

How to Interpret Negative AIC Values: A Simple Guide

Authored by
stats writer

December 5, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Interpret Negative AIC Values: A Simple Guide*.

PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=105526>

The Akaike Information Criterion: A Metric for Model Selection

The Akaike information criterion (AIC) stands as a cornerstone metric in statistical modeling, primarily utilized for comparing the relative quality of different statistical models fitted to the same set of data. Developed by statistician Hirotugu Akaike, AIC estimates the relative amount of information lost when a given model is used to represent the process that generated the data. Essentially, it serves as a measure of the trade-off between the complexity of a model and how well the model fits the observed data, helping practitioners select the optimal structure among several competing regression models.

In practice, when researchers or data scientists fit multiple candidate models--perhaps varying the predictor variables or the underlying distributions--they need a standardized, objective method to determine which model offers the best balance of parsimony and explanatory power. AIC provides this quantitative basis. The guiding principle is straightforward: among a set of candidate models, the model that minimizes the AIC score is deemed the superior choice, as it is estimated to lose the least amount of information.

Before diving into the common confusion surrounding negative scores, it is essential to revisit the fundamental calculation that defines this criterion. Understanding the components of the formula reveals precisely why AIC values are not constrained to be positive, thus demystifying the appearance of negative results in model comparison exercises. This foundational knowledge is crucial for proper interpretation.

Deconstructing the AIC Formula

The AIC is mathematically defined by the following expression:

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

This formula encapsulates the two primary drivers of the criterion: model complexity and goodness of fit. The first term, $2K$, acts as a penalty for increasing model complexity, while the second term, $-2\ln(L)$, quantifies the lack of fit to the data. Analyzing each component helps clarify the function of the criterion.

The variable **K** represents the number of estimated model parameters. This typically includes the intercept, the coefficients associated with the predictor variables, and often the variance of the errors. Multiplying K by two means that for every additional parameter introduced into the model, the AIC score increases by two points. This structure inherently discourages the use of overly complex models (those with too many parameters) unless the additional complexity is significantly justified by a better fit to the data. This penalty term is fundamental to the principle of parsimony in model selection.

The term $\ln(L)$, known as the log-likelihood, is a measure of the goodness of fit of the model. The **likelihood (L)** itself represents how probable the observed data are, given the model parameters. By taking the natural logarithm ($\ln$) and multiplying by -2, the term $-2\ln(L)$ transforms the likelihood into a measure of deviance or information loss. Crucially, a larger likelihood (meaning a better fitting model) results in a larger log-likelihood, and thus a smaller (more negative) value for the term $-2\ln(L)$. The balance between the complexity penalty ($2K$) and the fit measure ($-2\ln(L)$) determines the final AIC score.

The Role of Log-Likelihood in AIC Calculation

The log-likelihood, $\ln(L)$, is the driving force behind the magnitude of the AIC score and is the reason why AIC can frequently take on negative values. The **likelihood function** quantifies how well the estimated parameters explain the observed data. For many statistical distributions, especially those commonly used in regression, the likelihood values themselves tend to be very small, often numbers close to zero, reflecting the low probability of observing the exact data set under any specific parameter setting.

However, when we take the natural logarithm of the likelihood, $\ln(L)$, the resulting value is almost always negative because the natural logarithm of any number between 0 and 1 is negative. For instance, if the likelihood L is 0.001, then $\ln(L)$ is approximately -6.908. As the model fit improves--meaning the likelihood L increases towards 1 (its theoretical maximum)--the value of $\ln(L)$ approaches 0. In highly accurate models, the log-likelihood can even exceed 0, although this is uncommon in practice, particularly with complex real-world data.

Since the AIC formula uses -2 multiplied by $\ln(L)$, this component, $-2\ln(L)$, will typically be a large **positive** number when the model fit is poor (L is very small, $\ln(L)$ is a large negative number). Conversely, if the model fit is excellent (L is closer to 1), $\ln(L)$ will be a small negative number close to zero, or even slightly positive. In such cases of strong fit, the resulting term $-2\ln(L)$ becomes a small positive number, or even a negative number if $\ln(L)$ is positive. This variation in the log-likelihood term directly facilitates the possibility of negative AIC results.

Why Negative AIC Values Are Possible

The potential for negative AIC values is a direct consequence of the relationship between the number of parameters (K) and the strength of the model fit, specifically measured by the log-likelihood. Negative scores are not anomalous; they are simply a mathematical outcome when the measure of fit heavily outweighs the complexity penalty.

Recall the formula: $AIC = 2K - 2\ln(L)$. For the AIC score to be negative, the term $-2\ln(L)$ must be a negative number with a magnitude greater than the complexity penalty $2K$. Since K is always positive (representing the number of parameters), $2K$ is always positive. Therefore, the condition

for a negative AIC is that $2K < 2\ln(L)$. This inequality shows that the log-likelihood, $\ln(L)$, must be positive and large enough to overcome the positive penalty term $2K$.

