

Understanding Mean Absolute Error (MAE) vs. Root Mean Squared Error (RMSE) in Regression Analysis

Authored by
Mohammed Iooti

November 2, 2025

RECOMMENDED CITATION

Mohammed Iooti (2025). *Understanding Mean Absolute Error (MAE) vs. Root Mean Squared Error (RMSE) in Regression Analysis*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8460>

The Imperative Role of Error Metrics in Regression Analysis

[Regression models](#) are foundational tools in statistics and data science, utilized primarily to model and quantify the relationship between one or more [predictor variables](#) and a designated response variable. These powerful models strive to generate a mathematical representation that most accurately reflects the patterns observed in the underlying data structure, forming the bedrock of prediction and forecasting across various disciplines.

Once a regression model is trained, the most crucial subsequent phase is rigorous evaluation. Analysts must determine how effectively the model generalizes from the training data and forecasts unseen outcomes. This assessment requires quantifying the inherent gap between the model's predictions and the actual observed values--a discrepancy universally referred to as the error. Reliable measurement of this error is essential for selecting the optimal model for any given application.

Among the multitude of metrics available for quantifying model fit in [machine learning](#) and statistics, the [Mean Absolute Error \(MAE\)](#) and the [Root Mean Squared Error \(RMSE\)](#) stand out as the most widely adopted standards. While both metrics provide a single value representing the average magnitude of prediction errors, their distinct mathematical constructions lead to fundamentally different interpretations, particularly concerning their sensitivity to extreme values or data anomalies.

Mean Absolute Error (MAE): The Linear Approach to Accuracy

The Mean Absolute Error (MAE) offers a refreshingly robust and straightforward approach to measuring prediction accuracy. Formally, MAE is calculated by taking the average of the absolute differences between the actual observed values and the corresponding values predicted by the model. By utilizing absolute values, it avoids the cancellation of positive and negative errors, ensuring that all deviations contribute to the final score.

MAE is highly favored for its intuitive interpretability: it directly represents the average error magnitude expressed in the native units of the response variable. Crucially, the MAE structure is linear, meaning that the penalty for an error is directly proportional to its size. An error of 10 contributes exactly twice as much to the final metric as an error of 5. A low MAE score signifies that, on average, the model's predictions are tightly clustered around the true values.

This linearity grants MAE a significant advantage: its relative immunity to extreme values. Since large errors are not disproportionately magnified through squaring, MAE is considered far more robust to [outliers](#) compared to its squared counterpart, RMSE. This makes MAE the preferred metric when the dataset is known to contain noise or unavoidable measurement errors, and the priority is achieving the best overall prediction accuracy for the majority of observations.

The mathematical formulation for MAE is defined as follows:

$$\text{MAE} = 1/n * \sum |y_i - \hat{y}_i|$$

Where the variables are defined as:

\sum is the summation symbol, indicating the sum of all observations.

y_i is the observed actual value for the i th observation.

\hat{y}_i is the value predicted by the model for the i th observation.

n is the total sample size used for evaluation.

Deconstructing Root Mean Squared Error (RMSE): The Quadratic Penalty

The Root Mean Squared Error (RMSE) is another indispensable metric used extensively for [regression model](#) evaluation. Its calculation involves a three-stage process: first, calculating the squared differences between observed and predicted values; second, determining the average (mean) of these squared differences; and finally, taking the square root to return the metric to the original units of measurement.

The squaring operation embedded within the RMSE formula is the defining characteristic of this metric. This process ensures that large errors are heavily penalized, contributing exponentially more to the total error score than small errors. For instance, an error of 10 is penalized 100 times, while an error of 5 is penalized only 25 times. This characteristic makes RMSE acutely sensitive to the presence of significant deviations or [outliers](#) within the dataset.

Like MAE, RMSE is measured in the same units as the response variable, maintaining clear interpretability. However, due to the inherent squaring effect, RMSE will virtually always yield a higher numerical value than MAE for the identical set of predictions. A lower RMSE value indicates a superior model fit, particularly in domains where the primary objective is the minimization of substantial, potentially catastrophic, errors.

The mathematical formulation for RMSE is calculated as:

$$\text{RMSE} = \sqrt{\sum (y_i - \hat{y}_i)^2 / n}$$

Where the terms correspond to:

\sum is the summation symbol across all observations.

y_i is the observed actual value for the i th observation.

\hat{y}_i is the value predicted by the model for the i th observation.

n is the total sample size.

Visualizing the Difference: A Practical Calculation Example

To truly grasp the mechanical difference between these two metrics, we can examine a simple practical scenario. Imagine a regression model tasked with predicting the number of points scored by ten different players in a basketball game. We will compare the model's predictions against the actual observed scores to calculate the necessary errors for both MAE and RMSE.

