

Learning to Fit a Gamma Distribution to Data in R

Authored by
Mohammed loot

November 9, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning to Fit a Gamma Distribution to Data in R*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=14363>

This comprehensive tutorial is designed for analysts and data scientists looking to effectively model positive, continuous, and often skewed data using the statistical power of the [Gamma Distribution](#) within the [R programming language](#). The Gamma Distribution is a critical tool in fields like reliability engineering, finance, and queueing theory, as it is highly effective for modeling waiting times or the magnitude of independently and identically distributed components. Understanding how to accurately fit this distribution to real-world observations is a foundational skill in advanced statistical modeling.

In this guide, we will walk through the entire process, starting from generating synthetic data, installing the necessary tools, applying the powerful technique of maximum likelihood estimation, and finally, visually assessing the goodness of the fit. We emphasize using the appropriate packages and methods to ensure the resulting parameter estimates are both robust and meaningful for subsequent analysis and prediction.

Generating the Sample Data for Analysis

Before we can fit any distribution, we must establish a dataset. While real-world data would typically be imported, for demonstration purposes, we will generate a synthetic dataset that approximates a known [gamma distribution](#). This allows us to compare our estimated parameters against the true parameters used in the generation process, providing a clear validation of the fitting procedure.

We will utilize R's built-in functions to create fifty random values. The generating parameters are a **shape parameter** equal to 3 and a rate parameter (often inversely related to scale) of 10. Crucially, to make our synthetic data more representative of actual empirical observations, which are rarely perfectly distributed, we introduce a small amount of [Gaussian noise](#). This noise simulates the measurement errors or minor random fluctuations that are inherent in any real-world measurement system.

The following R code block demonstrates the generation and a quick preview of the initial values stored in the variable `z`. Note that the addition of Gaussian noise slightly perturbs the dataset, presenting a realistic challenge for the fitting algorithm to overcome.

```
#generate 50 random values that follow a gamma distribution with shape parameter = 3  
#and shape parameter = 10 combined with some gaussian noise  
z <- rgamma(50, 3, 10) + rnorm(50, 0, .02)
```

```
#view first 6 values
```

```
head(z)
```

```
0.07730 0.02495 0.12788 0.15011 0.08839 0.09941
```

Prerequisites: Installing and Loading the R Package

To perform high-quality distribution fitting in R, the standard base functions are often insufficient or overly complex for iterative fitting tasks. We rely on specialized libraries that streamline the estimation process. The package essential for this task is **fitdistrplus**, which provides a cohesive framework for fitting a wide variety of probability distributions to empirical data.

If you do not already have this library installed on your [R](#) environment, you must execute the installation command first. Once installed, the package must be loaded into the current R session using the `library()` function. This step makes all the powerful functions contained within **fitdistrplus**, especially the core function `fitdist()`, available for immediate use in your analysis.

The installation and loading procedure is standard practice in R and is shown below. Ensure you have an active internet connection if installation is required.

```
#install 'fitdistrplus' package if not already installed  
install.packages('fitdistrplus')
```

```
#load package  
library(fitdistrplus)
```

Understanding the fitdist Function Syntax

The primary function used for fitting distributions in this package is `fitdist()`. This function is highly versatile, designed to handle various distributions and estimation methods. Its general structure requires specifying the dataset, the name of the distribution being fitted, and the method used for parameter estimation. This modular design allows statisticians to quickly compare fits across different theoretical distributions or utilize different estimation techniques on the same data.

The general syntax for the `fitdist()` function is clearly defined, and understanding its arguments is key to successful implementation. The first argument is always the dataset, followed by the distribution choice (specified as a string, e.g., "gamma", "norm", "weibull"), and finally, the method of fitting. The flexibility of this syntax ensures that switching between modeling assumptions is straightforward.

For clarity and reference, the core structure of the function call is presented here:

```
fitdist(dataset, distr = "your distribution choice", method = "your method of fitting the data")
```

For our specific case, fitting the data in `z` to the [Gamma Distribution](#), we will select the most widely used and statistically robust estimation technique: [Maximum Likelihood Estimation](#) (MLE). MLE is

the preferred method because it provides parameter estimates that maximize the likelihood of observing the data we actually collected, offering superior asymptotic properties compared to simpler methods like moment matching.

