

Learning AIC: A Practical Guide to Calculating Akaike Information Criterion in R with Examples

Authored by
Mohammed loot

November 4, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning AIC: A Practical Guide to Calculating Akaike Information Criterion in R with Examples*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=9772>

Understanding the Akaike Information Criterion (AIC)

The [Akaike Information Criterion \(AIC\)](#) stands as a foundational metric in quantitative statistics, serving as an indispensable tool for [model selection](#). When researchers evaluate multiple competing [regression models](#) designed to explain a specific dataset, AIC provides a robust, relative measure of the quality of each statistical model. It helps determine which candidate model is most likely to be the best approximation of the underlying reality that generated the observed data.

The AIC was groundbreakingly developed by statistician Hirotugu Akaike. Conceptually, it quantifies the estimated information loss that occurs when a chosen statistical model is employed to represent the true process underlying the data. This metric is fundamentally crucial because it mandates a necessary trade-off: balancing the model's fidelity (or "goodness of fit," meaning how accurately it explains the variance in the data) against its inherent complexity (determined by the number of independent variables or parameters utilized).

A primary objective of the AIC is to mitigate the risk of [overfitting](#). Overfitting occurs when a model is excessively complex, fitting random noise or anomalies within the training data instead of capturing the underlying generalizable structure. The operational principle of the Akaike Criterion is straightforward: among a set of candidate models, the one exhibiting the **lowest AIC value** is mathematically considered the preferred choice, offering the most optimal balance between predictive accuracy and parsimony.

The Mathematical Foundation: AIC Formula Explained

To fully appreciate the selection process, it is essential to understand the concise mathematical definition of the standard AIC. This formula elegantly merges two critical measurements: the measure of model fit and the count of parameters required. The goal is to minimize the final AIC score, which is calculated based on the maximum [log-likelihood](#) achieved by the model and the complexity penalty imposed.

The standard formula used for calculating the Akaike Information Criterion is expressed as follows:

$$AIC = 2K - 2\ln(L)$$

The meaning of the constituent variables is paramount to interpreting the score, particularly the interplay between complexity (K) and fit (L).

K: This variable denotes the number of estimated **model parameters**, often referred to as the model's degrees of freedom. In the context of standard linear [regression analysis](#), K typically encompasses the intercept (β_0), the total number of predictor variables, and the error

variance term. Consequently, a model utilizing a single predictor variable would generally have a K value of 3.

$\ln(L)$: This term represents the maximum value attained by the model's [log-likelihood](#) function. The log-likelihood is a measure that quantifies the probability of observing the given data set under the assumption that the model is correct. Most specialized statistical environments, including the widely used [R programming language](#), automatically calculate this maximum likelihood estimate during the model fitting procedure.

The term K functions explicitly as a penalty mechanism. As the complexity of the model increases (i.e., K gets larger), the penalty grows, offsetting improvements in the log-likelihood that might result purely from adding extraneous variables. This design ensures that if two distinct models demonstrate an equally strong fit to the data (i.e., similar $\ln(L)$ values), the model with fewer parameters (smaller K) will yield a lower AIC score and will thus be scientifically preferred for its parsimony.

Implementing AIC Calculation in R

While the base installation of [R](#) provides a standard `AIC()` function, the systematic calculation and rigorous comparison of multiple candidate models are best executed using specialized packages tailored for model selection. For efficient analysis and reporting of AIC scores, we will leverage the powerful `aictab()` function, which is nested within the highly recommended [AICcmodavg](#) package. This package streamlines the process of generating comparative model selection tables.

A significant benefit of utilizing `aictab()` is its automatic calculation of the **AICc**, which is the Akaike Information Criterion corrected for small sample sizes. This correction is critically important and highly recommended in situations where the data is limited--specifically, when the ratio of the number of observations (n) to the number of estimated parameters (K) is less than 40 ($n/K < 40$). Employing the AICc ensures that the model selection process remains statistically reliable even when working with constrained datasets.

The subsequent section provides a comprehensive, step-by-step example demonstrating the practical application of this method in [R](#). The workflow involves three distinct steps: first, fitting all candidate [regression models](#) using standard functions; second, compiling these models into a single list; and finally, executing the powerful `aictab()` function to generate the formal, comparative model selection table.

Detailed Example: Applying AIC to the mtcars Dataset

For a practical demonstration of model selection using AICc, we will utilize the intrinsic [mtcars dataset](#), which is readily available within [R](#). Our objective is to construct three competing linear models aimed at predicting the dependent variable, miles per gallon (`mpg`), by using various

combinations of engine and vehicle performance characteristics, such as displacement, weight, and horsepower.

We define three distinct candidate models, each varying in complexity and the specific set of predictor variables included:

Model 1: This comprehensive model uses four predictors: `disp` (engine displacement), `hp` (gross horsepower), `wt` (vehicle weight), and `qsec` (quarter mile time).

Model 2: This simpler model incorporates only two predictors: `disp` and `qsec`.

Model 3: This second simple model uses the variables `disp` and `wt`.

The initial step involves fitting these linear models using R's standard `lm()` function. It is important to note that the AIC calculation depends entirely on the resulting maximum likelihood estimates generated during this fitting process.

