


# How to draw a covariance error ellipse?

Contents [\[hide\]](#) [\[hide\]](#)

- [1 Introduction](#)
- [2 Axis-aligned confidence ellipses](#)
- [3 Arbitrary confidence ellipses](#)
- [4 Source Code](#)
- [5 Conclusion](#)

## Introduction

In this post, I will show how to draw an error ellipse, a.k.a. confidence ellipse, for 2D normally distributed data. The error ellipse represents an iso-contour of the Gaussian distribution, and allows you to visualize a 2D confidence interval. The following figure shows a 95% confidence ellipse for a set of 2D normally distributed data samples. This confidence ellipse defines the region that contains 95% of all samples that can be drawn from the underlying Gaussian distribution.




**Figure 1.** 2D confidence ellipse for normally distributed data

In the next sections we will discuss how to obtain confidence ellipses for different confidence values (e.g. 99% confidence interval), and we will show how to plot these ellipses using Matlab or C++ code.

## Axis-aligned confidence ellipses

Before deriving a general methodology to obtain an error ellipse, let's have a look at the special case where the major axis of the ellipse is aligned with the X-axis, as shown by the following figure:



**Figure 2.** Confidence ellipse for uncorrelated Gaussian data

The above figure illustrates that the angle of the ellipse is determined by the covariance of the data. In this case, the covariance is zero, such that the data is uncorrelated, resulting in an axis-aligned error ellipse.

**Table 1.** Covariance matrix of the data shown in Figure 2

8.4213	0
0	0.9387

Furthermore, it is clear that the magnitudes of the ellipse axes depend on the [variance](#) of the data. In our case, the largest variance is in the direction of the X-axis, whereas the smallest variance lies in the direction of the Y-axis.

In general, the equation of an axis-aligned ellipse with a major axis of length  $2a$  and a minor axis of length  $2b$ , centered at the origin, is defined by the following equation:

$$(1) \left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 = 1$$

In our case, the length of the axes are defined by the standard deviations  $\sigma_x$  and  $\sigma_y$  of the data such that the equation of the error ellipse becomes:

$$(2) \left(\frac{x}{\sigma_x}\right)^2 + \left(\frac{y}{\sigma_y}\right)^2 = s$$

where  $s$  defines the scale of the ellipse and could be any arbitrary number (e.g.  $s=1$ ). The question is now how to choose  $s$ , such that the scale of the resulting ellipse represents a chosen confidence level (e.g. a 95% confidence level corresponds to  $s=5.991$ ).

Our 2D data is sampled from a multivariate Gaussian with zero covariance. This means that both the x-values and the y-values are normally distributed too. Therefore, the left hand side of equation (2) actually represents the sum of squares of independent normally distributed data samples. The sum of squared Gaussian data points is known to be distributed according to a so called [Chi-Square distribution](#). A Chi-Square distribution is defined in terms of ‘degrees of freedom’, which represent the number of unknowns. In our case there are two unknowns, and therefore two degrees of freedom.

Therefore, we can easily obtain the probability that the above sum, and thus  $s$  equals a specific value by calculating the Chi-Square likelihood. In fact, since we are interested in a confidence interval, we are looking for the probability that  $s$  is less than or equal to a specific value which can easily be obtained using the cumulative Chi-Square distribution. As statisticians are lazy people, we usually don’t try to calculate this probability, but simply look it up in a probability table: <https://people.richland.edu/james/lecture/m170/tbl-chi.html>.

For example, using this probability table we can easily find that, in the 2-degrees of freedom case:

$$P(s < 5.991) = 1 - 0.05 = 0.95$$

Therefore, a 95% confidence interval corresponds to  $s=5.991$ . In other words, 95% of the data will fall inside the ellipse defined as:

$$(3) \left(\frac{x}{\sigma_x}\right)^2 + \left(\frac{y}{\sigma_y}\right)^2 = 5.991$$


Similarly, a 99% confidence interval corresponds to  $s=9.210$  and a 90% confidence interval corresponds to  $s=4.605$ .

The error ellipse shown by figure 2 can therefore be drawn as an ellipse with a major axis length equal to  $2\sigma_x\sqrt{5.991}$  and the minor axis length to  $2\sigma_y\sqrt{5.991}$ .

## Arbitrary confidence ellipses

In cases where the data is not uncorrelated, such that a covariance exists, the resulting error ellipse will not be axis aligned. In this case, the reasoning of the above paragraph only holds if we temporarily define a new coordinate system such that the ellipse becomes axis-aligned, and then rotate the resulting ellipse afterwards.

In other words, whereas we calculated the variances  $\sigma_x$  and  $\sigma_y$  parallel to the x-axis and y-axis earlier, we now need to calculate these variances parallel to what will become the major and minor axis of the confidence ellipse. The directions in which these variances need to be calculated are illustrated by a pink and a green arrow in figure 1.



**Figure 1.** 2D confidence ellipse for normally distributed data

These directions are actually the directions in which the data varies the most, and are defined by the covariance matrix. The covariance matrix can be considered as a matrix that linearly transformed some original data to obtain the currently observed data. In a previous article about [eigenvectors and eigenvalues](#) we showed that the direction vectors along such a linear transformation are the eigenvectors of the transformation matrix. Indeed, the vectors shown by pink and green arrows in figure 1, are the eigenvectors of the covariance matrix of the data, whereas the length of the vectors corresponds to the eigenvalues.

The eigenvalues therefore represent the spread of the data in the direction of the eigenvectors. In other words, the eigenvalues represent the variance of the data in the direction of the eigenvectors. In the case of axis aligned error ellipses, i.e. when the covariance equals zero, the eigenvalues equal the variances of the covariance matrix and the eigenvectors are equal to the definition of the x-axis and y-axis. In the case of arbitrary correlated data, the eigenvectors represent the direction of the largest spread of the data, whereas the eigenvalues define how large this spread really is.


Thus, the 95% confidence ellipse can be defined similarly to the axis-aligned case, with the major axis of length  $2\sqrt{5.991\lambda_1}$  and the minor axis of length  $2\sqrt{5.991\lambda_2}$ , where  $\lambda_1$  and  $\lambda_2$  represent the eigenvalues of the covariance matrix.

To obtain the orientation of the ellipse, we simply calculate the angle of the largest eigenvector towards the x-axis:

$$(4) \quad \alpha = \arctan \frac{\mathbf{v}_1(y)}{\mathbf{v}_1(x)}$$

where  $\mathbf{v}_1$  is the eigenvector of the covariance matrix that corresponds to the largest eigenvalue.

Based on the minor and major axis lengths and the angle  $\alpha$  between the major axis and the x-axis, it becomes trivial to plot the confidence ellipse. Figure 3 shows error ellipses for several confidence values:



Confidence ellipses for normally distributed data

## Source Code

[Matlab source code](#)

[C++ source code \(uses OpenCV\)](#)

## Conclusion

In this article we showed how to obtain the error ellipse for 2D normally distributed data, according to a chosen confidence value. This is often useful when visualizing or analyzing data and will be of interest in a future article about [PCA](#).

Furthermore, source code samples were provided for Matlab and C++.