

Learning Pandas: Finding the Index of Minimum Values with `idxmin()`

Authored by
Mohammed loot

November 13, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning Pandas: Finding the Index of Minimum Values with `idxmin()`*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=23998>

In the demanding world of data analysis using [Python](#), the capacity to swiftly pinpoint specific data points within vast datasets is fundamental to deriving meaningful insights. When manipulating a [Pandas DataFrame](#), data scientists frequently encounter the need to determine the exact [index position](#) corresponding to the minimum value along a given dimension. This crucial task is optimally handled by leveraging one of Pandas' most powerful built-in methods: the [`idxmin\(\)` function](#). This function is specifically engineered for efficiency, designed to return the index label of the first occurrence of the minimum value, making it an indispensable tool for identifying data extremes and outliers.

Understanding how to properly apply [`idxmin\(\)`](#) allows analysts to move beyond simple aggregation, providing the necessary context to reference and manipulate the records that hold those extreme low values. Whether you are tracking the lowest performance metric in a sports dataset or finding the cheapest stock price in a time series, mastering this function significantly enhances data manipulation workflows.

Role and Importance of `idxmin()`

The core objective of the `idxmin()` method is to locate and report the [index label](#) (which might be a row label or a column label, depending on the axis chosen) associated with the lowest numerical entry found within a Pandas Series or DataFrame. This functionality provides a critical distinction from standard aggregation functions like `min()`, which merely return the minimum numerical value itself. By providing the crucial contextual location--the index--`idxmin()` enables subsequent operations that require reference to the entire record or row containing that minimum value.

The high performance of this function is derived from its optimization within the underlying Pandas and [NumPy](#) libraries, guaranteeing rapid execution even when dealing with extremely large datasets comprising millions of records. When executed against a [Pandas DataFrame](#), the result is a Series where the index of the output Series represents the axis across which the operation was performed, and the resulting values indicate the index labels where the minimums were identified.

This method is particularly valuable in scenarios where multiple records share the same minimum value. In such cases, `idxmin()` adheres to a consistent rule: it returns the index label corresponding to the first occurrence of that minimum value encountered during the traversal of the specified axis. This behavior ensures deterministic results, which is vital for reproducible data analysis.

Syntax and Essential Parameters of `DataFrame.idxmin()`

The `idxmin()` function utilizes an intuitive and flexible syntax, granting precise control over the search mechanism for locating minimum values across a data structure. Understanding the parameters is essential for successful implementation, particularly when dealing with missing

values or mixed [data types](#). The comprehensive syntax structure is defined as follows:

DataFrame.idxmin(axis=0, skipna=True, numeric_only=False)

Below is a detailed examination of the purpose and default behavior of each primary parameter:

axis: This argument determines the direction of the reduction operation. A value of **0** (the default setting) specifies a column-wise reduction, meaning the function iterates down each column and returns the row index label that holds the minimum value. Conversely, setting **1** specifies a row-wise reduction, where the function iterates across each row and returns the column index label associated with that row's minimum value.

skipna: This boolean parameter governs how missing data (NA or NaN values) are handled during the minimum search. By default, it is set to **True**, ensuring that missing values are ignored, allowing the calculation to proceed with available numerical data. If set to **False**, the presence of any NA value in the segment being analyzed may propagate an NA result, depending on the computational context.

numeric_only: When this is set to **True**, the operation is strictly limited to columns containing numerical [data types](#), such as **float**, **int**, or **boolean**. The default setting is **False**, which instructs the function to attempt the reduction across all columns. However, as we will see, this can lead to errors if incompatible string or object types are present.

Practical Demonstration: Setting Up Sample Data

To properly illustrate the practical application and key nuances of the `idxmin()` method, we will first establish a working example using a sample [Pandas DataFrame](#). This dataset simulates performance statistics for several basketball players, recording their team affiliation, points scored, and assists made during recent games.

This initial setup is crucial as it creates a realistic scenario involving mixed data types (strings for 'team' and integers for the metrics), allowing us to demonstrate how to correctly handle data selection before applying numerical reduction methods like `idxmin()`.

```
import pandas as pd
```

```
#create DataFrame
```

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

```
#view DataFrame
```

```
print(df)
```

```
team points assists  
0 A 12 8  
1 A 18 10  
2 B 18 11  
3 B 22 11  
4 C 30 7  
5 C 41 12  
6 C 12 8
```

Addressing Non-Numeric Data and the `TypeError`

A frequent hurdle encountered when applying aggregation or reduction functions like `idxmin()` across an entire `DataFrame` is the presence of non-numerical columns. Since the concept of identifying the index of the "minimum" requires a numerical comparison, applying this function to a column containing strings (such as the 'team' column in our example) will inevitably fail. Pandas and its underlying [NumPy](#) operations lack a defined mechanism for mathematically comparing and identifying the minimum index among categorical or textual [data types](#) in this context.

If we attempt to execute `idxmin()` on the entire `DataFrame`, including the non-numeric 'team' column, the following error is raised:

```
#find index of minimum value in each column  
df.idxmin()
```

```
TypeError: reduction operation 'argmin' not allowed for this dtype
```

The resulting [TypeError](#) confirms that the underlying reduction operation (`argmin`, native to NumPy) cannot be performed when the column's data type is incompatible with numerical comparison. This vital feedback loop highlights a core requirement in data processing: numerical reduction methods must be explicitly targeted only at the numerical subsets of the data structure.

Finding Minimum Indices Along Columns (The Default `Axis=0`)

To successfully utilize `idxmin()` to find the index of the minimum value within our statistical metrics, we must first subset the `DataFrame`, selecting only the numerical columns ('points' and 'assists'). By calling `idxmin()` on this filtered subset, we execute the default operation, instructing Pandas to traverse down each column (`axis=0`) and return the row index label that corresponds to the minimum value found in that column.

This column-wise application is the most common use case for `idxmin()`, enabling rapid

identification of the specific record associated with the lowest value across various performance indicators. The corrected and successful approach is demonstrated below:

```
#find index of minimum value in points and assists columns  
df].idxmin()
```

```
points 0  
assists 4  
dtype: int64
```

The output, presented as a Pandas Series, provides a clear and concise summary of where the minimum values reside. Based on our original DataFrame, we can interpret this result definitively: for the **points** column, the minimum value (12) first appears at [index position 0](#). Similarly, for the **assists** column, the minimum value (7) is located at [index position 4](#).

Locating Minimum Indices Across Rows (Axis=1)

While column-wise analysis is often prioritized, the `idxmin()` function is equally effective when applied horizontally across rows. By modifying the execution to include the parameter **axis=1**, we instruct Pandas to pivot the operation. Instead of scanning vertically, the function scans within each individual row, seeking the minimum numerical value and returning the corresponding column label (or column name) where that minimum is situated. This is extremely valuable for comparative analysis within a single observation or record.

As before, we must ensure that only the numerical columns are selected prior to execution. The adjustment for row-wise analysis is subtle but fundamentally changes the scope of the search:

```
#find index of minimum value in each row  
df].idxmin(axis=1)
```

```
0 assists  
1 assists  
2 assists  
3 assists  
4 assists  
5 assists  
6 assists  
dtype: object
```

The resulting Series provides the column index for every row in the subsetted DataFrame. In this

specific basketball performance example, the output consistently identifies **assists** as the minimum attribute for every player. This result occurs because, for every single row, the numerical value recorded under the 'assists' column is numerically smaller than the corresponding value in the 'points' column (e.g., in row 0, 8 assists is less than 12 points). This powerful application demonstrates the function's ability to quickly identify which attribute holds the lowest value on a per-record basis, offering immediate comparative insight.

Note: For a complete understanding of advanced usage, including behavior with various [data types](#) and edge cases involving NaN values, always consult the official [Pandas documentation](#) for the `idxmin()` function.

Additional Resources for Pandas Mastery

To further enhance your data analysis skills, the following tutorials cover other common and essential tasks within the Pandas library:

Featured Posts

[5 Statistical Biases to Avoid](#)

April 25, 2024

[5 Free Statistics Courses for Beginners](#)

April 19, 2024

[5 MIT Statistics Courses That Are Free](#)

April 18, 2024

[5 Free Books to Learn Statistics](#)

April 18, 2024

[How to Use the info\(\) Method in Pandas](#)

April 12, 2024

[How to Use pct_change\(\) in Pandas](#)

April 12, 2024