

Introduction to Statistical Methods

PH3010 Advanced Skills

RHUL Physics 2020/21

Abstract

In this module several important areas of statistical data analysis are explored, including parameter fitting and model testing with the method of least squares. This is done through a series of exercises that include fitting a polynomial and using the resulting covariance matrix of estimated parameters for error propagation, as well as analysis of historical data sets from Galileo and Ptolemy. Using measures of goodness-of-fit from the least squares fits, the level of agreement between the data and competing hypotheses is assessed.



1. Introduction

This module will provide an introduction to parameter fitting and model testing with the method of least squares, which are simple but powerful tools in the analysis of experimental data. The basic problem is illustrated in Fig. 1. Suppose we have measurements of a quantity y , each of which is carried out corresponding to a fixed value of some other variable x , called the control variable. Roughly speaking the goal is to find a smooth curve that passes close to the data points, called “fitting” a curve to the data.




Figure 1: Typical ingredients for the least-squares problem: measurements y with error bars σ are carried out at known values of a control variable x . The curve has been fitted to the data.

In Sec. 2 we will formulate this problem more precisely and state how the method of least squares is used to find the fitted curve. Then in Secs. 3, 4, 5 and 6 we will examine in greater detail several aspects of this problem: determining the fitted curve itself, quantifying its statistical uncertainty, testing whether the hypothesized functional form of the curve provides an adequate description of the data, and extending the procedure to the case of correlated data. Finally we present a set of exercises including analysis of historical data from Galileo and Ptolemy.

2. The method of least squares

Referring back to Fig. 1 we will suppose here that the x values are known very precisely, but the measured y values are subject to random fluctuations. For any given measurement we do not know exactly how far or in what direction the observed value has departed from its “true” value, but we will assume we have some measure σ of the typical size of the fluctuations. Suppose we have N measurements each labeled with an index i , and thus the data consist of the set of values (x_i, y_i, σ_i) with $i = 1, \dots, N$. These are displayed in Fig. 1 with the σ_i values as error bars.

Looking at the data in Fig. 1, one can believe that if it were possible to carry out the measurements without the random errors, i.e., in the case where σ is zero or negligibly small, then all of the (x, y) points would lie on a smooth curve. That is, there could be some function $f(x)$ that gives the “true” values, and the observed values differ from these as a result of the fluctuations.

More precisely we will treat the y_i as *independent random variables*. If one were to repeat the measurement of y_i many times under the same circumstances (i.e., at the same x_i), the observed y_i values would come out different each time, and would follow a distribution with a certain width. To characterize this width we give the standard deviation σ_i . Note that in

the real problem we only have one value y_i for a given x_i ; here the repeated measurements used to interpret the meaning of the standard deviation are purely hypothetical. The further statement that the set of y_i values are independent means that the measured value of any one of them has no influence on the fluctuations of any of the others.

If we only have the set of (x, y, σ) measurements, how can we figure out the curve $f(x)$? Given the available data, there is no way to give a definitive answer to this question. But we can *estimate* the curve using statistical methods. The approach we will take here requires that we hypothesize a formula for $f(x)$ that contains some adjustable parameters, i.e., constants whose values are not yet determined. We will usually use Greek letters for the parameters and may write them as a vector if we have more than one, e.g., $f(x; \boldsymbol{\theta})$ with $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)$. For example, we could suppose that $f(x)$ is a straight line, i.e.,

$$f(x; \boldsymbol{\theta}) = \theta_0 + \theta_1 x, \quad (1)$$

where here $\boldsymbol{\theta} = (\theta_0, \theta_1)$. The problem of estimating the curve is now reduced to estimating the unknown parameters. In this example, each possible choice for the values will yield a possible line, and the question is how to decide which point in the parameter space is best.

Suppose we pick an arbitrary point (θ_0, θ_1) , which corresponds to a given curve, as shown in Fig. 2. Our initial choice is evidently not very good, as the amount by which the data points miss the curve, $y_i - f(x_i; \boldsymbol{\theta})$, is quite large for most of the points.




Figure 2: The line corresponding to an arbitrarily chosen point in the parameter space (θ_0, θ_1) together with the data points. The thick vertical bars show the residuals $y_i - f(x_i; \boldsymbol{\theta})$.

Notice that for some of the points in Fig. 2, the difference or “residual” $y_i - f(x_i; \boldsymbol{\theta})$ is large but also the error bar σ_i is large. For these points one expects a bigger fluctuation in the measured value. Some of the measurements have small error bars, and so these should lie closer to the hypothesized curve. So to characterize how far a measured y_i is from the predicted value $f(x_i; \boldsymbol{\theta})$, we can measure the residual $y_i - f(x_i; \boldsymbol{\theta})$ in units of σ_i , i.e., we use

$$\frac{y_i - f(x_i; \boldsymbol{\theta})}{\sigma_i}, \quad (2)$$

sometimes called the *normalized* (or *standardized*) residual. The “best” curve should give small values for the normalized residuals, but these can be positive or negative for different data points, so we cannot simply minimize their sum. In the method of *Least Squares* (LS), we define the best parameter estimates to be those that minimize the sum of squares of the normalized residuals, a quantity called the chi-squared,

$$\chi^2(\boldsymbol{\theta}) = \sum_{i=1}^N \left(\frac{y_i - f(x_i; \boldsymbol{\theta})}{\sigma_i} \right)^2. \quad (3)$$

The values of the parameters that give the minimum chi-squared, called the *estimators* of the parameters, are written with hats, e.g., $\hat{\theta}_0$ and $\hat{\theta}_1$. Figure 3 shows the example of the fitted line, i.e., the curve $f(x; \boldsymbol{\theta})$ evaluated with the LS estimators for $\boldsymbol{\theta}$.




Figure 3: Least-squares fit of a straight line (see text).

One can derive the method of least squares from a more general starting point called the *method of maximum likelihood* in the special case where the data values y_i are statistically independent and each follow a Gaussian distribution (see, e.g., Ref. [2]). But it is equally valid to simply use the minimum of the $\chi^2(\boldsymbol{\theta})$ to define the estimators and it is important to emphasise that it is not possible in general to construct estimators that are “best” in every relevant sense. In practice it turns out that LS estimators have desirable properties in several respects, as we will see in examining the following questions:

1. How do we find the estimators, i.e., how do we minimize $\chi^2(\boldsymbol{\theta})$?
2. How do we quantify the statistical uncertainty in the estimated parameters that stems from the random fluctuations in the measurements, and how is this information used in an analysis problem, e.g., using error propagation?
3. How do we assess whether the hypothesized functional form $f(x; \boldsymbol{\theta})$ adequately describes the data?

3. Finding the minimum of $\chi^2(\boldsymbol{\theta})$

Depending on the problem there can be a variety of different methods employed to find the estimators, i.e., to minimize the quantity $\chi^2(\boldsymbol{\theta})$. We will illustrate using a simple example first in a case where we can write down the estimators in closed form in Sec. 3.1, and then by carrying out the minimization numerically in Sec. 3.2.

