place a column at a specific index. These methods provide robust ways to enhance your data structure, whether you are calculating derived metrics or merging external data.

However, before complex calculations or modifications begin, it is frequently necessary to isolate relevant data points. This article delves into the critical techniques used to generate a **new DataFrame** based on the columns of an **existing DataFrame**. Understanding how to efficiently subset and copy() data is crucial for maintaining data integrity and optimizing memory usage within your Python environment. We will explore three primary methods for creating these subsets, focusing on clarity, efficiency, and the vital practice of ensuring independent copies of the data.

Methods for Adding New Columns to a DataFrame

Although the core examples below focus on subsetting, it is valuable to quickly review the standard approaches for adding columns, as they represent the typical next step after creating a refined subset. Direct assignment is the simplest method: you simply assign a list or Series to a new column name, such as `df =` . This works seamlessly when the length of the assigned values matches the number of rows in the existing DataFrame. This method is highly efficient but lacks the functional chaining capabilities offered by other methods.

A more sophisticated and often preferred method, especially in modern Pandas workflows, is the use of DataFrame.assign(). This function is excellent because it returns a new object with all new columns added, without modifying the original DataFrame in place. This immutable approach is key for functional programming and avoiding unexpected side effects. You can assign multiple columns simultaneously by passing keyword arguments, making the code clean and readable, especially when defining complex calculated fields based on existing data.

Finally, for situations demanding precise control over column order, the DataFrame.insert() method is indispensable. Unlike direct assignment or `.assign()`, `.insert()` allows the user to specify the exact integer location (index) where the new column should be placed. This method is crucial when dealing with legacy systems or API requirements that demand specific column ordering, though it typically performs the modification in place, meaning it alters the original DataFrame directly.

Core Techniques for Creating a New DataFrame Subset

When analyzing large datasets, it is rarely necessary to work with every single column simultaneously. Creating a dedicated subset--a new `DataFrame` containing only the essential columns--streamlines memory usage and simplifies subsequent operations. There are three primary, reliable ways to achieve this column-based subsetting in `Pandas`, each tailored for a specific extraction scenario. These methods ensure that the resulting `DataFrame` is a dedicated entity, separate from the source data, which is essential for safe, isolated manipulations.

The choice between these methods depends entirely on the goal: do you need a handful of specific columns? Do you only require a single variable? Or is it easier to define your subset by listing the columns you wish to exclude? By mastering these three techniques, data practitioners gain fundamental control over data restructuring. It is important to remember that when using indexing techniques (like methods 1 and 2 below), the resulting object must be explicitly copied to guarantee independence, a concept we will detail shortly.

Here are the fundamental patterns for creating these new `DataFrames`:

Method 1: Create New DataFrame Using Multiple Columns from Old DataFrame (Index-based selection).

```
new_df = old_df.copy()
```

Method 2: Create New DataFrame Using One Column from Old DataFrame (Single column selection).

```
new_df = old_df.copy()
```

Method 3: Create New DataFrame Using All But One Column from Old DataFrame (Exclusion-based selection using the `.drop()` method).

```
new_df = old_df.drop('col1', axis=1)
```

Preparing the Environment and Sample Data

To clearly illustrate these subsetting methods, we will utilize a simple, synthetic dataset representative of sports statistics. This dataset includes columns for team identification, and player performance metrics like points, assists, and rebounds. Before executing any transformations, the initial step is always to import the `Pandas` library, typically aliased as `pd`, which provides the foundational structures like the `DataFrame` object itself. Once imported, we construct our source

data structure, `old_df`.

The sample data creation involves initializing the `DataFrame` constructor with a dictionary, where the keys represent the column names ('team', 'points', 'assists', 'rebounds') and the values are lists of corresponding data points. This setup allows for immediate, hands-on demonstration of how the subsetting techniques impact the structure of the data. Reviewing the original DataFrame output ensures that our starting point is correctly defined before we proceed with slicing and restructuring operations.