The table below presents the predicted points alongside the actual points scored, allowing for the calculation of the deviations required by both formulas:

Predicted Points (\hat{y}_i)	Actual points (y_i)
14	12
15	15
18	20
19	16
25	20
18	19
12	16
12	20
15	16
22	16

Applying the defined formulas to this sample yields an MAE of **3.2**. This result means that, on average, the absolute difference between the forecast and the actual score is 3.2 points. This metric provides a clear, average measure of error magnitude across the dataset.

For the exact same data, the RMSE is calculated to be **4**. While both values give an indication of the typical prediction difference, the RMSE's slightly higher value already hints at the underlying arithmetic difference, reflecting the penalization of larger errors through the squaring process.

The Critical Distinction: Sensitivity to Outliers and Error Cost

The fundamental philosophical distinction between RMSE and MAE lies in how they reflect the real-world cost associated with large prediction errors. Because RMSE relies on squared differences, it establishes a non-linear penalty structure. This makes RMSE profoundly more sensitive to large errors than MAE, which maintains a linear penalty structure.

This difference is pivotal when connecting the statistical metric to the business or domain objectives via the [cost function](#). If the consequence of being "off" by 20 units is deemed significantly more than twice as detrimental as being off by 10 units (a common scenario in engineering or finance where large failures are disproportionately costly), then RMSE is the superior choice because its squaring mechanism intrinsically reflects this exponential cost structure.

Conversely, if the error cost is strictly linear--meaning that an error of 20 is exactly twice as damaging as an error of 10--then MAE is the metric that accurately models the desired cost function, providing an unbiased average of the deviations.

To provide a dramatic illustration of RMSE's sensitivity, let us introduce a single, extreme [outlier](#) into our previous basketball prediction example, representing a failure to predict one player's score accurately:

Predicted Points (\hat{y}_i)	Actual points (y_i)
14	12
15	15
18	20
19	16
25	20
18	19
12	16
12	20
15	16
22	76

With the inclusion of this single extreme observation (where the actual score was 76, but the prediction was only 22), the calculated metrics change drastically:

MAE: 8

RMSE: 16.4356

While the MAE increased predictably, rising from 3.2 to 8, the RMSE experienced a massive surge, climbing from 4 to over 16.4. This exponential increase confirms that RMSE's use of squared differences heavily amplifies the penalty associated with the large deviation. This makes RMSE the definitive metric for model evaluation when the priority is detecting and severely

penalizing large prediction errors, even if they are infrequent.

Strategic Selection: Choosing the Right Metric for Your Model

In practical data science applications, the decision between MAE and RMSE must be dictated by the specific requirements of the domain and the level of tolerance for large errors. If the integrity or safety of a system relies heavily on avoiding extreme forecast failures--such as predicting structural loads in construction, forecasting severe weather events, or managing financial risk--RMSE is almost always the preferred choice. By optimizing for RMSE, the model is compelled to focus its effort on minimizing the worst-case scenarios.

Conversely, if the objective is to achieve the most stable and representative prediction across the entire dataset, especially when the data is noisy or contains unavoidable measurement anomalies, MAE provides a more robust performance indicator. MAE ensures that the model is not unduly skewed by a few extreme observations, resulting in better predictive performance for the majority of data points.

When evaluating multiple candidate [regression models](#), methodological consistency is non-negotiable. It is fundamentally critical to calculate and compare the exact same metric across all models under consideration. Comparing the MAE of one model against the RMSE of another model is statistically unsound and will inevitably lead to erroneous conclusions regarding which model is superior.

The standard procedure requires fitting several candidate models to the data, calculating one chosen metric (e.g., RMSE) for every model, and then selecting the model that yields the lowest value for that metric. This model is generally considered the best fit because its predictions are, on average, closest to the actual observed values based on the specific error minimization strategy (linear vs. quadratic penalty) chosen by the analyst.

Summary and Resources for Further Exploration

The ultimate decision between MAE and RMSE represents a choice between two distinct philosophies of error minimization. MAE employs a linear cost function, valuing proportional accuracy across all observations, which offers robustness against outliers. Conversely, RMSE employs a quadratic cost function, disproportionately penalizing large deviations, thereby prioritizing the mitigation of extreme, high-cost errors. A deep understanding of this mathematical difference is paramount for effective model selection and evaluation.

For further study and implementation, the following resources provide detailed guidance on calculating MAE using various statistical software packages:

The following tutorials explain how to calculate RMSE using different statistical software: