

Which metric, RMSE or R-Squared, should you use when evaluating a model's performance?

Authored by
stats writer

May 2, 2024

RECOMMENDED CITATION

stats writer (2024). *Which metric, RMSE or R-Squared, should you use when evaluating a model's performance?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=142010>

When evaluating a model's performance, it is important to consider both the RMSE (Root Mean Squared Error) and the R-Squared metric. These metrics provide different insights into the effectiveness of the model and should be used in conjunction with each other.

RMSE measures the average difference between the predicted values and the actual values, and a lower RMSE indicates a better fit. On the other hand, R-Squared measures the proportion of the variance in the dependent variable that is explained by the independent variables. A higher R-Squared indicates a better fit.

Therefore, the metric that should be used depends on the specific goals and requirements of the model. If the goal is to minimize errors and accurately predict values, RMSE should be the primary metric used. However, if the goal is to explain the relationship between variables and determine how well the model fits the data, R-Squared should be given more weight. In general, both metrics should be considered when evaluating a model's performance in order to gain a comprehensive understanding of its effectiveness.

RMSE vs. R-Squared: Which Metric Should You Use?

Regression models are used to quantify the relationship between one or more predictor variables and a response variable.

Whenever we fit a regression model, we want to understand how well the model "fits" the data. In other words, how well is the model able to use the values of the predictor variables to predict the value of the ?

Two metrics that statisticians often use to quantify how well a model fits a dataset are the root mean squared error (RMSE) and the R-squared (R²), which are calculated as follows:

RMSE: A metric that tells us how far apart the predicted values are from the observed values in a dataset, on average. The lower the RMSE, the better a model fits a dataset.

It is calculated as:

$$\text{RMSE} = \sqrt{\sum(P_i - O_i)^2 / n}$$

where:

Σ is a symbol that means "sum"
 P_i is the predicted value for the i th observation
 O_i is the observed value for the i th observation
 n is the sample size

R²: A metric that tells us the proportion of the variance in the response variable of a regression model that can be explained by the predictor variables. This value ranges from 0 to 1. The higher the R² value, the better a model fits a dataset.

It is calculated as:

$$R^2 = 1 - (\text{RSS}/\text{TSS})$$

where:

RSS represents the sum of squares of residuals
TSS represents the total sum of squares

RMSE vs. R2: Which Metric Should You Use?

When assessing how well a model fits a dataset, it's useful to calculate *both* the RMSE and the R2 value because each metric tells us something different.

On one hand, RMSE tells us the typical distance between the predicted value made by the regression model and the actual value.

On the other hand, R2 tells us how well the predictor variables can explain the variation in the response variable.

Price	Sq. Footage	# Bathrooms	# Bedrooms
\$ 459,000	3,240	4	5
\$ 394,000	3,200	3	6
\$ 285,000	2,500	4	4
\$ 245,000	2,634	2	3
\$ 356,000	2,800	3	3
...

Now suppose we'd like to use square footage, number of bathrooms, and number of bedrooms to predict

house price.

We can fit the following regression model:

$$\text{Price} = \beta_0 + \beta_1(\text{sq. footage}) + \beta_2(\text{\# bathrooms}) + \beta_3(\text{\# bedrooms})$$

Now suppose we fit this model and then calculate the following metrics to assess the goodness of fit of the model:

RMSE: 14,342 R²: 0.856

The RMSE value tells us that the average deviation between the predicted house price made by the model and the actual house price is \$14,342.

The R² value tells us that the predictor variables in the model (square footage, # bathrooms, and # bedrooms) are able to explain 85.6% of the variation in the house prices.

When determining if these values are "good" or not, we can compare these metrics to alternative models.

For example, suppose we fit another regression model

that uses a different set of predictor variables and calculate the following metrics for that model:

RMSE: 19,355 R²: 0.765

We can see that the RMSE value for this model is greater than the previous model. We can also see that the R² value for this model is less than the previous model. This tells us that this model fits the data worse than the previous model.

Summary

Here are the main points made in this article:

Both RMSE and R² quantify how well a regression model fits a dataset. The RMSE tells us how well a regression model can predict the value of the response variable in absolute terms while R² tells us how well a model can predict the value of the response variable in percentage terms. It's useful to calculate both the RMSE and R² for a given model because each metric gives us useful information.