Yet Another
Introduction to Numerical Methods

version August 19, 2025

D.V. Fedorov

ii

---

[1] http://en.wikipedia.org/wiki/GPL
[2] http://en.wikipedia.org/wiki/CC-BY-SA

Preface

This book evolved from lecture notes developed over several years of teaching numerical methods at the University of Aarhus. The book contains short descriptions of some of the most common numerical methods together with illustrational implementations of the discussed algorithms mostly in the C programming language. The implementations are solely for educational purposes, they are not thoroughly tested and are not tuned in any way other than to fit nicely on a single page of the book; some simple optimizations were often sacrificed to conciseness. For serious calculations one should use the GNU Scientific Library.

The content of the book is free as in freedom. You are permitted to copy and redistribute the book in original or modified form either gratis or for a fee. However, you must attribute the original author(s) and pass the same freedom to all recipients of your copies. See the GPL and CC-BY-SA licenses for more details.

2025
Dmitri Fedorov

Contents

Chapter 1

Interpolation

1.1 Introduction

In practice one often meets a situation where the function of interest, $f(x)$, is only represented by a discrete set of tabulated points,

$$\{x_i, y_i \doteq f(x_i) \mid i = 1 \ldots n\},$$

obtained, for example, by sampling, experimentation, or extensive numerical calculations.

Interpolation means constructing a (smooth) function, called *interpolating function* or *interpolant*, which passes exactly through the given points and hopefully approximates the tabulated function between the tabulated points. Interpolation is a specific case of *curve fitting* in which the fitting function must go exactly through the data points.

The interpolating function can be used for different practical needs like estimating the tabulated function between the tabulated points and estimating the derivatives/integrals involving the tabulated function.

1.2 Polynomial interpolation

Polynomial interpolation uses a polynomial as the interpolating function. Given a table of n points, $\{x_i, y_i\}$, where no two x_i are the same, one can construct a polynomial $P(x)$ of the order $n-1$ which passes exactly through the points: $P(x_i) = y_i$. This

polynomial can be intuitively written in the *Lagrange form*,

$$P(x) = \sum_{i=1}^{n} y_i \prod_{k \neq i}^{n} \frac{x - x_k}{x_i - x_k} \ .$$ (1.1)

The Lagrange interpolating polynomial always exists and is unique.

Table 1.1: Polynomial interpolation in C

```
double polinterp(int n, double *x, double *y, double z) {
    double s=0,p;
    for(int i=0;i<n;i++) {
        p=1; for(int k=0;k<n;k++) if(k!=i) p*=(z-x[k])/(x[i]-x[k]);
        s+=y[i]*p; }
    return s; }
```

Higher order interpolating polynomials are susceptible to the *Runge's phenomenon:* erratic oscillations close to the end-points of the interval, see Figure 1.1.
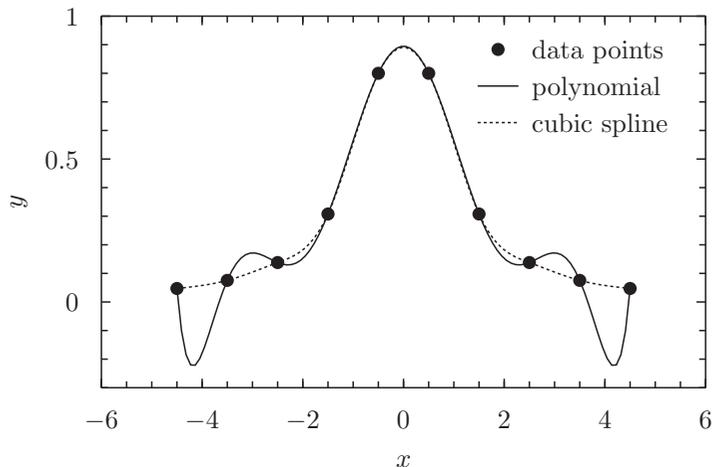
Figure 1.1: Lagrange interpolating polynomial, solid line, showing the Runge's phenomenon: large oscillations at the edges. For comparison the dashed line shows a cubic spline.

1.3 Spline interpolation

Spline interpolation uses a *piecewise polynomial*, $S(x)$, called *spline*, as the interpolating function,

$$S(x) = s_i(x) \text{ if } x \in [x_i, x_{i+1}] \Big|_{i=1,\dots,n-1} , \tag{1.2}$$

where $s_i(x)$ is a polynomial of a given order k. Spline interpolation avoids the problem of Runge's phenomenon. Originally, "spline" was a term for elastic rulers that were bent to pass through a number of predefined points. These were used to make technical drawings by hand.

The spline of the order $k \geq 1$ can be made continuous at the tabulated points,

$$s_i(x_i) = y_i , \ s_i(x_{i+1}) = y_{i+1} \Big|_{i=1,\dots,n-1} , \tag{1.3}$$

together with its $k-1$ derivatives,

$$\begin{aligned}
s_i'(x_{i+1}) &= s_{i+1}'(x_{i+1}) , \\
s_i''(x_{i+1}) &= s_{i+1}''(x_{i+1}) , \\
&\vdots \\
s_i^{(k-1)}(x_{i+1}) &= s_{i+1}^{(k-1)}(x_{i+1}) .
\end{aligned} \quad \Bigg| \ i = 1,\dots,n-2 \tag{1.4}$$

Continuity conditions (1.3) and (1.4) make $kn + n - 2k$ linear equations for the $(n-1)(k+1) = kn + n - k - 1$ coefficients of $n-1$ polynomials (1.2) of the order k. The missing $k-1$ conditions can be chosen (reasonably) arbitrarily.

The most popular is the cubic spline, where the polynomials $s_i(x)$ are of third order. The cubic spline is a continuous function together with its first and second derivatives. The cubic spline has a nice feature that it (sort of) minimizes the total curvature of the interpolating function. This makes the cubic splines look good.

Quadratic splines—continuous with the first derivative—are not nearly as good as cubic splines in most respects. In particular they might oscillate unpleasantly when a quick change in the tabulated function is followed by a period where the function is nearly a constant. Cubic splines are somewhat less susceptible to such oscillations.

Linear spline is simply a *polygon* drawn through the tabulated points.

1.3.1 Linear interpolation

Linear interpolation is a spline with linear polynomials. The continuity conditions (1.3) can be satisfied by choosing the spline in the (intuitive) form

$$s_i(x) = y_i + p_i(x - x_i) , \tag{1.5}$$

where

$$p_i = \frac{\Delta y_i}{\Delta x_i} \ , \quad \Delta y_i \doteq y_{i+1} - y_i \ , \quad \Delta x_i \doteq x_{i+1} - x_i \ . \tag{1.6}$$

The linear spline can be easily differentiated,

$$s_i'(x) = p_i \ , \tag{1.7}$$

and integrated,

$$\int_{x_i}^{x < x_{i+1}} s_i(t)dt = y_i(x - x_i) + p_i \frac{(x - x_i)^2}{2} \ . \tag{1.8}$$

Linear spline is continuous but its derivative is not.

Table 1.2: Linear interpolation in C

```
#include<assert.h>
double linterp(int n, double* x, double* y, double z){
    assert(n>1 && z>=x[0] && z<=x[n−1]);
    int i=0, j=n−1; /* binary search: */
    while(j−i>1){int m=(i+j)/2; if(z>x[m]) i=m; else j=m;}
    double dy=y[i+1]−y[i], dx=x[i+1]−x[i]; assert(dx>0);
    return y[i]+dy/dx*(z−x[i]);
}
```

Note that the search of the interval $[x_i \leq x \leq x_{i+1}]$ in an ordered array $\{x_i\}$ should be done with the *binary search* algorithm (also called *half-interval search*): the point x is compared to the middle element of the array, if it is less than the middle element, the algorithm repeats its action on the sub-array to the left of the middle element, if it is greater, on the sub-array to the right. When the remaining sub-array is reduced to two elements, the interval is found (see Table 1.2). The average number of operations for a binary search is $O(\log n)$.

1.3.2   Quadratic spline

Quadratic spline is made of second order polynomials, conveniently written in the form

$$s_i(x) = y_i + p_i(x - x_i) + c_i(x - x_i)(x - x_{i+1}) \Big|_{i=1,\dots,n-1} \ , \tag{1.9}$$

which identically satisfies the continuity conditions

$$s_i(x_i) = y_i \ , \ s_i(x_{i+1}) = y_{i+1} \Big|_{i=1,\dots,n-1} \ . \tag{1.10}$$

Substituting (1.9) into the derivative continuity condition,

$$s_i'(x_{i+1}) = s_{i+1}'(x_{i+1}) \Big|_{i=1,\dots,n-2} , \tag{1.11}$$

gives $n-2$ equations for $n-1$ unknown coefficients c_i,

$$p_i + c_i \Delta x_i = p_{i+1} - c_{i+1}\Delta x_{i+1} \Big|_{i=1,\dots,n-2} . \tag{1.12}$$

One coefficient can be chosen arbitrarily, for example $c_1 = 0$. The other coefficients can now be calculated recursively from (1.12),

$$c_{i+1} = \frac{1}{\Delta x_{i+1}} \left(p_{i+1} - p_i - c_i\Delta x_i \right) \Big|_{i=1,\dots,n-2} . \tag{1.13}$$

Alternatively, one can choose $c_{n-1} = 0$ and make the backward-recursion

$$c_i = \frac{1}{\Delta x_i} \left(p_{i+1} - p_i - c_{i+1}\Delta x_{i+1} \right) \Big|_{i=n-2,\dots,1} . \tag{1.14}$$

In practice, unless you know what your c_1 (or c_{n-1}) is, it is better to run both recursions and then average the resulting c's. This amounts to first running the forward-recursion from $c_1 = 0$ and then the backward recursion from $\frac{1}{2}c_{n-1}$.

The optimized form (1.9) of the quadratic spline can also be written in the ordinary form suitable for differentiation and integration,

$$s_i(x) = y_i + b_i(x - x_i) + c_i(x - x_i)^2 , \text{ where } b_i = p_i - c_i\Delta x_i . \tag{1.15}$$

An implementation of quadratic spline in C is listed in Table 1.3.2

1.3.3   Cubic spline

Cubic splines are made of third order polynomials,

$$s_i(x) = y_i + b_i(x - x_i) + c_i(x - x_i)^2 + d_i(x - x_i)^3 . \tag{1.16}$$

This form automatically satisfies the first half of continuity conditions (1.3): $s_i(x_i) = y_i$. The second half of continuity conditions (1.3), $s_i(x_{i+1}) = y_{i+1}$, and the continuity of the first and second derivatives (1.4) give a set of equations,

$$\begin{aligned}
y_i + b_i h_i + c_i h_i^2 + d_i h_i^3 &= y_{i+1} , \ i = 1, \dots, n-1 \\
b_i + 2c_i h_i + 3d_i h_i^2 &= b_{i+1} , \ i = 1, \dots, n-2 \\
2c_i + 6d_i h_i &= 2c_{i+1} , \ i = 1, \dots, n-2
\end{aligned} \tag{1.17}$$

Table 1.3: Quadratic spline in C

```c
#include <stdlib.h>
#include <assert.h>
typedef struct {int n; double *x, *y, *b, *c;} qspline;
qspline* qspline_alloc(int n,double* x,double* y){ //builds qspline
   qspline *s = (qspline*)malloc(sizeof(qspline));//spline
   s->b = (double*)malloc((n-1)*sizeof(double));   // b_i
   s->c = (double*)malloc((n-1)*sizeof(double));   // c_i
   s->x = (double*)malloc(n*sizeof(double));       // x_i
   s->y = (double*)malloc(n*sizeof(double));       // y_i
   s->n = n;  for(int i=0;i<n;i++){s->x[i]=x[i];  s->y[i]=y[i];}
   int i; double p[n-1], h[n-1];                    //VLA from C99
   for(i=0;i<n-1;i++){h[i]=x[i+1]-x[i];  p[i]=(y[i+1]-y[i])/h[i];}
   s->c[0]=0;                                       //recursion up:
   for(i=0;i<n-2;i++)s->c[i+1]=(p[i+1]-p[i]-s->c[i]*h[i])/h[i+1];
   s->c[n-2]/=2;                                    //recursion down:
   for(i=n-3;i>=0;i--)s->c[i]=(p[i+1]-p[i]-s->c[i+1]*h[i+1])/h[i];
   for(i=0;i<n-1;i++)s->b[i]=p[i]-s->c[i]*h[i];
   return s; }
double qspline_eval(qspline *s, double z){       //evaluates s(z)
   assert(z>=s->x[0] && z<=s->x[s->n-1]);
   int i=0, j=s->n-1;                             //binary search:
   while(j-i>1){int m=(i+j)/2; if(z>s->x[m]) i=m; else j=m;}
   double h=z-s->x[i];
   return s->y[i]+h*(s->b[i]+h*s->c[i]); }//inerpolating polynomial
void qspline_free(qspline *s){ //free the allocated memory
   free(s->x); free(s->y); free(s->b); free(s->c); free(s); }
```

where $h_i \doteq x_{i+1} - x_i$.

The set of equations (1.17) is a set of $3n-5$ linear equations for the $3(n-1)$ unknown coefficients $\{a_i, b_i, c_i \mid i = 1, \ldots, n-1\}$. Therefore two more equations should be added to the set to find the coefficients. If the two extra equations are also linear, the total system is linear and can be easily solved.

The spline is called *natural* if the extra conditions are given as vanishing second derivatives at the end-points,

$$S''(x_1) = S''(x_n) = 0 , \tag{1.18}$$

which gives

$$
\begin{aligned}
c_1 &= 0 , \\
c_{n-1} + 3d_{n-1}h_{n-1} &= 0 .
\end{aligned}
\tag{1.19}
$$

Solving the first two equations in (1.17) for c_i and d_i gives[1]

$$\begin{aligned}
c_i h_i &= -2b_i - b_{i+1} + 3p_i , \\
d_i h_i^2 &= b_i + b_{i+1} - 2p_i ,
\end{aligned} \qquad (1.20)$$

where $p_i \doteq \Delta y_i / h_i$. The natural conditions (1.19) and the third equation in (1.17) then produce the following tridiagonal system of n linear equations for the n coefficients b_i,

$$\begin{aligned}
2b_1 + b_2 &= 3p_1 , \\
b_i + \left(2\frac{h_i}{h_{i+1}} + 2\right) b_{i+1} + \frac{h_i}{h_{i+1}} b_{i+2} &= 3\left(p_i + p_{i+1}\frac{h_i}{h_{i+1}}\right) \Big|_{i=1,\ldots,n-2} , \\
b_{n-1} + 2b_n &= 3p_{n-1} ,
\end{aligned} \qquad (1.21)$$

or, in the matrix form,

$$\begin{pmatrix}
D_1 & Q_1 & 0 & 0 & \cdots \\
1 & D_2 & Q_2 & 0 & \cdots \\
0 & 1 & D_3 & Q_3 & \cdots \\
\vdots & \vdots & \ddots & \ddots & \ddots \\
\cdots & \cdots & 0 & 1 & D_n
\end{pmatrix}
\begin{pmatrix}
b_1 \\ \vdots \\ \vdots \\ \vdots \\ b_n
\end{pmatrix}
=
\begin{pmatrix}
B_1 \\ \vdots \\ \vdots \\ \vdots \\ B_n
\end{pmatrix} \qquad (1.22)$$

where the elements D_i at the main diagonal are

$$D_1 = 2 , \; D_{i+1} = 2\frac{h_i}{h_{i+1}} + 2 \Big|_{i=1,\ldots,n-2} , \; D_n = 2 , \qquad (1.23)$$

the elements Q_i at the above-main diagonal are

$$Q_1 = 1 , \; Q_{i+1} = \frac{h_i}{h_{i+1}} \Big|_{i=1,\ldots,n-2} , \qquad (1.24)$$

and the right-hand side terms B_i are

$$B_1 = 3p_1 , \; B_{i+1} = 3\left(p_i + p_{i+1}\frac{h_i}{h_{i+1}}\right) \Big|_{i=1,\ldots,n-2} , \; B_n = 3p_{n-1} . \qquad (1.25)$$

This system can be solved by one run of Gauss elimination and then a run of back-substitution. After a run of Gaussian elimination the system becomes

$$\begin{pmatrix}
\tilde{D}_1 & Q_1 & 0 & 0 & \cdots \\
0 & \tilde{D}_2 & Q_2 & 0 & \cdots \\
0 & 0 & \tilde{D}_3 & Q_3 & \cdots \\
\vdots & \vdots & \ddots & \ddots & \ddots \\
\cdots & \cdots & 0 & 0 & \tilde{D}_n
\end{pmatrix}
\begin{pmatrix}
b_1 \\ \vdots \\ \vdots \\ \vdots \\ b_n
\end{pmatrix}
=
\begin{pmatrix}
\tilde{B}_1 \\ \vdots \\ \vdots \\ \vdots \\ \tilde{B}_n
\end{pmatrix} , \qquad (1.26)$$

---

[1]introducing an auxiliary coefficient b_n.

where

$$\tilde{D}_1 = D_1 \, , \ \tilde{D}_i = D_i - Q_{i-1}/\tilde{D}_{i-1} \, \Big|_{i=2,\ldots,n} \, , \tag{1.27}$$

and

$$\tilde{B}_1 = B_1 \, , \ \tilde{B}_i = B_i - \tilde{B}_{i-1}/\tilde{D}_{i-1} \, \Big|_{i=2,\ldots,n} \, . \tag{1.28}$$

The triangular system (1.26) can be solved by a run of back-substitution,

$$b_n = \tilde{B}_n/\tilde{D}_n \, , \ b_i = (\tilde{B}_i - Q_i b_{i+1})/\tilde{D}_i \, \Big|_{i=n-1,\ldots,1} \, . \tag{1.29}$$

A C-implementation of cubic spline is listed in Table 1.3.3

1.3.4   Sub-splines

Sub-splines are—like splines—piecewise polynomials. However, unlike splines the sub-splines dispense with the demand of maximal differentiability of the spline—hence the name. Instead of maximal differentiability the sub-splines use one (or more) of their free parameters to achieve some other goals. Typically the sub-splines try to minimize the unpleasant wiggles that splines are prone to when the interpolated data contain a sub-set of points that are close to a line or an outlier.

Akima sub-spline [1] is an interpolating function in the form of a piecewise cubic polynomial, similar to the cubic spline,

$$S(x)\Big|_{x\in[x_i,x_{i+1}]} = a_i + b_i(x - x_i) + c_i(x - x_i)^2 + d_i(x - x_i)^3 \doteq A_i(x) \, . \tag{1.30}$$

However, unlike the cubic spline, Akima sub-spline dispenses with the demand of maximal differentiability of the spline—in this case, the continuity of the second derivative—hence the name *sub-spline*. Instead of achieving maximal differentiability Akima sub-splines try to reduce the wiggling which the ordinary splines are typically prone to (see Figure 1.2).

First let us note that the coefficients $\{a_i, b_i, c_i, d_i\}$ in eq. (1.30) are determined by the values of the derivatives $S_i' \doteq S'(x_i)$ of the sub-spline through the continuity conditions for the sub-spline and its first derivative,

$$A_i(x_i) = y_i, \ A_i'(x_i) = S_i', \ A_i(x_{i+1}) = y_{i+1}, \ A_i'(x_{i+1}) = S_{i+1}'. \tag{1.31}$$

Indeed, inserting (1.30) into (1.31) and solving for the coefficients gives

$$a_i = y_i, \ b_i = S_i', \ c_i = \frac{3p_i - 2S_i' - S_{i+1}'}{\Delta x_i}, \ d_i = \frac{S_i' + S_{i+1}' - 2p_i}{(\Delta x_i)^2}, \tag{1.32}$$

where $p_i \doteq \Delta y_i/\Delta x_i$, $\Delta y_i \doteq y_{i+1} - y_i$, $\Delta x_i \doteq x_{i+1} - x_i$.
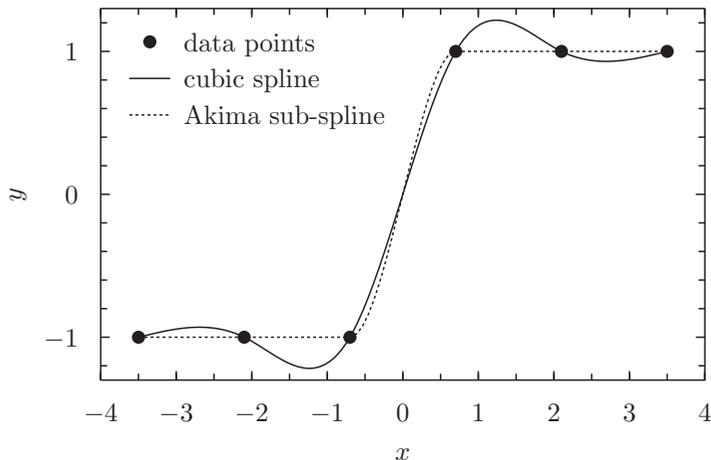
Figure 1.2: A cubic spline (solid line) showing the typical wiggles, compared to the Akima sub-spline (dashed line) where the wiggles are essentially removed.

In the ordinary cubic spline the derivatives S_i' are determined by the continuity condition of the second derivative of the spline. Sub-splines do without this continuity condition and can instead use the derivatives as free parameters to be chosen to satisfy some other condition.

Akima suggested to minimize the wiggling by choosing the derivatives as linear combinations of the nearest slopes,

$$S_i' = \frac{w_{i+1}p_{i-1} + w_{i-1}p_i}{w_{i+1} + w_{i-1}} \;, \quad \text{if} \quad w_{i+1} + w_{i-1} \neq 0 \;, \tag{1.33}$$

$$S_i' = \frac{p_{i-1} + p_i}{2} \;, \quad \text{if} \quad w_{i+1} + w_{i-1} = 0 \;, \tag{1.34}$$

where the weights w_i are given as

$$w_i = |p_i - p_{i-1}| \;. \tag{1.35}$$

The idea is that if three points lie close to a line, the sub-spline in this vicinity has to be close to this line. In other words, if $|p_i - p_{i-1}|$ is small, the nearby derivatives must be close to p_i.

The first two and the last two points need a special prescription, for example (naively) one can simply use

$$S_1' = p_1, \; S_2' = \frac{1}{2}p_1 + \frac{1}{2}p_2, \; S_n' = p_{n-1}, \; S_{n-1}' = \frac{1}{2}p_{n-1} + \frac{1}{2}p_{n-2}. \tag{1.36}$$

Table (1.5) shows a C-implementation of this algorithm.

1.4   Rational function interpolation

As the name suggests, the rational interpolation uses a rational function (a ratio of two polynomials) as the interpolant,

$$r_m^k(x) = \frac{p_0 + p_1 x + \cdots + p_k x^k}{q_0 + q_1 x + \cdots + q_m x^m} \, , \tag{1.37}$$

where $m > 0$. The rational interpolants are (generally) not susceptible to Runge phenomenon and are infinitely differentiable.