3.1. Finding LS estimators in closed form

The curve $f(x; \boldsymbol{\theta})$ that we are fitting to the data does not of course have to be a straight line; it can be any function. It turns out that the method of least squares is particularly easy for

problems where $f(x; \boldsymbol{\theta})$ is a linear function *of the parameters* (it does not matter whether it is a linear function of the control variable x). As an example one could have a polynomial of order M , i.e.,

$$f(x; \boldsymbol{\theta}) = \sum_{n=0}^M \theta_n x^n . \quad (4)$$

For this type of problem it is possible to find the values of the parameters that minimize the chi-squared in closed form, and we will show this here for the special case of a first-order polynomial, i.e., a straight line. Here the the function $\chi^2(\boldsymbol{\theta})$ takes on the form

$$\chi^2(\theta_0, \theta_1) = \sum_{i=1}^N \frac{(y_i - \theta_0 - \theta_1 x_i)^2}{\sigma_i^2} . \quad (5)$$

To find its minimum we set the derivatives of χ^2 with respect to both parameters equal to zero,

$$\frac{\partial \chi^2}{\partial \theta_0} = \sum_{i=1}^N \frac{-2(y_i - \theta_0 - \theta_1 x_i)}{\sigma_i^2} = 0 , \quad (6)$$

$$\frac{\partial \chi^2}{\partial \theta_1} = \sum_{i=1}^N \frac{-2x_i(y_i - \theta_0 - \theta_1 x_i)}{\sigma_i^2} = 0 . \quad (7)$$

These two equations can be rewritten using matrix notation as

$$\begin{pmatrix} \sum_{i=1}^N \frac{1}{\sigma_i^2} & \sum_{i=1}^N \frac{x_i}{\sigma_i^2} \\ \sum_{i=1}^N \frac{x_i}{\sigma_i^2} & \sum_{i=1}^N \frac{x_i^2}{\sigma_i^2} \end{pmatrix} \begin{pmatrix} \theta_0 \\ \theta_1 \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^N \frac{y_i}{\sigma_i^2} \\ \sum_{i=1}^N \frac{x_i y_i}{\sigma_i^2} \end{pmatrix} , \quad (8)$$

which has the general form

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} \theta_0 \\ \theta_1 \end{pmatrix} = \begin{pmatrix} e \\ f \end{pmatrix} . \quad (9)$$

By comparing Eqs. (8) and (9) we can read off the values of a , b , c , d , e and f . Then using the fact that a 2×2 matrix A and its inverse A^{-1} can be written

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} , \quad A^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} , \quad (10)$$

and applying A^{-1} to both sides of Eq. (9), we can find solutions

$$\hat{\theta}_0 = \frac{de - bf}{ad - bc} , \quad (11)$$

$$\hat{\theta}_1 = \frac{af - ec}{ad - bc} . \quad (12)$$

Here we write the solutions with hats because these are the estimators for the parameters θ_0 and θ_1 .

We can extract a few lessons from this example. First, we see that the solutions can be obtained by solving a system of linear equations using standard matrix techniques. Furthermore, the data values y_i enter into the estimators through the quantities e and f , and therefore from Eqs. (11) and (12) we see that the estimators $\hat{\theta}_0$ and $\hat{\theta}_1$ are linear functions of the y_i . This remains true for an arbitrary number of parameters, as long as the function that we are fitting, $f(x; \boldsymbol{\theta})$, is linear in the parameters (see, e.g., Ref. [2]). Finally we see that although we can write down expressions for the solutions in closed form, they are somewhat involved even for the simple case of two parameters. If the number of parameters is higher then this approach may become impractical, and in any case if $f(x; \boldsymbol{\theta})$ is nonlinear in any of the parameters then we must use other methods to find the minimum of $\chi^2(\boldsymbol{\theta})$.

3.2. Numerical minimization of $\chi^2(\boldsymbol{\theta})$

In most problems of practical interest it is too difficult or impossible to find formulas for the estimators in closed form as done in the previous section. Rather, we must find the minimum of $\chi^2(\boldsymbol{\theta})$ (called the *objective function*) numerically. The study of algorithms used to carry out this task constitutes an entire branch of applied mathematics called *mathematical optimization*.

Here we will not go into the details of these algorithms other than to mention some of their general principles. The basic idea is to evaluate the objective function $\chi^2(\boldsymbol{\theta})$ at different points in the parameter space and in this way to try to figure out where it has its minimum. Some algorithms make use of the derivatives of the objective function with respect to the parameters if these are known. One must start at a certain point in the parameter space, say, $\boldsymbol{\theta}_{\text{start}}$. Sometimes default values can be used for this; in other problems one may need to have a rough preliminary estimate.

For example, starting with some initial guess for the parameters, hold all but one of them constant, say, θ_0 , and determine the value of θ_0 that minimizes the χ^2 . This can be done by finding the solution to $\partial\chi^2/\partial\theta_0 = 0$ using, e.g., the Newton–Raphson method. Then hold all parameters constant at the newly found point except the next one, say, θ_1 , and determine its value so as to minimize the χ^2 . After cycling through all of the parameters one will have moved to a point closer to the minimum. Repeat the procedure as many times as necessary until the change in the parameters from one iteration to the next drops below some threshold, or alternatively, until the change in the χ^2 between iterations falls below a specified value.

In fact the most successful algorithms do not minimize by varying separately individual parameters but rather by moving along certain well-chosen directions in the parameter space, e.g., in the direction where the objective function is decreasing most rapidly (called the method of steepest descent) or using so-called conjugate gradient directions. More details on these types of algorithms can be found in Chapter 10 of Ref. [1].

Here we will show a simple example of numerical minimization using routines in Python. Let us suppose we have the data points shown above in Fig. 1 and we want to fit a straight line. A simple but relatively flexible tool in Python is the routine `curve_fit`, which is part of the package `scipy.optimize`. To use this we first need to import the relevant packages:

```
import numpy as np
from scipy.optimize import curve_fit
```

We need to define the function that we will be fitting to the data. For this example suppose it is a straight line:

```
def func(x, *theta):
    theta0, theta1 = theta
    return theta0 + theta1*x
```

Then we need to have the data values (x_i, y_i, σ_i) in the form of NumPy arrays. In practice we would probably read the data values in from a file, but for this example let us suppose we have the arrays defined by

```
x = np.array([1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0])
y = np.array([2.7, 3.9, 5.5, 5.8, 6.5, 6.3, 7.7, 8.5, 8.7])
sig = np.array([0.3, 0.5, 0.7, 0.6, 0.4, 0.3, 0.7, 0.8, 0.5])
```

To carry out the least-squares fit we usually want to supply initial values for the parameters. If we have reasonable guesses then these can be used; if they are not supplied then the program will take the initial values equal to 1.0. Looking at the graph we see that the slope and y intercept are both of order unity, so we define

```
p0 = np.array([1.0, 1.0])
```

In practice the result should be insensitive to the initial values, but for some special problems one may need to choose this point with some care. To actually fit the curve we use

```
thetaHat, cov = curve_fit(func, x, y, p0, sig, absolute_sigma=True)
```

Notice that it is necessary here to include the argument `absolute_sigma=True` in order for the statistical errors of the estimated parameters to have the interpretation that we will require. This results in $\hat{\theta}_0 = 2.26 \pm 0.29$ and $\hat{\theta}_1 = 0.741 \pm 0.057$, as represented by the fitted curve shown earlier in Fig. 3. Here the indicated statistical errors correspond to the standard deviations of the estimators, as described in the next section.

4. Statistical errors of the fitted parameters

In this section we will examine the statistical uncertainties in the estimated parameters, which stem from the random fluctuations in the original measurements. First in Sec. 4.1 we will discuss in some detail how to interpret the statistical errors and how to find them using the Monte Carlo method. In Sec. 4.2 we will show a much simpler way to estimate the errors using derivatives of the function $\chi^2(\boldsymbol{\theta})$, and in Sec. 4.3 we show how to use the statistical uncertainties with error propagation.

