

# How to Easily Compare Three Columns in Pandas

Authored by  
**stats writer**

November 28, 2025

## RECOMMENDED CITATION

stats writer (2025). *How to Easily Compare Three Columns in Pandas*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=101067>

The ability to efficiently compare data across multiple columns is a fundamental requirement in professional [data analysis](#). When working with large datasets, identifying rows where three or more specific attributes match exactly often forms the basis for validation, filtering, or quality control. This guide focuses on leveraging the strengths of the [Pandas](#) library in [Python](#) to perform precise three-column comparisons within a single [DataFrame](#). To do this, we primarily utilize the powerful row-wise application function coupled with Python's efficient chained comparison operators, which returns a [Boolean](#) result indicating equality.

## Introduction to Column Comparison in Pandas

Pandas, a cornerstone of the [Python](#) data science stack, provides unparalleled tooling for manipulating tabular data. While simple column comparisons can be managed using functions like `.eq()`, complex requirements--such as checking if three distinct columns are simultaneously identical--demand a more generalized solution. Identifying these fully matched rows is crucial for data auditing, ETL process verification, and ensuring transactional consistency across different recorded fields, making the comparison technique we discuss here invaluable.

We will demonstrate how to create a highly readable and performant solution that generates a new status column based on the congruence of the specified columns. This method bypasses the need for multiple logical AND operators often associated with element-wise comparisons (like `(A==B) & (B==C)`), favoring the more concise and memory-efficient row-wise application provided by `.apply()`.

The core technique centers around the `DataFrame.apply()` method, which facilitates running custom [Python](#) functions, often defined via a [lambda](#) expression, across the rows (`axis=1`) of the [DataFrame](#). This approach allows us to exploit the inherent efficiency of Python's equality chaining (`A == B == C`) to return a single, definitive [Boolean](#) outcome for each observation, confirming whether all three columns align perfectly.

## Implementing the Generic Comparison Syntax

To implement the comparison, we construct a new column--conventionally named `all_matching`--which will store the outcome of our row-wise check. This ensures that the original data remains untouched while providing a clear flag for subsequent filtering or reporting based on data consistency.

The following basic syntax demonstrates the use of `.apply()` to compare values in three placeholder columns (`col1`, `col2`, `col3`). The power of this approach lies in the simplicity of the [lambda](#) function, which leverages Python's ability to handle transitive equality checks seamlessly and efficiently for high-volume [Pandas](#) operations.

This syntax creates a new column called **all\_matching** that returns a value of **True** if all of the columns have matching values, otherwise it returns **False**.

```
df = df.apply(lambda x: x.col1 == x.col2 == x.col3, axis = 1)
```

This syntax creates a new column called **all\_matching** that returns a value of **True** if all of the columns have matching values, otherwise it returns **False**.

## Setting Up the Practical Example in Pandas

To solidify our understanding, we will now proceed with a concrete example. We begin by importing the [Pandas](#) library and defining a sample [DataFrame](#). This DataFrame is crucial as it contains varied data, allowing us to observe cases of perfect matches, partial matches, and complete mismatches across the three columns 'A', 'B', and 'C'.

This setup mimics a common scenario where raw data is loaded and requires an immediate quality check before further processing in the [data analysis](#) pipeline. The columns 'A', 'B', and 'C' represent distinct observations or measurements that ideally should align perfectly based on the underlying business rules or data integrity constraints.

The following example shows how to use this syntax in practice. Suppose we have the following pandas DataFrame with three columns:

```
import pandas as pd
```

```
#create DataFrame
```

```
df = pd.DataFrame({'A': ,  
'B': ,  
'C': })
```

```
#view DataFrame
```

```
print(df)
```

```
A B C
```

```
0 4 4 4
```

```
1 0 2 0
```

```
2 3 3 3
```

```
3 3 5 5
```

```
4 6 6 5
```

```
5 8 4 10
```

```
6 7 7 7
```

```
7 9 7 9
8 12 12 12
```

## Executing the Comparison and Updating the DataFrame

With our DataFrame initialized, we now execute the core comparison logic. We pass the row-wise comparison `lambda` function to `.apply()`, targeting the 'A', 'B', and 'C' columns specifically. This is a crucial step in transforming raw data into meaningful intelligence, identifying exactly which rows conform to the strict three-way equality condition.

The resulting new column, `all_matching`, is appended to the DataFrame. Its values reflect the definitive answer for each row: `True` if A equals B equals C; `False` otherwise. This method is highly optimized within the `Pandas` ecosystem, ensuring rapid processing even for millions of records by performing the comparison in a vectorized manner across the specified axis.

We can use the following code to create a new column called `all_matching` that returns `True` if all three columns match in a given row and `False` if they do not:

```
#create new column that displays whether or not all column values match
```

```
df = df.apply(lambda x: x.A == x.B == x.C, axis = 1)
```

```
#view updated DataFrame
```

```
print(df)
```

```
A B C all_matching
```

```
0 4 4 4 True
```

```
1 0 2 0 False
```

```
2 3 3 3 True
```

```
3 3 5 5 False
```

```
4 6 6 5 False
```

```
5 8 4 10 False
```

```
6 7 7 7 True
```

```
7 9 7 9 False
```

```
8 12 12 12 True
```

## Interpreting the Boolean Output

The final DataFrame clearly shows the results of the comparison. The `all_matching` column, being a `Boolean` series, is immediately useful for subsequent operations, such as filtering the DataFrame to select only the perfectly matched records (i.e., where `all_matching == True`).

A value of `True` confirms that for that row index, the value in column A is the same as B, and the value in B is the same as C. A value of `False` indicates that the condition `A=B=C` was not met, signaling a potential inconsistency or mismatch that needs further investigation. This simple Boolean flag transforms raw equality checks into actionable insights.

The new column called **`all_matching`** shows whether or not the values in all three columns match in a given row.

For example:

All three values match in the first row (4, 4, 4), so **`True`** is returned.

Not every value matches in the second row (0, 2, 0) because `0 != 2`, so **`False`** is returned.

In row 4, `A=6` and `B=6`, but `C=5`, illustrating that even if two columns match, the absence of a third match results in **`False`**.

## Conclusion and Advanced Applications

This methodology provides a highly effective and Pythonic solution for verifying three-way equality checks in Pandas. By utilizing the `.apply(lambda x: x.A == x.B == x.C, axis=1)` pattern, we streamline complex comparison logic into a single, declarative line of code, significantly improving the clarity and maintainability of our Python scripts for data analysis.

This robust foundation allows data professionals to move quickly from identifying inconsistencies to addressing them. Furthermore, the approach is highly flexible; if you needed to compare four or five columns, the syntax would simply extend (e.g., `x.A == x.B == x.C == x.D`). If the comparison needed to handle nuanced data types or required custom logic, the function passed to `.apply()` could be easily modified.

And so on.