One popular family of rational function interpolants is the so called *univariate barycentric interpolation*. One example of this can be illustrated as follows. Suppose the table to interpolate contains only two points, (x_0, y_0) and (x_1, y_1). Then the (linear) interpolant can be intuitively written as

$$F(x) = \frac{x_1 - x}{x_1 - x_0} y_0 + \frac{x - x_0}{x_1 - x_0} y_1 \, . \tag{1.38}$$

In order to generalize this approach to larger tables one can rewrite it as

$$F(x) = \frac{(x - x_1)y_0 - (x - x_0)y_1}{-(x_1 - x_0)} = \frac{(x - x_1)y_0 - (x - x_0)y_1}{(x - x_1) - (x - x_0)} \, . \tag{1.39}$$

Now we divide both the numerator and denominator by $(x - x_0)(x - x_1)$,

$$F(x) = \frac{\dfrac{1}{x - x_0} y_0 - \dfrac{1}{x - x_1} y_1}{\dfrac{1}{x - x_0} - \dfrac{1}{x - x_1}} = \frac{\displaystyle\sum_{i=0}^{1} \dfrac{(-1)^i}{x - x_i} y_i}{\displaystyle\sum_{i=0}^{1} \dfrac{(-1)^i}{x - x_i}} \tag{1.40}$$

Now, the generalization to a table with $n + 1$ points seems as easy as changing the summation limit from 1 to n (first suggested by Berrut in 1988 [5]),

$$B_1(x) = \frac{\displaystyle\sum_{i=0}^{n} \dfrac{(-1)^i}{x - x_i} y_i}{\displaystyle\sum_{i=0}^{n} \dfrac{(-1)^i}{x - x_i}} \, . \tag{1.41}$$

In order to see that B_1 indeed goes through the tabulated points we can calculate the limit,

$$\lim_{x \to x_i} B_1(x) = \frac{\dfrac{(-1)^i}{x - x_i} y_i}{\dfrac{(-1)^i}{x - x_i}} = y_i \, . \tag{1.42}$$

One can show that $B_1(x)$ is indeed a rational function (of the order of at most n by n) by multiplying both the numerator and the denominator by $\prod_i (x - x_i)$. One can also show that $B_1(x)$ has no poles on the real axis as the denominator never vanishes for real arguments. Indeed suppose that $x \in [x_0, x_1]$. The denominator is then given as

$$\sum_{i=0}^{n} \frac{(-1)^i}{x - x_i} = \left(\frac{1}{x - x_0} \right) + \left(\frac{1}{x_1 - x} - \frac{1}{x_2 - x} \right) + \left(\frac{1}{x_3 - x} - \ldots \right) + \ldots \quad (1.43)$$

which is always larger than zero since every term in parentheses is larger than zero (the proof is similar if x falls in any other sub-interval).

Berrut has also suggested a slightly different rational interpolant,

$$B_2(x) = \frac{\dfrac{1}{x - x_0} y_0 + 2 \sum\limits_{i=1}^{n-1} \dfrac{(-1)^i}{x - x_i} y_i + \dfrac{(-1)^n}{x - x_n} y_n}{\dfrac{1}{x - x_0} + 2 \sum\limits_{i=1}^{n-1} \dfrac{(-1)^i}{x - x_i} + \dfrac{(-1)^n}{x - x_n}} . \quad (1.44)$$

Berrut interpolants are as slow to evaluate as the Lagrange interpolating polynomial—$O(n)$ operations—but are more stable, see the illustration on Figure (1.3).

1.5 Multivariate interpolation

Interpolation of a function in more than one variable is called *multivariate interpolation*. The function of interest is represented as a set of discrete points in a multidimensional space. The points may or may not lie on a regular grid.

1.5.1 Nearest-neighbor interpolation

Nearest-neighbor interpolation approximates the value of the function at a non-tabulated point by the value at the nearest tabulated point, yielding a piecewise-constant interpolating function. It can be used for both regular and irregular grids.

1.5.2 Piecewise-linear interpolation

Piecewise-linear interpolation is used to interpolate functions of two variables tabulated on irregular grids. The tabulated 2D region is triangulated – subdivided into a set of non-intersecting triangles whose union is the original region. Inside each triangle the interpolating function $S(x, y)$ is taken in the linear form,

$$S(x, y) = a + bx + cy , \quad (1.45)$$

where the three constants are determined by the three conditions that the interpolating function is equal the tabulated values at the three vertexes of the triangle.

Figure 1.3: Polynomial interpolant showing the Runge's phenomenon (large oscillations at the edges) compared to rational Berrut interpolant where the oscillations are significantly reduced.
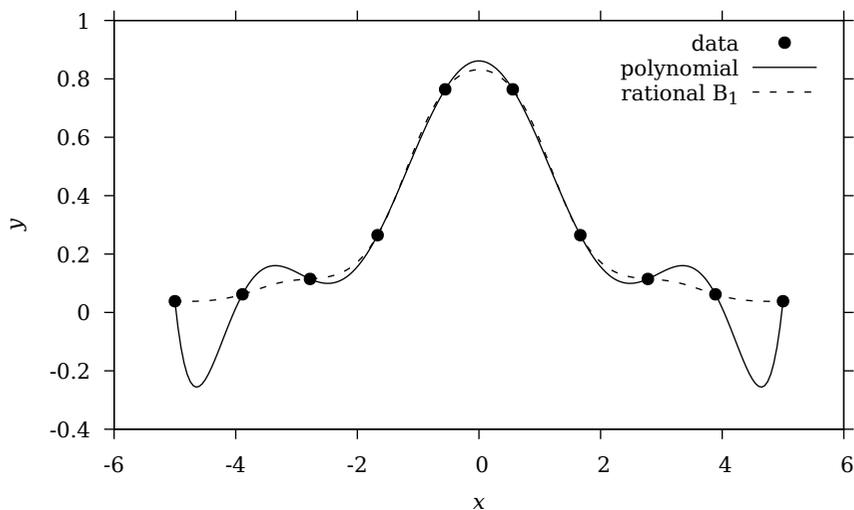
1.5.3   Bi-linear interpolation

Bi-linear interpolation is used to interpolate functions of two variables tabulated on regular rectilinear 2D grids. The interpolating function $B(x,y)$ inside each of the grid rectangles is taken as a bilinear function of x and y,

$$B(x, y) = a + bx + cy + dxy \, , \tag{1.46}$$

where the four constants a, b, c, d are obtained from the four conditions that the interpolating function is equal the tabulated values at the four nearest tabulated points (which are the vertexes of the given grid rectangle).

Table 1.4: Cubic spline in C

```c
#include<stdlib.h>
#include<assert.h>
#include<stdio.h>
typedef struct {int n; double *x,*y,*b,*c,*d;} cubic_spline;
cubic_spline* cubic_spline_alloc(int n, double *x, double *y)
{// builds natural cubic spline
    cubic_spline* s = (cubic_spline*)malloc(sizeof(cubic_spline));
    s->x = (double*)malloc(n*sizeof(double));
    s->y = (double*)malloc(n*sizeof(double));
    s->b = (double*)malloc(n*sizeof(double));
    s->c = (double*)malloc((n-1)*sizeof(double));
    s->d = (double*)malloc((n-1)*sizeof(double));
    s->n = n;  for(int i=0;i<n;i++){s->x[i]=x[i];  s->y[i]=y[i];}
    double h[n-1],p[n-1];  // VLA
    for(int i=0;i<n-1;i++){h[i]=x[i+1]-x[i];  assert(h[i]>0);}
    for(int i=0;i<n-1;i++) p[i]=(y[i+1]-y[i])/h[i];
    double D[n], Q[n-1], B[n];// building the tridiagonal system:
    D[0]=2;  for(int i=0;i<n-2;i++)D[i+1]=2*h[i]/h[i+1]+2;  D[n-1]=2;
    Q[0]=1;  for(int i=0;i<n-2;i++)Q[i+1]=h[i]/h[i+1];
    for(int i=0;i<n-2;i++)B[i+1]=3*(p[i]+p[i+1]*h[i]/h[i+1]);
    B[0]=3*p[0];  B[n-1]=3*p[n-2];  //Gauss elimination :
    for(int i=1;i<n;i++){ D[i]-=Q[i-1]/D[i-1];  B[i]-=B[i-1]/D[i-1];  }
    s->b[n-1]=B[n-1]/D[n-1];  //back-substitution :
    for(int i=n-2;i>=0;i--) s->b[i]=(B[i]-Q[i]*s->b[i+1])/D[i];
    for(int i=0;i<n-1;i++){
        s->c[i]=(-2*s->b[i]-s->b[i+1]+3*p[i])/h[i];
        s->d[i]=(s->b[i]+s->b[i+1]-2*p[i])/h[i]/h[i];
    }
    return s;
}
double cubic_spline_eval(cubic_spline *s,double z){
        assert(z>=s->x[0] && z<=s->x[s->n-1]);
        int i=0, j=s->n-1;// binary search for the interval for z :
        while(j-i>1){int m=(i+j)/2;  if(z>s->x[m]) i=m;  else j=m;  }
        double h=z-s->x[i];// calculate the inerpolating spline :
        return s->y[i]+h*(s->b[i]+h*(s->c[i]+h*s->d[i]));
}
void cubic_spline_free(cubic_spline *s){ //free the allocated memory
    free(s->x);free(s->y);free(s->b);free(s->c);free(s->d);free(s);}
```

Table 1.5: Akima sub-spline in C

```c
#include<assert.h>
#include<stdlib.h>
#include<math.h>
typedef struct {int n; double *x,*y,*b,*c,*d;} akima_spline;
akima_spline* akima_spline_alloc(int n, double *x, double *y){
    assert(n>2); double h[n-1],p[n-1]; /* VLA */
    for(int i=0;i<n-1;i++){h[i]=x[i+1]-x[i]; assert(h[i]>0);}
    for(int i=0;i<n-1;i++) p[i]=(y[i+1]-y[i])/h[i];
    akima_spline *s = (akima_spline*)malloc(sizeof(akima_spline));
    s->x = (double*)malloc(n*sizeof(double));
    s->y = (double*)malloc(n*sizeof(double));
    s->b = (double*)malloc(n*sizeof(double));
    s->c = (double*)malloc((n-1)*sizeof(double));
    s->d = (double*)malloc((n-1)*sizeof(double));
    s->n = n; for(int i=0;i<n;i++){s->x[i]=x[i]; s->y[i]=y[i];}
    s->b[0]   =p[0];    s->b[1]   =(p[0]+p[1])/2;
    s->b[n-1]=p[n-2]; s->b[n-2]=(p[n-2]+p[n-3])/2;
    for(int i=2;i<n-2;i++){
        double w1=fabs(p[i+1]-p[i]), w2=fabs(p[i-1]-p[i-2]);
        if(w1+w2==0) s->b[i]=(p[i-1]+p[i])/2;
        else s->b[i]=(w1*p[i-1]+w2*p[i])/(w1+w2);
    }
    for(int i=0;i<n-1;i++){
        s->c[i]=(3*p[i]-2*s->b[i]-s->b[i+1])/h[i];
        s->d[i]=(s->b[i+1]+s->b[i]-2*p[i])/h[i]/h[i];
    }
    return s;
}
double akima_spline_eval(akima_spline *s, double z){
    assert(z>=s->x[0] && z<=s->x[s->n-1]);
    int i=0, j=s->n-1;
    while(j-i>1){int m=(i+j)/2; if(z>s->x[m]) i=m; else j=m;}
    double h=z-s->x[i];
    return s->y[i]+h*(s->b[i]+h*(s->c[i]+h*s->d[i]));
}
void akima_spline_free(akima_spline *s){
    free(s->x);free(s->y);free(s->b);free(s->c);free(s->d);free(s);}
```

Chapter 2

Systems of linear equations

2.1 Introduction

A *system of linear equations* (or *linear system*) is a collection of linear equations involving the same set of unknown variables. A general system of n linear equations with m unknowns can be written as

$$
\begin{cases}
A_{11}x_1 & + & A_{12}x_2 & + & \cdots & + & A_{1m}x_m & = & b_1 \\
A_{21}x_1 & + & A_{22}x_2 & + & \cdots & + & A_{2m}x_m & = & b_2 \\
\vdots & & \vdots & & & & \vdots & & \vdots \\
A_{n1}x_1 & + & A_{n2}x_2 & + & \cdots & + & A_{nm}x_m & = & b_n
\end{cases} , \tag{2.1}
$$

where $x_1, x_2, \ldots, x_m$ are the unknown variables, $A_{11}, A_{12}, \ldots, A_{nm}$ are the (constant) coefficients, and $b_1, b_2, \ldots, b_n$ are the (constant) right-hand side terms.

The system can be equivalently written in the matrix form,

$$
\mathrm{A}\mathbf{x} = \mathbf{b} , \tag{2.2}
$$

where $\mathrm{A} \doteq \{A_{ij}\}$ is the $n \times m$ matrix of the coefficients, $\mathbf{x} \doteq \{x_j\}$ is the size-m column-vector of the unknown variables, and $\mathbf{b} \doteq \{b_i\}$ is the size-n column-vector of right-hand side terms.

A solution to a linear system is a set of values for the variables $\mathbf{x}$ which satisfies all equations.

Systems of linear equations occur quite regularly in applied mathematics. Therefore computational algorithms for finding solutions of linear systems are an important part of numerical methods. A system of non-linear equations can often be approximated by a linear system – a helpful technique (called *linearization*) in creating a mathematical model of an otherwise a more complex system.

If $m = n$ the matrix A is called *square*. A square system has a unique solution if A is invertible.

2.2   Triangular systems

An efficient algorithm to solve numerically a square system of linear equations is to transform the original system into an equivalent *triangular system*,

$$T\mathbf{y} = \mathbf{c} \ , \tag{2.3}$$

where T is a *triangular matrix* – a special kind of square matrix where the matrix elements either below (upper triangular) or above (lower triangular) the main diagonal are zero.

Indeed, an upper triangular system $U\mathbf{y} = \mathbf{c}$ can be easily solved by *back-substitution*,

$$y_i = \frac{1}{U_{ii}} \left(c_i - \sum_{k=i+1}^{n} U_{ik} y_k \right), \ i = n, n-1, \ldots, 1 \ , \tag{2.4}$$

where one first computes $y_n = b_n/U_{nn}$, then substitutes *back* into the previous equation to solve for y_{n-1}, and repeats through y_1.

Here is a Csharp implementation of the in-place[1] back-substitution:

```
static void backsub(matrix U, vector c){
    for(int i=c.size-1; i>=0; i--){
        double sum=0;
        for(int k=i+1; k<c.size; k++) sum+=U[i,k]*c[k];
        c[i]=(c[i]-sum)/U[i,i]; } }
```

For a lower triangular system $L\mathbf{y} = \mathbf{c}$ the equivalent procedure is called *forward-substitution*,

$$y_i = \frac{1}{L_{ii}} \left(c_i - \sum_{k=1}^{i-1} L_{ik} y_k \right), \ i = 1, 2, \ldots, n \ . \tag{2.5}$$

2.3   Reduction to triangular form

Popular algorithms for reducing a square system of linear equations to a triangular form are *LU-decomposition* and *QR-decomposition*.

---

[1] here *in-place* means the right-hand side $\mathbf{c}$ is replaced by the solution $\mathbf{y}$.

2.3.1 QR-decomposition

QR-decomposition is a factorization of a matrix into a product of an orthogonal matrix Q, such that $Q^{\mathsf{T}}Q = 1$, where $^{\mathsf{T}}$ denotes transposition, and a right triangular matrix R, such that

$$A = QR . \tag{2.6}$$

QR-decomposition can be used to convert (by multiplying with Q^{T} from the left) a linear system $A\mathbf{x} = \mathbf{b}$ into the triangular form,

$$R\mathbf{x} = Q^{\mathsf{T}}\mathbf{b} , \tag{2.7}$$

which can be solved directly by back-substitution.

QR-decomposition can also be performed on non-square matrices with few long columns. Generally speaking a rectangular $n \times m$ matrix A can be represented as a product, $A = QR$, of an orthogonal $n \times m$ matrix Q, $Q^{\mathsf{T}}Q = 1$, and a right-triangular $m \times m$ matrix R.

QR-decomposition of a matrix can be computed using several methods, such as Gram-Schmidt orthogonalization, Householder transformation [17], or Givens rotation [13].

Gram-Schmidt orthogonalization

Gram-Schmidt orthogonalization is an algorithm for orthogonalization of a set of vectors in a given inner product space. It takes a linearly independent set of vectors $A = \{\mathbf{a}_1, \ldots, \mathbf{a}_m\}$ and generates an orthogonal set $Q = \{\mathbf{q}_1, \ldots, \mathbf{q}_m\}$ which spans the same subspace as A. The algorithm is given as

```
for i = 1 to m :
    qᵢ ← aᵢ/‖aᵢ‖
    for j = i + 1 to m :  aⱼ ← aⱼ − ⟨qᵢ|aⱼ⟩ qᵢ
```

where $\langle \mathbf{a}|\mathbf{b} \rangle$ is the inner product of two vectors, and $\|\mathbf{a}\| \doteq \sqrt{\langle \mathbf{a}|\mathbf{a} \rangle}$ is the vector's norm. This variant of the algorithm, where all remaining vectors $\mathbf{a}_j$ are made orthogonal to $\mathbf{q}_i$ as soon as the latter is calculated, is considered to be numerically stable and is referred to as *stabilized* or *modified*.

Stabilized Gram-Schmidt orthogonalization can be used to compute QR-decomposition of a matrix A by orthogonalization of its column-vectors $\mathbf{a}_i$ with the inner product

$$\langle \mathbf{a}|\mathbf{b} \rangle = \mathbf{a}^{\mathsf{T}}\mathbf{b} \equiv \sum_{k=1}^{n}(\mathbf{a})_k(\mathbf{b})_k , \tag{2.8}$$

where n is the length of column-vectors $\mathbf{a}$ and $\mathbf{b}$, and $(\mathbf{a})_k$ is the kth element of the column-vector,

$$\text{for } i = 1 \text{ to } m :$$
$$R_{ii} = \sqrt{\mathbf{a}_i^\mathsf{T}\mathbf{a}_i} \; ; \quad \mathbf{q}_i = \mathbf{a}_i/R_{ii}$$
$$\text{for } j = i+1 \text{ to } m :$$
$$R_{ij} = \mathbf{q}_i^\mathsf{T}\mathbf{a}_j \; ; \quad \mathbf{a}_j = \mathbf{a}_j - \mathbf{q}_i R_{ij} \; .$$

After orthogonalization the matrices $Q = \{\mathbf{q}_1 \ldots \mathbf{q}_m\}$ and R are the sought orthogonal and right-triangular factors of matrix A. A Csharp implementation might look like this,

```csharp
matrix Q=A.copy(), R=new matrix(m,m);
for(int i=0;i<m;i++){
   R[i,i]=Q[i].norm(); /* Q[i] points to the i-th columb */
   Q[i]/=R[i,i];
   for(int j=i+1;j<m;j++){
      R[i,j]=Q[i].dot(Q[j]);
      Q[j]-=Q[i]*R[i,j]; } }
```

The factorization is unique under requirement that the diagonal elements of R are positive. For a $n \times m$ matrix the complexity of the algorithm is $O(m^2 n)$.

Gram-Schmidt decomposition with column pivoting

Pivoted decomposition differs from the ordinary Gram-Schmidt in that at each iteration it takes the largest of the remaining columns and thus introduces the permutation matrix P,

$$AP = QR \; , \tag{2.9}$$

that is (generally) chosen so that the diagonal elements of the R-matrix are decreasing,

$$|R_{11}| \geq |R_{22}| \geq \cdots \geq |R_{mm}| \; . \tag{2.10}$$

Pivoted QR-decomposition can be used when matrix A is rank deficient or its rank is in doubt. With exact arithmetics if $\text{rank}(A) = k$ then the sub-matrix of R with rows and columns from $k+1$ to m would be zero. Numerical determination of rank requires a criterion for deciding when a small diagonal element of R should be treated as zero – a practical choice that depends on both the matrix and the application.

Householder transformation

A square matrix H of the form

$$H = 1 - \frac{2}{\mathbf{u}^\mathsf{T}\mathbf{u}}\mathbf{u}\mathbf{u}^\mathsf{T} \tag{2.11}$$

is called *Householder matrix*, where the vector $\mathbf{u}$ is called a *Householder vector*. Householder matrices are symmetric and orthogonal,

$$H^\mathsf{T} = H \; , \; H^\mathsf{T}H = 1 \; . \tag{2.12}$$

The transformation induced by the Householder matrix on a given vector $\mathbf{a}$,

$$\mathbf{a} \to \mathrm{H}\mathbf{a} \,, \tag{2.13}$$

is called a *Householder transformation* or *Householder reflection*. The transformation changes the sign of the affected vector's component in the $\mathbf{u}$ direction, or, in other words, makes a reflection of the vector about the hyper-plane perpendicular to $\mathbf{u}$, hence the name.

Householder transformation can be used to zero selected components of a given vector $\mathbf{a}$. For example, one can zero all components but the first one, such that

$$\mathrm{H}\mathbf{a} = \gamma \mathbf{e}_1 \,, \tag{2.14}$$

where γ is a number and $\mathbf{e}_1$ is the unit vector in the first direction. The factor γ can be easily calculated,

$$\|\mathbf{a}\|^2 \doteq \mathbf{a}^{\mathsf{T}}\mathbf{a} = \mathbf{a}^{\mathsf{T}}\mathrm{H}^{\mathsf{T}}\mathrm{H}\mathbf{a} = (\gamma\mathbf{e}_1)^{\mathsf{T}}(\gamma\mathbf{e}_1) = \gamma^2 \,, \tag{2.15}$$

$$\Rightarrow \ \gamma = \pm\|a\| \,. \tag{2.16}$$

To find the Householder vector, we notice that

$$\mathbf{a} = \mathrm{H}^{\mathsf{T}}\mathrm{H}\mathbf{a} = \mathrm{H}^{\mathsf{T}}\gamma\mathbf{e}_1 = \gamma\mathbf{e}_1 - \frac{2(\mathbf{u})_1}{\mathbf{u}^{\mathsf{T}}\mathbf{u}}\mathbf{u} \,, \tag{2.17}$$

$$\Rightarrow \ \frac{2(\mathbf{u})_1}{\mathbf{u}^{\mathsf{T}}\mathbf{u}}\mathbf{u} = \gamma\mathbf{e}^1 - \mathbf{a} \,, \tag{2.18}$$

where $(\mathbf{u})_1$ is the first component of the vector $\mathbf{u}$. One usually chooses $(\mathbf{u})_1 = 1$ (for the sake of the possibility to store the other components of the Householder vector in the zeroed elements of the vector $\mathbf{a}$) and stores the factor

$$\frac{2}{\mathbf{u}^{\mathsf{T}}\mathbf{u}} \equiv \tau \tag{2.19}$$

separately. With this convention one readily finds τ from the first component of equation (2.18),

$$\tau = \gamma - (\mathbf{a})_1 \,. \tag{2.20}$$

where $(\mathbf{a})_1$ is the first element of the vector $\mathbf{a}$. For the sake of numerical stability the sign of γ has to be chosen opposite to the sign of $(\mathbf{a})_1$,

$$\gamma = -\mathrm{sign}\left((\mathbf{a})_1\right)\|\mathbf{a}\| \,. \tag{2.21}$$

Finally, the Householder reflection, which zeroes all component of a vector $\mathbf{a}$ but the first, is given as

$$\mathrm{H} = 1 - \tau\mathbf{u}\mathbf{u}^{\mathsf{T}} \,, \ \tau = -\mathrm{sign}((\mathbf{a})_1)\|\mathbf{a}\| - (\mathbf{a})_1 \,, \ (\mathbf{u})_1 = 1 \,, \ (\mathbf{u})_{i>1} = -\frac{1}{\tau}(\mathbf{a})_i \,. \tag{2.22}$$

Now, a QR-decomposition of an $n \times n$ matrix A by Householder transformations can be performed in the following way:

1. Build the size-n Householder vector $\mathbf{u}_1$ which zeroes the sub-diagonal elements of the first column of matrix A, such that

$$
H_1 A =
\begin{bmatrix}
\star & \star & \cdots & \star \\
0 & & & \\
\vdots & & A_1 & \\
0 & & &
\end{bmatrix},
\tag{2.23}
$$

   where $H_1 = 1 - \tau_1 \mathbf{u}_1 \mathbf{u}_1^\mathsf{T}$ and where $\star$ denotes (generally) non-zero matrix elements. In practice one does not build the matrix H_1 explicitly, but rather calculates the matrix $H_1 A$ in-place, consecutively applying the Householder reflection to columns the matrix A, thus avoiding computationally expensive matrix-matrix operations. The zeroed sub-diagonal elements of the first column of the matrix A can be used to store the elements of the Householder vector $\mathbf{u}_1$ while the factor τ_1 has to be stored separately in a special array. This is the storage scheme used by LAPACK and GSL.

2. Similarly, build the size-$(n-1)$ Householder vector $\mathbf{u}_2$ which zeroes the sub-diagonal elements of the first column of matrix A_1 from eq. (2.23). With the transformation matrix H_2 defined as

$$
H_2 =
\begin{bmatrix}
1 & 0 & \cdots & & 0 \\
0 & & & & \\
\vdots & & 1 - \tau_2 \mathbf{u}_2 \mathbf{u}_2^\mathsf{T} & \\
0 & & & &
\end{bmatrix}.
\tag{2.24}
$$

   the two transformations together zero the sub-diagonal elements of the two first columns of matrix A,

$$
H_2 H_1 A =
\begin{bmatrix}
\star & \star & \star & \cdots & \star \\
0 & \star & \star & \cdots & \star \\
0 & 0 & & & \\
\vdots & \vdots & & A_3 & \\
0 & 0 & & &
\end{bmatrix},
\tag{2.25}
$$

3. Repeating the process zero the sub-diagonal elements of the remaining columns.

For column k the corresponding Householder matrix is

$$H_k = \left[\begin{array}{c|c} I_{k-1} & 0 \\ \hline 0 & 1 - \tau_k \mathbf{u}_k \mathbf{u}_k^{\mathsf{T}} \end{array}\right] , \tag{2.26}$$

where I_{k-1} is an identity matrix of size $k-1$, $\mathbf{u}_k$ is the size-(n-k+1) Householder vector that zeroes the sub-diagonal elements of matrix A_{k-1} from the previous step. The corresponding transformation step is

$$H_k \dots H_2 H_1 A = \left[\begin{array}{c|c} R_k & \star \\ \hline 0 & A_k \end{array}\right] , \tag{2.27}$$

where R_k is a size-k right-triangular matrix.

After $n-1$ steps the matrix A will be transformed into a right triangular matrix,

$$H_{n-1} \cdots H_2 H_1 A = R . \tag{2.28}$$

4. Finally, introducing an orthogonal matrix $Q = H_1^{\mathsf{T}} H_2^{\mathsf{T}} \dots H_{n-1}^{\mathsf{T}}$ and multiplying eq. (2.28) by Q from the left, we get the sought QR-decomposition,

$$A = QR . \tag{2.29}$$

In practice one does not build explicitly the Q matrix but rather applies the successive Householder reflections stored during the decomposition.

Givens rotations

A Givens rotation is a transformation in the form

$$A \rightarrow G(p, q, \theta)A , \tag{2.30}$$

where A is the object to be transformed—matrix of vector—and $G(p, q, \theta)$ is the Givens rotation matrix (also known as Jacobi rotation matrix): an orthogonal matrix in the form

$$G(p, q, \theta) = \begin{bmatrix} 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & & \vdots & & \vdots \\ 0 & \cdots & \cos\theta & \cdots & \sin\theta & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots & & \vdots \\ 0 & \cdots & -\sin\theta & \cdots & \cos\theta & \cdots & 0 \\ \vdots & & \vdots & & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & 0 & \cdots & 1 \end{bmatrix} \begin{array}{l} \\ \\ \leftarrow \text{row } p \\ \\ \leftarrow \text{row } q \\ \\ \end{array} . \tag{2.31}$$

When a Givens rotation matrix $G(p, q, \theta)$ multiplies a vector $\mathbf{x}$, only elements x_p and x_q are affected. Considering only these two affected elements, the Givens rotation is given explicitly as

$$\begin{bmatrix} x'_p \\ x'_q \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x_p \\ x_q \end{bmatrix} = \begin{bmatrix} x_p\cos\theta + x_q\sin\theta \\ -x_p\sin\theta + x_q\cos\theta \end{bmatrix}. \tag{2.32}$$

Apparently the rotation can zero the element x'_q, if the angle θ is chosen as

$$\tan\theta = \frac{x_q}{x_p} \quad \Rightarrow \quad \theta = \mathrm{atan2}(x_q, x_p). \tag{2.33}$$

A sequence of Givens rotations,

$$G = \prod_{n \geq q > p = 1}^{m} G(p, q, \theta_{qp}), \tag{2.34}$$

(where $n \times m$ is the dimension of the matrix A) can zero all elements of a matrix below the main diagonal if the angles θ_{qp} are chosen to zero the elements with indices q, p of the partially transformed matrix just before applying the matrix $G(p, q, \theta_{qp})$. The resulting matrix is obviously the R-matrix of the sought QR-decomposition of the matrix A where $G = Q^\mathsf{T}$.

In practice one does not explicitly builds the G matrix but rather stores the θ angles in the places of the corresponding zeroed elements of the original matrix:

```
for(int p=0;p<A.size2;p++){
   for(int q=p+1;q<A.size1;q++){
      double theta=Atan2(A[q,p],A[p,p]);
      double c=Cos(theta),s=Sin(theta);
      for(int k=q;k<A.size2;k++){
         double xp=A[p,k], xq=A[q,k];
         A[p,k]= xp*c+xq*s;
         A[q,k]=-xp*s+xq*c; }
      A[q,p]=theta; } }
```

When solving the linear system $A\mathbf{x} = \mathbf{b}$ one transforms it into the equivalent triangular system $R\mathbf{x} = G\mathbf{b}$ where one calculates $G\mathbf{b}$ by successively applying the individual Givens rotations with the stored θ-angles:

```
for(int p=0;p<G.size2;p++){
   for(int q=p+1;q<G.size1;q++){
      double theta = G[q,p];
      double c=Cos(theta),s=Sin(theta);
      double bp=b[p], bq=b[q];
      b[p]=+bp*c+bq*s;
      b[q]=-bp*s+bq*c; } }
```

The triangular system $R\mathbf{x} = G\mathbf{b}$ is then solved by the ordinary back-substitution.

If one needs to build the Q-matrix explicitly, one uses

$$Q_{ij} = \mathbf{e}_i^\mathsf{T} Q \mathbf{e}_j = \mathbf{e}_j^\mathsf{T} Q^\mathsf{T} \mathbf{e}_i , \qquad (2.35)$$

where $\mathbf{e}_i$ is the unit vector in the direction i and where again one can use the successive rotations to calculate $Q^\mathsf{T} \mathbf{e}_i$,

Since each Givens rotation only affects two rows of the matrix it is possible to apply a set of rotations in parallel. Givens rotations are also more efficient on sparse matrices.

2.3.2 LU-decomposition

LU-decomposition is a factorization of a square matrix A into a product of a lower triangular matrix L and an upper triangular matrix U,

$$A = LU . \qquad (2.36)$$

The linear system $A\mathbf{x} = \mathbf{b}$ after LU-decomposition of the matrix A becomes $LU\mathbf{x} = \mathbf{b}$ and can be solved by first solving $L\mathbf{y} = \mathbf{b}$ for $\mathbf{y}$ and then $U\mathbf{x} = \mathbf{y}$ for $\mathbf{x}$ with two runs of forward and backward substitutions.

If A is an $n \times n$ matrix, the condition (2.36) is a set of n^2 equations,

$$\sum_{k=1}^{n} L_{ik}U_{kj} = A_{ij} \big|_{i,j=1\ldots n} , \qquad (2.37)$$

for $n^2 + n$ unknown elements of the triangular matrices L and U. The decomposition is thus not unique.

Usually the decomposition is made unique by providing extra n conditions e.g. by the requirement that the elements of the main diagonal of the matrix L are equal one,

$$L_{ii} = 1 , \ i = 1\ldots n . \qquad (2.38)$$

The system (2.37) with the extra conditions (2.38) can then be easily solved row after row using the *Doolittle's algorithm*,

> for $i = 1\ldots n$:
> $L_{ii} = 1$
>   for $j = i\ldots n$: $U_{ij} = A_{ij} - \sum_{k<i} L_{ik}U_{kj}$
>   for $j = i+1\ldots n$: $L_{ji} = \frac{1}{U_{ii}} \left(A_{ji} - \sum_{k<j} L_{jk}U_{ki} \right)$

In a slightly different *Crout's algorithm* it is the matrix U that has unit diagonal elements,

```
for  i = 1 ... n  :
    U_ii = 1
    for  j = i ... n  :   L_ji = A_ji − Σ_{k<i} L_jk U_ki
    for  j = i + 1 ... n  :   U_ij = (1/L_ii) ( A_ji − Σ_{k<j} L_jk U_ki )
```

Without a proper ordering (permutations) in the matrix, the factorization may fail. For example, it is easy to verify that $A_{11} = L_{11}U_{11}$. If $A_{11} = 0$, then at least one of L_{11} and U_{11} has to be zero, which implies either L or U is singular, which is impossible if A is non-singular. This is however only a procedural problem. It can be removed by simply reordering the rows of A so that the first element of the permuted matrix is nonzero (or, even better, the largest in absolute value among all elements of the column below the diagonal). The same problem in subsequent factorization steps can be removed in a similar way. Such algorithm is referred to as *partial pivoting*. It requires an extra integer array to keep track of row permutations.

2.3.3 Cholesky decomposition

The Cholesky decomposition of a Hermitian positive-definite matrix A is a decomposition in the form

$$A = \mathrm{LL}^{\dagger} , \tag{2.39}$$

where L is a lower triangular matrix with real and positive diagonal elements, and $\mathrm{L}^{\dagger}$ is the conjugate transpose of L.

For real symmetric positive-definite matrices the decomposition reads

$$A = \mathrm{LL}^{\mathsf{T}} , \tag{2.40}$$

where L is real.

The decomposition can be calculated using the following in-place algorithm,

$$L_{jj} = \sqrt{A_{jj} - \sum_{k=1}^{j-1} L_{jk}^2} , \; L_{ij} = \frac{1}{L_{jj}} \left(A_{ij} - \sum_{k=1}^{j-1} L_{ik}L_{jk} \right)\Bigg|_{i>j} . \tag{2.41}$$

The expression under the square root is always positive if A is real and positive-definite.

When applicable, the Cholesky decomposition is about twice as efficient as LU-decomposition for solving systems of linear equations.

2.4 Determinant of a matrix

LU- and QR-decompositions allow $O(n^3)$ calculation of the determinant of a square matrix. Indeed, for the LU-decomposition,

$$\det A = \det LU = \det L \det U = \det U = \prod_{i=1}^{n} U_{ii} \; . \tag{2.42}$$

For the Gram-Schmidt QR-decomposition

$$\det A = \det QR = \det Q \det R \; . \tag{2.43}$$

Since Q is an orthogonal matrix $(\det Q)^2 = 1$,

$$|\det A| = |\det R| = \left| \prod_{i=1}^{n} R_{ii} \right| \; . \tag{2.44}$$

With Gram-Schmidt method one arbitrarily assigns positive sign to diagonal elements of the R-matrix thus removing from the R-matrix the memory of the original sign of the determinant.

However with Givens rotation method the determinant of the individual rotation matrix—and thus the determinant of the total rotation matrix—is equal one, therefore for a square matrix A the QR-decomposition A = GR via Givens rotations allows calculation of the determinant with the correct sign,

$$\det A = \det R \equiv \prod_{i=1}^{n} R_{ii} \tag{2.45}$$

2.5 Matrix inverse

The inverse A^{-1} of a square $n \times n$ matrix A can be calculated by solving n linear equations

$$A\mathbf{x}_i = \mathbf{e}_i \Big|_{i=1,\dots,n} \; , \tag{2.46}$$

where $\mathbf{e}_i$ is the unit-vector in the i-direction: a column where all elements are equal zero except for the element number i which is equal one. Thus the set of columns $\{\mathbf{e}_i\}_{i=1,\dots,n}$ form the identity matrix. The matrix made of columns $\mathbf{x}_i$ is apparently the inverse of A.

Chapter 3

Ordinary least-squares problem

3.1 Introduction

An *ordinary least-squares problem* (also called *linear least-squares problem*) is the problem of minimization of the (squared) norm of a vector (hence the name, least squares) where the minimization parameters enter the vector linearly. Reformulation of a given problem in terms of an ordinary least-squares problem is often called the least-squares method. Ordinary least-squares problems can be usually solved using linear algebra methods – a great asset that ensures their extensive applications in science end engineering.

In signal processing least-squares methods are used for smoothing, prediction, deconvolution, error recovery, and de-clipping. In physics it is used, in particular, in fitting a theoretical model to experimental data with uncertainties, and in deconvolution of the detected signal from the detector response.

3.2 Overdetermined linear systems

A system of linear equations is considered *overdetermined* if there are more equations than unknown variables. If all equations of an overdetermined system are linearly independent, the system has no exact solution. However, it is usually possible to find an approximate solution to an overdetermined system using the least-squares method.

Consider a linear system

$$\mathbf{A}\mathbf{c} = \mathbf{b} \,, \tag{3.1}$$

where A is a tall $n \times m$ matrix, $\mathbf{c}$ is an m-component vector of unknown variables and $\mathbf{b}$ is an n-component vector of the right-hand side terms. If the number of equations n is larger than the number of unknowns m, the system is overdetermined and generally has no solution. However, it is still possible to find an approximate solution — the one where $A\mathbf{c}$ is only approximately equal $\mathbf{b}$ — in the sense that the Euclidean norm of the difference between $A\mathbf{c}$ and $\mathbf{b}$ is minimized,

$$\mathbf{c} : \ \min_{\mathbf{c}} \|A\mathbf{c} - \mathbf{b}\|^2 \ . \tag{3.2}$$

The problem (3.2) is an example of an ordinary least-squares problem. The vector $\mathbf{c}$ that minimizes $\|A\mathbf{c} - \mathbf{b}\|^2$ is usually called the *least-squares solution.*

Theoretically, the solution to this minimization problem is given by the equation

$$\frac{\partial}{\partial \mathbf{c}^{\mathsf{T}}} \left((\mathbf{c}^{\mathsf{T}} A^{\mathsf{T}} - \mathbf{b}^{\mathsf{T}})(A\mathbf{c} - \mathbf{b}) \right) = 2 \left((A^{\mathsf{T}} A)\mathbf{c} - A^{\mathsf{T}} \mathbf{b} \right) = 0 \ , \tag{3.3}$$

with the solution

$$\mathbf{c} = (A^{\mathsf{T}} A)^{-1} A^{\mathsf{T}} \mathbf{b} \ , \tag{3.4}$$

where $(A^{\mathsf{T}} A)^{-1} A^{\mathsf{T}}$ is the left pseudo-inverse of the tall matrix A,

$$(A^{\mathsf{T}} A)^{-1} A^{\mathsf{T}} A = \mathbf{1} \ . \tag{3.5}$$

In practice however instead of calculating the pseudo-inverse one should use QR or SVD decompositions as described below.

3.2.1 Least-squares solution via QR-decomposition

The linear least-squares problem can be solved by QR-decomposition. The matrix A is factorized as $A = QR$, where Q is $n \times m$ matrix with orthonormal columns, $Q^{\mathsf{T}} Q = 1$, and R is an $m \times m$ upper triangular matrix. The matrix Q is a semi-orthonormal matrix whose columns span the range (column space) of matrix A.

The matrix QQ^{T} is the projector on the range of matrix A and $(1 - QQ^{\mathsf{T}})$ is the corresponding orthogonal projector. Indeed,

$$(1 - QQ^{\mathsf{T}})QQ^{\mathsf{T}} = QQ^{\mathsf{T}}(1 - QQ^{\mathsf{T}}) = 0 \ . \tag{3.6}$$

The vector $A\mathbf{c} - \mathbf{b}$ can be represented as a sum of two orthogonal components, the one within the range of the matrix and the orthogonal one,

$$A\mathbf{c} - \mathbf{b} = QQ^{\mathsf{T}}(A\mathbf{c} - \mathbf{b}) + (1 - QQ^{\mathsf{T}})(A\mathbf{c} - \mathbf{b}) \ . \tag{3.7}$$

The Euclidean norm $\|A\mathbf{c} - \mathbf{b}\|^2$ is then given as the sum of the norms of the two orthogonal components,

$$\|A\mathbf{c} - \mathbf{b}\|^2 = \|QQ^{\mathsf{T}}(A\mathbf{c} - \mathbf{b})\|^2 + \|(1 - QQ^{\mathsf{T}})(A\mathbf{c} - \mathbf{b})\|^2 \ . \tag{3.8}$$

The second term,

$$\|(1 - QQ^\mathsf{T})(A\mathbf{c} - \mathbf{b})\|^2 = \|(1 - QQ^\mathsf{T})\mathbf{b}\|^2 \,, \tag{3.9}$$

is the norm of the range-orthogonal component of the right-hand-side $\mathbf{b}$: it is independent of the variables $\mathbf{c}$ and can not be reduced by their variations. However, the first term

$$\|QQ^\mathsf{T}(A\mathbf{c} - \mathbf{b})\|^2 = \|(R\mathbf{c} - Q^\mathsf{T}\mathbf{b})\|^2 \,, \tag{3.10}$$

can be reduced down to zero by solving the $m \times m$ system of linear equations

$$R\mathbf{c} = Q^\mathsf{T}\mathbf{b} \,. \tag{3.11}$$

The system is right-triangular and can be readily solved by back-substitution.

Thus the solution to the ordinary least-squares problem (3.2) is given by the solution of the triangular system (3.11).

3.2.2   Least-squares solution via SVD

Under the *thin singular value decomposition* we shall understand a representation of a tall $n \times m$ $(n > m)$ matrix A in the form

$$A = USV^\mathsf{T} \,, \tag{3.12}$$

where U is a semi-orthonormal $(U^\mathsf{T}U = 1)$ $n \times m$ matrix that spans the range of matrix A, S is a square $m \times m$ diagonal matrix with non-negative real numbers (called singular values) on the diagonal, and V is a square $m \times m$ orthonormal matrix $(V^\mathsf{T}V = VV^\mathsf{T} = 1)$.

Singular value decomposition can too be used to solve the least squares problem $A\mathbf{c} \approx \mathbf{b}$ the same way as with QR-decomposition only using the U matrix for the projection of the equation on the column-space of the matrix. Specifically,

$$U^\mathsf{T}(A\mathbf{c} - \mathbf{b}) = 0 \;\Rightarrow\; SV^\mathsf{T}\mathbf{c} = U^\mathsf{T}\mathbf{b} \;\Rightarrow\; \mathbf{c} = VS^{-1}U^\mathsf{T}\mathbf{b} \,, \tag{3.13}$$

where one should not build the S^{-1} matrix but should rather divide by the corresponding singular values the components of the vector it acts upon.

Notice that $VS^{-1}U^\mathsf{T}$ is the pseudo-inverse A^- of the matrix A. Therefore the solution to the least squares problem can be (theoretically) written as

$$\mathbf{c} = A^-\mathbf{b} \,. \tag{3.14}$$

3.3   Signal smoothing

Least squares method can be use to smooth a noisy signal – a sequence of numbers represented by a vector, $\mathbf{y}$. The idea is to find a new signal, $\mathbf{x}$, that is similar to the noisy signal $\mathbf{y}$ but smoother. The smoothness of the signal can be measured by the smallness of its second derivative (indeed a signal with vanishing second derivative is a straight line – the smoothest possible signal). The smoothers based on minimization of the second derivative are the most common in the literature although other derivatives can also be used.

The second derivative of a discrete signal $\mathbf{x}$ can be approximated by its second-order difference, $D\mathbf{x}$, where the matrix D is given as (pay attention to the first and the last rows)

$$ D = \begin{pmatrix} 1 & -2 & 1 & & & & \\ 1 & -2 & 1 & & & & \\ & 1 & -2 & 1 & & & \\ & & & \vdots & & & \\ & & & & 1 & -2 & 1 \\ & & & & 1 & -2 & 1 \end{pmatrix} . \tag{3.15} $$

The smooth signal $\mathbf{x}$ can then be obtained as a solution to the following least-squares problem,

$$ \mathbf{x} : \min_{\mathbf{x}} \left(\|\mathbf{x} - \mathbf{y}\|^2 + \lambda \|D\mathbf{x}\|^2 \right) , \tag{3.16} $$

where $\lambda > 0$ is the smoothing parameter: when $\lambda \to 0$ there is no smoothing and $\mathbf{x}$ converges to $\mathbf{y}$; when $\lambda \to \infty$ the second derivative is infinitely penalized and $\mathbf{x}$ converges to a linear fit to the data.

The minimum of the form (3.16) is located where its partial derivatives vanish,

$$ \frac{\partial}{\partial \mathbf{x}^\mathsf{T}} \left((\mathbf{x} - \mathbf{y})^\mathsf{T}(\mathbf{x} - \mathbf{y}) + \lambda \mathbf{x}^\mathsf{T} D^\mathsf{T} D \mathbf{x} \right) = 2\left((I + \lambda D^\mathsf{T} D)\mathbf{x} - \mathbf{y} \right) = 0 , \tag{3.17} $$

(where I is the identity matrix) which gives the (least-squares) solution

$$ \mathbf{x} = (I + \lambda D^\mathsf{T} D)^{-1} \mathbf{y} . \tag{3.18} $$

Of course in practice one should not explicitly calculate the inverse matrix but rather solve the linear equation

$$ (I + \lambda D^\mathsf{T} D)\mathbf{x} = \mathbf{y} . \tag{3.19} $$

The matrix in this linear equation is *banded* (has only few non-zero diagonals) therefore in order to make the method efficient one has to take advantage of this fact (rather than use generic linear solvers).

And example of smoothing using this algorithm is given at Figure (3.1).

Figure 3.1: Noisy signal smoothing using the least-squares algorithm (3.16).
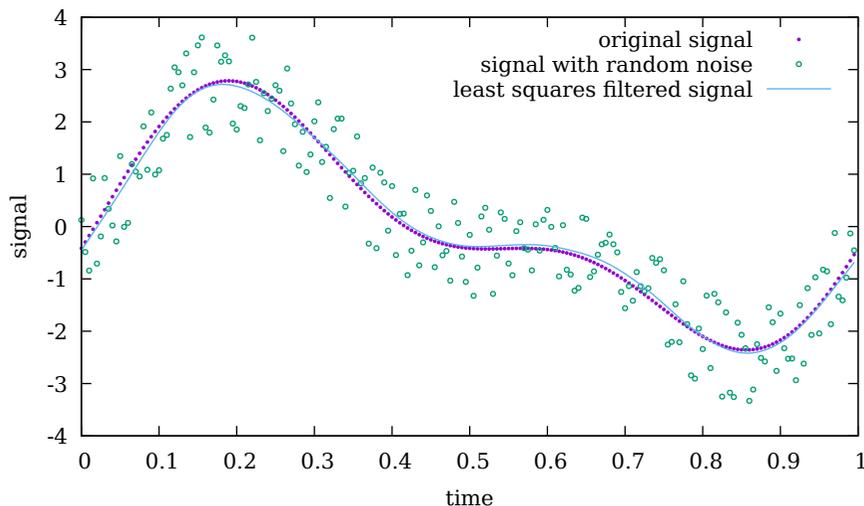
3.4 Signal extrapolation (prediction)

Least-squares method can also be used to find patterns (correlations) in a given signal (sequence of numbers) and then use these patterns for signal extrapolation.