Fit three models predicting MPG using the mtcars dataset

```
model1 <- lm(mpg ~ disp + hp + wt + qsec, data = mtcars)
```

```
model2 <- lm(mpg ~ disp + qsec, data = mtcars)
```

```
model3 <- lm(mpg ~ disp + wt, data = mtcars)
```

After fitting the models, we proceed to load the necessary library, organize our candidate models into a list structure, and assign descriptive names for clarity in the output table. Finally, we invoke the `aictab()` function to execute the comparative analysis, which will present the crucial AICc values for all three models side-by-side.

library(AICcmodavg)

```
# Define list of candidate models
```

```
models <- list(model1, model2, model3)
```

```
# Specify corresponding model names
```

```
mod.names <- c('disp.hp.wt.qsec', 'disp.qsec', 'disp.wt')
```

```
# Calculate AICc for each model
```

```
aictab(cand.set = models, modnames = mod.names)
```

Model selection based on AICc:

```
K AICc Delta_AICc AICcWt Cum.Wt LL
disp.hp.wt.qsec 6 162.43 0.00 0.83 0.83 -73.53
disp.wt 4 165.65 3.22 0.17 1.00 -78.08
disp.qsec 4 173.32 10.89 0.00 1.00 -81.92
```

Interpreting the AICc Model Selection Output

The output generated by the `aictab()` function is automatically organized in ascending order based on the **AICc** score, ensuring that the model with the strongest empirical support is immediately presented first. Interpreting the various columns of this table provides a comprehensive overview of how each model performs relative to its competitors, allowing for nuanced [model selection](#) decisions.

The critical components of the output table are defined as follows:

K: Represents the number of **parameters** estimated in the model. For instance, Model 1 has $K=6$, indicating its greater complexity compared to Models 2 and 3 ($K=4$).

AICc: The corrected Akaike Information Criterion. This is the central metric; the model with the minimum AICc value is deemed the most parsimonious and supported model.

Delta_AICc: This difference score measures how much worse a model is compared to the absolute best model (which has a Delta_AICc of 0.00). Generally, models with a Delta_AICc value greater than 10 are considered to have negligible support relative to the top model.

AICcWt: Known as the **Akaike Weight**, this column provides the probability that the current model is the true best approximating model, given the candidate set. In our example, Model 1's weight of 0.83 suggests an 83% relative probability of being the strongest candidate.

Cum.Wt: The **Cumulative Weight**, calculated by summing the AICc weights sequentially. This metric helps practitioners identify the total evidence supporting a specific subset of the models.

LL: The maximum [log-likelihood](#) achieved by the model, reflecting the raw fit quality to the observed data.

Based on the results, Model 1 (`disp.hp.wt.qsec`) clearly emerged as the superior choice. Its AICc value of 162.43 is the lowest, and its high Akaike Weight confirms its strong empirical support. Although this model is the most complex ($K=6$), the substantial improvement in its fit (as reflected in the log-likelihood) effectively compensated for the complexity penalty, validating its selection over the simpler models.

Conclusion and Next Steps in Regression Analysis

The outcome of the AICc comparison unequivocally established Model 1 as the optimal candidate for predicting vehicle miles per gallon (`mpg`) based on the supplied data. This model successfully achieved the most effective compromise--the hallmark of successful [model selection](#)--by maximizing the explained variance while receiving the lowest penalty for its complexity. The functional structure of this validated model is explicitly defined as:

$$\text{mpg} = \beta_0 + \beta_1(\text{disp}) + \beta_2(\text{hp}) + \beta_3(\text{wt}) + \beta_4(\text{qsec})$$

With the selection phase complete, the statistical procedure shifts focus to the meticulous evaluation of the final model's diagnostics. This subsequent phase requires fitting the chosen model and carefully examining core statistical indicators. These indicators include the overall significance (p-value), the adjusted R-squared value (which quantifies explanatory power), and most importantly, the detailed interpretation of the individual [beta coefficients](#).

The analysis of the magnitude and directional sign of the **beta coefficients** is fundamental for articulating the precise, quantitative relationship between each predictor variable (displacement, horsepower, weight, and quarter-mile time) and the response variable (MPG). This thorough interpretation provides actionable insights and completes the rigorous cycle of statistical model validation.

Alternative and Advanced Model Selection Techniques

While the Akaike Information Criterion (AIC/AICc) is a powerful standard for balancing model fit and complexity, sophisticated researchers often employ complementary criteria to ensure the utmost robustness in their [model selection](#) process. Combining multiple techniques strengthens the generalizability and reliability of the final chosen model.

These robust alternatives and supplementary methods include:

The **Bayesian Information Criterion (BIC)**: Similar in goal to AIC, BIC also uses the maximum [log-likelihood](#) but applies a significantly heavier penalty for increasing the number of parameters (\$K\$). This heavier penalty tends to favor more parsimonious models than AIC, especially with large datasets.

Cross-validation Methods: Techniques such as k-fold validation are used to rigorously assess a model's out-of-sample predictive performance. This provides crucial evidence regarding how well the model generalizes to new, unseen data, offering protection against subtle forms of [overfitting](#).

Model Averaging: This advanced technique involves deriving parameter estimates not just from the single "best" model, but by averaging the estimates across a set of high-performing models, weighted by their respective AICc weights (AICcWt). This approach often yields more stable and reliable parameter estimates.

Mastering these various techniques ensures that researchers can navigate the complexities of statistical modeling, moving beyond simple comparison to achieve truly robust and reliable analytical results.