4.1. Meaning of statistical errors and estimation by Monte Carlo method

First let us consider carefully the statement that the estimators $\hat{\boldsymbol{\theta}}$ (e.g., $\hat{\theta}_0$ and $\hat{\theta}_1$ in the example above) have statistical errors. What this means is that if we were to repeat the entire experiment with a new, statistically independent set of measurements y_i with $i = 1, \dots, N$, then the fluctuations in the data would in general be different from before and therefore the numerical values of the estimates $\hat{\boldsymbol{\theta}}$ would also differ. Equivalently we can say that the estimators $\hat{\boldsymbol{\theta}}$ are functions of random variables y_1, \dots, y_N , and they are therefore random variables themselves.

Although we may only carry out the real experiment once, we can simulate it many times with the Monte Carlo method. Figure 4 shows the outcomes of many independent

simulations of the same set of measurements of the straight-line fit described above, each of which is characterized by two fitted parameters, $\hat{\theta}_0$ and $\hat{\theta}_1$.




Figure 4: Estimates $\hat{\theta}_0$ and $\hat{\theta}_1$ from repeated simulations of the straight-line fit, each based on a statistically independent set of measurements.

Figures 5 show the projections of the points from the scatter plot onto the $\hat{\theta}_0$ and $\hat{\theta}_1$ axes, which give histograms of the corresponding quantities. The widths, or more precisely, the standard deviations of these distributions are used to quantify the statistical errors in the parameters, which we can write as $\sigma_{\hat{\theta}_0}$ and $\sigma_{\hat{\theta}_1}$ (note the hats).




Figure 5: Projections of the scatter plot from Fig. 4 onto (a) the $\hat{\theta}_0$ and (b) the $\hat{\theta}_1$ axes.

Returning to the scatter plot in Fig. 4, we see another interesting feature, namely, that the cloud of points is tilted. This shows that if the value of, say, $\hat{\theta}_0$ comes out higher than average, then there is a higher probability for the value of $\hat{\theta}_1$ to be lower than its average. That is, the quantities $\hat{\theta}_0$ and $\hat{\theta}_1$ are *correlated*, in this case the correlation being negative.¹

Before we quantify the statistical errors in the estimators we need to review briefly how to characterise the typical amount of fluctuation in random variables. Recall that to do this for a single random variable, say, u , one defines the *variance* $V[u]$ as the mean of its square minus the square of its mean, i.e.,

$$V[u] = \langle u^2 \rangle - \langle u \rangle^2 . \quad (13)$$

¹Often one hears statements such as “the parameters are correlated”, but this is not strictly correct. Here we are treating the parameters θ_0 and θ_1 as unknown constants, not as random variables. It is the estimators, $\hat{\theta}_0$ and $\hat{\theta}_1$, which are functions of the random data, that vary upon repetition of the experiment, and therefore we should say that the *estimators* of the parameters are correlated.

Often we use the notation $V[u] = \sigma_u^2$, and the square root of the variance, σ_u , is the standard deviation. In a similar way, to characterise the amount of variation in a pair of random variables, say, u and v , we define the *covariance* as the mean of their product minus the product of the means,

$$\text{cov}[u, v] = \langle uv \rangle - \langle u \rangle \langle v \rangle . \quad (14)$$

If u and v have some units such as kg or cm, then the covariance has the units given by their product. It is often convenient to use the dimensionless *correlation coefficient* ρ , defined as the covariance of the two variables divided by the product of their standard deviations,

$$\rho = \frac{\text{cov}[u, v]}{\sigma_u \sigma_v} . \quad (15)$$

One can show that the correlation coefficient for any pair of variables lies in the range $-1 \leq \rho \leq 1$. The extreme values of $\rho = \pm 1$ corresponding to a 100% positive (+1) or negative (-1) correlation (i.e., perfect *anticorrelation*). The corresponding scatter plot is then infinitesimally thin and tilted up ($\rho = +1$) or down ($\rho = -1$). If $\rho = 0$ then the variables are uncorrelated, and their scatter plot is not tilted.

If one has more than two random variables, e.g., a set of m estimators $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \dots, \hat{\theta}_m)$, then there is a covariance for each pair, and these numbers can be arranged in a *covariance matrix*,

$$U_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j] . \quad (16)$$

From the definition of covariance it is easy to check that the matrix is square and symmetric. The diagonal elements simply give the variances

$$U_{ii} = \text{cov}[\hat{\theta}_i, \hat{\theta}_i] = \sigma_{\hat{\theta}_i}^2 , \quad (17)$$

and so by taking their square roots one finds the standard deviations $\sigma_{\hat{\theta}_i}$. The set of correlation coefficients also form a matrix ρ_{ij} , given by

$$\rho_{ij} = \frac{U_{ij}}{\sigma_{\hat{\theta}_i} \sigma_{\hat{\theta}_j}} . \quad (18)$$

It is easy to verify that the diagonal elements are $\rho_{ii} = 1$ for all i , which is to say that a random variable is always 100% positively correlated with itself.

From the scatter plot of the simulated estimates $\hat{\theta}_0$ and $\hat{\theta}_1$ we see that they have a negative correlation. Using standard formulas or NumPy routines to estimate covariance from a sample of observed values (see, e.g., Refs. [2, 3]), we find

$$\sigma_{\hat{\theta}_0} = 0.29 , \quad (19)$$

$$\sigma_{\hat{\theta}_1} = 0.057 , \quad (20)$$

$$\text{cov}[\hat{\theta}_0, \hat{\theta}_1] = -0.0142 , \quad (21)$$

$$\rho = -0.86 . \quad (22)$$

4.2. Statistical errors from derivatives of $\chi^2(\boldsymbol{\theta})$

The Monte Carlo technique described in the previous section provides estimates of the covariance of the estimators but it requires a certain amount of work to set up the simulations and work out the required numbers. The covariance matrix can also be determined directly from the single real data set, with no need of simulated data, by using the second derivatives of $\chi^2(\boldsymbol{\theta})$ with respect to the parameters.

Derivation of the required formulas is nontrivial and goes beyond the scope of this module. We claim here without proof that the inverse of the covariance matrix for a set of estimators $U_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]$ can be estimated by

$$U_{ij}^{-1} = \frac{1}{2} \left(\frac{\partial^2 \chi^2}{\partial \theta_i \partial \theta_j} \right)_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} . \quad (23)$$

In general this formula is an approximation but it is exact or at least close to being so in most cases of practical interest. When using the Python routines described above, the `curve_fit` function finds the minimum of $\chi^2(\boldsymbol{\theta})$ numerically and then by taking small steps about the minimum, it finds the matrix of second derivatives to get U^{-1} . It then inverts this and returns the covariance matrix U .

For the example of the straight line shown above, `curve_fit` returns the covariance matrix

$$U = \begin{pmatrix} 0.08537 & -0.01438 \\ -0.01438 & 0.003275 \end{pmatrix} , \quad (24)$$

where the first row/column corresponds to the estimator $\hat{\theta}_0$ and the second to $\hat{\theta}_1$. By taking the square root of the diagonal elements one finds the standard deviations $\sigma_{\hat{\theta}_0} = 0.29$ and $\sigma_{\hat{\theta}_1} = 0.057$, and combining with the covariance one finds the correlation coefficient $\rho_{01} = -0.86$. All of these numbers agree to good level of precision with the values found using the Monte Carlo calculation in Sec. 4.1.