Suppose the signal $\{x_i\}_{i=1\ldots N}$ is correlated such that the term x_k can be predicted by the n preceding terms,

$$x_k = x_{k-1}a_n + x_{k-2}a_{n-1} + \cdots + x_{k-n}a_1 . \tag{3.20}$$

Applying this anzats (often called *linear prediction*) to the subset $\{x_{n+1}, x_{n+2}, \ldots, x_N\}$ of the signal gives the system of equations

$$
\begin{pmatrix}
x_1 & x_2 & \ldots & x_n \\
x_2 & x_3 & \ldots & x_{n+1} \\
\vdots & \vdots & \vdots & \vdots \\
\vdots & \vdots & \vdots & \vdots \\
x_{N-n} & x_{N-n+1} & \ldots & x_{N-1}
\end{pmatrix}
\begin{pmatrix}
a_1 \\
a_2 \\
\vdots \\
a_n
\end{pmatrix}
=
\begin{pmatrix}
x_{n+1} \\
x_{n+2} \\
\vdots \\
\vdots \\
x_N
\end{pmatrix} . \tag{3.21}
$$

This is an overdetermined system the least-squares solution of which determines the correlation parameters $\{a_1, \ldots, a_n\}$ which allow extrapolation of the sequence beyond

Figure 3.2: Signal extrapolation (prediction) using the linear prediction anzats (3.20) with $n = 6$.
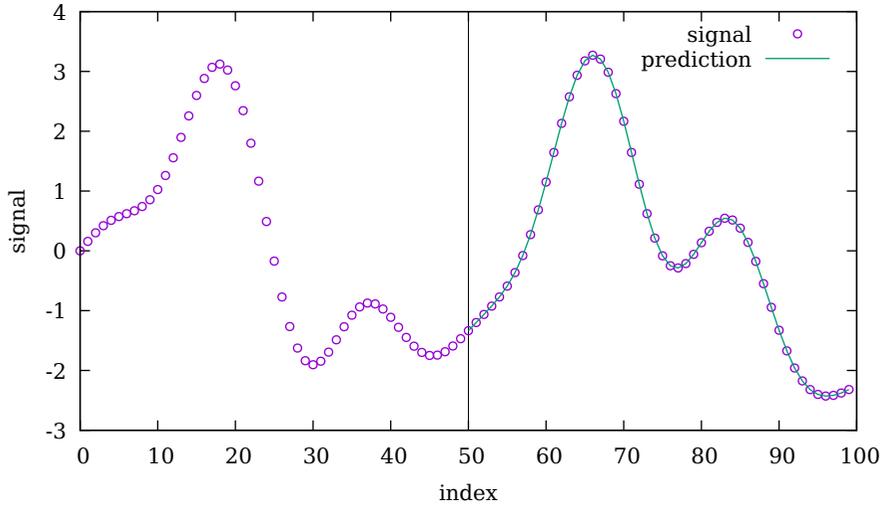
the last term, x_N,

$$
\begin{aligned}
x_{N+1} &= a_1 x_{N-n+1} + a_2 x_{N-n+2} + \cdots + a_n x_N & (3.22) \\
x_{N+2} &= a_1 x_{N-n+2} + a_2 x_{N-n+3} + \cdots + a_n x_{N+1} & (3.23) \\
&\vdots \quad . & (3.24)
\end{aligned}
$$

An illustration of this method is given at figure (3.2) where the six correlation parameters where determined from the first 50 values of the signal and then extrapolated to the next 50 values (and compared with the exact values). The signal is

$$
x_k = 2 \sin\left(0.9 \frac{2\pi k}{N-1} \right) - \sin\left(2.1 \frac{2\pi k}{N-1} \right) + \frac{1}{2} \sin\left(3.1 \frac{2\pi k}{N-1} \right) , \qquad (3.25)
$$

where $N = 50$.

3.5   Missing data recovery (error concealment)

Sometimes one can have a signal at hand where several entries are lost due to some sort of data corruption, for example transmission errors (packet loss). Among the

Figure 3.3: Least-squares missing data recovery
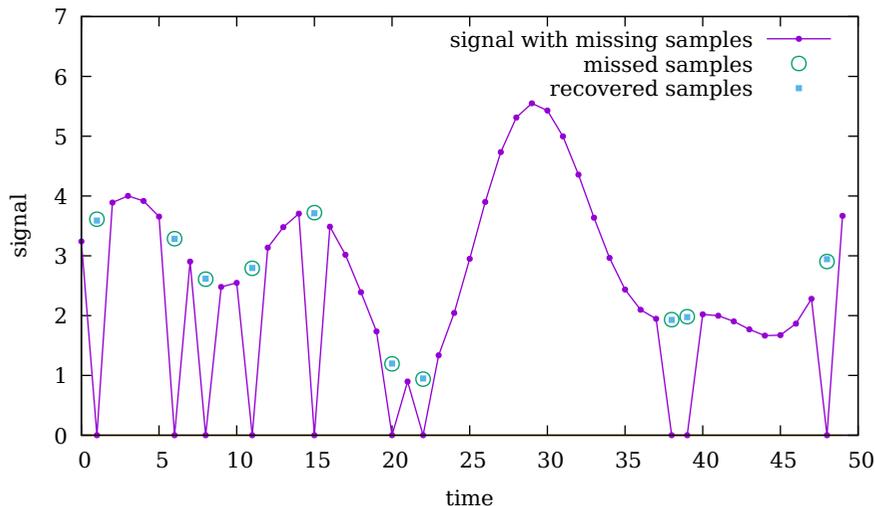
techniques to recover the missing data in a corrupted signal (often referred to as *error concealment* techniques) there is one based on the least-squares method.

Let us assume that in a received signal, $\mathbf{y}$, n entries (out of the total N) at positions $m_1, m_2, \ldots, m_n$ have been lost and have been replaced with zeros. If one knew the missing entries, $\{z_1 \ldots z_n\} \doteq \mathbf{z}$, the recovered signal $\mathbf{x}$ would be given as

$$\mathbf{x} = \mathbf{y} + \mathrm{M}\mathbf{z} , \tag{3.26}$$

where M is the matrix that inserts the elements of $\mathbf{z}$ into the missing positions of $\mathbf{y}$. It is an $N \times n$ matrix with all elements equal zero except for the elements

$$\mathrm{M}_{m_k,k}\big|_{k=1,\ldots,n} = 1 . \tag{3.27}$$

Like in signal smoothing the unknown entries $\mathbf{z}$ can be estimated from the condition that the recovered signal is smooth,

$$\mathbf{z} : \min_{\mathbf{z}} \|\mathrm{D}\mathbf{x}\|^2 = \min_{\mathbf{z}} \|\mathrm{D}(\mathbf{y} + \mathrm{M}\mathbf{z})\|^2 . \tag{3.28}$$

The solution is given by the least-squares solution to the overdetermined system

$$\mathrm{DM}\mathbf{z} = -\mathrm{D}\mathbf{y} . \tag{3.29}$$

Figure 3.4: Least-squares data declipping

3.6   Signal declipping

A problem very similar to missing data recovery is *signal declipping*. Here the corrupted signal misses those samples, the absolute value of which are larger than certain cutoff (for example, the detector threshold). The recovery is similar to missing data, however in this particular case one minimizes the third derivative, rather than second, forcing the declipped signal to be in the form of a parabola. The corresponding matrix is given by the third order finite difference formulae,

$$
D = \begin{pmatrix}
-1 & 3 & -3 & 1 & & & \\
 & -1 & 3 & -3 & 1 & & \\
-1/2 & 1 & 0 & -1 & 1/2 & & \\
 & -1/2 & 1 & 0 & -1 & 1/2 & \\
 & & & \vdots & & & \\
 & & -1/2 & 1 & 0 & -1 & 1/2 \\
 & & -1 & 3 & -3 & 1 & \\
 & & & -1 & 3 & -3 & 1
\end{pmatrix} .
\tag{3.30}
$$

Figure (3.4) shows an example of least-squares declipping.

3.7 Ordinary least-squares curve fitting

Ordinary least-squares curve fitting is a problem of fitting n (experimental) data points $\{x_i, y_i \pm \Delta y_i\}_{i=1,\ldots,n}$, where Δy_i are experimental errors, by a linear combination, $F_{\mathbf{c}}$, of m functions $\{f_k(x)\}_{k=1,\ldots,m}$,

$$F_{\mathbf{c}}(x) = \sum_{k=1}^{m} c_k f_k(x) , \tag{3.31}$$

where the coefficients c_k are the fitting parameters.

The objective of the least-squares fit is to minimize the square deviation, called χ^2, between the fitting function $F_{\mathbf{c}}(x)$ and the experimental data [4],

$$\chi^2 = \sum_{i=1}^{n} \left(\frac{F(x_i) - y_i}{\Delta y_i} \right)^2 . \tag{3.32}$$

where the individual deviations from experimental points are weighted with their inverse errors in order to promote contributions from the more precise measurements.

Minimization of χ^2 with respect to the coefficient c_k in (3.31) is apparently equivalent to the least-squares problem (3.2) where

$$A_{ik} = \frac{f_k(x_i)}{\Delta y_i} , \ b_i = \frac{y_i}{\Delta y_i} . \tag{3.33}$$

3.7.1 Variances and correlations of fitting parameters

Suppose δy_i is a small deviation of the measured value of the physical observable at hand from its exact value. The corresponding deviation δc_k of the fitting coefficient is then given as

$$\delta c_k = \sum_i \frac{\partial c_k}{\partial y_i} \delta y_i . \tag{3.34}$$

In a good experiment the deviations δy_i are statistically independent and distributed normally with the standard deviations Δy_i. The deviations (3.34) are then also distributed normally with *variances*

$$\langle \delta c_k \delta c_k \rangle = \sum_i \left(\frac{\partial c_k}{\partial y_i} \Delta y_i \right)^2 = \sum_i \left(\frac{\partial c_k}{\partial b_i} \right)^2 . \tag{3.35}$$

The standard errors in the fitting coefficients are then given as the square roots of variances,

$$\Delta c_k = \sqrt{\langle \delta c_k \delta c_k \rangle} = \sqrt{\sum_i \left(\frac{\partial c_k}{\partial b_i} \right)^2} . \tag{3.36}$$

The variances are diagonal elements of the *covariance matrix*, Σ, made of *covariances*,

$$\Sigma_{kq} \equiv \langle \delta c_k \delta c_q \rangle = \sum_i \frac{\partial c_k}{\partial b_i} \frac{\partial c_q}{\partial b_i} \ . \tag{3.37}$$

Covariances $\langle \delta c_k \delta c_q \rangle$ are measures of to what extent the coefficients c_k and c_q change together if the measured values y_i are varied. The normalized covariances,

$$\frac{\langle \delta c_k \delta c_q \rangle}{\sqrt{\langle \delta c_k \delta c_k \rangle \langle \delta c_q \delta c_q \rangle}} \tag{3.38}$$

are called *correlations*.

Using $\mathbf{c} = A^- \mathbf{b}$ the covariance matrix can be calculated as

$$\Sigma = \left(\frac{\partial \mathbf{c}}{\partial \mathbf{b}} \right) \left(\frac{\partial \mathbf{c}}{\partial \mathbf{b}} \right)^\mathsf{T} = A^- A^{-\mathsf{T}} = (A^\mathsf{T} A)^{-1} \ . \tag{3.39}$$

The square roots of the diagonal elements of this matrix provide the estimates of the errors $\Delta \mathbf{c}$ of the fitting coefficients,

$$\Delta c_k = \left. \sqrt{\Sigma_{kk}} \right|_{k=1\ldots m} , \tag{3.40}$$

and the (normalized) off-diagonal elements provide the estimates of their correlations.

With SVD the covariance matrix (3.39) can be calculated as

$$\Sigma = (A^\mathsf{T} A)^{-1} = (VS^2 V^\mathsf{T})^{-1} = VS^{-2} V^T \ . \tag{3.41}$$

With QR-decomposition the covariance matrix (3.39) can be calculated as

$$\Sigma = (A^\mathsf{T} A)^{-1} = (R^\mathsf{T} R)^{-1} = R^{-1} (R^{-1})^\mathsf{T} \ . \tag{3.42}$$

Table 3.7.1 shows how a Csharp implementation of the ordinary least squares fit via QR decomposition could look like. An illustration of a fit is shown on Figure 3.5 where a polynomial is fitted to a set of data.

```
static (vector,matrix) lsfit
(Func<double,double>[] fs, vector x, vector y, vector dy){
    int n = x.size , m=fs.Length;
    var A = new matrix(n,m);
    var b = new vector(n);
    for(int i=0;i<n;i++){
        b[i]=y[i]/dy[i];
        for(int k=0;k<m;k++)A[i,k]=fs[k](x[i])/dy[i];
        }
    vector c = A.solve(b); // solves ||A*c-b||->min
    matrix AI = A.inverse(); // calculates pseudoinverse
    matrix Σ = AI*AI.T;
    return (c, Σ);
    }
```

Table 3.1: A Csharp implemetation of the ordinary least-squares fit.
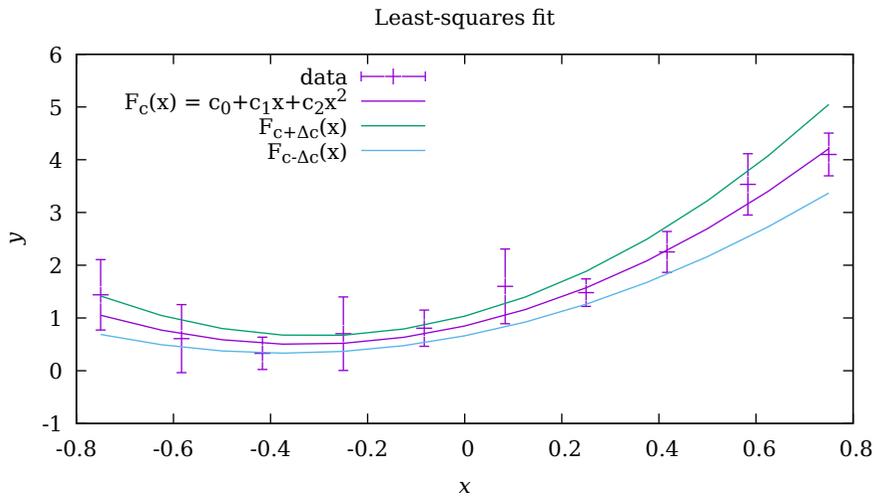
Figure 3.5: Ordinary least squares fit of $F_{\mathbf{c}}(x) = c_1 + c_2 x + c_3 x^2$ to a set of data. Shown are fits with optimal coefficiens $\mathbf{c}$ as well as with $\mathbf{c} + \Delta\mathbf{c}$ and $\mathbf{c} - \Delta\mathbf{c}$.

Chapter 4

Eigenvalues and eigenvectors

4.1 Introduction

A non-zero column-vector $\mathbf{v}$ is called the *eigenvector* of a matrix A with the *eigenvalue* λ, if

$$A\mathbf{v} = \lambda\mathbf{v} \, . \tag{4.1}$$

If an $n \times n$ matrix A is real and symmetric, $A^\mathsf{T} = A$, then it has n real eigenvalues $\lambda_1, \ldots, \lambda_n$, and its (orthogonalized) eigenvectors $V = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ form a full basis,

$$VV^\mathsf{T} = V^\mathsf{T}V = \mathbf{1} \, , \tag{4.2}$$

in which the matrix is diagonal,

$$V^\mathsf{T}AV = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & & \vdots \\ \vdots & & \ddots & \\ 0 & \cdots & & \lambda_n \end{bmatrix} \equiv D \, . \tag{4.3}$$

Matrix diagonalization means finding all eigenvalues and (optionally) eigenvectors of a matrix. Once all eigenvalues and eigenvectors are found, the *Eigenvalue Decomposition* (EVD) of the matrix is given as

$$A = VDV^\mathsf{T} \, . \tag{4.4}$$

Eigenvalues and eigenvectors enjoy a multitude of applications in different branches of science and technology.

4.2   Similarity transformations

Orthogonal transformations,

$$A \to Q^{\mathsf{T}}AQ \ , \tag{4.5}$$

where $Q^T Q = \mathbf{1}$, and, generally, similarity transformations,

$$A \to S^{-1}AS \ , \tag{4.6}$$

preserve eigenvalues and eigenvectors. Therefore one of the strategies to diagonalize a matrix is to apply a sequence of similarity transformations (also called rotations) which (iteratively) turn the matrix into diagonal form.

4.2.1   Jacobi eigenvalue algorithm

Jacobi eigenvalue algorithm is an iterative method to calculate the eigenvalues and eigenvectors of a real symmetric matrix by a sequence of Jacobi rotations.

Jacobi rotation is an orthogonal transformation which zeroes a pair of the off-diagonal elements of a (real symmetric) matrix A,

$$A \to A' = J(p,q)^{\mathsf{T}}AJ(p,q) \ : \ A'_{pq} = A'_{qp} = 0 \ . \tag{4.7}$$

The orthogonal matrix $J(p,q)$ which eliminates the element A_{pq} is called the Jacobi rotation matrix. It is equal identity matrix except for the four elements with indices pp, pq, qp, and qq,

$$J(p,q) = \begin{bmatrix} 1 & & & & & & \\ & \ddots & & & & 0 & \\ & & \cos\theta & \cdots & \sin\theta & & \\ & & \vdots & \ddots & \vdots & & \\ & & -\sin\theta & \cdots & \cos\theta & & \\ & 0 & & & & \ddots & \\ & & & & & & 1 \end{bmatrix} \begin{matrix} \\ \\ \leftarrow \text{row } p \\ \\ \leftarrow \text{row } q \\ \\ \\ \end{matrix} \ . \tag{4.8}$$

Or explicitly,

$$\begin{aligned} J(p,q)_{ij} &= \delta_{ij} \ \forall \ ij \notin \{pq,qp,pp,qq\} \ ; \\ J(p,q)_{pp} &= \cos\theta = J(p,q)_{qq} \ ; \\ J(p,q)_{pq} &= \sin\theta = -J(p,q)_{qp} \ . \end{aligned} \tag{4.9}$$

After a Jacobi rotation, $A \to A' = J^\mathsf{T} A J$, the matrix elements of A' become

$$
\begin{aligned}
A'_{ij} &= A_{ij} \ \forall \ i \neq p, q \wedge j \neq p, q \\
A'_{pi} &= A'_{ip} = cA_{pi} - sA_{qi} \ \forall \ i \neq p, q \ ; \\
A'_{qi} &= A'_{iq} = sA_{pi} + cA_{qi} \ \forall \ i \neq p, q \ ; \\
A'_{pp} &= c^2 A_{pp} - 2scA_{pq} + s^2 A_{qq} \ ; \\
A'_{qq} &= s^2 A_{pp} + 2scA_{pq} + c^2 A_{qq} \ ; \\
A'_{pq} &= A'_{qp} = sc(A_{pp} - A_{qq}) + (c^2 - s^2)A_{pq} \ ,
\end{aligned}
\tag{4.10}
$$

where $c \equiv \cos\theta$, $s \equiv \sin\theta$. The angle θ is chosen such that after rotation the matrix element A'_{pq} is zeroed,

$$
\tan(2\theta) = \frac{2A_{pq}}{A_{qq} - A_{pp}} \ \Rightarrow \ A'_{pq} = 0 \ , \theta = \frac{1}{2}\mathrm{atan2}(2A_{pq}, A_{qq} - A_{pp}) \ ,
\tag{4.11}
$$

where the atan2 correctly deals with the cases where one or two arguments are equal zero.

A side effect of zeroing a given off-diagonal element A_{pq} by a Jacobi rotation is that other off-diagonal elements are changed. Namely, the elements of the rows and columns with indices p and q. However, after the Jacobi rotation the sum of squares of all off-diagonal elements is reduced. The algorithm repeatedly performs rotations until the off-diagonal elements become sufficiently small.

The convergence of the Jacobi method can be proved for two strategies for choosing the order in which the elements are zeroed:

1. *Classical method*: with each rotation the largest of the remaining off-diagonal elements is zeroed.

2. *Cyclic method*: the off-diagonal elements are zeroed in strict order, e.g. row after row.

Although the classical method allows the least number of rotations, it is typically slower than the cyclic method since searching for the largest element is an $O(n^2)$ operation. The count can be reduced by keeping an additional array with indexes of the largest elements in each row. Updating this array after each rotation is only an $O(n)$ operation.

A *sweep* is a sequence of Jacobi rotations applied to all non-diagonal elements. Typically the method converges after a small number of sweeps. The operation count is $O(n)$ for a Jacobi rotation and $O(n^3)$ for a sweep.

The typical convergence criterion is that the diagonal elements have not changed after a sweep. Other criteria can also be used, like the sum of absolute values of the off-diagonal elements is small, $\sum_{i<j} |A_{ij}| < \epsilon$, where ϵ is the required accuracy, or the largest off-diagonal element is small, $\max |A_{i<j}| < \epsilon$.

Eigenvectors

Once the matrix A is diagonalized,

$$Q^\mathsf{T} AQ = D , \tag{4.12}$$

the matrix Q becomes the matrix of eigenvectors. The latter can therefore be calculated as $V = \mathbf{1}J_0 J_2...$, where J_i are the successive Jacobi matrices. At each iteration the update of the V-matrix is given as

$$
\begin{aligned}
V_{ij} &\to V_{ij} , \ j \neq p,q \\
V_{ip} &\to cV_{ip} - sV_{iq} \\
V_{iq} &\to sV_{ip} + cV_{iq}
\end{aligned} \tag{4.13}
$$

Alternatively, if only one (or few) eigenvector $\mathbf{v}_k$ is needed, one can instead solve the (singular) system $(A - \lambda_k)\mathbf{v} = 0$.

Ordering of eigenvalues

Suppose the matrix element A_{pq} to be zeroed by a Jacobi rotation is already zero. Then the function

$$\theta = \frac{1}{2}\mathrm{atan2}(2A_{pq}, A_{qq} - A_{pp}) \tag{4.14}$$

will return 0 if $A_{qq} > A_{pp}$ or $\pi/2$ if $A_{qq} < A_{pp}$. In the first case no rotation will take place, while in the second case the rotation with $\theta = \pi/2$ will be applied. The latter will exchange the matrix elements A_{qq} and A_{pp}. That is, the diagonal elements will be arranged in accending order.

4.2.2   QR/QL algorithm

An orthogonal transformation of a real symmetric matrix, $A \to Q^T AQ = RQ$, where Q is from the QR-decomposition of A, partly turns the matrix A into diagonal form. Successive iterations eventually make it diagonal. If there are degenerate eigenvalues there will be a corresponding block-diagonal sub-matrix.

For convergence properties it is of advantage to use *shifts*: instead of QR[A] we do QR[A − s**1**] and then A → RQ + s**1**. The shift s can be chosen as A$_{nn}$. As soon as an eigenvalue is found the matrix is deflated, that is, the corresponding row and column are crossed out.

Accumulating the successive transformation matrices Q$_i$ into the total matrix Q = Q$_1$ ... Q$_N$, such that $Q^T AQ = \Lambda$, gives the eigenvectors as columns of the Q matrix.