The following code block sets up our environment and defines the base DataFrame we will manipulate throughout the subsequent examples:

```
import pandas as pd
```

```
#create DataFrame
```

```
old_df = pd.DataFrame({'team': ,  
'points': ,  
'assists': ,  
'rebounds': })
```

```
#view DataFrame
```

```
print(old_df)
```

Example 1: Selecting Multiple Columns for a New DataFrame

The most common method for creating a subset involves explicitly listing the required column names using double brackets (list of column names). This technique is highly readable and ensures that only the named columns are included in the new structure. In this example, we aim to isolate the core performance statistics--specifically 'points' and 'rebounds'--while excluding auxiliary data like 'team' and 'assists'. The use of double square brackets, `[]`, is mandatory because the outer brackets perform the indexing operation on the DataFrame, and the inner brackets define the list of column keys we are selecting.

Immediately following the column selection, we chain the powerful `copy()` method. This step is non-negotiable for index-based selection methods. By invoking `.copy()`, we force `Pandas` to allocate new memory for the data and metadata of `new_df`. Without this explicit command, the resulting `new_df` might be a "view" of the original `old_df`, meaning any modification to `new_df` could inadvertently alter the data in `old_df`, leading to the dreaded `SettingWithCopyWarning` and potentially corrupting the source data. The output confirms that the new DataFrame retains only the specified columns and maintains the original row indexing.

The following code demonstrates the creation of `new_df` based on a selection of multiple columns, followed by verification of the structure and type:

#create new DataFrame from existing DataFrame

```
new_df = old_df.copy()
```

```
#view new DataFrame
```

```
print(new_df)
```

```
points rebounds
```

```
0 18 11
```

```
1 22 8
```

```
2 19 10
```

```
3 14 6
```

```
4 14 6
```

```
5 11 7
```

```
6 20 9
```

```
7 28 12
```

```
#check data type of new DataFrame
```

```
type(new_df)
```

```
pandas.core.frame.DataFrame
```

Notice that this new DataFrame, `new_df`, successfully contains only the **points** and **rebounds** columns, confirming the selective extraction process was executed correctly. This technique is fundamental for creating focused analytical datasets.

Importance of Using the `.copy()` Method

In Pandas, when you subset a DataFrame using indexing (like `df[1]`), the resulting object can sometimes be a "view" of the original data rather than a complete, independent "copy." A view means that the underlying data in memory is shared between the original and the subset. If you later modify a value in the subset, that change might unintentionally propagate back to the original DataFrame, which is almost always undesirable in a robust data pipeline. This ambiguity arises because Pandas tries to optimize performance by avoiding unnecessary data duplication.

To eliminate this potential source of error and confusion, it is critically important to explicitly call the `copy()` function immediately after the selection operation, as demonstrated in the examples above. The `.copy()` method forces the creation of a deep copy, meaning that the new DataFrame, `new_df`, receives its own dedicated block of memory for both its structure and its data values. This

ensures complete isolation: modifications made to `new_df` will have no effect whatsoever on the original `old_df`, guaranteeing data integrity and predictability in complex data manipulation scripts.

Note: It's crucial to use the `copy()` function when creating the new DataFrame via index selection (Methods 1 and 2) so that we avoid any potential `SettingWithCopyWarning` issues and guarantee that the new DataFrame is truly independent. Failure to do so can lead to subtle but serious bugs in data pipelines where data mutation is involved.

Example 2: Extracting a Single Column into a New DataFrame

While Method 1 focused on extracting multiple columns, a slightly different approach is required when the goal is to create a new DataFrame consisting of only a single column, yet preserving the structure of a two-dimensional DataFrame object. If one were to use single square brackets (e.g., `old_df`), the result would be a one-dimensional Pandas Series, not a DataFrame. To maintain the DataFrame structure, even for a single column, we must still use the list notation within the indexing brackets.