4.3. Using the covariance of estimators with error propagation

Suppose we have carried out a least-squares fit and found a set of estimators $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \dots, \hat{\theta}_m)$ and their covariance matrix $U_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]$. Very often we then form some function u of the estimators as the quantity of interest, e.g., $u(\hat{\boldsymbol{\theta}}) = (\hat{\theta}_1 \hat{\theta}_2 - 2\hat{\theta}_3)^2$. Because the estimators $\hat{\theta}_i$ are themselves random variables, a function of them is also a random variable with a certain variance. To characterise the statistical uncertainty in the function, we need to “propagate” the covariance of the $\hat{\theta}_i$ through to u . An analogous procedure is carried out with more than one function, e.g., u and v , whereby one finds their variances as well as the covariance $\text{cov}[u, v]$.

One way to carry out this procedure would be to simulate the entire experiment many times, just as in Sec. 4.1. Each time we determine the estimates $\hat{\boldsymbol{\theta}}$, with these evaluate the function $u(\hat{\boldsymbol{\theta}})$, and record its value, e.g., in a histogram. By repeating a sufficient number of times one obtains the distribution of u and thus its standard deviation can be found.

Although the Monte Carlo procedure is robust in the sense that it is essentially guaranteed to produce a meaningful result, it may require some effort to set up. So it is useful to have a simpler approximate method, called (linear) *error propagation*. This is based on a linear

approximation to the function in the neighborhood of the estimates actually found. For a pair of functions $u(\hat{\theta})$ and $v(\hat{\theta})$ one finds for the covariance (see, e.g, Ref. [2]),

$$\text{cov}[u, v] \approx \sum_{i,j=1}^m \frac{\partial u}{\partial \hat{\theta}_i} \frac{\partial v}{\partial \hat{\theta}_j} U_{ij}, \quad (25)$$

where $U_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]$ is the covariance matrix of the estimators of the parameters and the derivatives are evaluated at the estimated values of the parameters. In the case where there is only one function, say, u , its variance is given by $\sigma_u^2 = \text{cov}[u, u]$, i.e.,

$$\sigma_u^2 \approx \sum_{i,j=1}^m \frac{\partial u}{\partial \hat{\theta}_i} \frac{\partial u}{\partial \hat{\theta}_j} U_{ij}. \quad (26)$$

In the special case where the covariance may be diagonal, i.e., $U_{ij} = \delta_{ij} \sigma_{\hat{\theta}_i}^2$, then one can carry out one of the two sums to find

$$\sigma_u^2 \approx \sum_{i=1}^m \left(\frac{\partial u}{\partial \hat{\theta}_i} \right)^2 \sigma_{\hat{\theta}_i}^2. \quad (27)$$

It should be emphasised, however, that estimators for parameters are correlated in general, i.e., the covariance matrix usually has nonzero off-diagonal elements, and thus one must use the full covariance in propagating the errors.

The approximation on which Eqs. (25), (26) and (27) are based is that the functions should be linear in the neighborhood of the estimated values, where the size of the relevant neighborhood is determined by the standard deviations of the corresponding estimators. Suppose, for example, we have a function of the form $u = \hat{\theta}_1 / \hat{\theta}_2$, and that $\hat{\theta}_2 = 10 \pm 1$. The function is (exactly) linear in $\hat{\theta}_1$, and although it is nonlinear in $\hat{\theta}_2$, this is approximately linear if we only consider the small region around 10 ± 1 . So in this case linear error propagation should be a good approximation. If, on the other hand, we have $\hat{\theta}_2 = 10 \pm 10$, then the function is very nonlinear if we consider this larger range. In this case linear error propagation can give wildly wrong results. The Monte Carlo approach, however, could be used in such a case to infer the correct distribution of the function u .

5. Goodness-of-fit from the minimum of $\chi^2(\theta)$

In the example from Sec. 2 we hypothesized that the data shown in Fig. 3 could be described by a straight line. Having carried out the fit we then found the values of the parameters for the slope and y -intercept. But how do we know that this hypothesis is valid? Perhaps we should have used a curve with more adjustable parameters, such as the third-order polynomial fitted to the same data in Fig. 1. Or perhaps the data come out as in Fig. 6 below; in this case the hypothesis of a straight line is clearly not very good. Beyond a visual comparison of the fitted curve with the data, how can we quantify the “goodness of fit”?

Before we go any further we should head off a common misunderstanding. One might think that if the hypothesized function is bad, then this should be reflected in large values of the statistical errors for the fitted parameters. This is not true. The level of agreement between data and hypothesis and the size of statistical errors are in general very separate questions, and in the particular example shown here they are completely unrelated. In fact,




Figure 6: Fit resulting in a poor level of agreement between the hypothesized fit function and the data.

to obtain the data points shown in Fig. 6, the points from the earlier example (Fig. 3) were simply shifted up or down; their error bars were left unchanged. And as a result, the standard deviations of the estimators of the parameters shown in Fig. 6 are exactly the same as before.

The reason for this perhaps surprising behaviour lies in the meaning of the statistical errors. What they reflect is the level with which the parameter estimates would fluctuate upon repetition of the experiment with a new set of measured y_i . It could be that each repetition results in poor agreement, as found in Fig. 6, but that the fitted parameters do not vary much for each repetition, so the standard deviations of the estimators may be small. What we would like to quantify, however, is something different, namely the level to which a straight line provides an adequate description of the data.

One way to address this question is to look at the value of the χ^2 at its minimum, i.e., at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$,

$$\chi_{\min}^2 = \sum_{i=1}^N \frac{(y_i - f(x_i; \hat{\boldsymbol{\theta}}))^2}{\sigma_i^2}. \quad (28)$$

From the form of this equation we can immediately see that if the fitted function is in good agreement with the data, then the contribution of each term in the sum is small and therefore χ_{\min}^2 should be small, so this statistic can be used to reflect the goodness-of-fit. But we should not be surprised to see some nonzero value even if the fit function is correct, because of the statistical fluctuations in the data. To make a quantitative statement about the goodness-of-fit we need to know how large of a χ_{\min}^2 to expect in the case where the hypothesized fit function is correct.

Recall that the standard deviation σ_i reflects the typical (more precisely, the root mean square) level of departure of the measurement y_i from its mean. That is to say, if the fitted value $f(x_i; \hat{\boldsymbol{\theta}})$ were equal to the true mean of y_i , then one would expect $|y_i - f(x_i; \hat{\boldsymbol{\theta}})|$ to be similar to σ_i and therefore each normalized residual (cf. Eq. (2)) should be of order unity,

$$\left| \frac{y_i - f(x_i; \hat{\boldsymbol{\theta}})}{\sigma_i} \right| \sim 1. \quad (29)$$

Since there are N terms in the sum, one would naively expect χ_{\min}^2 to be roughly equal to N . In fact this cannot be quite right, since we have adjusted some number of parameters, so the fitted values $f(x_i; \hat{\boldsymbol{\theta}})$ get pulled even closer to y_i than would be obtained from the “true” function. For example, if one had N data points and $m = N$ parameters, then one would be

able to adjust their values so that the curve goes exactly through each of the points and one would find $\chi_{\min}^2 = 0$.

