

Learning to Resolve the R Warning: “glm.fit: algorithm did not converge

Authored by
Mohammed loot

November 2, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Resolve the R Warning: “glm.fit: algorithm did not converge*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8472>

When conducting advanced statistical modeling using the [R programming language](#), data scientists and statisticians frequently rely on the `glm()` function to fit models belonging to the family of [Generalized Linear Models](#) (GLMs). However, a common and potentially misleading warning that arises during this process, particularly when utilizing [logistic regression](#) for binary outcomes, is the dreaded message:

glm.fit: algorithm did not converge

This warning signifies a failure in the standard iterative optimization procedure used by the `glm.fit` function. Rather than reaching stable, finite values, the algorithm is unable to locate reliable [maximum likelihood estimates](#) (MLEs) for the model coefficients. Understanding the root cause of this failure is paramount for producing statistically valid results and trustworthy models.

The core statistical problem driving this convergence issue is frequently a condition known as [perfect separation](#), or complete separation. This article serves as a comprehensive guide, providing a detailed explanation of why separation occurs, how it affects the estimation process, and outlining powerful, robust methodologies for effectively resolving the issue in practical data analysis scenarios.

Understanding the `glm.fit: algorithm did not converge` Warning

The process of fitting a [Generalized Linear Model](#) relies heavily on iterative optimization techniques. Specifically, the `glm.fit` function in R typically employs the [Iteratively Reweighted Least Squares](#) (IRLS) algorithm. This method works by iteratively updating the coefficient estimates, gradually moving toward the point where the log-likelihood function is maximized. The algorithm is deemed successful, or "converged," when the magnitude of the change in coefficient estimates between successive iterations falls below a predefined, infinitesimally small tolerance threshold.

When the algorithm fails to converge, it is a clear signal that the model is attempting to estimate coefficients whose true values are infinitely large--either positive or negative. In mathematical terms, the standard technique of [maximum likelihood estimation](#) breaks down because the likelihood function does not have a finite maximum point in the parameter space. Instead, the function continues to increase as the coefficients approach infinity, preventing the IRLS algorithm from settling on stable, usable values.

It is crucial to note that R will often output numerical coefficient estimates even when the convergence warning is present. However, these coefficients are inherently unreliable, as they are merely the values reached when the iteration limit was hit or when numerical overflow occurred. Consequently, any associated standard errors, confidence intervals, and P-values derived from

these unreliable estimates are invalid. Proceeding with interpretation or prediction based on a model that has failed to converge is a significant statistical error that must be corrected by addressing the underlying cause.

The Root Cause: Perfect Separation (Complete Separation)

The overwhelming majority of convergence failures in [logistic regression](#) models can be traced back to the statistical phenomenon known as perfect, or complete, separation. This condition arises specifically in binary classification contexts when one or more predictor variables (or a linear combination of them) possess the ability to perfectly divide the response variable (Y) into its distinct categories (usually coded as 0 and 1). In essence, a simple rule based on the predictor flawlessly predicts the outcome for every single data point in the sample.

Consider the implications of perfect separation on the logistic link function. If separation exists, the probability of the positive outcome ($P(Y=1)$) approaches 1 for observations on one side of the separation point, and approaches 0 for observations on the other side. To mathematically model this deterministic, non-stochastic relationship--a boundary that is infinitely sharp--the corresponding logistic regression coefficients must tend toward infinity (either positive or negative).

Since the goal of [maximum likelihood estimation](#) is to find the coefficients that maximize the likelihood of observing the data, and the likelihood continues to increase indefinitely as the coefficients grow, the MLEs for the separating variables become undefined. The model's attempt to fit an ideal, perfect boundary forces these coefficient values to extreme limits, resulting directly in the convergence failure warning and unreliable estimates. Partial separation, a related but less severe issue where a predictor strongly predicts the outcome but not perfectly, can also lead to convergence issues or extremely inflated standard errors.

Practical Example: Reproducing the Warning in R

To solidify the concept of [perfect separation](#), we can construct a synthetic dataset in R where the predictor variable x flawlessly determines the binary outcome y . We define a simple scenario: all y values are 0 when x is less than 1, and all y values are 1 when x is 1 or greater. This design immediately establishes an unambiguous separation point, which should trigger the convergence warning when fitted using `glm()`.

```
#create data frame
```

```
df <- data.frame(x=c(.1, .2, .3, .4, .5, .6, .7, .8, .9, 1, 1, 1.1, 1.3, 1.5, 1.7),  
y=c(0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 1, 1, 1, 1))
```

```
#attempt to fit logistic regression model
```

```
glm(y~x, data=df, family="binomial")
```

```
Call: glm(formula = y ~ x, family = "binomial", data = df)
```

Coefficients:

(Intercept) x

-409.1 431.1

Degrees of Freedom: 14 Total (i.e. Null); 13 Residual

Null Deviance: 20.19

Residual Deviance: 2.468e-09 AIC: 4

Warning messages:

1: glm.fit: algorithm did not converge

2: glm.fit: fitted probabilities numerically 0 or 1 occurred

The output clearly displays the primary warning: `glm.fit: algorithm did not converge`. Furthermore, the accompanying secondary warning, `glm.fit: fitted probabilities numerically 0 or 1 occurred`, explicitly confirms the statistical scenario. This secondary message is the functional definition of separation within the R environment, indicating that the model attempted to assign perfect probabilities (0 or 1) to certain observations.