If only one (or few) eigenvector $\mathbf{v}_k$ is needed one can instead solve the (singular) system $(A - \lambda_k)\mathbf{v} = 0$.

Tridiagonalization.

Each iteration of the QR/QL algorithm is an $O(n^3)$ operation. On a tridiagonal matrix it is only $O(n)$. Therefore the effective strategy is first to make the matrix tridiagonal and then apply the QR/QL algorithm. Tridiagonalization of a matrix is a non-iterative operation with a fixed number of steps.

4.3 Eigenvalues of updated matrix

In practice it happens quite often that the matrix A to be diagonalized is given in the form of a diagonal matrix, D, plus an update matrix, W,

$$A = D + W, \tag{4.15}$$

where the update W is a simpler, in a certain sense, matrix which allows a more efficient calculation of the updated eigenvalues, as compared to general diagonalization algorithms.

The most common updates are

- symmetric rank-1 update,

$$W = \mathbf{u}\mathbf{u}^\mathsf{T}, \tag{4.16}$$

  where $\mathbf{u}$ is a columnt-vector;

- symmetric rank-2 update,

$$W = \mathbf{u}\mathbf{v}^\mathsf{T} + \mathbf{v}\mathbf{u}^\mathsf{T}; \tag{4.17}$$

- symmetric row/column update – a special case of rank-2 update,

$$W = \begin{bmatrix} 0 & \dots & u_1 & \dots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ u_1 & \dots & u_p & \dots & u_n \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \dots & u_n & \dots & 0 \end{bmatrix} \equiv \mathbf{e}(p)\mathbf{u}^\mathsf{T} + \mathbf{u}\mathbf{e}(p)^\mathsf{T}, \tag{4.18}$$

  where $\mathbf{e}(p)$ is the unit vector in the p-direction.

4.3.1 Rank-1 update

We assume that a size-n real symmetric matrix A to be diagonalized is given in the form of a diagonal matrix plus a rank-1 update,

$$A = D + \sigma\mathbf{u}\mathbf{u}^\mathsf{T}, \tag{4.19}$$

where D is a diagonal matrix with diagonal elements $\{d_1, \ldots, d_n\}$ and $\mathbf{u}$ is a given vector. The diagonalization of such matrix can be done in $O(m^2)$ operations, where $m \leq n$ is the number of non-zero elements in the update vector $\mathbf{u}$, as compared to $O(n^3)$ operations for a general diagonalization [14].

The eigenvalue equation for the updated matrix reads

$$\left(D + \sigma \mathbf{u}\mathbf{u}^\mathsf{T}\right) \mathbf{q} = \lambda \mathbf{q} \,, \tag{4.20}$$

where λ is an eigenvalue and $\mathbf{q}$ is the corresponding eigenvector. The equation can be rewritten as

$$(D - \lambda 1)\, \mathbf{q} + \sigma \mathbf{u}\mathbf{u}^\mathsf{T}\mathbf{q} = 0 \,. \tag{4.21}$$

Multiplying from the left with $\mathbf{u}^\mathsf{T} \left(D - \lambda 1\right)^{-1}$ gives

$$\mathbf{u}^\mathsf{T}\mathbf{q} + \mathbf{u}^\mathsf{T} \left(D - \lambda 1\right)^{-1} \sigma \mathbf{u}\mathbf{u}^\mathsf{T}\mathbf{q} = 0 \,. \tag{4.22}$$

Finally, dividing by $\mathbf{u}^\mathsf{T}\mathbf{q}$ leads to the (scalar) *secular equation* (or *characteristic equation*) in λ,

$$1 + \sum_{i=1}^{m} \frac{\sigma u_i^2}{d_i - \lambda} = 0 \,, \tag{4.23}$$

where the summation index counts the m non-zero components of the update vector $\mathbf{u}$. The m roots of this equation determine the (updated) eigenvalues[1].

Finding a root of a rational function requires an iterative technique, such as the Newton-Raphson method. Therefore diagonalization of an updated matrix is still an iterative procedure. However, each root can be found in $O(1)$ iterations, each iteration requiring $O(m)$ operations. Therefore the iterative part of this algorithm — finding all m roots — needs $O(m^2)$ operations.

Finding roots of this particular secular equation can be simplified by utilizing the fact that its roots are bounded by the eigenvalues d_i of the matrix D. Indeed if we denote the roots as $\lambda_1, \lambda_2, \ldots, \lambda_n$ and assume that $\lambda_i \leq \lambda_{i+1}$ and $d_i \leq d_{i+1}$, it can be shown that

1. if $\sigma \geq 0$,

$$d_i \leq \lambda_i \leq d_{i+1} \,, \qquad i = 1, \ldots, n-1 \,, \tag{4.24}$$
$$d_n \leq \lambda_n \leq d_n + \sigma \mathbf{u}^\mathsf{T}\mathbf{u} \,; \tag{4.25}$$

2. if $\sigma \leq 0$,

$$d_{i-1} \leq \lambda_i \leq d_i \,, \qquad i = 2, \ldots, n \,, \tag{4.26}$$
$$d_1 + \sigma \mathbf{u}^\mathsf{T}\mathbf{u} \leq \lambda_1 \leq d_1 \,. \tag{4.27}$$

---

[1]Multiplying this equation by $\prod_{i=1}^{m}(d_i - \lambda)$ leads to an equivalent polynomial equation of the order m, which has exactly m roots.

4.3.2 Symmetric row/column update

The matrix A to be diagonalized is given in the form

$$
A = D + \mathbf{e}(p)\mathbf{u}^\mathsf{T} + \mathbf{u}\mathbf{e}(p)^\mathsf{T} =
\begin{bmatrix}
d_1 & \cdots & u_1 & \cdots & 0 \\
\vdots & \ddots & \vdots & \ddots & \vdots \\
u_1 & \cdots & d_p & \cdots & u_n \\
\vdots & \ddots & \vdots & \ddots & \vdots \\
0 & \cdots & u_n & \cdots & d_n
\end{bmatrix},
\tag{4.28}
$$

where D is a diagonal matrix with diagonal elements $\{d_i | i = 1, \ldots, n\}$, $\mathbf{e}(p)$ is the unit vector in the p-direction, and $\mathbf{u}$ is a given update vector where the p-th element can be assumed to equal zero, $u_p = 0$, without loss of generality. Indeed, if the element is not zero, one can simply redefine $d_p \rightarrow d_p + 2u_p$, $u_p \rightarrow 0$.

The eigenvalue equation for matrix A is given as

$$
(D - \lambda)\mathbf{x} + \mathbf{e}(p)\mathbf{u}^\mathsf{T}\mathbf{x} + \mathbf{u}\mathbf{e}(p)^\mathsf{T}\mathbf{x} = 0 \,,
\tag{4.29}
$$

where $\mathbf{x}$ is an eigenvector and λ is the corresponding eigenvalue. The component number p of this vector-equation reads

$$
(d_p - \lambda)x_p + \mathbf{u}^\mathsf{T}\mathbf{x} = 0 \,,
\tag{4.30}
$$

while the component number $k \neq p$ reads

$$
(d_k - \lambda)x_k + u_k x_p = 0 \,,
\tag{4.31}
$$

Dividing the last equation by $(d_k - \lambda)$, multiplying from the left with $\sum_{k=1}^{n} u_k$, substituting $\mathbf{u}^\mathsf{T}\mathbf{x}$ using equation (4.30) and dividing by x_p gives the secular equation,

$$
-(d_p - \lambda) + \sum_{k \neq p}^{n} \frac{u_k^2}{d_k - \lambda} = 0 \,,
\tag{4.32}
$$

which determines the updated eigenvalues.

4.3.3 Symmetric rank-2 update

A symmetric rank-2 update can be represented as two consecutive rank-1 updates,

$$
\mathbf{u}\mathbf{v}^\mathsf{T} + \mathbf{v}\mathbf{u}^\mathsf{T} = \mathbf{a}\mathbf{a}^\mathsf{T} - \mathbf{b}\mathbf{b}^\mathsf{T} \,,
\tag{4.33}
$$

where

$$
\mathbf{a} = \frac{1}{\sqrt{2}}(\mathbf{u} + \mathbf{v}) \,, \ \mathbf{b} = \frac{1}{\sqrt{2}}(\mathbf{u} - \mathbf{v}) \,.
\tag{4.34}
$$

The eigenvalues can then be found by applying the rank-1 update method twice.

4.4   Singular Value Decomposition

Singular Value Decomposition (SVD) is a factorization of matrix A in the form

$$A = USV^\mathsf{T} , \tag{4.35}$$

where S is a diagonal matrix, U is (generally) a semi-orthogonal matrix ($U^\mathsf{T}U = 1$), and V is an orthogonal matrix ($V^\mathsf{T}V = VV^\mathsf{T} = 1$).

The elements of the diagonal matrix S are called the *singular values* of matrix A. Singular values can always be chosen non-negative by changing the signs of the corresponding columns of matrix U. The columns of the matrices U and V are called the (correspondingly left and right) singular vectors.

4.4.1   Applications of SVD

Column space, null space, and rank

The SVD of a matrix A provides a representation of the *column space* (also called *range*) and the *null space* of the matrix: the columns of the V matrix corresponding to vanishing singular values span the null space while the columns of the U matrix corresponding to non-vanishing singular values span the range of the matrix.

The rank of the matrix is given by the number of non-vanishing singular values.

Pseudoinverse

An $n \times m$ matrix A has its *pseudoinverse $m \times n$* matrix A^- if the following equation is satisfied,

$$AA^-A = A . \tag{4.36}$$

Given the SVD of the matrix, $A = USV^\mathsf{T}$, its pseudoinverse is

$$A^- = VS^-U^\mathsf{T} , \tag{4.37}$$

where S^- is the pseudoinverse of the matrix S which is formed by replacing non-vanishing elements of S with their inverses.

Least squares solution

The least sqaures solution to an overdetermined system $Ax = b$ is given via the matrix' pseudoinverse as

$$x = A^-b . \tag{4.38}$$

Solution to homogenious systems

Solutions to a homogenious system $Ax = 0$ are given by the null-space of the matrix, that is, by any linear combination of the columns of the matrix V corresponding to vanishing singular values of matrix A.

Lower-rank matrix approximation

Sometimes one needs to approximate a matrix A with a smaller matrix $\tilde{A}$ of a given rank r. The SVD of A provides one such approximation,

$$\tilde{A} = \tilde{U}\tilde{S}\tilde{V}^{\mathsf{T}} , \tag{4.39}$$

where $\tilde{S}$ is a diagonal matrix which contains only the r largest singular values, and where $\tilde{U}$ and $\tilde{V}^{\mathsf{T}}$ contain only the corresponding singular vectors. It minimizes the *Frobenius norm* of the difference between the matrix and its approximation.

4.4.2 Relation to eigenvalues of $A^{\mathsf{T}}A$

Singular values are equal the square roots of the eigenvalues of the real symmetrix matrix $A^{\mathsf{T}}A$. Indeed, by construction the matrix $A^{\mathsf{T}}A$ is positive definite, therefore its eigenvalues are not-negiative and its eigenvalue decomposition can be written as

$$A^{\mathsf{T}}A = VS^2V^{\mathsf{T}} \tag{4.40}$$

where S is the diagonal matrix of square roots of $A^{\mathsf{T}}A$, and V is the matrix of corresponding eigenvectors. Then the singular value decomposition of matrix A can be formally written as

$$A = USV^{\mathsf{T}} , \ U = AVS^{-1} . \tag{4.41}$$

4.4.3 Calculation of SVD

While it is theoretically possible to calculate the SVD of a matrix A by diagonalizing $A^{\mathsf{T}}A$, it is not recommended in practice for several reasons including the potential numerical issues related to squaring small singular values leading to rounding errors. Generally it is more efficient to use algorithms specifically designed for SVD computations.

The most widely used algorithm, called Golub-Reinsch, is a generalization of the QR/QL algorithm for matrix diagonalization. However, like QR/QL in matrix diagonalization, it is a bit tedious and tricky to implement for a learner, therefore we shall discuss here two other, classical, methods which might not be as fast as Golub-Reinsch, but they are stable, precise, and easy to implement.

One-sided Jacobi SVD algorithm

This is an iterative algorithm [10] where the given matrix A is gradually transformed into a matrix with orhogonal columns with the elementary iteration given as

$$A \leftarrow AJ(p, q, \theta) \,, \tag{4.42}$$

where the (rotation) angle θ is chosen such that the columns of the matrix A with the indices p and q become (after the rotation) orthogonal[2]. The indices (p, q) are swept cyclically, $(p = 1 \ldots m, q = p + 1, \ldots, m)$, where m is the number of columns.

After the algoritm has converged, that is, the transformed matrix A has orthogonal columns, the SVD of the original matrix is recovered as follows: the matrix V is the accumulation of the rotation matrices J, the matrix U is given by normalizing the columns of the transformed matrix A, and the diagonal matrix S contains the norms of the columns of the transformed matrix A.

Two-sided Jacobi SVD algorithm

This is also an iterative algorithm where a (square, else see below) matrix A is transformed ino a diagonal matrix, where in the elementary iteration one first applies a Givens rotation to symmetrize a pair of off-diagonal elements of the matrix and then applies a Jacobi transformation to eliminate these off-diagonal elements,

$$A \rightarrow J^\mathsf{T} G^\mathsf{T} A J \,. \tag{4.44}$$

Just like in the Jacobi eigenvalue algorithm the iterations are performed in cyclic sweeps over all non-diagonal elements of the matrix.

For a 2×2 matrix the two-sided Jacobi SVD transformation is given as following: first, one applies a Givens rotation to symmetrize two off-diagonal elements,

$$A = \begin{bmatrix} w & x \\ y & z \end{bmatrix} \rightarrow G^\mathsf{T} A = \begin{bmatrix} \cos\varphi & -\sin\varphi \\ \sin\varphi & \cos\varphi \end{bmatrix} \begin{bmatrix} w & x \\ y & z \end{bmatrix} = \begin{bmatrix} a & b \\ b & c \end{bmatrix} \,, \tag{4.45}$$

where the rotation angle $\varphi = \operatorname{atan2}(x - y, w + z)$; and, second, one makes the usual Jacobi transformation to eliminate the off-diagonal elements b,

$$\begin{aligned} GA &\rightarrow J^\mathsf{T} G^\mathsf{T} A J \\ &= \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} = \begin{bmatrix} d_1 & 0 \\ 0 & d_1 \end{bmatrix} . \end{aligned} \tag{4.46}$$

---

[2]The angle is given as

$$\theta = \frac{1}{2} \operatorname{atan2}(2A_p^\mathsf{T} A_q, A_q^\mathsf{T} A_q - A_p^\mathsf{T} A_p) \,, \tag{4.43}$$

where A_p and A_q are the corresponding columns of the matris A.

The matrices U and V are accumulated (from identity matrices) as

$$U \leftarrow UGJ\,, \tag{4.47}$$
$$V \leftarrow VJ\,. \tag{4.48}$$

If the matrix A is a tall $n \times m$ non-square matrix $(n > m)$, the first step should be the QR-decomposition,

$$A = QR\,, \tag{4.49}$$

where Q is the $n \times m$ orthogonal matrix and R is a square triangular $m \times m$ matrix.

The second step is the normal SVD of the square matrix R,

$$R = U'DV^{\mathsf{T}}\,. \tag{4.50}$$

Now the SVD of the original matrix A is given as

$$A = UDV^{\mathsf{T}}\,, \tag{4.51}$$

where

$$U = QU'\,. \tag{4.52}$$

Chapter 5

Power iteration methods and Krylov subspaces

5.1 Power iteration

Power method is an iterative method to calculate an eigenvalue and the corresponding eigenvector of a (real symmetrix) matrix A using the *power iteration*

$$\mathbf{x}_{i+1} = A\mathbf{x}_i \ . \tag{5.1}$$

The iteration converges to the eigenvector with the largest eigenvalue. Indeed, according to the spectral theorem for self-conjugate operators the eigenvectors $\mathbf{v}_i$,

$$A\mathbf{v}_i = \lambda_i \mathbf{v}_i, \tag{5.2}$$

of a real symmetric matrix A form an orthogonal basis such that any vector $\mathbf{x}_0$ can be represented as a linear combination of the eigenvectors,

$$\mathbf{x}_0 = \sum_k c_k \mathbf{v}_k \ . \tag{5.3}$$

Acting on $\mathbf{x}_0$ with the matrix A gives

$$A^i \mathbf{x}_0 = \sum_k \lambda_k^i c_k \mathbf{v}_k \ . \tag{5.4}$$

Thus the contribution from the eigenvector corresponding to the largest eigenvalue is amplified at each iteration with the factor

$$\left(\frac{\lambda_{\text{largest}}}{\lambda_{\text{next largest}}} \right)^i \ . \tag{5.5}$$

The eigenvalue can be estimated using the *Rayleigh quotient*,

$$\lambda[\mathbf{x}_i] = \frac{\mathbf{x}_i^T A \mathbf{x}_i}{\mathbf{x}_i^T \mathbf{x}_i} = \frac{\mathbf{x}_{i+1}^T \mathbf{x}_i}{\mathbf{x}_i^T \mathbf{x}_i}. \tag{5.6}$$

5.2 Inverse iteration

Alternatively, the *inverse power iteration* with the inverse matrix,

$$\mathbf{x}_{i+1} = A^{-1} \mathbf{x}_i , \tag{5.7}$$

converges to the smallest (in the absolute value) eigenvalue of matrix A.

Finally, the *shifted inverse power iteration*,

$$\mathbf{x}_{i+1} = (A - s1)^{-1} \mathbf{x}_i , \tag{5.8}$$

where 1 signifies the identity matrix of the same size as A, converges to the eigenvalue closest to the given number s.

The *inverse iteration method* is a refinement of the inverse power method where the trick is not to invert the matrix in (5.8) but rather solve the linear system

$$(A - s1)\mathbf{x}_{i+1} = \mathbf{x}_i \tag{5.9}$$

using e.g. QR-decomposition.

The better approximation s to the sought eigenvalue is chosen, the faster convergence one gets. However, incorrect choice of s can lead to slow convergence or to the convergence to a different eigenvector. In practice the method is usually used when good approximation for the eigenvalue is known, and hence one needs only few (quite often just one) iteration.

One can update the estimate for the eigenvalue using the Rayleigh quotient $\lambda[\mathbf{x}_i]$ after each iteration and get faster convergence for the price of $O(n^3)$ operations per QR-decomposition; or one can instead make more iterations (with $O(n^2)$ operations per iteration) using the same matrix $(A - s1)$. The optimal strategy is probably an update after several iterations.

5.3 Krylov subspaces

When calculating an eigenvalue of a matrix A using the power method, one starts with an initial random vector $\mathbf{b}$ and then computes iteratively the sequence $A\mathbf{b}, A^2\mathbf{b}, \ldots, A^{n-1}\mathbf{b}$ normalising and storing the result in $\mathbf{b}$ on each iteration. The sequence converges to the eigenvector of the largest eigenvalue of A.

The set of vectors

$$\mathcal{K}_n = \left\{ \mathbf{b}, A\mathbf{b}, A^2\mathbf{b}, \dots, A^{n-1}\mathbf{b} \right\} , \tag{5.10}$$

where $n < \mathrm{rank}(A)$, is called the order-n *Krylov matrix*, and the subspace spanned by these vectors is called the order-n *Krylov subspace* [2]. The vectors are not orthogonal but can be made so e.g. by Gram-Schmidt orthogonalisation.

Krylov subspaces are the basis of several successful iterative methods in numerical linear algebra, in particular: Arnoldi and Lanczos methods for finding one (or a few) eigenvalues of a matrix; and GMRES (Generalised Minimum RESidual) method for solving systems of linear equations.

These methods are particularly suitable for large sparse matrices as they avoid matrix-matrix operations but rather multiply vectors by matrices and work with the resulting vectors and matrices in Krylov subspaces of modest sizes.

The Krylov subspaces can be expected to approximate the subspace of the extreme (the highest and the lowest) eigenvalues of the given matrix.

5.4 Arnoldi iteration

Arnoldi iteration is an algorithm where the order-n orthogonalised Krylov matrix Q_n for a given matrix A is built using stabilised Gram-Schmidt process [3]:

```
start with a set Q = {q₁} where q₁ is a random normalised vector;
repeat for k = 2 to n :
    make a new vector qₖ = Aqₖ₋₁
    orthogonalise qₖ to all vectors qᵢ ∈ Q storing qᵢ†qₖ → hᵢ,ₖ₋₁
    normalise qₖ storing ‖qₖ‖ → hₖ,ₖ₋₁
    add qₖ to the set Q
```

By construction the matrix H_n made of the elements h_{jk} is an upper Hessenberg matrix,

$$H_n = \begin{bmatrix} h_{1,1} & h_{1,2} & h_{1,3} & \cdots & h_{1,n} \\ h_{2,1} & h_{2,2} & h_{2,3} & \cdots & h_{2,n} \\ 0 & h_{3,2} & h_{3,3} & \cdots & h_{3,n} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & h_{n,n-1} & h_{n,n} \end{bmatrix} , \tag{5.11}$$

which is a partial orthogonal reduction of A into Hessenberg form,

$$H_n = Q_n^\dagger A Q_n . \tag{5.12}$$

The matrix H_n can be viewed as a representation of A in the Krylov subspace $\mathcal{K}_n$. The eigenvalues and eigenvectors of the matrix H_n approximate the extreme eigenvalues of matrix A.

Since H_n is a Hessenberg matrix of modest size one can relatively easily apply to it the standard algorithms of linear algebra.

In practice if the size n of the Krylov subspace becomes too large the method is restarted.

5.5   Lanczos iteration

Lanczos iteration is Arnoldi iteration for Hermitian matrices [21], in which case the Hessenberg matrix H_n of Arnoldi method becomes a tridiagonal matrix T_n.

The Lanczos algorithm thus reduces the original hermitian $N \times N$ matrix A into a smaller $n \times n$ tridiagonal matrix T_n by an orthogonal projection onto the order-n Krylov subspace. The eigenvalues and eigenvectors of a tridiagonal matrix of a modest size can be easily found by e.g. the QR-diagonalisation method.

In practice the Lanczos method is not very stable due to round-off errors leading to quick loss of orthogonality. The eigenvalues of the resulting tridiagonal matrix may then not be a good approximation to the original matrix. Library implementations fight the stability issues by trying to prevent the loss of orthogonality and/or to recover the orthogonality after the basis is generated.

5.6   Generalised minimum residual (GMRES)

GMRES is an iterative method for the numerical solution of a system of linear equations,

$$A\mathbf{x} = \mathbf{b} \,, \tag{5.13}$$

where the exact solution is approximated by the least-squares solution that minimises the residual $\|A\mathbf{x} - \mathbf{b}\|^2$ in the Krylov subspace $\mathcal{K}_n$ of matrix A.

The original equation is projected on the Krylov subspace,

$$A\mathbf{x} = \mathbf{b} \quad \Rightarrow \quad (Q_n^\dagger A Q_n)(Q_n^\dagger \mathbf{x}) = (Q_n^\dagger \mathbf{b}) \,, \tag{5.14}$$

which gives a linear system of size-n,

$$H_n \tilde{\mathbf{x}} = \tilde{\mathbf{b}} \,, \tag{5.15}$$

where

$$H_n = Q_n^\dagger A Q_n \,, \quad \tilde{\mathbf{b}} = Q_n^\dagger \mathbf{b} \,. \tag{5.16}$$

The Hessenberg equation (5.15) can be easily solved by one run of Gauss elimination and then a run of back-substitution (like in cubic splines). The approximate solution to the original equation is then given as

$$\mathbf{x} \approx Q_n \tilde{\mathbf{x}} \, . \tag{5.17}$$

Chapter 6

Ordinary differential equations

6.1 Introduction

Ordinary differential equations (ODE) are generally defined as differential equations in one variable where the highest order derivative enters linearly. Such equations invariably arise in many different contexts throughout mathematics (and science generally) as soon as changes in the system at hand are considered, usually with respect to variations of certain parameters.