The minimized value χ_{\min}^2 is of course a function of the data and as such it is itself a random variable. If one were to repeat the experiment many times, its value would follow a certain distribution. One can show that if the hypothesised form of the function $f(x; \boldsymbol{\theta})$ is correct, and if the y_i are Gaussian distributed, and if certain additional conditions are satisfied that usually hold in practice, then the statistic χ_{\min}^2 will follow a probability distribution function (pdf) called, not surprisingly, the chi-squared distribution. Calling the variable χ_{\min}^2 here z to simplify the notation, the chi-squared pdf is given by

$$f_{\chi^2}(z; n) = \frac{1}{2^{n/2}\Gamma(n/2)} z^{n/2-1} e^{-z/2}. \quad (30)$$

Here the parameter n is called the “number of degrees of freedom” and takes on positive integer values, $n = 1, 2, \dots$. If there are N measurements y_1, \dots, y_N and m fitted parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)$, then the number of degrees of freedom is given by $n = N - m$.

The distribution of χ_{\min}^2 for the example of the straight-line fit from repeated simulations of the measurement is shown as a histogram in Fig. 7. The fit had $N = 9$ measured y values and $m = 2$ fitted parameters, so the relevant number of degrees of freedom is $n = 7$. Also shown on the plot as a solid curve is the chi-squared pdf for 7 degrees of freedom.




Figure 7: Distribution of the value of χ_{\min}^2 from Monte Carlo simulations of the measurement and fit of a straight line. The solid curve shows the chi-squared pdf for $n = 7$ degrees of freedom.

The mean and standard deviation of a chi-squared distributed variable for n degrees of freedom are given by

$$\langle z \rangle = n, \quad (31)$$

$$\sigma_z = \sqrt{2n}. \quad (32)$$

Since the mean is equal to the number of degrees of freedom, we can immediately get a good idea of the goodness-of-fit simply by comparing the χ_{\min}^2 obtained to $n = N - m$, the number of measurements minus number of fitted parameters. Sometimes this is given as a ratio, χ_{\min}^2/n , called the *chi-squared per degree of freedom*. If this ratio comes out close to unity, the fit is “good”; if it is very large, the fit is bad, and if it is much less than one then it means that the data points have come out closer to the hypothesised curve than what one would expect, given the random fluctuations that should normally be present.

For the fit shown in Fig. 6, there are 9 measurements and 2 fitted parameters, so if the straight-line hypothesis is true, one would expect a value of χ_{\min}^2 of around 7. In the fit shown, a value $\chi_{\min}^2 = 20.9$ was found, or a chi-squared per degree of freedom of almost 3. This can be contrasted with the fit shown in Fig. 3, which has $\chi_{\min}^2 = 8.2$, for the same number of degrees of freedom.

The fact that the value of χ_{\min}^2 in Fig. 6 came out quite high does not prove that the straight-line hypothesis is wrong. It can happen, as a result of the random fluctuations in the data, that the value of χ_{\min}^2 comes out sometimes high and sometimes low. After all it is a random variable, as illustrated by its distribution shown in Fig. 7. We can take a step further in quantifying the goodness-of-fit by giving the probability, under assumption of the hypothesis used (here a straight line), to find a value of χ_{\min}^2 as high as the one found or higher. This can be found by integrating the chi-squared pdf from the value of χ_{\min}^2 observed to infinity,

$$p = \int_{\chi_{\min}^2}^{\infty} f_{\chi^2}(z; n) dz, \quad (33)$$

where $f_{\chi^2}(z; n)$ is the chi-squared pdf for n degrees of freedom from Eq. (30). The probability p obtained from Eq. (33) is called the p -value of the hypothesis. If it comes out very small, it means that there is only a very small probability, if the hypothesis is true, to find data at least as incompatible with it or more so. We may therefore be tempted to think of the p -value as something like the probability that the hypothesis is true, but this is not a correct statement. There is an important distinction between these quantities, which we will not consider in detail here (further discussion can be found, e.g., in Ref. [2]). For now we will simply emphasize that the p -value is a measure of compatibility between the data and a given hypothesis, and if it comes out very small then the hypothesis is disfavoured.

For the straight-line fit shown in Fig. 3, we had $\chi_{\min}^2 = 8.2$ and 7 degrees of freedom. Using these values in Eq. (33) and carrying out the integral with the Python routine `scipy.stats.chi2.sf` gives a p -value of 0.31. Thus if the straight-line hypothesis is correct, there is a substantial probability to get an even higher value of χ_{\min}^2 , and so we would not reject the hypothesis in this case. But in the fit in Fig. 6 we found $\chi_{\min}^2 = 20.9$, corresponding to a p -value of 0.0039. It is still possible that the straight-line hypothesis is correct, but if it is, there is only a 4 in a thousand chance to see a χ_{\min}^2 at least as high as the one we found.

If we find a very low p -value (or equivalently, χ_{\min}^2 substantially greater than the number of degrees of freedom), then we may wish to try another hypothesis. Using the same data as in Fig. 6 we can try, for example, a second-order polynomial,

$$f(x; \boldsymbol{\theta}) = \theta_0 + \theta_1 x + \theta_2 x^2. \quad (34)$$

The fit using this function is shown in Fig. 8. It gives $\chi_{\min}^2 = 3.5$ for 6 degrees of freedom, corresponding to a p -value of 0.75, and thus a much better level of agreement between the data and hypothesis.

6. Least squares with correlated measurements

In the method of least squares as described above, the quantity $\chi^2(\boldsymbol{\theta})$, whose minimum defines the estimators $\hat{\boldsymbol{\theta}}$, is given by




Figure 8: Fit to the same data values as in Fig. 6 but here using a fit function with three adjustable parameters.

$$\chi^2(\boldsymbol{\theta}) = \sum_{i=1}^N \frac{(y_i - f(x_i; \boldsymbol{\theta}))^2}{\sigma_i^2}. \quad (35)$$

As mentioned in Sec. 2, one can show that this recipe follows from a more general principle, namely, the method of maximum likelihood, for the special case where the measured quantities y_i are independent and Gaussian distributed. In practice the measured quantities we deal with very often satisfy these two criteria, and this is an important reason why the method of least squares is so widely used.

It is not uncommon, however, that measured values are Gaussian but they are correlated, and hence not independent. That is, each measurement does not simply have a variance $V[y_i]$, but the set of N measurements is described by an $N \times N$ covariance matrix $V_{ij} = \text{cov}[y_i, y_j]$. The elements of this matrix are related to the correlation coefficients through

$$V_{ij} = \sigma_i \sigma_j \rho_{ij}. \quad (36)$$

Equation (36) includes the special case of $i = j$, i.e., the diagonal elements have a correlation coefficient $\rho_{ii} = 1$ and the corresponding elements $V_{ii} = \sigma_i^2$ are the variances. In such a problem, one can generalise the method of least squares by replacing the sum of squares given above by

$$\chi^2(\boldsymbol{\theta}) = \sum_{i,j=1}^N (y_i - f(x_i; \boldsymbol{\theta})) V_{ij}^{-1} (y_j - f(x_j; \boldsymbol{\theta})), \quad (37)$$

where V^{-1} is the inverse of the covariance matrix. As in the uncorrelated case, this version of $\chi^2(\boldsymbol{\theta})$ can be derived from the method of maximum likelihood where the data follow a *multivariate* Gaussian distribution (see, e.g., [2]).