Therefore, to extract just the 'points' column as a new DataFrame, the syntax remains `old_df['points']`. This ensures that the resulting object has columns and rows, maintaining the expected data structure for future operations like merging or advanced indexing that require two dimensions. Just like in the previous example, the explicit use of the `copy()` function is essential here to guarantee that `new_df` is an independent entity, decoupled from the memory space of `old_df`, preventing unwanted data contamination.

The following code shows how to create a new DataFrame using only the 'points' column from the original data, verifying that the output maintains the DataFrame type:

```
#create new DataFrame from existing DataFrame
```

```
new_df = old_df[['points']].copy()
```

```
#view new DataFrame
```

```
print(new_df)
```

```
points
```

```
0 18
```

```
1 22
```

```
2 19
```

```
3 14
```

```
4 14
```

```
5 11
```

```
6 20
```

7 28

```
#check data type of new DataFrame
```

```
type(new_df)
```

```
pandas.core.frame.DataFrame
```

Example 3: Excluding Specific Columns Using `.drop()`

Often, it is easier to define the desired subset by listing the columns you do *not* want, rather than the dozens of columns you do want. For this exclusion-based selection, the `DataFrame.drop()` method is the ideal tool. This method allows you to remove specified rows or columns from a DataFrame, returning the resulting object without the removed elements. A key benefit of using `.drop()` for column removal is that, by default, it already returns a new DataFrame copy, mitigating the need for an explicit `.copy()` call unless performing complex, multi-stage assignments.

To use `DataFrame.drop()` to remove columns, two essential arguments must be provided: the label(s) of the column(s) to remove (e.g., `'points'`) and the `axis` parameter. Setting `axis=1` (or `axis='columns'`) tells Pandas that the provided label refers to a column name, as opposed to `axis=0` (or `axis='index'`), which is the default and refers to row labels. By dropping the `'points'` column, we create a new DataFrame containing all other original columns, which in our sample dataset includes `'team'`, `'assists'`, and `'rebounds'`.

This method is particularly valuable when performing iterative data cleaning or when dealing with highly sparse datasets where only a few problematic or irrelevant columns need to be discarded. The resulting DataFrame, `new_df`, is a clean slate containing the majority of the original data structure, ready for analysis or further modification.

The following code demonstrates how to use the `DataFrame.drop()` method to exclude the `'points'` column:

```
#create new DataFrame from existing DataFrame
```

```
new_df = old_df.drop('points', axis=1)
```

```
#view new DataFrame
```

```
print(new_df)
```

```
team assists rebounds
```

```
0 A 5 11
```

```
1 A 7 8
```

```
2 A 7 10
```

```
3 A 9 6
```

```
4 B 12 6
```

```
5 B 9 7
```

```
6 B 9 9
```

```
7 B 4 12
```

```
#check data type of new DataFrame
```

```
type(new_df)
```

```
pandas.core.frame.DataFrame
```

Notice that this new DataFrame contains all of the columns from the original DataFrame *except* the **points** column, demonstrating the efficacy of the exclusion method.

Summary and Best Practices for Data Subsetting

Mastering the creation of new DataFrames through subsetting is foundational to efficient data analysis in Pandas. Whether you are adding new calculated columns using DataFrame.assign() or restructuring data by selecting specific fields, the underlying principle is to ensure data integrity and memory efficiency. The three methods explored--selection by multiple columns, selection by single column (using list notation), and exclusion using DataFrame.drop()--provide a comprehensive toolkit for handling different data extraction requirements.

A crucial best practice, reiterated throughout these examples, is the importance of using the copy() method when performing index-based selections (Methods 1 and 2). This safeguards the original dataset from unintended modifications caused by the ambiguous "view vs. copy" behavior in Pandas. By making explicit copies, developers ensure that their data pipelines are robust, predictable, and free from the pitfalls of shared memory access, resulting in cleaner and more maintainable code.

Ultimately, the choice of method depends on context: use indexing (and .copy()) when you know exactly which few columns you need to retain, and use .drop() when you only need to discard a few columns from a large set. These techniques, combined with robust column addition methods, enable complex data transformations necessary for modern analytical workloads.