Executing the Maximum Likelihood Estimation (MLE) Fit

We are now ready to execute the parameter estimation using the [fitdistrplus](#) package. By specifying `distr = "gamma"` and `method = "mle"`, we instruct R to find the optimal **shape parameter** and rate parameter for the Gamma Distribution that best describes the observed dataset `z`. This process involves complex numerical optimization, which `fitdistr()` handles internally.

The resulting object, stored in the variable `fit`, contains not just the parameter estimates but also detailed information about the fitting process, including the achieved log-likelihood value and the standard errors associated with the estimates. Viewing the summary of this object is crucial, as it provides the numerical results of our statistical model. The summary output includes the estimated shape and rate parameters, which are central to defining our fitted distribution, along with their respective standard errors, indicating the precision of the estimates.

```
#fit our dataset to a gamma distribution using mle
```

```
fit <- fitdistr(z, distr = "gamma", method = "mle")
```

```
#view the summary of the fit
```

```
summary(fit)
```

The execution of the summary command produces a detailed output, typically displayed directly in the R console. This output confirms the method used ([MLE](#)), the parameters estimated (shape and rate), and critical goodness-of-fit metrics. We can observe how closely the estimated parameters align with the true parameters (shape=3, rate=10) we used to generate the synthetic data, serving as a powerful validation of the fitting procedure.

```
Fitting of the distribution ' gamma ' by maximum likelihood
Parameters :
      estimate Std. Error
shape 2.550018  0.4803947
rate  9.711618  2.0216362
Loglikelihood: 26.51357  AIC:  -49.02714  BIC:  -45.2031
Correlation matrix:
      shape    rate
shape 1.0000000 0.9049895
rate  0.9049895 1.0000000
```

Visualizing the Goodness of Fit

While numerical outputs like the log-likelihood or AIC/BIC values provide quantitative measures of fit quality, visualization is indispensable for truly assessing how well the theoretical [Gamma Distribution](#) aligns with the empirical data. The `plot()` function, when applied directly to the `fit` object generated by [fitdistrplus](#), automatically generates four key diagnostic plots.

These plots offer different perspectives on the fit:

The **Density Plot** compares the fitted distribution's theoretical density curve against the empirical histogram of the data.

The **Cumulative Distribution Function (CDF) Plot** compares the theoretical CDF to the empirical CDF.

The **QQ Plot (Quantile-Quantile Plot)** compares the quantiles of the data against the theoretical quantiles of the fitted distribution. A good fit results in points aligning closely along the diagonal line.





The **PP Plot (Probability-Probability Plot)** compares the empirical probabilities against the fitted theoretical probabilities.

Analyzing these visual tools allows the statistician to identify any systematic deviations or outliers that might suggest the Gamma Distribution is not the most appropriate model, or that the data contains underlying structures not captured by the current model. For our generated data, we expect these plots to show a very strong alignment, confirming the success of the [MLE](#) estimation.

To generate these informative plots, we use the following simple command:

```
#produce plots  
plot(fit)
```

Executing this command yields the four comprehensive diagnostic charts, providing immediate visual feedback on the quality of our distribution fit.

Empirical and theoretical dens.**Q-Q plot****Empirical and theoretical CDFs****P-P plot**

Complete Code Summary for Reproducibility

For ease of implementation and reproducibility, the entire sequence of R commands required to generate the data, load the necessary package, fit the [Gamma Distribution](#) using MLE, and visualize the results is provided below. This comprehensive script encapsulates all the steps discussed in this tutorial, allowing users to run the analysis seamlessly in their own [R](#) environment. This unified approach is highly beneficial for integrating distribution fitting into larger data processing pipelines.

Remember that accurate distribution fitting is a cornerstone of probabilistic modeling. Mastering the

use of tools like [fitdistrplus](#) in R ensures that your subsequent statistical inferences, such as confidence intervals or hypothesis tests, are based on reliable and well-justified distributional assumptions derived from the observed data.

#install 'fitdistrplus' package if not already installed

```
install.packages('fitdistrplus')
```

```
#load package
```

```
library(fitdistrplus)
```

```
#generate 50 random values that follow a gamma distribution with shape parameter = 3
```

```
#and shape parameter = 10 combined with some gaussian noise
```

```
z <- rgamma(50, 3, 10) + rnorm(50, 0, .02)
```

```
#fit our dataset to a gamma distribution using mle
```

```
fit <- fitdist(z, distr = "gamma", method = "mle")
```

```
#view the summary of the fit
```

```
summary(fit)
```

```
#produce plots to visualize the fit
```

```
plot(fit)
```