If we use Eq. (37) for the case where the measurements are in fact uncorrelated, then the covariance matrix is diagonal, i.e., $V_{ij} = \delta_{ij} \sigma_i^2$, and thus its inverse is simply $V_{ij}^{-1} = \delta_{ij} / \sigma_i^2$. Substituting this expression into Eq. (37) allows one to carry out one of the two sums, and the formula for $\chi^2(\boldsymbol{\theta})$ becomes again a sum of squares as in Eq. (35).

7. Exercises

Here we give several exercises that illustrate the methods described above. Exercise 1 is similar to the straight-line fit shown above and requires an extended error analysis. Exercises 2 and 3 use historical data on projectile motion from Galileo and on refraction of light from Ptolemy. The goodness-of-fit is analyzed to choose between competing hypotheses. Exercise 4 explores algorithms that can be used to find the minimum of $\chi^2(\boldsymbol{\theta})$ numerically.

Exercise 1: Polynomial fit and error analysis

Consider the following set of (x, y, σ) data points:

```
x = np.array([1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0])
y = np.array([2.7, 3.9, 5.5, 5.8, 6.5, 6.3, 7.7, 8.5, 8.7])
sig = np.array([0.3, 0.5, 0.7, 0.6, 0.4, 0.3, 0.7, 0.8, 0.5])
```

1(a): Using these data carry out the least-squares fit of an M th order polynomial, i.e., with $M + 1$ adjustable parameters, for $M = 1, 2, 3$.

1(b): For each fit, use the error propagation formula Eq. (26) to find the standard deviation of the fitted function σ_f as a function of x . Note that to do this you will need to compute the derivatives of $f(x; \hat{\boldsymbol{\theta}})$ with respect to the components of $\hat{\boldsymbol{\theta}}$. Display the fitted curve plus-or-minus one standard deviation as a shaded band, and extend the x axis to at least 20. (The shaded band can be made with the function `matplotlib.fill_between`.) The result for $M = 1$ is shown in Fig. 9. Note how the size of the error band increases when one goes to x values outside the region where data are available; investigate how this behaviour changes as the order of the polynomial is increased.




Figure 9: Result of least-squares fit with the one-sigma error shown as a shaded band.

1(c): Consider the fit with $M = 3$ and define the difference

$$\Delta_{ab}(\hat{\boldsymbol{\theta}}) = f(a; \hat{\boldsymbol{\theta}}) - f(b; \hat{\boldsymbol{\theta}}). \quad (38)$$

Using error propagation, find the standard deviation of $\Delta_{ab}(\hat{\boldsymbol{\theta}})$ for $a = 5$ and $b = 6, 10, 20$. Compare these values you find with the standard deviation of f that you plotted for this fit as a shaded band evaluated at both a and b .

Exercise 2: Analysis of Galileo's ball and ramp data

In this exercise we consider an experiment performed by Galileo using a ball and an inclined ramp as shown in Fig. 10. The ball starts at a height h above the edge of the ramp, and its trajectory is forced to be horizontal before it falls over the edge. The horizontal distance d from the edge to the point of impact is measured for different values of h . Five data points obtained by Galileo in 1608 are shown in Table 1 (from Ref. [4]).




Figure 10: The configuration Galileo's ball and ramp experiment.

Table 1: Galileo's data on horizontal distance before impact d for five values of the starting height h . The units are punti (points); one punto is slightly less than 1 mm.

h	d
1000	1500
828	1340
800	1328
600	1172
300	800

We will assume the heights h are known with negligible error, and that the horizontal distances d have uncertainties of $\sigma = 15$ punti (points); one punto is slightly less than 1 mm. It is not actually known what the measurement uncertainties were, but 1–2% is plausible. In addition, we know that if $h = 0$, then the horizontal distance d will be zero, i.e. if the ball is started at the very edge of the ramp, it will fall straight down to the floor.

We could proceed by applying the known laws of mechanics to the system and deriving the relationship between d and h . For purposes of this exercise, however, we will pretend we do not yet know Newton's laws (as was the case for Galileo), and we will simply try different hypotheses and compare their predictions with the data.

Consider the following three hypotheses for the functional relationship between d and h : a linear relation with a single free parameter α ,

$$d = \alpha h , \tag{39}$$

a quadratic relation with two parameters, α and β ,

$$d = \alpha h + \beta h^2 , \tag{40}$$

and something which is a nonlinear function of the parameters,

$$d = \alpha h^\beta . \tag{41}$$

The goals of this exercise are to investigate these three hypotheses. More specifically, you should write a Python program to do the following:

2(a) Carry out a least-squares fit of the parameters for each hypothesis. Find the estimated parameter values, their standard deviations and if relevant their covariance and correlation coefficient.

2(b) Produce plots of the fitted curves along with the data points and their errors.

2(c) Quantify the level of agreement between the data and each hypothesis with both a p -value and the chi-squared per degree of freedom. Discuss which of the hypotheses are disfavoured and which are preferred.

2(d) Use Newton's Laws to derive the relation between d and h and predict the values of the parameters. Discuss how you would compare this prediction to the fit results obtained above. What parameters of the system can be related using your theoretical prediction to the fitted parameters? What important experimental uncertainties could render such a comparison difficult?

Bonus question: From the relation between d and h that you derive in 2(d), one can predict the height of the table in punti. Find the height and using error propagation determine its statistical error from those of the fitted parameters. Compare your result to the value that Galileo recorded in his notebook. (See Appendix A and look where it says "altezza della tavola".)

Exercise 3: Analysis of refraction data from Ptolemy

The astronomer Claudius Ptolemy performed experiments on the refraction of light using a circular copper disc submerged to its centre in water, as illustrated in Fig. 11. Angles of refraction θ_r for 8 values of the angle of incidence θ_i obtained by Ptolemy around 140 A.D. are shown in Table 2 (from Ref. [5]).

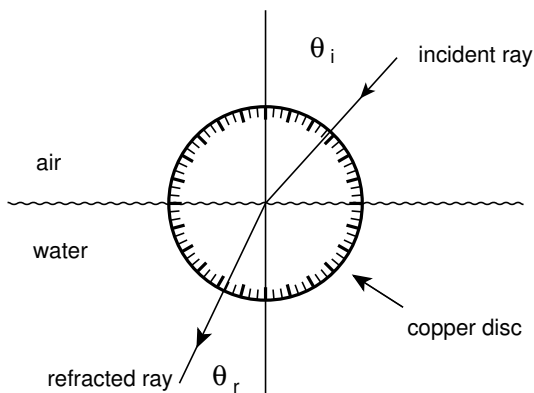


Figure 11: The apparatus used by Ptolemy to investigate the refraction of light.

For purposes of this exercise we will take the angles of incidence to be known with negligible error and treat the angles of reflection as independent Gaussian-distributed measurements

Table 2: Angles of incidence and refraction (in degrees).

θ_i	θ_r
10	8
20	$15\frac{1}{2}$
30	$22\frac{1}{2}$
40	29
50	35
60	$40\frac{1}{2}$
70	$45\frac{1}{2}$
80	50

with standard deviations of $\sigma = \frac{1}{2}^\circ$. (This is a reasonable guess given that the angles are reported to the nearest half degree. Note that we can absorb an error in θ_i into an effective error in θ_r .)

