

# Learning to Calculate Averages in MongoDB with Aggregation Pipelines

Authored by  
**Mohammed loot**

October 31, 2025

## RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Calculate Averages in MongoDB with Aggregation Pipelines*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=6777>

## Introduction to Averaging Data in MongoDB

Calculating the average value of a specific field is a foundational requirement in virtually all forms of data analysis, providing immediate and valuable statistical insights into large datasets. Within the NoSQL environment of [MongoDB](#), this complex operation is executed with high efficiency using the powerful, multi-stage [Aggregation Pipeline](#). This comprehensive guide is designed to clarify the mechanics of aggregation and walk you through two essential methods for computing averages: deriving a single global average across an entire [collection](#), and calculating detailed grouped averages based on unique field values.

Whether your goal is to gain a rapid statistical overview of system performance metrics or to conduct a granular analysis of average sales figures segmented by region, the [Aggregation Pipeline](#) provides the necessary flexibility and robust processing capabilities. By mastering the application of key operators, specifically [\\$group](#) for structuring data and [\\$avg](#) for computation, you can transform raw data stored in your [documents](#) into precise, actionable statistical measures.

The subsequent sections of this article will provide clear theoretical explanations, reinforced by practical, ready-to-implement code examples. These demonstrations will illustrate how to construct and execute effective aggregation queries, ensuring you can confidently and accurately derive meaningful average values from any [document](#) structure within your [collection](#).

## Understanding the MongoDB Aggregation Pipeline

Before attempting specific averaging techniques, it is vital to establish a strong understanding of the [MongoDB Aggregation Pipeline](#) framework. This framework functions much like a standard manufacturing assembly line: it accepts input data (the [documents](#) from a collection), processes them through a sequence of defined stages, and ultimately returns computed, aggregated results. Each stage in the pipeline performs a specialized operation, such as filtering, projecting, sorting, or grouping.

For the specific task of calculating averages, two primary aggregation operators form the core of the required pipeline logic. These operators work in tandem to first categorize the data and then perform the mathematical calculation:

**\$group**: This is a fundamental stage operator that collects all input [documents](#) and segments them based on a specified identifier expression defined in the `_id` field. After grouping, it applies accumulator expressions (like `$avg`) to the documents contained within each new group.

**\$avg**: This is an accumulator operator used exclusively within the [\\$group](#) stage. Its function is straightforward: it calculates the arithmetic mean (average) of all numeric values produced by the specified expression within that particular group.

The overall syntax for initiating this process involves calling the `db.collection.aggregate()` [method](#). This method accepts a JavaScript array containing one or more aggregation stages. For most averaging scenarios, the pipeline will consist of a single, powerful `$group` stage, as demonstrated in the practical examples provided later in this article.

## Method 1: Calculating a Global Average

The simplest form of data aggregation is computing a single, overall average for a numerical field across every [document](#) stored in a given [collection](#). This global average provides a quick, high-level summary statistic, often used as a benchmark or baseline for further analysis.

To achieve this global calculation using the [Aggregation Pipeline](#), we utilize the `$group` aggregation stage, but with a unique configuration for the grouping key. We set the `_id` field to `null`. By specifying `_id: null`, the pipeline is instructed to treat \*all\* incoming [documents](#) as belonging to a single, unified group. This effectively bypasses category-based grouping and allows the `$avg` accumulator to calculate one single, comprehensive average for the entire collection.

The following syntax represents the structure required to calculate a global average, where `$valueField` is replaced by the actual field name you wish to average:

```
db.collection.aggregate()
```

Upon execution, this query will return exactly one [document](#). This result [document](#) will contain the single computed average, typically aliased as `avg_val`, and will feature `_id: null`, confirming its status as the global summary statistic for the entire dataset within the specified [collection](#).

## Method 2: Calculating Averages by Group

While a global average is useful, real-world data analysis often demands a more granular approach, requiring averages to be broken down by distinct categories, such as product type, geographic region, or user status. [MongoDB](#) expertly handles this requirement by allowing you to define a grouping key based on the value of a specific field within your [documents](#). This method facilitates detailed comparative analysis between different segments of your data.

To implement this powerful grouped aggregation, we modify the definition of the `_id` parameter within the `$group` stage. Instead of using `null`, we set `_id` to reference the field containing the category information, for instance, `_id: "$groupField"`. This directive tells the [Aggregation Pipeline](#) to process the collection, creating a new group for every unique value found in `groupField`. Subsequently, the `$avg` accumulator calculates the average specifically for the subset of documents belonging to that distinct group.