A positive log-likelihood occurs only when the likelihood L itself is greater than 1 (since $\ln(1) = 0$). While likelihood is often constrained to be between 0 and 1 in specific contexts (like probability mass functions), in generalized contexts involving density estimates, the likelihood can, mathematically, exceed 1, especially when dealing with normalized probability density functions evaluated at certain data points. This phenomenon results in $\ln(L)$ being positive.

Suppose we have a model with 7 parameters ($K=7$) and an exceptionally high log-likelihood of 70 (meaning the model fits the data extremely well). We would calculate the AIC as:

$$\text{AIC} = 2 * 7 - 2 * 70 = 14 - 140 = \mathbf{-126}$$

This result, **-126**, is a large negative value, indicating a model that fits the data substantially better than the complexity penalty would suggest. This calculation demonstrates definitively that negative AIC values are not only possible but are expected when models exhibit a strong explanatory fit relative to the model complexity.

Interpreting Negative AIC Values: Focus on Relativity

The most critical concept to grasp when dealing with AIC is that its **absolute value is fundamentally meaningless**. Whether the score is +50, -10, or -500, the sign (positive or negative) provides no intrinsic information about the quality of the model in isolation. AIC is designed purely as a comparative metric used to rank competing models.

The simple and definitive interpretation is this: **The lower the AIC value, the better the fit of the model relative to the other models being considered**. This rule holds true regardless of the sign. Negative AIC values are interpreted exactly the same way as positive ones. They simply indicate a lower score on the comparative scale.

Consider the practical scenario where two regression models are compared. If Model 1 has an AIC value of -56.5 and Model 2 has an AIC value of -103.3, Model 2 is the superior choice. Even though both values are negative, -103.3 is lower than -56.5, signifying that Model 2 has a better balance of fit and parsimony, or equivalently, less estimated information loss. The sign of the AIC score is merely a consequence of the chosen scale of the log-likelihood, which, as previously discussed, is often mathematically arbitrary due to constant shifts.

Practical Comparison Using Delta AIC

While we use the lowest AIC value to identify the "best" model, a more rigorous and informative approach involves calculating the difference between the AIC scores, known as the **Delta AIC**

(Δ_i). The Delta AIC allows researchers to quantify the evidence supporting one model over another and to rank all candidate models based on their proximity to the minimum score.

The Delta AIC for any given model i is calculated by subtracting the minimum AIC value found among the candidate set ($\text{AIC}_{\min}$) from the model's own AIC value (AIC_i):

$$\Delta_i = \text{AIC}_i - \text{AIC}_{\min}$$

The model with the minimum AIC score naturally has a Delta AIC of 0. All other models will have positive Delta AIC values. These differences are highly interpretable, indicating how much worse a given model is compared to the best model. For instance, a model with an AIC difference between 0 and 2 is considered to have substantial empirical support, whereas a model with a difference of 10 or more is generally considered to have very little support and is unlikely to be the true generating model.

This relative metric highlights why the absolute magnitude, including the negative or positive sign, is irrelevant. When calculating Δ_i , the arbitrary additive constant inherent in the definition of the log-likelihood cancels out, leaving only the meaningful comparative difference. This process converts potentially large, negative, and confusing AIC scores into a simple, positive, interpretable scale starting at zero, thereby facilitating clear decision-making in model selection.

Authoritative Perspectives on AIC Interpretation

The principle that only relative AIC values matter is consistently reinforced across standard statistical literature and official documentation. Leading textbooks on statistical modeling emphasize the non-absolute nature of the AIC scale, regardless of whether the results are positive or negative.

As noted by experts in statistical inference:

Usually, AIC is positive; however, it can be shifted by any additive constant, and some shifts can result in negative values of AIC... It is not the absolute size of the AIC value, it is the relative values over the set of models considered, and particularly the differences between AIC values, that are important.

Another authoritative perspective confirms that, similar to raw likelihood measures, the absolute value of AIC is largely arbitrary, being determined by constants that depend on the data set and the choice of probability distribution function. Therefore, AIC can only be validly used to compare models fitted on **identical samples**.

The objective remains focused: identifying the best model--the one with the smallest AIC value--which corresponds to the least estimated information loss relative to the true underlying process. In

conclusion, the presence of negative AIC values should never be a cause for concern or misinterpretation. These values are a natural mathematical outcome of a strong model fit relative to the complexity penalty. Focus solely on the hierarchy of the scores: the lowest AIC score, positive or negative, always designates the relatively best model among the candidates analyzed.

ARABPSYCHOLOGY.COM