Notice the sheer magnitude of the estimated coefficients (`-409.1` and `431.1`). These extremely large values are the model's desperate attempt to fit the infinitely steep slope required to perfectly separate the data. Since the predictor variable `x` perfectly separates the response variable `y`--where `y=0` for all `x < 1` and `y=1` for all `x >= 1`--the standard maximum likelihood approach fails entirely.

A Comparative Example: Modeling Without Separation

To appreciate a properly functioning model, it is helpful to contrast the previous example with a dataset that exhibits overlap, meaning the predictor variable is not able to perfectly separate the outcome. By introducing minor changes to the response variable `y`, we create a scenario of non-perfect separation. This overlap prevents the coefficients from needing to tend toward infinity, thereby allowing the [IRLS](#) algorithm to converge normally and locate stable [maximum likelihood estimates](#).

The following R code demonstrates the fitting process for a revised dataset where the predictor variable is intentionally corrupted to prevent perfect partitioning of the binary response variable. This successfully avoids the computational instability and the resulting convergence error:

```
#create data frame
```

```
df <- data.frame(x=c(.1, .2, .3, .4, .5, .6, .7, .8, .9, 1, 1, 1.1, 1.3, 1.5, 1.7),
```

```
y=c(0, 1, 0, 1, 0, 0, 0, 0, 0, 1, 1, 1, 1, 1, 1))
```

```
#fit logistic regression model  
glm(y~x, data=df, family="binomial")
```

```
Call: glm(formula = y ~ x, family = "binomial", data = df)
```

```
Coefficients:
```

```
(Intercept) x  
-2.112 2.886
```

```
Degrees of Freedom: 14 Total (i.e. Null); 13 Residual
```

```
Null Deviance: 20.73
```

```
Residual Deviance: 16.31 AIC: 20.31
```

In this modified scenario, we do not receive any warning message from R. The resulting coefficients (-2.112 and 2.886) are finite, stable, and statistically meaningful. This successful outcome confirms that the model converged because the data contains enough inherent uncertainty and noise--or overlap--to prevent the predictor variable from achieving [perfect separation](#), thus allowing the likelihood function to reach a well-defined maximum.

Strategies for Handling Complete Separation

When the `glm.fit: algorithm did not converge` warning signals complete separation, the standard approach of ignoring the warning and utilizing the inflated coefficients is statistically indefensible. The appropriate remediation strategy depends heavily on the context of the analysis and whether the separation is believed to be an anomaly specific to the sample size or a genuine, deterministic relationship existing in the population.

Method 1: Implement Penalized Regression Techniques (Recommended Statistical Solution).

If the primary goal is to obtain regularized, stable, and finite coefficient estimates suitable for inference, hypothesis testing, or prediction, the statistically superior method is the application of [penalized regression](#). These methods modify the standard likelihood function by adding a penalty term that biases the coefficients toward zero. This crucial step prevents the coefficients from spiraling toward infinity, successfully resolving the convergence failure.

Common implementations include [Lasso logistic regression](#) (which applies an L1 penalty) or [Elastic-Net regularization](#) (which combines L1 and L2 penalties). These techniques are accessible in R through packages like `glmnet` or `brms`. By shrinking the inflated coefficients, penalized methods provide statistically robust estimates that account for the strong predictive power of the separating variable without breaking the optimization algorithm.

Method 2: Remove the Separating Variable or Collapse Categories.

If the separation is caused by a single categorical predictor with sparse cell counts, one solution is to remove that predictor entirely from the model, assuming it is not theoretically essential. Alternatively, if the predictor is categorical, combining or collapsing sparsely populated categories (e.g., merging rare outcomes into an "Other" category) can introduce necessary overlap, thereby resolving the separation issue. This is a data-driven approach that modifies the input features rather than the fitting procedure.

Method 3: Utilize the Deterministic Relationship for Prediction.

In cases where the perfect separation is highly plausible in the true population (i.e., the relationship genuinely appears deterministic), the simplest practical solution is to use the predictor variable itself to model the outcome directly, bypassing the need for a probabilistic [GLM](#). For instance, if you observe that the response variable y is always equal to 0 when the predictor x is less than 1, you can establish this deterministic rule as part of your prediction mechanism. This approach validates the strong relationship without attempting to force an unstable probabilistic model onto perfectly predictable data.

Conclusion and Further Resources

The warning `glm.fit: algorithm did not converge` is not merely a technical glitch; it is a vital statistical diagnostic that signals the failure of standard maximum likelihood estimation due to [perfect separation](#). Recognizing this limitation guides the data analyst toward more sophisticated and appropriate modeling strategies, such as penalized regression. Addressing the cause of non-convergence ensures that the final model is both stable and statistically defensible.

For researchers interested in deepening their knowledge of generalized linear models and advanced modeling strategies in R, the following resources are recommended:

Detailed documentation on the `glm()` function in the [R programming language](#).

Tutorials focusing on the implementation of [Lasso](#) and Ridge regression using the `glmnet` package.

Statistical theory explaining the principles of [maximum likelihood estimation](#) when dealing with boundary cases.