The general syntax below demonstrates how to calculate an average partitioned by the values held in `$groupField`, using the [db.collection.aggregate\(\) method](#):

### **db.collection.aggregate()**

Unlike the global average, this query will typically return multiple result [documents](#). Each output document corresponds to a unique value found in the `groupField` (now represented by the `_id`) and provides the calculated average value (`avg_val`) specific to that category. This method is fundamental for comparative analysis and business intelligence reporting.

## **Setting Up Our Example Data**

To effectively demonstrate both the global and grouped average aggregation methods, we will utilize a small, focused dataset. We will create a sample [collection](#) named `teams`, which is designed to store basic statistics related to different basketball team performances, specifically tracking points scored and rebounds achieved.

The following commands will insert five distinct [documents](#) into the `teams` [collection](#). These documents represent individual game statistics or player records, providing the numerical data necessary for our averaging calculations:

```
db.teams.insertOne({team: "Mavs", points: 30, rebounds: 8})
db.teams.insertOne({team: "Mavs", points: 30, rebounds: 12})
db.teams.insertOne({team: "Spurs", points: 20, rebounds: 7})
db.teams.insertOne({team: "Spurs", points: 25, rebounds: 5})
db.teams.insertOne({team: "Spurs", points: 25, rebounds: 9})
```

This dataset structure allows us to perform two types of calculations: first, determining the average points across all records regardless of the team (the global average), and second, calculating the average points specifically broken down by the two distinct teams, "Mavs" and "Spurs" (the grouped average).

## **Practical Example 1: Global Average of Points**

Our first practical application demonstrates Method 1: calculating the overall average of the `points` field across the entire `teams` collection. This operation ignores the `team` categorization and computes a single, definitive average value that summarizes the performance across all recorded statistics.

We execute the [db.collection.aggregate\(\) method](#) using a pipeline containing a single

`$group` stage. By setting `_id: null`, we force all five input **documents** into one group. The `$avg` accumulator then calculates the average of the `$points` field within this unified group:

```
db.teams.aggregate()
```

Executing this query against our sample data yields the following result:

```
{ _id: null, avg_val: 26 }
```

The resulting **document** confirms that the overall average value for the `points` field across all records is **26**. The presence of `_id: null` in the output is the standard indicator for a global aggregation result. We can verify this computation by manually summing the points ( $30 + 30 + 20 + 25 + 25 = 130$ ) and dividing by the total count of documents (5), which confirms the accuracy of the **MongoDB** aggregation.

## Practical Example 2: Grouped Average of Points by Team

Moving to a more insightful level of analysis, we now apply Method 2 to calculate the average `points` for each distinct team recorded in the `teams` **collection**. This requires partitioning the data based on the values in the `team` field before performing the average calculation.

To achieve this segregation, the `db.collection.aggregate()` **method** is called with a single `$group` stage, where the `_id` is explicitly set to reference the grouping field: `_id: "$team"`. This configuration instructs the pipeline to generate a separate average for every unique team name encountered. The `$avg` accumulator then processes only the points belonging to the documents in its respective group:

```
db.teams.aggregate()
```

Executing this grouped query provides the following structured results, offering a clear comparative breakdown of average performance:

```
{ _id: 'Spurs', avg_val: 23.333333333333332 }  
{ _id: 'Mavs', avg_val: 30 }
```

These results reveal specific insights that the global average masked. We can clearly observe the following performance metrics:

The **Spurs** team has an average points score of approximately **23.33** (derived from the sum of 70 divided by 3 entries).

The **Mavs** team maintains a perfect average of **30** points (derived from the sum of 60 divided by 2 entries).

This type of grouped aggregation is crucial for identifying key differences and trends within subsets of your data, providing far greater analytical depth than a single global statistic can offer.

## Conclusion and Further Learning

The [MongoDB Aggregation Pipeline](#) stands as an exceptionally versatile and powerful framework central to effective data analysis within the [MongoDB](#) ecosystem. Calculating averages, whether globally using `_id: null` or by specific categories using a field reference in the `_id`, demonstrates a fundamental capability of this pipeline. By mastering the application of the [db.collection.aggregate\(\) method](#) in conjunction with the `$group` and `$avg` operators, developers and analysts can efficiently transform raw [documents](#) into precise, meaningful statistical insights.

These techniques form the bedrock for more advanced aggregation queries involving multiple stages, such as filtering data using `$match` before grouping it, or sorting the resulting averages using `$sort`. For those seeking to deepen their expertise and explore the full range of aggregation possibilities, we strongly recommend consulting the official [MongoDB documentation for \\$avg](#) and other related aggregation operators.

## Additional Resources

To further expand your knowledge of [MongoDB](#) and its powerful features, explore these related tutorials on common aggregation and data manipulation operations: