

How to Find the Minimum Value Across Multiple Columns in Pandas

Authored by
stats writer

November 24, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Find the Minimum Value Across Multiple Columns in Pandas*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=100202>

Data analysis often requires us to calculate statistics not just vertically (column-wise), but horizontally (row-wise). When working with the [Pandas](#) library in [Python](#), a common requirement is identifying the lowest observation within a specific subset of columns for every single row in a [DataFrame](#). This operation is essential for tasks ranging from identifying performance bottlenecks to finding the most conservative estimate across several metrics.

To achieve this row-wise minimum calculation, we leverage the built-in [min\(\) function](#). The key to performing this operation successfully is understanding and correctly applying the `axis` parameter. By setting the `axis` parameter to 1, we instruct Pandas to compute the minimum value horizontally, thereby scanning across the specified columns for each individual record. This methodology efficiently returns a Pandas Series object containing the minimum result for every row. This is incredibly useful for finding the lowest statistical value observed for each distinct entry across your chosen metrics.

Core Methods for Calculating Row Minimums

There are primarily two highly efficient methods for extracting the minimum value across specific columns in a Pandas [DataFrame](#). Both approaches rely on indexing the desired columns and then applying the [min\(\) function](#) with the crucial argument `axis=1`. The choice between the methods generally depends on whether you need the resulting minimum values as a standalone Series or integrated directly back into the original DataFrame as a new column for persistent storage and further analysis.

The first method provides a quick calculation, yielding a Pandas Series that aligns perfectly with the index of the original DataFrame. This is useful for temporary checks or filtering steps. The second method, however, is often preferred when preparing data for robust modeling or analysis, as it keeps the derived metric contained within the primary data structure, minimizing the risk of misalignment or errors during subsequent processing steps. Understanding the implications of the `axis` argument is central to mastering these row-wise computations, as it dictates the direction of aggregation.

Method 1: Finding Minimum Value Across Multiple Columns (Series Output)

This approach is the most direct way to generate the minimum values across a subset of columns. We select the columns of interest using standard bracket notation, effectively creating a temporary DataFrame slice, and chain the `min()` method immediately afterward. Specifying `axis=1` tells Pandas to iterate over the column headers (the axis indexed by 1), calculating the minimum value horizontally across those selected columns, rather than the row indices (axis indexed by 0), which is the default behavior.

The resulting calculation efficiently generates a new Series. This Series maintains the indexing

structure of the original DataFrame, ensuring perfect alignment between the calculated minimum and its corresponding row context. This intermediate result is excellent for quick validation, performing statistical checks, or when the minimum calculation is only a transient step in a larger data pipeline where the final output is not intended to modify the original DataFrame structure.

```
df].min(axis=1)
```

Method 2: Adding New Column Containing Minimum Value Across Multiple Columns (DataFrame Modification)

Data scientists frequently need to incorporate newly calculated statistics directly into their existing data structure. This ensures that the derived metric is permanently associated with the corresponding record, creating a rich, feature-engineered dataset. To achieve this permanent integration, we utilize the exact same calculation used in Method 1, but we assign its output to a new column name within the DataFrame using standard assignment syntax.

The advantage of this method is the creation of a persistent feature, which can be easily used for subsequent filtering, visualization, or machine learning model training without needing to rejoin or align separate data structures. By assigning the output of the `min(axis=1)` operation to `df`, we leverage Pandas' optimized indexing to ensure that the minimum value is correctly inserted into the corresponding row across the entire dataset.

```
df = df].min(axis=1)
```

This approach solidifies data integrity and simplifies the process of tracking derived metrics alongside original observations, making the resulting DataFrame ready for complex, multi-stage analytical pipelines.

Demonstration Setup: Sample Pandas DataFrame Creation

To properly illustrate these two powerful techniques, we will utilize a sample Pandas DataFrame representing hypothetical player statistics in a sports context. This dataset includes measurements for `points`, `rebounds`, and `assists` across several unique players (A through G). Understanding the structure and numerical nature of the source data is the foundation for accurate analysis and successful application of row-wise functions.

We begin by importing the necessary library, aliasing it as `pd`, and then constructing the data structure using the `pd.DataFrame()` constructor. Notice the clear separation of categorical data (`player`) and the numerical metrics (`points`, `rebounds`, `assists`) that will be the focus of our minimum value calculation. The initialization ensures that we have a well-defined structure to

practice indexing and aggregation.