Ordinary differential equations can be generally reformulated as (coupled) systems of first-order ordinary differential equations,

$$\mathbf{y}'(x) = \mathbf{f}(x, \mathbf{y}) \,, \tag{6.1}$$

where $\mathbf{y}' \doteq d\mathbf{y}/dx$, and the variables $\mathbf{y}$ and the *right-hand side function* $\mathbf{f}(x, \mathbf{y})$ are understood as column-vectors. For example, a second order differential equation in the form

$$u'' = g(x, u, u') \tag{6.2}$$

can be rewritten as a system of two first-order equations,

$$\begin{cases} y_1' = y_2 \\ y_2' = g(x, y_1, y_2) \end{cases} , \tag{6.3}$$

using the variable substitution $y_1 = u$, $y_2 = u'$.

In practice ODEs are usually supplemented with boundary conditions which pick out a certain class or a unique solution of the ODE. In the following we shall mostly

consider the *initial value problem*: an ODE with the boundary condition in the form of an initial condition at a given point a,

$$\mathbf{y}(a) = \mathbf{y}_0 \,. \tag{6.4}$$

The problem is then to find the value of the solution $\mathbf{y}$ at some other point b.

Finding a solution to an ODE is often referred to as *integrating* the ODE.

An integration algorithm typically advances the solution from the initial point a to the final point b in a number of discrete steps

$$\{x_0 \doteq a, x_1, \ldots, x_{n-1}, x_n \doteq b\}. \tag{6.5}$$

An efficient algorithm tries to integrate an ODE using as few steps as possible under the constraint of the given accuracy goal. For this purpose the algorithm should continuously adjust the step-size during the integration, using few larger steps in the regions where the solution is smooth and perhaps many smaller steps in more treacherous regions.

Typically, an adaptive step-size ODE integrator is implemented as two routines. One of them—called *driver*—monitors the local errors and tolerances and adjusts the step-sizes. To actually perform a step the driver calls a separate routine—the *stepper*—which advances the solution by one step, using one of the many available algorithms, and estimates the local error. The GNU Scientific Library, GSL, implements about a dozen of different steppers and a tunable adaptive driver.

In the following we shall discuss several of the popular driving algorithms and stepping methods for solving initial-value ODE problems.

6.2   Adaptive step-size control

Let *tolerance* τ be the maximal accepted error consistent with the required accuracy to be achieved in the integration of an ODE. Suppose the integration is done in n steps of size h_i such that $\sum_{i=1}^{n} h_i = b - a$. Under assumption that the errors at the integration steps are random and statistically uncorrelated, the local tolerance τ_i for the step i has to scale as the square root of the step-size,

$$\tau_i = \tau \sqrt{\frac{h_i}{b-a}}. \tag{6.6}$$

Indeed, if the local error e_i on the step i is less than the local tolerance, $e_i \leq \tau_i$, the total error E will be consistent with the total tolerance τ,

$$E \approx \sqrt{\sum_{i=1}^{n} e_i^2} \leq \sqrt{\sum_{i=1}^{n} \tau_i^2} = \tau \sqrt{\sum_{i=1}^{n} \frac{h_i}{b-a}} = \tau \,. \tag{6.7}$$

The current step h_i is accepted if the local error e_i is smaller than the local tolerance τ_i, after which the next step is attempted with the step-size adjusted according to the following empirical prescription [11],

$$h_{i+1} = h_i \times \left(\frac{\tau_i}{e_i}\right)^{\text{Power}} \times \text{Safety}, \tag{6.8}$$

where Power ≈ 0.25 and Safety ≈ 0.95.

If the local error is larger than the local tolerance the step is rejected and a new step is attempted with the step-size adjusted according to the same prescription (6.8).

One simple prescription for the local tolerance τ_i and the local error e_i to be used in (6.8) is

$$\tau_i = (\epsilon\|\mathbf{y}_i\| + \delta)\sqrt{\frac{h_i}{b-a}} \;,\;\; e_i = \|\delta\mathbf{y}_i\| \;, \tag{6.9}$$

where δ and ϵ are the required absolute and relative precision and $\delta\mathbf{y}_i$ is the estimate of the integration error at the step i.

A more elaborate prescription considers components of the solution separately,

$$(\tau_i)_k = \left(\epsilon|(\mathbf{y}_i)_k| + \delta\right)\sqrt{\frac{h_i}{b-a}} \;,\;\; (\mathbf{e}_i)_k = |(\delta\mathbf{y}_i)_k| \;, \tag{6.10}$$

where the index k runs over the components of the solution. In this case the step acceptance criterion also becomes component-wise: the step is accepted, if

$$\forall k :\; (\mathbf{e}_i)_k < (\tau_i)_k \;. \tag{6.11}$$

The factor τ_i/e_i in the step adjustment formula (6.8) is then replaced by

$$\frac{\tau_i}{e_i} \to \min_k \frac{(\tau_i)_k}{(\mathbf{e}_i)_k} \;. \tag{6.12}$$

Yet another refinement is to include the derivatives $\mathbf{y}'$ of the solution into the local tolerance estimate, either overall,

$$\tau_i = \left(\epsilon\alpha\|\mathbf{y}_i\| + \epsilon\beta\|\mathbf{y}_i'\| + \delta\right)\sqrt{\frac{h_i}{b-a}} \;, \tag{6.13}$$

or component-wise,

$$(\tau_i)_k = \left(\epsilon\alpha|(\mathbf{y}_i)_k| + \epsilon\beta|(\mathbf{y}_i')_k| + \delta\right)\sqrt{\frac{h_i}{b-a}} \;. \tag{6.14}$$

The weights α and β are chosen by the user.

Table 6.1 shows an implementation of the discussed driver in Python.

Table 6.1: An implementation of an ODE driver in Python

```
def driver(f,a,ya,b,h=0.125,acc=0.01,eps=0.01) :
    x=a; y=ya; xlist=[x]; ylist=[y]
    while True :
        if x>=b  : return (xlist,ylist) # return the path
        if x+h>b : h=b−x
        (yh,dy) = stepper(f,x,y,h) # stepper returns y(x+h) and δy
        tol = (acc+eps*yh.norm())*math.sqrt(h/(b−a))
        err = dy.norm()
        if err<tol :
            x+=h; y=yh; xlist.append(x); ylist.append(y)
        if err>0 : h*=min( (tol/err)**0.25*0.95 , 2 )
        else     : h*=2
```

6.3   Error estimate

In an adaptive step-size algorithm the stepping routine must provide an estimate of the integration error, upon which the driver bases its strategy to determine the optimal step-size for a user-specified accuracy goal.

A stepping method is generally characterized by its *order*: a method has order p if it can integrate exactly an ODE where the solution is a polynomial of order p. In other words, for small h the error of the order-p method is $O(h^{p+1})$.

There are two popular methods to estimate the error: step doubling, where one compares the results from the full step and from two half-steps, and two-orders method where one compares the results from steppers of two different orders.

6.3.1   Step doubling (Runge's principle)

For sufficiently small steps the error δy of an integration step for a method of a given order p can be estimated by comparing the solution $\mathbf{y}_{\text{full\_step}}$, obtained with one full-step integration, against a potentially more precise solution, $\mathbf{y}_{\text{two\_half\_steps}}$, obtained with two consecutive half-step integrations,

$$\delta\mathbf{y} = \frac{\mathbf{y}_{\text{full\_step}} - \mathbf{y}_{\text{two\_half\_steps}}}{2^p - 1} \, . \tag{6.15}$$

where p is the order of the algorithm used. Indeed, if the step-size h is small, we can assume

$$\delta\mathbf{y}_{\text{full\_step}} \quad = \quad Ch^{p+1} \, , \tag{6.16}$$

$$\delta\mathbf{y}_{\text{two\_half\_steps}} \quad = \quad 2C\left(\frac{h}{2}\right)^{p+1} = \frac{Ch^{p+1}}{2^p} \, , \tag{6.17}$$

where $\delta\mathbf{y}_{\text{full\_step}}$ and $\delta\mathbf{y}_{\text{two\_half\_steps}}$ are the errors of the full-step and two half-steps integrations, and C is an unknown constant. The two can be combined as

$$
\begin{aligned}
\mathbf{y}_{\text{full\_step}} - \mathbf{y}_{\text{two\_half\_steps}} &= \delta\mathbf{y}_{\text{full\_step}} - \delta\mathbf{y}_{\text{two\_half\_steps}} \\
&= \frac{Ch^{p+1}}{2^p}(2^p - 1)\,,
\end{aligned}
\tag{6.18}
$$

from which it follows that

$$
\frac{Ch^{p+1}}{2^p} = \frac{\mathbf{y}_{\text{full\_step}} - \mathbf{y}_{\text{two\_half\_steps}}}{2^p - 1}\,.
\tag{6.19}
$$

One has, of course, to take the potentially more precise $\mathbf{y}_{\text{two\_half\_steps}}$ as the approximation to the solution $\mathbf{y}$. Its error is then given as

$$
\delta\mathbf{y}_{\text{two\_half\_steps}} = \frac{Ch^{p+1}}{2^p} = \frac{\mathbf{y}_{\text{full\_step}} - \mathbf{y}_{\text{two\_half\_steps}}}{2^p - 1}\,,
\tag{6.20}
$$

which had to be demonstrated. This prescription is often referred to as the *Runge's principle*.

One drawback of the Runge's principle is that the full-step and the two half-step calculations generally do not share evaluations of the right-hand side function $\mathbf{f}(x, \mathbf{y})$, and therefore many extra evaluations are needed to estimate the error.

6.3.2 Different orders

An alternative prescription for error estimation is to make the same step-size integration using two methods of *different orders*, with the difference between the two solutions providing the estimate of the error. If the lower order method mostly uses the same evaluations of the right-hand side function—in which case it is called *embedded* in the higher order method—the error estimate does not need additional evaluations.

Predictor-corrector methods are naturally of embedded type: the correction—which generally increases the order of the method—itself can serve as the estimate of the error.

6.4 Runge-Kutta steppers

The Runge-Kutta steppers approximate the solution of the differential equation using polynomials. The coefficients of the polynomials are determined by sampling the right-hand-side, $\mathbf{f}(x, \mathbf{y})$, of the differential equation at certain points (mostly) within the step.

6.4.1 Linear stepper

The linear approximation is given by the first two terms of the Taylor expansion of the solution at the point x_i,

$$\mathbf{y}(x) \approx \mathbf{p}_1(x) = \mathbf{y}_i + \mathbf{y}'_i(x - x_i) \,, \tag{6.21}$$

where the derivative $\mathbf{y}'_i$ is determined from the equation itself, $\mathbf{y}'_i = \mathbf{f}(x, \mathbf{y}_i)$. In this parlance the values of the derivative are often denoted as $\mathbf{k}$ with index, for example the derivative $\mathbf{y}'_i$ is often called $\mathbf{k}_0$,

$$\mathbf{p}_1(x) = \mathbf{y}_i + \mathbf{k}_0(x - x_i) \,. \tag{6.22}$$

With this notation the linear Runge-Kutta stepper (called the Euler's rule) is given as

$$\mathbf{k}_0 = \mathbf{f}(x, \mathbf{y}_i) \tag{6.23}$$
$$\mathbf{y}_{i+1} \approx \mathbf{p}_1(x + h) = \mathbf{y}_i + \mathbf{k}_0 h \,. \tag{6.24}$$

6.4.2 Quadratic steppers and Butcher tableaux

One can improve upon the $\mathbf{p}_1$-approximation by adding a second order term,

$$\mathbf{y}(x) \approx \mathbf{p}_2(x) = \mathbf{p}_1(x) + \mathbf{c}(x - x_i)^2 \,, \tag{6.25}$$

where the coefficient $\mathbf{c}$ can be found by matching the polynomial $\mathbf{p}_2$ against our differential equation at a certain point z within the step,

$$\mathbf{p}'_2(z) = \mathbf{f}(z, \mathbf{p}_1(z)) \,. \tag{6.26}$$

This condition should be viewed in the sense of Taylor expansion — that is, in the same order of h — that's why we used $\mathbf{p}_1$ polynomial at the right-hand-side. This gives

$$\mathbf{c} = \frac{\mathbf{f}(z, \mathbf{p}_1(z)) - \mathbf{p}'_1(z)}{2(z - x_i)} \,. \tag{6.27}$$

One usually identifies the sampling point z by the corresponding fraction of the step, $z = x_i + \alpha h$. Written with the α the $\mathbf{c}$-coefficient is given as

$$\mathbf{c} = \frac{\mathbf{f}(x_i + \alpha h, \mathbf{y}_i + \alpha \mathbf{k}_0 h)) - \mathbf{k}_0}{2\alpha h} = \frac{\mathbf{k}_1 - \mathbf{k}_0}{2\alpha h} \,, \tag{6.28}$$

where

$$\mathbf{k}_1 \doteq \mathbf{f}(x_i + \alpha h, \mathbf{y}_i + \alpha \mathbf{k}_0 h) \,. \tag{6.29}$$

Finally, the second order approximating polynomial is given as

$$\mathbf{p}_2(x) = \mathbf{y}_i + \mathbf{k}_0(x - x_i) + (\mathbf{k}_1 - \mathbf{k}_0)\frac{(x - x_i)^2}{2\alpha h} \ . \tag{6.30}$$

The value of the function at the end of the step, $\mathbf{y}_{i+1} \approx \mathbf{p}_2(x_i + h)$, is then given as

$$\mathbf{y}_{i+1} = \mathbf{y}_i + \left(1 - \frac{1}{2\alpha}\right)\mathbf{k}_0 h + \frac{1}{2\alpha}\mathbf{k}_1 h \ . \tag{6.31}$$

Summarizing, the (generic) second order Runge-Kutta stepper is given as

$$\begin{aligned}
\mathbf{k}_0 &= \mathbf{f}(x_i, \mathbf{y}_i) & (6.32) \\
\mathbf{k}_1 &= \mathbf{f}(x_i + \alpha h, \mathbf{y}_i + \alpha \mathbf{k}_0 h) & (6.33) \\
\mathbf{y}_{i+1} &= \mathbf{y}_i + \left(1 - \frac{1}{2\alpha}\right)\mathbf{k}_0 h + \frac{1}{2\alpha}\mathbf{k}_1 h \ . & (6.34)
\end{aligned}$$

It is customary in this business to represent the Runge-Kutta steppers with the so called Butcher tableaux which collect the step-sizes and the coefficient of the polynomials in a table. For example, the Butcher tableau of out generic second order stepper is given as

	step sizes	coefficients before $\mathbf{k}_i h$	
$\mathbf{k}_0$	0		
$\mathbf{k}_1$	α	α	
$\mathbf{y}_{i+1}$		$\left(1 - \frac{1}{2\alpha}\right)$	$\frac{1}{2\alpha}$

$$\tag{6.35}$$

The intermediate step α can be chosen arbitrarily (or from a certain condition) the stepper is of second order for any α. Here are some popular second order methods – midpoint, Heun's, and Ralston's.

0			0			0	
1/2	1/2		1	1		2/3	2/3
	0	1		1/2	1/2		1/4	3/4

$$\tag{6.36}$$

6.4.3 Cubic approximation

One can further increase the order of the approximating polynomial,

$$\mathbf{p}_3(x) = \mathbf{p}_2(x) + \mathbf{d}(x - x_i)^3 \ , \tag{6.37}$$

where d can be found from the matching condition at a (possibly) different point z_2 within the step,

$$\mathbf{p}_3'(z_2) = \mathbf{f}(z_2, \mathbf{p}_2(z_2)) \ . \tag{6.38}$$

This gives

$$\mathbf{d} = \frac{\mathbf{f}(z_2, \mathbf{p}_2(z_2)) - \mathbf{p}_2'(z_2)}{3(z_2 - x_i)^2} \ . \tag{6.39}$$

Taking $z_2 = x_i + \beta h$ and denoting

$$\mathbf{k}_2 \doteq \mathbf{f}\left(z_2, \mathbf{p}_2(z_2)\right) = \mathbf{f}\left(x_i + \beta h, \mathbf{y}_i + \left(\beta - \frac{\beta^2}{2\alpha}\right)\mathbf{k}_0 h + \frac{\beta^2}{2\alpha}\mathbf{k}_1 h\right) \tag{6.40}$$

we get

$$\mathbf{d} = \frac{\mathbf{k}_2 - \left(\mathbf{k}_0 + (\mathbf{k}_1 - \mathbf{k}_0)\frac{\beta}{\alpha}\right)}{3\beta^2 h^2} \ , \tag{6.41}$$

and

$$\mathbf{p}_3(x) = \mathbf{y}_i + \mathbf{k}_0(x - x_i) + \frac{\mathbf{k}_1 - \mathbf{k}_0}{2\alpha h}(x - x_i)^2$$

$$+ \frac{\mathbf{k}_2 - \left(\mathbf{k}_0 + (\mathbf{k}_1 - \mathbf{k}_0)\frac{\beta}{\alpha}\right)}{3\beta^2 h^2}(x - x_i)^3 \ . \tag{6.42}$$

This gives the following (generic) cubic Runge-Kutta stepper,

$$
\begin{aligned}
\mathbf{k}_0 &= \mathbf{f}(x_i, \mathbf{y}_i) & (6.43)\\
\mathbf{k}_1 &= \mathbf{f}(x_i + \alpha h, \mathbf{y}_i + \alpha \mathbf{k}_0 h) & (6.44)\\
\mathbf{k}_2 &= \mathbf{f}\left(x_i + \beta h, \mathbf{y}_i + \left(\beta - \frac{\beta^2}{2\alpha}\right)\mathbf{k}_0 h + \frac{\beta^2}{2\alpha}\mathbf{k}_1 h\right) & (6.45)\\
\mathbf{y}_{i+1} &= \mathbf{y}_i + \left(1 - \frac{1}{2\alpha} - \frac{1}{3\beta^2} + \frac{1}{3\alpha\beta}\right)\mathbf{k}_0 h + \left(\frac{1}{2\alpha} - \frac{1}{3\alpha\beta}\right)\mathbf{k}_1 h &\\
&\quad + \left(\frac{1}{3\beta^2}\right)\mathbf{k}_2 h \ . & (6.46)
\end{aligned}
$$

The corresponding Butcher tableau is

$$
\begin{array}{c|ccc}
0 & & & \\
\hline
\alpha & \alpha & & \\
\beta & \beta - \frac{\beta^2}{2\alpha} & \frac{\beta^2}{2\alpha} & \\
\hline
p_3 & 1 - \frac{1}{2\alpha} - \frac{1}{3\beta^2} + \frac{1}{3\alpha\beta} & \frac{1}{2\alpha} - \frac{1}{3\alpha\beta} & \frac{1}{3\beta^2} \\
p_2 & 1 - \frac{1}{2\alpha} & \frac{1}{2\alpha} & 0
\end{array}
\tag{6.47}
$$

The intermediates steps α and β can be chosen largely arbitrarily.

Here are several popular third order Runge-Kutta steppers, Heun's, Ralston's, and 8/15th,

$$
\begin{array}{c|ccc}
0 \\
1/3 & 1/3 \\
2/3 & 0 & 2/3 \\
\hline
& 1/4 & 0 & 3/4
\end{array}
\qquad
\begin{array}{c|ccc}
0 \\
1/2 & 1/2 \\
3/4 & 0 & 3/4 \\
\hline
& 2/9 & 3/9 & 4/9
\end{array}
\qquad
\begin{array}{c|ccc}
0 \\
8/15 & 8/15 \\
2/3 & 1/4 & 5/12 \\
\hline
& 1/4 & 0 & 3/4
\end{array}
\tag{6.48}
$$

And, in principle, the coefficients in the rows can be also reshuffled somewhat (preferable shifting the weight to the right) with the condition that the sum of the coefficients in the row must be equal one.

6.4.4 Higher order steppers

Using this approach one can easily further increase the order of the approximating polynomial by adding higher powers of $(x - x_i)$ and fixing the coefficients in front by an additional sampling of the right-hand-side at certain points within the step. This led to a bunch of steppers of different orders with different sampling points. For example, here is the classic Runge-Kutta fourth-order method,

$$
\begin{array}{c|cccc}
0 \\
1/2 & 1/2 \\
1/2 & 0 & 1/2 \\
1 & 0 & 0 & 1 \\
\hline
& 1/6 & 1/3 & 1/3 & 1/6
\end{array}
\quad .
\tag{6.49}
$$

One can find a list of different Runge-Kutta methods in Wikipedia.

6.4.5 Embedded methods with error estimates

The embedded Runge-Kutta methods in addition to advancing the solution by one step also produce an estimate of the local error of the step. This is done by having two methods in the tableau, one with a certain order p and another one with order $p-1$. The difference between the two methods gives the estimate of the local error. If the lower order method uses the same **k**-values as the higher order method, it is called *embedded*. Embedded methods allow effective estimate of the error without extra evaluations of the right-hand-side.

Table 6.2: Embedded midpoint/Euler method with error estimate

```
public static (vector,vector) rkstep12
    (Func<double,vector,vector> f, double x, vector y, double h)
{
vector k0 = f(x,y);                 /* embedded lower order formula (Euler) */
vector k1 = f(x+h/2,y+k0*(h/2));/* higher order formula (midpoint) */
vector yh = y+k1*h;                 /* y(x+h) estimate */
vector dy = (k1−k0)*h;              /* error estimate */
return (yh,dy);
}
```

Since the embedded rule uses the same **k**'s the Butcher's tableau for this kind of method is simply extended by one row to give the coefficients of the lower order rule.

The simplest embedded methods are Heun-Euler method, and midpoint-Euler method,

$$
\begin{array}{c|cc}
0 & & \\
1 & 1 & \\
\hline
 & 1/2 & 1/2 \\
 & 1 & 0
\end{array}
\ ,\quad
\begin{array}{c|cc}
0 & & \\
1/2 & 1/2 & \\
\hline
 & 0 & 1 \\
 & 1 & 0
\end{array}
\ ,
\tag{6.50}
$$

which both combine methods of orders 2 and 1. Table (6.2) shows a C# implementation of the embedded midpoint/Euler method with error estimate.

Here are simple embedded methods of orders 3 and 2,

$$
\begin{array}{c|ccc}
0 & & & \\
1/2 & 1/2 & & \\
3/4 & 0 & 3/4 & \\
\hline
 & 2/9 & 3/9 & 4/9 \\
 & 0 & 1 & 0
\end{array}
\ ,\quad
\begin{array}{c|ccc}
0 & & & \\
1/3 & 1/3 & & \\
2/3 & 0 & 2/3 & \\
\hline
 & 1/3 & 0 & 2/3 \\
 & 1/4 & 0 & 3/4
\end{array}
\ ,
\tag{6.51}
$$

The *Bogacki-Shampine method* [6] also combines methods of orders 3 and 2,

$$
\begin{array}{c|cccc}
0 & & & & \\
1/2 & 1/2 & & & \\
3/4 & 0 & 3/4 & & \\
1 & 2/9 & 1/3 & 4/9 & \\
\hline
 & 2/9 & 1/3 & 4/9 & 0 \\
 & 7/24 & 1/4 & 1/3 & 1/8
\end{array}
\qquad (6.52)
$$

Bogacki and Shampine argue that their method has better stability properties and actually outperforms higher order methods at lower accuracy goal calculations. This method has the FSAL—First Same As Last—property: the value $\mathbf{k}_4$ at one step equals $\mathbf{k}_1$ at the next step; thus only three function evaluations are needed per step.

The Runge-Kutta-Fehlberg method [12]—called *RKF45*—implemented in the renowned `rkf45` Fortran routine, has two methods of orders 5 and 4:

$$
\begin{array}{c|cccccc}
0 & & & & & & \\
1/4 & 1/4 & & & & & \\
3/8 & 3/32 & 9/32 & & & & \\
12/13 & 1932/2197 & -7200/2197 & 7296/2197 & & & \\
1 & 439/216 & -8 & 3680/513 & -845/4104 & & \\
1/2 & -8/27 & 2 & -3544/2565 & 1859/4104 & -11/40 & \\
\hline
 & 16/135 & 0 & 6656/12825 & 28561/56430 & -9/50 & 2/55 \\
 & 25/216 & 0 & 1408/2565 & 2197/4104 & -1/5 & 0
\end{array}
$$

6.5 Implicit methods

Instead of the forward Euler method one could employ the backward Euler method where the derivative is approximated as

$$
y'(x) \approx \frac{y(x) - y(x - h)}{h} \; , \qquad (6.53)
$$

which gives the following (backward Euler) stepper,

$$
y_{x+h} = y_x + hf(x + h, y_{x+h}) \; . \qquad (6.54)
$$

The backward Euler methods is an *implicit* method: one has to solve the above equation to find y_{x+h}. It generally costs time to solve this equation numerically – a disadvantage as compared to explicit methods. However implicit methods are usually more stable for stiff (difficult) equations where a larger step h can be used as compared to explicit methods.

Just like with explicit methods one can devise higher-order implicit methods, for example, the implicit Heun's method (trapezoidal rule),

$$y_{x+h} = y_x + h\frac{1}{2}\big(f(x, y_x) + f(x + h, y_{x+h})\big) \ . \tag{6.55}$$

6.6 Multistep methods

Multistep methods try to use the information about the function gathered at the previous steps. They are generally not *self-starting* as there are no previous steps at the start of the integration. The first step must be done with a one-step method like Runge-Kutta.

A number of multistep methods have been devised (and named after different mathematicians); we shall only consider a few simple ones here to get the idea of how it works.

6.6.1 Two-step method

Given the previous point, $(x_{i-1}, \mathbf{y}_{i-1})$, in addition to the current point $(x_i, \mathbf{y}_i)$, the sought function $\mathbf{y}$ can be approximated in the vicinity of the point x_i as a second order polynomial,

$$\mathbf{y}(x) \approx \mathbf{p}_2(x) = \mathbf{y}_i + \mathbf{y}'_i \cdot (x - x_i) + \mathbf{c} \cdot (x - x_i)^2, \tag{6.56}$$

where $\mathbf{y}'_i = \mathbf{f}(x_i, \mathbf{y}_i)$ and the coefficient $\mathbf{c}$ can be found from the condition

$$\mathbf{p}_2(x_{i-1}) = \mathbf{y}_{i-1} \ , \tag{6.57}$$

which gives

$$\mathbf{c} = \frac{\mathbf{y}_{i-1} - \mathbf{y}_i + \mathbf{y}'_i \cdot (x_i - x_{i-1})}{(x_i - x_{i-1})^2} \ . \tag{6.58}$$

The value $\mathbf{y}_{i+1}$ of the function at the next point, $x_{i+1} \doteq x_i + h$, can now be estimated as $\mathbf{y}_{i+1} = \mathbf{p}_2(x_{i+1})$ from (6.56).

The error of this second-order two-step stepper can be estimated by a comparison with the first-order Euler's step, which is given by the linear part of (6.56). The correction term $\mathbf{c}h^2$ can serve as the error estimate,

$$\delta\mathbf{y} = \mathbf{c}h^2 \ . \tag{6.59}$$

6.6.2   Two-step method with extra evaluation

One can further increase the order of the approximation (6.56) by adding a third order term,

$$\mathbf{y}(x) \approx \mathbf{p}_3(x) = \mathbf{p}_2(x) + \mathbf{d} \cdot (x - x_i)^2 (x - x_{i-1}) \,. \tag{6.60}$$

The coefficient $\mathbf{d}$ can be found from the matching condition at a certain point t inside the interval,

$$\mathbf{p}_3'(t) = \mathbf{f}(t, \mathbf{p}_2(t)) \,, \tag{6.61}$$

where $x_i < t < x_i + h$. This gives

$$\mathbf{d} = \frac{\mathbf{f}(t, \mathbf{p}_2(t)) - \mathbf{y}_i' - 2\mathbf{c} \cdot (t - x_i)}{2(t - x_i)(t - x_{i-1}) + (t - x_i)^2} \,. \tag{6.62}$$

The error estimate at the point $x_{i+1} \doteq x_0 + h$ is again given as the difference between the higher and the lower order methods,

$$\delta\mathbf{y} = \mathbf{p}_3(x_{i+1}) - \mathbf{p}_2(x_{i+1}) \,. \tag{6.63}$$

6.7   Predictor-corrector methods

Predictor-corrector methods use extra iterations to improve the solution. The algorithm proceeds in two steps: first, the predictor step calculates a rough approximation of $\mathbf{y}(x + h)$. Second, the corrector step refines the initial approximation. The corrector step can be repeated in the hope that this achieves an even better approximation to the true solution.

For example, the two-point Runge-Kutta method (**??**) is as actually a predictor-corrector method, as it first calculates the *prediction*, $\mathbf{y}^p$, for $\mathbf{y}(x_{i+1})$,

$$\mathbf{y}^p = \mathbf{y}_i + h\mathbf{f}(x_i, \mathbf{y}_i) \,, \tag{6.64}$$

and then uses this prediction in the *correction* step to calculate the corrected value, $\mathbf{y}^c$,

$$\mathbf{y}^c = \mathbf{y}_i + h\frac{1}{2}\left(\mathbf{f}(x_i, \mathbf{y}_i) + \mathbf{f}(x_{i+1}, \mathbf{y}^p)\right) \,. \tag{6.65}$$

6.7.1   Two-step method with correction

Similarly, one can use the two-step approximation (6.56) as a predictor, and then improve it by one order with a correction step, namely

$$\check{\mathbf{y}}(x) = \bar{\mathbf{y}}(x) + \check{\mathbf{d}} \cdot (x - x_i)^2 (x - x_{i-1}). \tag{6.66}$$

The coefficient $\check{\mathbf{d}}$ can be found from the condition $\check{\mathbf{y}}'(x_{i+1}) = \bar{\mathbf{f}}_{i+1}$, where $\bar{\mathbf{f}}_{i+1} \doteq \mathbf{f}(x_{i+1}, \bar{\mathbf{y}}(x_{i+1}))$,

$$\check{\mathbf{d}} = \frac{\bar{\mathbf{f}}_{i+1} - \mathbf{y}_i' - 2\mathbf{c} \cdot (x_{i+1} - x_i)}{2(x_{i+1} - x_i)(x_{i+1} - x_{i-1}) + (x_{i+1} - x_i)^2}. \tag{6.67}$$

Equation (6.66) gives a better estimate, $\mathbf{y}_{i+1} = \check{\mathbf{y}}(x_{i+1})$, of the sought function at the point x_{i+1}. In this context the formula (6.56) serves as *predictor*, and (6.66) as *corrector*. The difference between the two gives an estimate of the error.

This method is equivalent to the two-step method with an extra evaluation where the extra evaluation is done at the full step.

Chapter 7

Numerical integration

7.1 Introduction

The term *numerical integration* refers to a broad family of algorithms to compute a numerical approximation to a definite (Riemann) integral.

Generally, the integral is approximated by a weighted sum of function values within the domain of integration,

$$\int_a^b f(x)dx \approx \sum_{i=1}^n w_i f(x_i) \,. \tag{7.1}$$

Expression (7.1) is often referred to as *quadrature* (*cubature* for multidimensional integrals) or *rule*. The abscissas x_i (also called *nodes*) and the weights w_i of a quadrature are usually optimized—using one of a large number of different strategies—to suit a particular class of integration problems.

For a given numerical integration problem the choice of the quadrature algorithm depends on several factors, in particular on the integrand. Different classes of integrands generally require different quadratures for the most effective calculation.

A popular numerical integration library is QUADPACK [24]. It includes general purpose routines—like QAGS, based on an adaptive Gauss–Kronrod quadrature with acceleration—as well as a number of specialized routines. The GNU scientific library [11] (GSL) implements most of the QUADPACK routines and in addition includes a modern general-purpose adaptive routine CQUAD based on Clenshaw-Curtis quadratures [16].

In the following we shall consider some of the popular numerical integration algorithms.

7.2   Rectangle and trapezium rules

In mathematics, the *Reimann integral* is generally defined in terms of *Riemann sums* [26]. If the integration interval $[a, b]$ is partitioned into n subintervals,

$$a = t_0 < t_1 < t_2 < \cdots < t_n = b \; . \tag{7.2}$$

the Riemann sum is defined as

$$\sum_{i=1}^{n} f(x_i) \Delta x_i \; , \tag{7.3}$$

where $x_i \in [t_{i-1}, t_i]$ and $\Delta x_i = t_i - t_{i-1}$. Geometrically a Riemann sum can be interpreted as the area of a collection of adjacent rectangles with widths Δx_i and heights $f(x_i)$.

The Rieman integral is defined as the limit of a Riemann sum as the *mesh*—the length of the largest subinterval—of the partition approaches zero. Specifically, the number denoted as

$$\int_a^b f(x) dx \tag{7.4}$$

is called the Riemann integral, if for any $\epsilon > 0$ there exists $\delta > 0$ such that for any partition (7.2) with $\max \Delta x_i < \delta$ we have

$$\left| \sum_{i=1}^{n} f(x_i) \Delta x_i - \int_a^b f(x) dx \right| < \epsilon \; . \tag{7.5}$$

A definite integral can be interpreted as the net signed area bounded by the graph of the integrand.

Now, the n-point *rectangle quadrature* is simply the Riemann sum (7.3),

$$\int_a^b f(x) dx \approx \sum_{i=1}^{n} f(x_i) \Delta x_i \; , \tag{7.6}$$

where the node x_i is often (but not always) taken in the middle of the corresponding subinterval, $x_i = t_{i-1} + \frac{1}{2}\Delta x_i$, and the subintervals are often (but not always) chosen equal, $\Delta x_i = (b - a)/n$. Geometrically the n-point rectangle rule is an approximation to the integral given by the area of a collection of n adjacent equal rectangles whose heights are determined by the values of the function (at the middle of the rectangle).

An n-point *trapezium rule* uses instead a collection of trapezia fitted under the graph,

$$\int_a^b f(x) dx \approx \sum_{i=1}^{n} \frac{f(t_{i-1}) + f(t_i)}{2} \Delta x_i \; . \tag{7.7}$$

Importantly, the trapezium rule is the average of two Riemann sums,

$$\sum_{i=1}^{n} \frac{f(t_{i-1}) + f(t_i)}{2} \Delta x_i = \frac{1}{2} \sum_{i=1}^{n} f(t_{i-1}) \Delta x_i + \frac{1}{2} \sum_{i=1}^{n} f(t_i) \Delta x_i \ . \tag{7.8}$$

Rectangle and trapezium quadratures both have the important feature of closely following the very mathematical definition of the integral as the limit of the Riemann sums. Therefore—disregarding the round-off errors—these two rules cannot fail if the integral exists.

For certain partitions of the interval the rectangle and trapezium rules coincide. For example, for the nodes

$$x_i = a + (b - a) \frac{i - \frac{1}{2}}{n} \ , \quad i = 1, \dots, n \tag{7.9}$$

both rules give the same quadrature with equal weights, $w_i = (b - a)/n$,

$$\int_a^b f(x) dx \approx \frac{b - a}{n} \sum_{i=1}^{n} f\left(a + (b - a) \frac{i - \frac{1}{2}}{n}\right) \ . \tag{7.10}$$

Rectangle and trapezium quadratures are rarely used on their own—because of the slow convergence—but they often serve as the basis for more advanced quadratures, for example adaptive quadratures and variable transformation quadratures considered below.

7.3 Quadratures with regularly spaced abscissas

A quadrature (7.1) with n predefined nodes x_i has n free parameters: the weights w_i. A set of n parameters can generally be tuned to satisfy n conditions. The archetypal set of conditions in quadratures is that the quadrature integrates exactly a set of n functions,

$$\{\phi_1(x), \dots, \phi_n(x)\} \ . \tag{7.11}$$

This leads to a set of n equations,

$$\sum_{i=1}^{n} w_i \phi_k(x_i) = I_k \Big|_{k=1,\dots,n} \ , \tag{7.12}$$

where the integrals

$$I_k \doteq \int_a^b \phi_k(x) dx \tag{7.13}$$

are assumed to be known. Equations (7.12) are linear in w_i and can be easily solved.

Since integration is a linear operation, the quadrature will then also integrate exactly any linear combination of functions (7.11).

A popular choice for predefined nodes is a *closed* set—that is, including the end-points of the interval—of evenly spaced abscissas,

$$x_i = a + \frac{i-1}{n-1}(b-a)\Big|_{i=1,\dots,n}. \tag{7.14}$$

However, in practice it often happens that the integrand has an integrable singularity at one or both ends of the interval. In this case one can choose an *open set* of equidistant nodes,

$$x_i = a + \frac{i-\frac{1}{2}}{n}(b-a)\Big|_{i=1,\dots,n}. \tag{7.15}$$

The set of functions to be integrated exactly is generally chosen to suite the properties of the integrands at hand: the integrands must be well represented by linear combinations of the chosen functions.

7.3.1   Classical quadratures

Suppose the integrand can be well represented by the first few terms of its Taylor series,

$$f(x) = \sum_{k=0}^{\infty} \frac{f^{(k)}(a)}{k!}(x-a)^k, \tag{7.16}$$

where $f^{(k)}$ is the k-th derivative of the integrand. This is often the case for analytic—that is, infinitely differentiable—functions. For such integrands one can obviously choose polynomials

$$\{1, x, x^2, \dots, x^{n-1}\} \tag{7.17}$$

as the set of functions to be integrated exactly.

This leads to the so called *classical quadratures*: quadratures with regularly spaced abscissas and *polynomials* as exactly integrable functions.

An n-point classical quadrature integrates exactly the first n terms of the function's Taylor expansion (7.16). The x^n order term will not be integrated exactly and will lead to an error of the quadrature. Thus the error E_n of the n-point classical quadrature is on the order of the integral of the x^n term in (7.16),

$$E_n \approx \int_a^b \frac{f^{(n)}(a)}{n!}(x-a)^n dx = \frac{f^{(n)}(a)}{(n+1)!}h^{n+1} \propto h^{n+1}, \tag{7.18}$$

where $h = b - a$ is the length of the integration interval. A quadrature with the error of the order h^{n+1} is often called a *degree-n* quadrature.

Table 7.1: Maxima script to calculate analytically the weights of an n-point classical quadrature with predefined abscissas in the interval $[0, 1]$.

```
n: 8; xs: makelist((i-1)/(n-1),i,1,n); /* nodes: adapt to your needs */
ws: makelist(concat(w,i),i,1,n);
ps: makelist(x^i,i,0,n-1); /* polynomials */
fs: makelist(buildq([i:i,ps:ps],lambda([x],ps[i])),i,1,n);
integ01: lambda([f],integrate(f(x),x,0,1));
Is: maplist(integ01,fs); /* calculate the integrals */
eq: lambda([f],lreduce("+",maplist(f,xs)*ws));
eqs: maplist(eq,fs)-Is; /* build equations */
solve(eqs,ws); /* solve for the weights */
```

If the integrand is smooth enough and the length h is small enough a classical quadrature with not so large n can provide a good approximation for the integral. However, for large n the weights of classical quadratures tend to have alternating signs, which leads to large round-off errors, which in turn negates the potentially higher accuracy of the quadrature. Again, if the integrand violates the assumption of Taylor expansion—for example by having an integrable singularity inside the integration interval—the higher order quadratures may perform poorly.

Classical quadratures are mostly of historical interest nowadays. Alternative methods—such as quadratures with optimized abscissas, adaptive, and variable transformation quadratures—are more stable and accurate and are normally preferred to classical quadratures.

Classical quadratures with equally spaced abscissas—both closed and open sets—are generally referred to as *Newton-Cotes quadratures*. An interested reader can generate Newton–-Cotes quadratures of any degree n using the Maxima script in Table (7.1).

7.4 Quadratures with optimized abscissas

In quadratures with optimized abscissas not only the weights w_i but also the abscissas x_i are chosen optimally. The number of free parameters is thus $2n$ and one can choose a set of $2n$ functions,

$$\{\phi_1(x), \ldots, \phi_{2n}(x)\} , \tag{7.19}$$

to be integrated exactly. This gives a system of $2n$ equations, linear in w_i and non-linear in x_i,

$$\sum_{i=1}^{n} w_i \phi_k(x_i) = I_k \Big|_{k=1,\ldots,2n} , \tag{7.20}$$

where again

$$I_k \doteq \int_a^b \phi_k(x)dx \ . \tag{7.21}$$

The weights and abscissas of the quadrature can be determined by solving this system of equations[1].

Although quadratures with optimized abcissas are generally of much higher order, $2n - 1$ compared to $n - 1$ for non-optimal abscissas, the optimal points generally can not be reused at the next iteration in an adaptive algorithm.

7.4.1   Gauss quadratures

Gauss quadratures deal with a slightly more general form of integrals,

$$\int_a^b \omega(x)f(x)dx \ , \tag{7.23}$$

where $\omega(x)$ is a positive weight function. For $\omega(x) = 1$ the problem is the same as considered above. Popular choices of the weight function include $\omega(x) = (1 - x^2)^{\pm 1/2}$, $\exp(-x)$, $\exp(-x^2)$ and others. The idea is to represent the integrand as a product $\omega(x)f(x)$ such that all the difficulties go into the weight function $\omega(x)$ while the remaining factor $f(x)$ is smooth and well represented by polynomials.

An N-point *Gauss quadrature* is a quadrature with optimized abcissas,

$$\int_a^b \omega(x)f(x)dx \approx \sum_{i=1}^N w_i f(x_i) \ , \tag{7.24}$$

which integrates exactly a set of $2N$ polynomials of the orders $1, \ldots, 2N - 1$ with the given weight $\omega(x)$.

Fundamental theorem

There is a theorem stating that there exists a set of polynomials $p_n(x)$, orthogonal on the interval $[a, b]$ with the weight function $\omega(x)$,

$$\int_a^b \omega(x)p_n(x)p_k(x) \propto \delta_{nk} \ . \tag{7.25}$$

---

[1]Here is, for example, an $n = 2$ quadrature with optimized abscissas,

$$\int_{-1}^1 f(x)dx \approx f\left(-\sqrt{\tfrac{1}{3}}\right) + f\left(+\sqrt{\tfrac{1}{3}}\right) \ . \tag{7.22}$$

Now, one can prove that the optimal nodes for the N-point Gauss quadrature are the roots of the polynomial $p_N(x)$,

$$p_N(x_i) = 0 \ . \tag{7.26}$$

The idea behind the proof is to consider the integral

$$\int_a^b \omega(x)q(x)p_N(x)dx = 0 \ , \tag{7.27}$$

where $q(x)$ is an arbitrary polynomial of degree less than N. The quadrature should represent this integral exactly,

$$\sum_{i=1}^N w_i q(x_i)p_N(x_i) = 0 \ . \tag{7.28}$$

Apparently this is only possible if x_i are the roots of p_N.

Calculation of nodes and weights

A neat algorithm—usually refered to as Golub-Welsch algorithm [15]—for calculation of the nodes and weights of a Gauss quadrature is based on the symmetric form of the three-term reccurence relation for orthogonal polynomials,

$$xp_{n-1}(x) = \beta_n p_n(x) + \alpha_n p_{n-1}(x) + \beta_{n-1}p_{n-2}(x) \ , \tag{7.29}$$

where $p_{-1}(x) \doteq 0$, $p_1(x) \doteq 1$, and $n = 1,\ldots,N$. This reccurence relation can be written in the matrix form,

$$x\mathbf{p}(x) = \mathbf{J}\mathbf{p}(x) + \beta_N p_N(x)\mathbf{e}_N \ , \tag{7.30}$$

where $\mathbf{p}(x) \doteq \{p_0(x),\ldots,p_{N-1}(x)\}^T$, $\mathbf{e}_N = \{0,\ldots,0,1\}^T$, and the tridiagonal matrix $\mathbf{J}$ — usually refered to as *Jacobi matrix* or *Jacobi operator* — is given as

$$\mathbf{J} = \begin{pmatrix} \alpha_1 & \beta_1 & & & \\ \beta_1 & \alpha_2 & \beta_2 & & \\ & \beta_2 & \alpha_3 & \beta_3 & \\ & & \ddots & \ddots & \\ & & & \beta_{N-1} & \alpha_N \end{pmatrix} \ . \tag{7.31}$$

Substituting the roots x_i of p_N — that is, the set $\{x_i \mid p_N(x_i) = 0\}$ — into the matrix equation (7.30) leads to eigenvalue problem for the Jacobi matrix,

$$\mathbf{J}\mathbf{p}(x_i) = x_i\mathbf{p}(x_i) \,. \tag{7.32}$$

Thus, the nodes of an N-point Gauss quadrature (the roots of the polynomial p_N) are the eigenvalues of the Jacobi matrix J and can be calculated by a standard diagonalization[2] routine.

The weights can be obtained considering N integrals,

$$\int_a^b \omega(x)p_n(x)dx = \delta_{n0}\int_a^b \omega(x)dx \,, \ n = 0, \ldots, N-1 \,. \tag{7.33}$$

Applying our quadrature gives the matrix equation,

$$\mathrm{P}\mathbf{w} = \mathbf{e}_1 \int_a^b \omega(x)dx \,, \tag{7.34}$$

where $\mathbf{w} \doteq \{w_1, \ldots, w_N\}^T$, $\mathbf{e}_1 = \{1, 0, \ldots, 0\}^T$, and

$$\mathrm{P} \doteq \begin{pmatrix} p_0(x_1) & \ldots & p_0(x_N) \\ p_1(x_1) & \ldots & p_1(x_N) \\ \ldots & \ldots & \ldots \\ p_{N-1}(x_1) & \ldots & p_{N-1}(x_N) \end{pmatrix} \,. \tag{7.35}$$

Equation (7.34) is linear in w_i and can be solved directly. However, if diagonalization of the Jacobi matrix provided the normalized eigenvectors, the weigths can be readily obtained using the following method.

The matrix P apparently consists of non-normalized column eigenvectors of the matrix J. The eigenvectors are orthogonal and therefore $\mathrm{P}^T\mathrm{P}$ is a diagonal matrix with positive elements. Multiplying (7.34) by P^T and then by $(\mathrm{P}^T\mathrm{P})^{-1}$ from the left gives

$$\mathbf{w} = (\mathrm{P}^T\mathrm{P})^{-1}\mathrm{P}^T\mathbf{e}_1 \int_a^b \omega(x)dx \,. \tag{7.36}$$

From $p_0(x) = 1$ it follows that $\mathrm{P}^T\mathbf{e}_1 = \{1, \ldots, 1\}^T$ and therefore

$$w_i = \frac{1}{(\mathrm{P}^T\mathrm{P})_{ii}} \int_a^b \omega(x)dx \,. \tag{7.37}$$

---

[2]A symmetric tridiagonal matrix can be diagonalized very effectively using the QR/RL algorithm.

