

# Learn How to Replace NaN Values in Pandas with Data from Another Column

Authored by  
**Mohammed looti**

October 29, 2025

## RECOMMENDED CITATION

Mohammed looti (2025). *Learn How to Replace NaN Values in Pandas with Data from Another Column*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=5372>

## The Critical Challenge of Missing Data in Pandas

In the specialized field of [Pandas](#)-based data analysis and manipulation, encountering **missing data** is not merely a possibility--it is an inevitability. These informational voids can severely compromise the integrity, accuracy, and eventual utility of statistical models and reports if they are not addressed with careful precision. Within the [Pandas](#) ecosystem, these missing values are conventionally identified by the marker [NaN \(Not a Number\)](#). This special floating-point representation, inherited directly from the underlying [NumPy](#) library, signifies that a value is either undefined, unrepresentable, or simply absent.

The origins of missing data are diverse, ranging from preventable human errors, such as incorrect data entry or corrupted file transfers, to systemic issues inherent in the collection process, like non-responses in surveys or sensor failures. Effective data cleaning requires more than just acknowledging these gaps; it demands a strategic pipeline for remediation. Standard approaches often involve outright removal of incomplete records (dropping rows or columns), or statistical [imputation](#), where missing entries are replaced by constants, or the mean, median, or mode derived from the available data. However, these universal methods often discard valuable context or introduce statistical bias.

This comprehensive guide delves into a superior, contextually rich technique: utilizing the existing, valid data residing in a secondary column to fill the [NaN](#) entries in a primary target column. This approach is profoundly powerful when your dataset naturally includes fallback or secondary data fields that contain relevant information designed to complete the primary column's record. This ensures that the imputation process maintains the contextual integrity of each individual observation.

## Introducing the `fillna()` Method for Contextual Imputation

The [fillna\(\) method](#) is arguably one of the most essential functions in the [Pandas](#) library for managing data quality. It is designed specifically to handle and replace missing (**NaN**) values within both [DataFrames](#) and [Series](#) objects. While its basic applications involve filling missing slots with a simple scalar value (e.g., replacing all NaNs with 0 or a calculated mean), or employing sequence-based filling strategies like forward-fill (`ffill`) or backward-fill (`bfill`), its true sophistication emerges when it is supplied with another [Series](#) (i.e., another column) as its primary argument.

This capacity for column-to-column replacement transforms the [fillna\(\) method](#) from a simple replacement tool into an advanced mechanism for conditional, contextual [imputation](#). Instead of imposing a statistical average across all rows, this technique allows data from a reliable backup column to be slotted into the precise locations where the primary data is deficient. This is

particularly valuable in scenarios involving heterogeneous data sources, such as merging customer records where one source provides a primary identifier and a secondary source offers a reliable, alternative identifier for the same record.

The key functional distinction here is the reliance on index alignment. When one [Series](#) calls [fillna\(\)](#) and passes a second [Series](#), Pandas ensures that the replacement value is taken from the corresponding row index. This guarantees that the imputation is performed on a row-by-row basis, preserving the contextual link between the primary and secondary data points. This contextual maintenance is critical for maintaining data fidelity and ensuring that the resulting dataset is logical and consistent.

## Mechanism: How Pandas Aligns Columns for Replacement

Understanding the internal mechanics of index alignment is essential for effectively leveraging the column-to-column [fillna\(\) method](#). A [DataFrame](#) is inherently built upon aligned data structures, meaning that every [Series](#) (column) within it shares a common, immutable index. When you execute the replacement operation, [Pandas](#) does not rely on the sequential position of rows but rather on this shared index.

The process unfolds as follows: the target column (the one calling [fillna\(\)](#)) iterates through its elements, identifying every instance of **NaN**. For each identified missing value, Pandas retrieves the index label of that position. It then uses that exact same index label to look up the corresponding value in the source column (the column passed as the argument). If the value retrieved from the source column at that index is non-missing, it successfully replaces the **NaN** in the target column. If the source column also contains a missing value at that specific index, the target column's **NaN** is preserved.

This index-based alignment provides robust safety and flexibility. It means that even if you were to rearrange the rows of your [DataFrame](#), the column-to-column [fillna\(\)](#) operation would still correctly match the values based on their unique index identifiers, eliminating the risk of mismatched replacements. This mechanism contrasts sharply with operations that rely solely on positional indexing, which can fail if the data is sorted or filtered differently between steps.

## Practical Implementation: Step-by-Step `fillna()` Syntax

The syntax for executing this precise imputation technique is surprisingly concise, reflecting the efficiency of the [Pandas](#) library. The core operation is performed using a single, powerful line of code:

```
df = df.fillna(df)
```

To fully appreciate the mechanism, we must dissect the roles of the three main components in this statement. Firstly, **df** on the left side of the assignment operator (=) designates the destination. This is the column within your [DataFrame](#), conventionally named `df`, that will be updated in place. The resulting [Series](#) returned by the right side of the equation will overwrite the original contents of `col1`, now containing fewer, or zero, **NaN** values.

Secondly, the sequence `df.fillna(...)` represents the function call. We are explicitly calling the [fillna\(\) method](#) directly on the target column, `col1`. This initiates the process of scanning `col1` for all missing entries. Importantly, this operation generates a \*new\* [Series](#) with the replacements applied, which is why the assignment back to `df` is necessary to persist the changes.

Finally, (`col2`) is the vital argument supplied to [fillna\(\)](#). This source column instructs [Pandas](#) on where to find the replacement data. As discussed, the alignment between `col1` and `col2` is strictly governed by their shared index. If `col1` has a **NaN** at index 5, the value from `col2` at index 5 is used to fill the gap. This clarity of purpose makes the column-to-column fill method far more deterministic and reliable than blanket imputation techniques.

## Demonstration: Filling NaNs with a Real-World Example

To provide a concrete understanding of this imputation technique, we will now walk through a practical example, setting up a sample [Pandas DataFrame](#) with deliberately introduced missing values and then applying the solution.

The foundational step involves importing the necessary libraries: [NumPy](#) is required to generate the specific `np.nan` values that Pandas recognizes as missing data, and [Pandas](#) is used for the [DataFrame](#) construction and manipulation. We create a simple DataFrame representing two competing teams, where the primary column, `team1`, has missing entries:

```
import numpy as np
import pandas as pd
```

```
#create DataFrame with some NaN values
df = pd.DataFrame({'team1': ,
'team2': })
```

```
#view DataFrame
df
```

```
team1 team2
0 Mavs Spurs
1 NaN Lakers
```

2 Nets Kings  
3 Hawks Celtics  
4 NaN Heat  
5 Jazz Magic

The visualization of the initial [DataFrame](#) clearly shows that `team1` is incomplete at index 1 and index 4. Our goal is to use the reliable and complete data present in the `team2` column to fill these precise gaps. The [fillna\(\) method](#) is applied directly to `team1`, using `team2` as the source of replacement values.

Executing the imputation line yields the following result:

```
#fill NaNs in team1 column with corresponding values in team2 column
```

```
df = df.fillna(df)
```

```
#view updated DataFrame
```

```
df
```

```
team1 team2  
0 Mavs Spurs  
1 Lakers Lakers  
2 Nets Kings  
3 Hawks Celtics  
4 Heat Heat  
5 Jazz Magic
```

The resulting [DataFrame](#) confirms the successful operation. At index 1, where `team1` was **NaN**, it is now populated by 'Lakers' (the value from `team2` at index 1). Similarly, at index 4, the missing entry is replaced by 'Heat'. This example perfectly illustrates the precision and effectiveness of using column-to-column imputation to ensure data completeness while respecting row-level context.

## Advanced Techniques and Handling Edge Cases

While the basic column-to-column imputation is powerful, data professionals must be aware of advanced techniques and specific edge cases to ensure robust data cleaning pipelines.

One common scenario involves using **multiple fallback columns**. If you have a hierarchy of preference--say, you want to fill `col1` first with `col2`, and then any remaining gaps with `col3`--you can easily chain the [fillna\(\) method](#) calls. The expression `df = df.fillna(df).fillna(df)`

executes the fills sequentially. The first call replaces `col1`'s NaNs using `col2`. The result of this first operation is then used as the starting point for the second `fillna()` call, which uses `col3` to fill any gaps that still remain (i.e., where both `col1` and `col2` were missing).

A crucial edge case involves **concurrent missing values**. If the target column (`col1`) has a **NaN** at a specific index, and the source column (`col2`) also has a **NaN** at that \*exact same index\*, the target column will retain the **NaN** after the operation. The [fillna\(\) method](#) is designed only to replace missing values with non-missing values from the argument [Series](#). Therefore, robust data cleaning might require preliminary steps to ensure the fallback columns are as complete as possible.

Finally, **data type compatibility** is a practical consideration. [DataFrames](#) will automatically attempt type coercion when mixing data types. For example, if you fill a numeric column (like `float64`, which holds NaNs) with values from a string column, the resulting column will typically be converted to the `object` dtype (string). It is essential to verify the data types using `df.dtypes` before and after imputation to ensure that the resulting column type is suitable for subsequent analysis or modeling steps.

## Conclusion and Summary of Best Practices

Leveraging the [fillna\(\) method](#) with a [Series](#) argument represents a sophisticated and highly effective strategy for handling missing data in a [Pandas DataFrame](#). This technique transcends the limitations of generalized imputation methods by providing precise, contextually grounded replacements based on existing, related data within the dataset.

The primary benefit of this approach lies in its reliance on index alignment, guaranteeing that replacement values are accurately matched to the corresponding records, thereby enhancing the overall completeness and reliability of the data. As a best practice, always evaluate the logical relationship between the target and source columns before implementation. Only use a fallback column if its values genuinely serve as meaningful and accurate substitutes for the missing information in the primary column. Mastering this technique is indispensable for any data professional seeking to build clean, high-quality datasets for advanced statistical analysis and machine learning.

## Additional Resources

For a comprehensive understanding of the [fillna\(\) function](#) and its various parameters, refer to the official [Pandas documentation](#). This resource offers in-depth details on all available options, including different fill methods and handling of specific data types.

The following tutorials explain how to perform other common operations in [Pandas](#):

[How to add new columns to a DataFrame](#)

[How to select rows and columns in a DataFrame](#)

[How to calculate descriptive statistics in a DataFrame](#)