The resulting DataFrame, `df`, provides seven distinct observations, indexed from 0 to 6. Our goal in the following examples is strictly to find the minimum value achieved by each player when comparing only their `points` and `rebounds` attributes, thus demonstrating how to effectively select a subset of columns for localized calculation, deliberately ignoring the `assists` column.

```
import pandas as pd
```

```
#create DataFrame
```

```
df = pd.DataFrame({'player': ,  
'points': ,  
'rebounds': ,  
'assists': })
```

```
#view DataFrame
```

```
print(df)
```

```
player points rebounds assists
```

```
0 A 28 5 10
```

```
1 B 17 6 13
```

```
2 C 19 4 7
```

```
3 D 14 7 8
```

```
4 E 23 14 4
```

```
5 F 26 12 5
```

```
6 G 5 9 8
```

Example 1: Generating a Series of Row Minimums

In this first practical example, we implement Method 1 by focusing specifically on the `points` and `rebounds` columns. The precise objective is to calculate the lower of the two values for every player (row) and present this information as a new Series object. This is achieved by indexing the two columns using double brackets and then applying the `min()` function, confirming that `axis=1` is set to ensure the calculation scans horizontally.

Executing the following command yields a Series whose index aligns perfectly with the player rows in the original DataFrame. It is critical to recognize that this operation does not modify the DataFrame itself; it simply returns a computed result based on the specified columns.

```
#find minimum value across points and rebounds columns
```

```
df].min(axis=1)
```

```
0 5
1 6
2 4
3 7
4 14
5 12
6 5
dtype: int64
```

Interpreting the Resulting Minimum Series

The output Series provides a clear, concise summary of the minimum performance metric between points and rebounds for each row. Interpreting these results confirms that Pandas correctly performs the element-wise comparison across the columns defined by the list .

We can verify the output by comparing the original data:

For index 0 (Player A, 28 points, 5 rebounds), the minimum value is correctly identified as **5**. This shows the calculation works across disparate values.

For index 1 (Player B, 17 points, 6 rebounds), the minimum value is **6**.

For index 2 (Player C, 19 points, 4 rebounds), the minimum value is **4**.

This calculation is crucial for analysts seeking to identify the "weakest link" or baseline performance metric for each observation when multiple data points measure similar concepts. The consistency and speed of the vectorized operation underpin the efficiency of using the Pandas [library](#).

Example 2: Adding the Minimum Value as a New Feature Column

This example demonstrates the second method, which involves appending the calculated minimum directly into the existing [DataFrame](#) as a new feature column. This transformation is highly valuable for feature engineering, as it permanently links the derived summary statistic (the minimum value) to its source record.

To implement Method 2, we reuse the exact expression `df].min(axis=1)` and assign its output to the new column name `min_points_rebs`. This assignment operation automatically handles the necessary indexing and data alignment, creating a clean, updated DataFrame structure ready for further use.

```
#add new column that contains min value across points and rebounds columns
```

```
df = df.min(axis=1)
```

```
#view updated DataFrame
```

```
print(df)
```

```
player points rebounds assists min_points_rebs
```

```
0 A 28 5 10 5
```

```
1 B 17 6 13 6
```

```
2 C 19 4 7 4
```

```
3 D 14 7 8 7
```

```
4 E 23 14 4 14
```

```
5 F 26 12 5 12
```

```
6 G 5 9 8 5
```

Analyzing the Updated DataFrame Structure

The new column titled `min_points_rebs` now contains the minimum value across the points and rebounds columns for each row in the `DataFrame`. This successful integration confirms the power and simplicity of vectorized row-wise operations in Pandas.

For instance, examining Row 4 (Player E), which had 23 points and 14 rebounds, the newly created `min_points_rebs` column correctly registers the minimum as **14**. This derived column can now be seamlessly utilized for complex queries, such as filtering for players whose baseline metric (the minimum of points/rebounds) falls below a specific performance threshold, or visualizing the distribution of these minimum values relative to other features like assists.

Using the `min()` function combined with `axis=1` is the most optimized method for this calculation. It avoids the performance overhead associated with manual iteration methods like `apply()` or loops, especially when dealing with large datasets, ensuring fast and scalable data processing.

Conclusion: Mastering Row-Wise Operations in Pandas

Finding the minimum value across multiple columns in a `DataFrame` is a straightforward yet powerful operation made simple by the `Pandas` library. By correctly specifying the columns for comparison and setting the `axis` parameter to 1 within the `min()` method, data analysts can quickly and efficiently derive row-specific summary statistics.

Whether the objective is to generate a transient Series for immediate use (Method 1) or integrate a new persistent column for long-term feature storage (Method 2), these techniques are essential components of robust data preparation and exploratory data analysis (EDA). Mastering the

concept of axis manipulation in Pandas--understanding the difference between vertical (`axis=0`, the default) and horizontal (`axis=1`) operations--is crucial for efficient and high-performance data manipulation in modern data science workflows.

The ability to perform these row-wise aggregations allows for sophisticated feature creation and targeted data reduction, enabling analysts to focus on the most important metrics derived from complex raw data.

ARABPSYCHOLOGY.COM