Until the discovery of the correct law of refraction (see below), a commonly used hypothesis was

$$\theta_r = \alpha\theta_i, \quad (42)$$

although it is reported that Ptolemy preferred the formula

$$\theta_r = \alpha\theta_i - \beta\theta_i^2. \quad (43)$$

3(a): Find the least-squares estimates of the parameters for both hypotheses and determine the minimized χ^2 . Comment on the goodness-of-fit for both hypotheses. Is it plausible that all of the data values are based on actual measurements?²

The law of refraction discovered by the Persian mathematician and physicist Ibn Sahl in the 10th century and rediscovered by others including Snell in 1621 is

$$\theta_r = \sin^{-1} \left(\frac{\sin \theta_i}{r} \right), \quad (44)$$

where $r = n_r/n_i$ is the ratio of indices of refraction of the two media.

3(b): Determine the least-squares estimate for r and find value of the minimized χ^2 . Comment on the validity of the Gaussian assumption for θ_r with $\sigma = \frac{1}{2}^\circ$.

For all of the fits, report the parameter estimates, their standard deviations, and where relevant the covariance and correlation coefficient. Make plots of the fitted curves and the data.

Exercise 4: Algorithms for numerical minimization

In most applications we use a specialised software package such as `scipy.optimize` to minimize an objective function like $\chi^2(\boldsymbol{\theta})$. It is nevertheless an interesting exercise to create such

²See R. Feynman, R. Leighton and M. Sands, *The Feynman Lectures on Physics*, Vol. I, Addison-Wesley, Menlo Park, 1963, Section 26-2.

a routine yourself. The goal here is to implement a minimization algorithm in Python and to compare its results to those obtained from `curve_fit`.

Because the mathematics of function minimization can be treated separately from the method of least squares, here we will change notation and suppose that we have an objective function $f(\mathbf{x})$ of an n -dimensional argument $\mathbf{x} = (x_1, \dots, x_n)$. The goal is to find the point in \mathbf{x} -space that gives the minimum of $f(\mathbf{x})$. We will assume in doing so that the analyst supplies a starting point \mathbf{x}_s , which can be based on an initial guess.

All methods that achieve this goal require that one be able to evaluate the objective function f at an arbitrary point \mathbf{x} . If, in addition, one is able to evaluate the derivatives of f with respect to the components of \mathbf{x} then algorithms can be used that will find the minimum much faster. In some cases one may know the derivatives in closed form; in other situations one may have to estimate them numerically. For the functions we have considered in the examples above it is relatively simple to calculate the first and second derivatives of $\chi^2(\boldsymbol{\theta})$ with respect to the parameters, and so for the method we implement here we will suppose that these are available.

One simple strategy to minimize $f(\mathbf{x})$ is to consider all of the components of \mathbf{x} as fixed except one. Starting from the point \mathbf{x}_0 , we vary only this chosen component to minimize f . We then repeat this procedure for all of the components of \mathbf{x} in turn.

Consider first the function f as depending only on one non-constant component, which here for simplicity we will call x , without a subscript. Suppose that at a particular step of the algorithm the value of x is x_0 , i.e., here the index refers to a step of the algorithm, not to the component. We can expand $f(x)$ in a Taylor series to second order about its initial value x_0 , i.e.,

$$f(x) \approx f(x_0) + f'(x_0)(x - x_0) + \frac{1}{2!}f''(x_0)(x - x_0)^2 . \quad (45)$$

To minimize this function with respect to x , we set its derivative equal to zero,

$$f'(x) \approx f'(x_0) + f''(x_0)(x - x_0) = 0 . \quad (46)$$

Solving for x gives

$$x = x_0 - \frac{f'(x_0)}{f''(x_0)} . \quad (47)$$

This x is used to update the corresponding component of \mathbf{x} , which in turn gives an updated estimate of the position of the minimum of $f(\mathbf{x})$. Having updated one component, we then cycle through all of the remaining ones and minimize with respect to each while holding the rest constant.

After carrying out minimizations with respect to all of the components we can evaluate the objective function and see how much it changed from its value at the previous full step. The procedure is repeated until the change in $f(\mathbf{x})$ from one iteration to the next falls below a desired threshold, which determines the final numerical accuracy.

One can easily generalise this algorithm to minimize not with respect to the individual coordinates, but rather along an arbitrary direction in \mathbf{x} -space. If we have available the derivatives of f , i.e.,

$$\nabla f = \left(\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_n} \right), \quad (48)$$

then we can generalise the procedure described above to minimize f along the direction of $-\nabla f$, i.e., the direction in which the function is decreasing most rapidly. This is called the *method of steepest descent*.

Although the method of steepest descent should eventually find the minimum of the objective function, in many cases it turns out to be faster to choose what are called *conjugate gradient* directions; more details can be found in Chapter 10 of Ref. [1].

4(a): The goal of this exercise is to create a Python program that will find the minimum of the objective function $\chi^2(\boldsymbol{\theta})$ numerically. Implement the algorithm with minimization along the coordinate directions as described above and test it by fitting a polynomial to the data shown in Sec. 7. Use different orders for the polynomial and thus different numbers of fitted parameters. Compare the results you obtain with those from the routine `curve_fit` and comment on any differences.

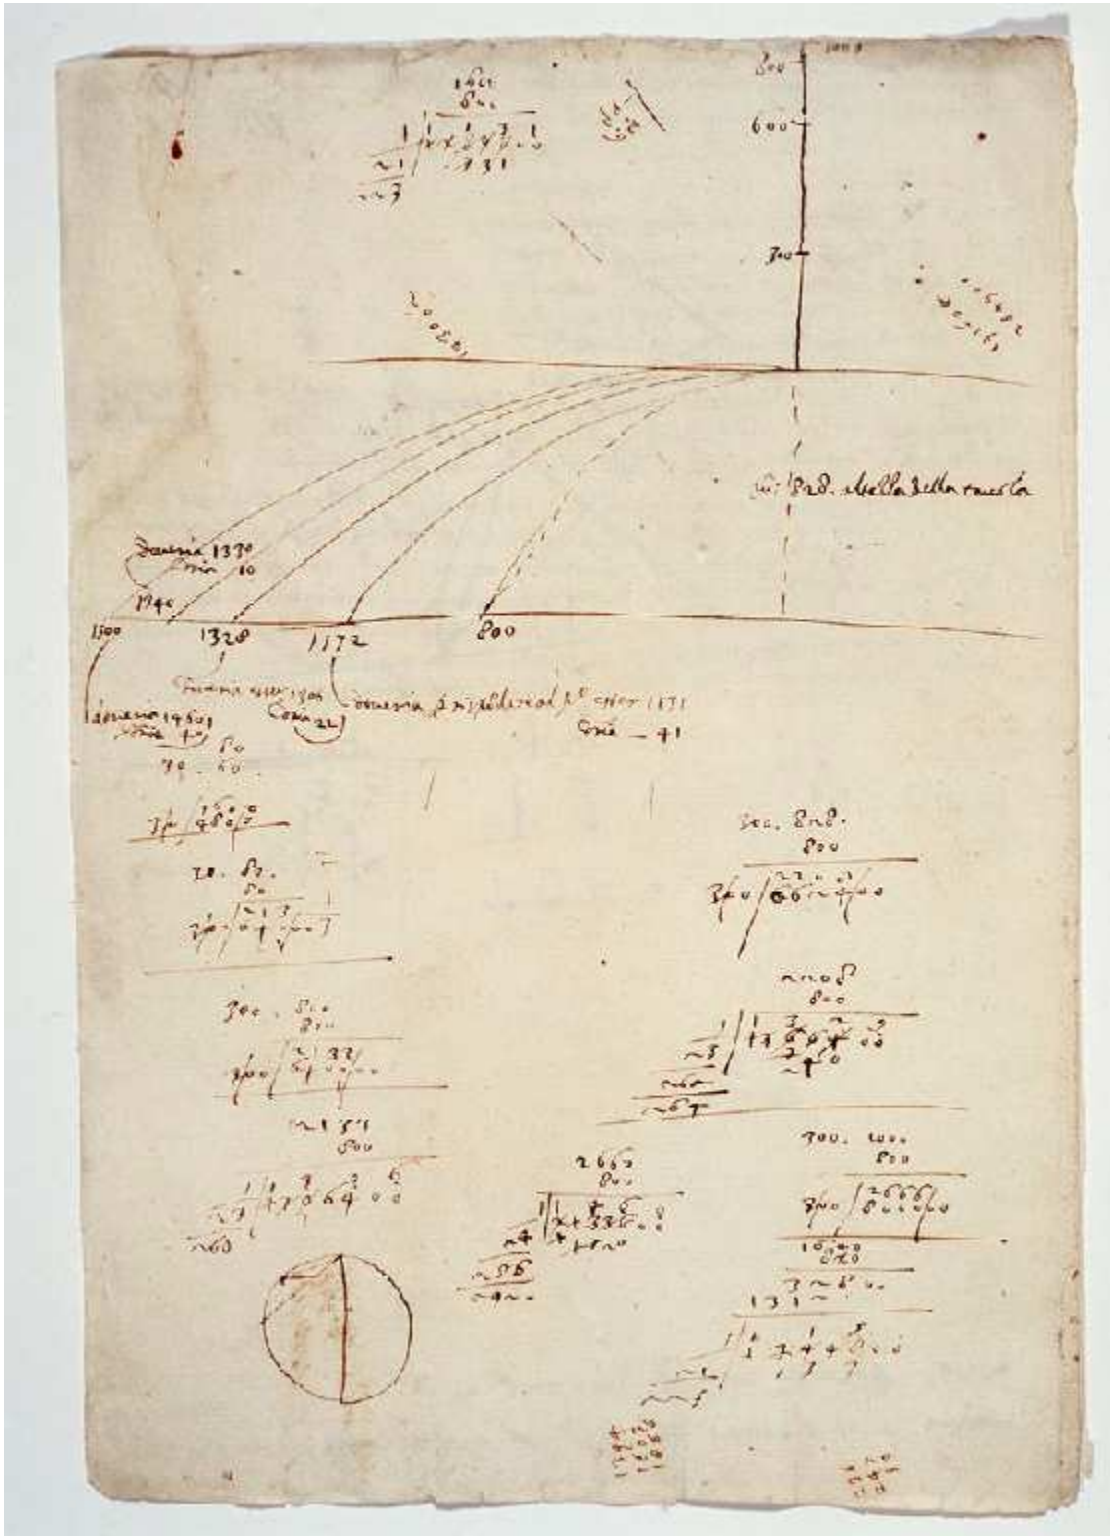
4(b): Extend exercise 4(a) by using the method of steepest descent.

References

- [1] S. Brandt, *Statistical and Computational Methods in Data Analysis*, Springer, New York (1997).
- [2] G. Cowan, *Statistical Data Analysis*, Oxford University Press, 1998.
- [3] RHUL Physics Department, *Mathematical Formula Sheet for Physics* (2017); available on moodle.
- [4] Stillman Drake and James Maclachlan, Galileo's discovery of the parabolic trajectory, *Scientific American* **232** (March 1975) 102; Stillman Drake, *Galileo at Work*, University of Chicago Press, Chicago (1978).
- [5] Olaf Pedersen and Mogens Pihl, *Early Physics and Astronomy: A Historical Introduction*, MacDonald and Janes, London, 1974.
- [6] Galileo Galilei, Manuscript *f.116*, Biblioteca Nazionale Centrale di Firenze, www.bncf.firenze.sbn.it.

A. Galileo's notes on the ball and ramp experiment

Manuscript f.116 from Galileo's notebooks showing the ball and ramp experiment (from [6]).



B. Python code

Program `simpleFit.py` for polynomial fits.

```
1  # simpleFit.py
2  # G. Cowan / RHUL Physics / October 2017
3  # Simple program to illustrate least-squares fitting with curve_fit
4
5  import matplotlib
6  import matplotlib.pyplot as plt
7  import numpy as np
8  from scipy.optimize import curve_fit
9
10 # define fit function
11 def func(x, *theta):
12     theta0, theta1 = theta
13     return theta0 + theta1*x
14
15 # set data values
16 x = np.array([1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0])
17 y = np.array([2.7, 3.9, 5.5, 5.8, 6.5, 6.3, 7.7, 8.5, 8.7])
18 sig = np.array([0.3, 0.5, 0.7, 0.6, 0.4, 0.3, 0.7, 0.8, 0.5])
19
20 # set default parameter values and do the fit
21 p0 = np.array([1.0, 1.0])
22 thetaHat, cov = curve_fit(func, x, y, p0, sig, absolute_sigma=True)
23
24 # Retrieve minimized chi-squared, etc.
25 numPoints = len(x)
26 numPar = len(p0)
27 ndof = numPoints - numPar
28 chisq = sum((y - func(x, *thetaHat))/sig)**2)
29 print ("chisq = ", chisq, " ", ndof = " ", ndof)
30
31 # Print fit parameters and covariance matrix
32 print ("\n", "Fitted parameters and standard deviations:")
33 sigThetaHat = np.sqrt(np.diag(cov))
34 for i in range(len(thetaHat)):
35     print ("thetaHat[" , i, " ] = ", thetaHat[i], " +- ", sigThetaHat[i])
36
37 print ("\n", "i, j, cov[i,j], rho[i,j]:")
38 for i in range(len(thetaHat)):
39     for j in range(len(thetaHat)):
40         rho = cov[i][j] / (sigThetaHat[i]*sigThetaHat[j])
41         print (i, " ", j, " ", cov[i][j], " ", rho)
42
43 # Set up plot
44 matplotlib.rcParams.update({'font.size':18}) # set all font sizes
45 plt.clf()
46 fig, ax = plt.subplots(1,1)
47 plt.gcf().subplots_adjust(bottom=0.15)
48 plt.gcf().subplots_adjust(left=0.15)
49 plt.errorbar(x, y, yerr=sig, xerr=0, color='black', fmt='o', label='data')
50 plt.xlabel(r'$x$')
51 plt.ylabel(r'$y$', labelpad=10)
52 xMin = 0
53 xMax = 10
54 yMin = 0
55 yMax = 10
56 plt.xlim(xMin, xMax)
57 plt.ylim(yMin, yMax)
58 xPlot = np.linspace(xMin, xMax, 100) # enough points for a smooth curve
59 fit = func(xPlot, *thetaHat)
60 plt.plot(xPlot, fit, 'red', linewidth=2, label='fit result')
61
62 # Tweak legend
63 handles, labels = ax.get_legend_handles_labels()
64 handles = [handles[1], handles[0]]
65 labels = [labels[1], labels[0]]
66 handles = [handles[0][0], handles[1]] # turn off error bar for data in legend
```



```
67 plt.legend(handles, labels, loc='lower right', fontsize=14, frameon=False)
68
69 # Make and store plot
70 plt.show()
71 plt.savefig("simpleFit.pdf", format='pdf')
```

G. Cowan
RHUL Physics
Version 3.0 / September 2020