Let the matrix V be the set of the normalized column eigenvectors of the matrix J. The matrix V is then connected with the matrix P through the normalization equation,

$$V = \sqrt{(P^T P)^{-1}} P \,. \tag{7.38}$$

Therefore, again taking into account that $p_0(x) = 1$, equation (7.37) can be written as

$$w_i = (V_{1i})^2 \int_a^b \omega(x) dx \,. \tag{7.39}$$

Example: Gauss-Legendre quadrature

Gauss-Legendre quadrature deals with the weight $\omega(x) = 1$ on the interval $[-1, 1]$. The associated polynomials are Legendre polynomials $\mathcal{P}_n(x)$, hence the name. Their reccurence relation is usually given as

$$(2n - 1) x \mathcal{P}_{n-1}(x) = n \mathcal{P}_n(x) + (n - 1) \mathcal{P}_{n-2}(x) \,. \tag{7.40}$$

Rescaling the polynomials (preserving $p_0(x) = 1$) as

$$\sqrt{2n + 1} \mathcal{P}_n(x) = p_n(x) \tag{7.41}$$

reduces this reccurence relation to the symmetric form (7.29),

$$x p_{n-1}(x) = \frac{1}{2} \frac{1}{\sqrt{1 - (2n)^{-2}}} p_n(x) + \frac{1}{2} \frac{1}{\sqrt{1 - (2(n-1))^{-2}}} p_{n-2}(x) \,. \tag{7.42}$$

Correspondingly, the coefficients in the matrix J are

$$\alpha_n = 0 \,, \ \beta_n = \frac{1}{2} \frac{1}{\sqrt{1 - (2n)^{-2}}} \,. \tag{7.43}$$

The problem of finding the nodes and the weights of the N-point Gauss-Legendre quadrature is thus reduced to the eigenvalue problem for the Jacobi matrix with coefficients (7.43).

As an illustration of this algorithm Table (7.2) shows an Octave function which calculates the nodes and the weights of the N-point Gauss-Legendre quadrature and then integrates a given function.

7.4.2 Gauss-Kronrod quadratures

Generally, the error of a numerical integration is estimated by comparing the results from two rules of different orders. However, for ordinary Gauss quadratures the nodes

Table 7.2: An Octave function that calculates the nodes and weights of the N-point Gauss-Legendre quadrature and then integrates a given function.

```
function Q = gauss_legendre(f,a,b,N)
beta = .5./sqrt(1-(2*(1:N-1)).^(-2));  % reccurence relation
J = diag(beta,1) + diag(beta,-1);      % Jacobi matrix
[V,D] = eig(J);                        % diagonalization of J
x = diag(D);  [x,i] = sort(x);         % sorted nodes
w = V(1,i).^2*2;                       % weights
Q = w*f((a+b)/2+(b-a)/2*x)*(b-a)/2;    % integral
endfunction;
```

for two rules of different orders almost never coinside. This means that one can not reuse the points of the lower order rule when calculating the hihger order rule.

Gauss-Kronrod algorithm [20] remedies this inefficiency. The points inherited from the lower order rule are reused in the higher order rule as predefined nodes (with n weights as free parameters), and then m more optimal points are added (m abscissas and m weights as free parameters). The order of the method is $n + 2m - 1$. The lower order rule becomes *embedded*—that is, it uses a subset of the nodes—into the higher order rule. On the next iteration the procedure is repeated.

Patterson [23] has tabulated nodes and weigths for several sequences of embedded Gauss-Kronrod rules.

7.5   Adaptive quadratures

Higher order quadratures suffer from round-off errors as the weights w_i generally have alternating signs. Again, using high order polynomials is dangerous as they typically oscillate wildly and may lead to Runge's phenomenon. Therefore, if the error of the quadrature is yet too large for a quadrature with sufficiently large n, the best strategy is to subdivide the interval in two and then use the quadrature on the half-intervals. Indeed, if the error is of the order h^k, the subdivision would lead to reduced error, $2\left(h/2\right)^k < h^k$, if $k > 1$.

An *adaptive quadrature* is an algorithm where the integration interval is subdivided into adaptively refined subintervals until the given accuracy goal is reached.

Adaptive algorithms are usually built on pairs of quadrature rules – a higher order rule,

$$Q = \sum_i w_i f(x_i), \tag{7.44}$$

where w_i are the weights of the higher order rule and Q is the higher order estimate of

the integral, and a lower order rule,

$$q = \sum_i v_i f(x_i), \tag{7.45}$$

where v_i are the weights of the lower order rule and q is the the lower order estimate of the integral. The difference between the higher order rule and the lower order rule gives an estimate of the error,

$$\delta Q = |Q - q| . \tag{7.46}$$

The integration result is accepted, if the error δQ is smaller than tolerance,

$$\delta Q < \delta + \epsilon |Q| , \tag{7.47}$$

where δ is the absolute accuracy goal and ϵ is the relative accuracy goal of the integration.

If the error estimate is larger than tolerance, the interval is subdivided into two half-intervals and the procedure applies recursively to subintervals with the same relative accuracy goal ϵ and rescaled absolute accuracy goal $\delta/\sqrt{2}$.

The points x_i are usually chosen such that the two quadratures use the same points, and that the points can be reused in the subsequent recursive steps. The reuse of the function evaluations made at the previous step of adaptive integration is very important for the efficiency of the algorithm. The equally-spaced abscissas naturally provide for such a reuse.

As an example, Table 7.3 shows an implementation of the described algorithm using

$$x_i \;=\; \left\{ \frac{1}{6}, \frac{2}{6}, \frac{4}{6}, \frac{5}{6} \right\} \text{ (easily reusable points)} , \tag{7.48}$$

$$w_i \;=\; \left\{ \frac{2}{6}, \frac{1}{6}, \frac{1}{6}, \frac{2}{6} \right\} \text{ (trapezium rule)} , \tag{7.49}$$

$$v_i \;=\; \left\{ \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4} \right\} \text{ (rectangle rule)} . \tag{7.50}$$

During recursion the function values at the points #2 and #3 are inherited from the previous step and need not to be recalculated.

The points and weights are cited for the rescaled integration interval $[0, 1]$. The transformation of the points and weights to the original interval $[a, b]$ is given as

$$\begin{aligned} x_i &\;\rightarrow\; a + (b - a)x_i , \\ w_i &\;\rightarrow\; (b - a)w_i . \end{aligned} \tag{7.51}$$

This implementation calculates directly the Riemann sums and can therefore deal with integrable singularities, although rather inefficiently.

More efficient adaptive routines keep track of the subdivisions of the interval and the local errors [16]. This allows detection of singularities and switching in their vicinity to specifically tuned quadratures. It also allows better estimates of local and global errors.

Here is an embedded 8-point open quadrature,

$$x_i = \left\{ \frac{1}{12}, \frac{2}{12}, \frac{4}{12}, \frac{5}{12}, \frac{7}{12}, \frac{8}{12}, \frac{10}{12}, \frac{11}{12} \right\},$$ (7.52)

$$w_i = \left\{ \frac{4738}{19845}, \frac{-59}{567}, \frac{5869}{13230}, \frac{-74}{945}, w_4, w_3, w_2, w_1 \right\},$$ (7.53)

$$v_i = \left\{ \frac{208}{945}, \frac{-7}{135}, \frac{209}{630}, 0, v_4, v_3, v_2, v_1 \right\}.$$ (7.54)

Table 7.3: Recursive adaptive integrator in Python

```
def adapt(f,a,b,acc=0.01,eps=0.01,f2=math.nan,f3=math.nan) :
    f1 = f(a+(b-a)/6)
    f4 = f(a+5*(b-a)/6);
    if math.isnan(f2) : # first call
        f2 = f(a+2*(b-a)/6)
        f3 = f(a+4*(b-a)/6)
    Q = (2*f1+f2+f3+2*f4)/6*(b-a)
    q = (f1+f4+f2+f3)/4*(b-a);
    tolerance = acc+eps*abs(Q)
    error = abs(Q-q);
    if error < tolerance : return Q;
    else :
        Q1 = adapt(f,a,(a+b)/2,acc/math.sqrt(2.),eps,f1,f2)
        Q2 = adapt(f,(a+b)/2,b,acc/math.sqrt(2.),eps,f3,f4)
        return Q1+Q2
```

7.6   Variable transformation quadratures

The idea behind *variable transformation quadratures* is to apply the given quadrature—either with optimimized or regularly spaced nodes—not to the original integral, but to a variable transformed integral [22],

$$\int_a^b f(x)dx = \int_{t_a}^{t_b} f\big(g(t)\big)g'(t)dt \approx \sum_{i=1}^N w_i f\big(g(t_i)\big)g'(t_i),$$ (7.55)

where the transformation $x = g(t)$ is chosen such that the transformed integral better suits the given quadrature. Here g' denotes the derivative and $[t_a, t_b]$ is the corresponding interval in the new variable.

For example, the Gauss-Legendre quadrature assumes the integrand can be well represented with polynomials and performs poorly on integrals with integrable singularities like

$$I = \int_0^1 \frac{1}{2\sqrt{x}} dx \, . \tag{7.56}$$

However, a simple varibale transformation $x = t^2$ removes the singularity,

$$I = \int_0^1 dt \, , \tag{7.57}$$

and the Gauss-Legendre quadrature for the transformed integral gives exact result. The price is that the transformed quadrature performs less effectively on smooth functions.

Some of the popular variable transformation quadratures are Clenshaw-Curtis [8], based on the transformation

$$\int_{-1}^1 f(x)dx = \int_0^\pi f(\cos\theta) \sin\theta d\theta \, , \tag{7.58}$$

and "tanh-sinh" quadrature [22], based on the transformation

$$\int_{-1}^1 f(x)dx = \int_{-\infty}^\infty f\left(\tanh\left(\frac{\pi}{2}\sinh(t)\right)\right) \frac{\pi}{2} \frac{\cosh(t)}{\cosh^2\left(\frac{\pi}{2}\sinh(t)\right)} dt \, . \tag{7.59}$$

and another one,

$$\int_0^1 f(x)dx = \int_0^1 f\left(3t^2 - 2t^3\right) 6(t - t^2)dt \, . \tag{7.60}$$

Generally, the equally spaced trapezium/rectangle rule is used after the transformation.

7.7 Infinite intervals

One way to calculate an integral over infinite interval is to transform it by a variable sustitution into an integral over a finite interval. The latter can then be evaluated by ordinary integration methods. Table 7.4 lists several of such transformation.

Table 7.4: Variable transformations reducing infinite interval integrals into integrals over finite intervals.

$$\int_{-\infty}^{+\infty} f(x)dx = \int_{-1}^{+1} f\left(\frac{t}{1-t^2}\right) \frac{1+t^2}{(1-t^2)^2} dt \ , \tag{7.61}$$

$$\int_{-\infty}^{+\infty} f(x)dx = \int_{0}^{1} \left(f\left(\frac{1-t}{t}\right) + f\left(-\frac{1-t}{t}\right)\right) \frac{dt}{t^2} \ , \tag{7.62}$$

$$\int_{a}^{+\infty} f(x)dx = \int_{0}^{1} f\left(a + \frac{t}{1-t}\right) \frac{1}{(1-t)^2} dt \ , \tag{7.63}$$

$$\int_{a}^{+\infty} f(x)dx = \int_{0}^{1} f\left(a + \frac{1-t}{t}\right) \frac{dt}{t^2} \ , \tag{7.64}$$

$$\int_{-\infty}^{b} f(x)dx = \int_{-1}^{0} f\left(b + \frac{t}{1+t}\right) \frac{1}{(1+t)^2} dt \ , \tag{7.65}$$

$$\int_{-\infty}^{b} f(x)dx = \int_{0}^{1} f\left(b - \frac{1-t}{t}\right) \frac{dt}{t^2} \ . \tag{7.66}$$

Chapter 8

Monte Carlo integration

8.1 Introduction

Monte Carlo integration is a quadrature (cubature) where the nodes are chosen randomly [28]. Typically no assumption is made about the smoothness of the integrand, not even that it is continuous.

Monte Carlo algorithms are particularly suited for multi-dimensional integrations where one of the problems is that the integration region, Ω, might have a quite complicated boundary which can not be easily described by simple functions. On the other hand, it is usually much easier to find out whether a given point lies within the integration region or not. Therefore a popular strategy is to create an auxiliary rectangular volume, V, which encompasses the integration volume Ω, and an auxiliary function which coincides with the integrand inside the volume Ω and is equal zero outside. Then the integral of the auxiliary function over the auxiliary volume is equal the original integral.

However, the auxiliary function is generally non-continuous at the boundary; thus ordinary quadratures—that assume continuity of the integrand—are bound to have difficulties here. One the contrary the Monte-Carlo quadratures will do just as good (or as bad) as with continuous integrands.

A typical implementation of a Monte Carlo algorithm integrates the given function over a rectangular volume, specified by the coordinates of its "lower-left" and "upper-right" vertices, assuming the user has provided the encompassing volume with the auxiliary function.

Plain Monte Carlo algorithm distributes points uniformly throughout the integration region using uncorrelated pseudo-random sequences of points.

Adaptive algorithms, such as VEGAS and MISER, distribute points non-uniformly

in an attempt to reduce integration error using correspondingly *importance* and *stratified* sampling.

Yet another strategy to reduce the error is to use correlated quasi-random sequences.

The GNU Scientific Library, GSL, implements a plain Monte Carlo integration algorithm; a stratified sampling algorithm, MISER; an importance sampling algorithm, VEGAS; and a number of quasi-random generators.

8.2   Plain Monte Carlo sampling

Plain Monte Carlo is a quadrature with random abscissas and equal weights,

$$\int_V f(\mathbf{x})dV \approx w \sum_{i=1}^{N} f(\mathbf{x}_i) \, , \tag{8.1}$$

where $\mathbf{x}$ is a point in the multi-dimensional integration space. One free parameter, w, allows one condition to be satisfied: the quadrature must integrate exactly a constant function. This gives $w = V/N$,

$$\int_V f(\mathbf{x})dV \approx \frac{V}{N} \sum_{i=1}^{N} f(\mathbf{x}_i) \doteq V\langle f \rangle \, . \tag{8.2}$$

Under the assumptions of the *central limit theorem* the error of the integration can be estimated as

$$\text{error} = V \frac{\sigma}{\sqrt{N}} \, , \tag{8.3}$$

where σ is the variance of the sample,

$$\sigma^2 = \langle f^2 \rangle - \langle f \rangle^2 \, . \tag{8.4}$$

The familiar $1/\sqrt{N}$ convergence of a random walk process is quite slow: to reduce the error by a factor 10 requires 100-fold increase in the number of sample points.

Expression (8.3) provides only a statistical estimate of the error, which is not a strict bound; random sampling may not uncover all the important features of the function, resulting in an underestimate of the error.

A simple implementation of the plain Monte Carlo algorithm is shown in Table 8.1.

Table 8.1: Plain Monte Carlo integrator

```
def plainmc(f, a, b, N):
    dim=len(a)
    V=1; for i in range(dim): V*=a[i]-b[i]
    sum=0; sum2=0
    for i in range(N):
        x=[random.uniform(a[i],b[i]) for i in range(dim)]
        fx=f(x)
        sum+=fx; sum2+=fx*fx
    mean = sum/N, variance2 = sum2/N-mean*mean
    return (mean*V, sqrt(variance2/N)*V;
```

8.3   Importance sampling

Suppose the points are distributed not uniformly but with some density $\rho(\mathbf{x})$. That is, the number of points Δn in the volume ΔV around point $\mathbf{x}$ is given as

$$\Delta n = \frac{N}{V}\rho(\mathbf{x})\Delta V, \tag{8.5}$$

where ρ is normalised such that $\int_V \rho dV = V$.
    The estimate of the integral is then given as

$$\int_V f(\mathbf{x})dV \approx \sum_{i=1}^{N} f(\mathbf{x}_i)\Delta V_i = \sum_{i=1}^{N} f(\mathbf{x}_i)\frac{V}{N\rho(\mathbf{x}_i)} = V\left\langle \frac{f}{\rho} \right\rangle , \tag{8.6}$$

where

$$\Delta V_i = \frac{V}{N\rho(x_i)} \tag{8.7}$$

is the volume-per-point at the point x_i.
    The corresponding variance is now given by

$$\sigma^2 = \left\langle \left(\frac{f}{\rho}\right)^2 \right\rangle - \left\langle \frac{f}{\rho} \right\rangle^2 . \tag{8.8}$$

Apparently if the ratio f/ρ is close to a constant, the variance is reduced.
    It is tempting to take $\rho = |f|$ and sample directly from the integrand. However in practice evaluations of the integrand are typically expensive. Therefore a better strategy is to build an approximate density in the product form, $\rho(x, y, \ldots, z) = \rho_x(x)\rho_y(y)\ldots\rho_z(z)$, and then sample from this approximate density. A popular routine of this sort is called VEGAS.

8.4   Stratified sampling

Stratified sampling is a generalisation of the recursive adaptive integration algorithm to random quadratures in multi-dimensional spaces.

Table 8.2: Recursive stratified sampling algorithm

```
sample N random points with plain Monte Carlo;
estimate the average and the error;
IF the error is acceptable :
   RETURN the average and the error;
ELSE :
   FOR EACH dimension :
      subdivide the volume in two along the dimension;
      estimate the sub—variances in the two sub—volumes;
   pick the dimension with the largest sub—variance;
   subdivide the volume in two along this dimension;
   dispatch two recursive calls to each of the sub—volumes;
   estimate the grand average and grand error;
   RETURN the grand average and grand error;
```

The ordinary "dividing by two" strategy does not work for multi-dimensional integrations as the number of sub-volumes grows way too fast to keep track of. Instead one estimates along which dimension a subdivision should bring the most dividends and only subdivides along this dimension. Such strategy is called *recursive stratified sampling*. A simple variant of this algorithm is presented in Table 8.2.

In a stratified sample the points are concentrated in the regions where the variance of the function is largest, see the illustration in Figure 8.1.

The naive implementation in Table 8.2 keeps throwing points until the given tolerance is achieved. This might be problematic for multidimensional integrals where for the given tolerance one might need excessively many points. The GSL's implementation instead takes the number of points as the argument and then returns the estimate of the integral and the estimate of the error.

8.5   Quasi-random (low-discrepancy) sampling

Pseudo-random sampling has high discrepancy[1]: it typically creates regions with high density of points and other regions with low density of points, see an illustration on Figure 8.2 (left). With pseudo-random sampling there is a finite probability that all the N points would fall into one half of the region and none into the other half.

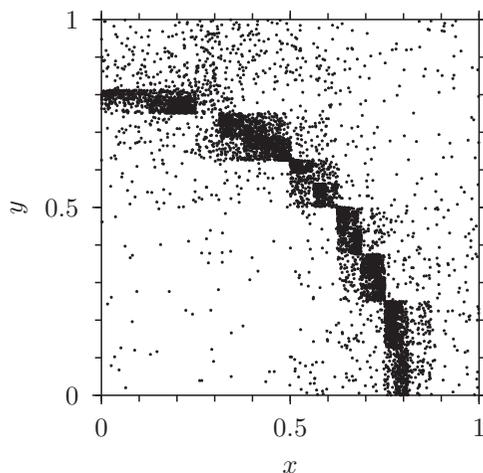[1]discrepancy is a measure of how unevenly the points are distributed over the region.

Figure 8.1: Stratified sample of a discontinuous function, $f(x, y) = 1$ if $x^2 + y^2 < 0.8^2$ otherwise $f(x, y) = 0$, built with the algorithm in Table 8.2.

Quasi-random sequences avoid this phenomenon by distributing points in a highly correlated manner with a specific requirement of low discrepancy, see Figure 8.2 for an example. Quasi-random sampling is like a computation on a grid where the grid constant must not be known in advance as the grid is ever gradually refined and the points are always distributed uniformly over the region. The computation can be stopped at any time.

By placing points more evenly than at random, the quasi-random sequences try to improve on the $1/\sqrt{N}$ convergence rate of pseudo-random sampling.

The central limit theorem does not apply in this case as the points are not statistically independent. Therefore the variance can not be used as an estimate of the error. The error estimation is actually not trivial. In practice one can employ two different sequences and use the difference in the resulting integrals as an error estimate.

8.5.1 Van der Corput and Halton sequences

A *van der Corput sequence* is a low-discrepancy sequence over the unit interval. It is constructed by reversing the base-b representation of the sequence of natural numbers $(1, 2, 3, \dots)$. For example, the decimal van der Corput sequence begins as

$$0.1, 0.2, 0.3, \dots, 0.8, 0.9, 0.01, 0.11, 0.21, 0.31, \dots, 0.91, 0.02, 0.12, \dots . \tag{8.9}$$

In a base-b representation a natural number n with s digits $\{d_i \mid i = 1 \dots s, \ 0 \le$
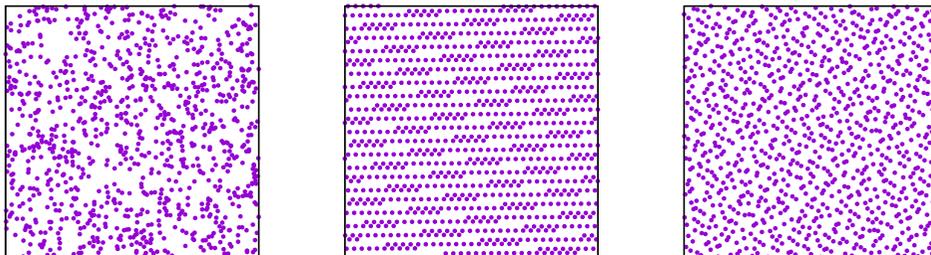
Figure 8.2: Typical distributions of pseudo-random points (left), and quasi-random low-discrepancy points: additive (center) and base-2/3 Halton (right) sequences. The first thousand points are plotted in each case.

$d_i < b\}$ is given as

$$n = \sum_{k=1}^{s} d_k b^{k-1} \ . \tag{8.10}$$

The corresponding base-b van der Corput number $q_b(n)$ is then given as

$$q_b(n) = \sum_{k=1}^{s} d_k b^{-k} \ . \tag{8.11}$$

Here is a Python implementation of this algorithm,

```
def corput(n,b=2):
    q=0; bk=1/b
    while(n>0): q += n % b * bk; n //= b; bk /= b;
    return q
```

The van der Corput numbers of any base are uniformly distributed over the unit interval. They also form a dense set in the unit interval: there exists a subsequence of the van der Corput sequence which converges to any given real number in $[0, 1]$.

The Halton sequence is a generalization of the van der Corput sequence to d-dimensional spaces. One chooses a set of coprime bases $b_1, \ldots, b_d$ and then for each dimension i generates a van der Corput sequence with its own base b_i. The n-th Halton d-dimentional point $\mathbf{x}$ in the unit volume is then given as

$$\mathbf{x}_{b_1,\ldots,b_d}(n) = \{q_{b_1}(n), \ldots, q_{b_d}(n)\} \ . \tag{8.12}$$

Here is a Python implementation which calls the `corput` function listed above,

```
class Halton:
    def __init__(self,dim) : self.bases=prime_numbers(dim)
    def x(self,n) : return [corput(n,b) for b in self.bases]
```

where a (short) list of n prime numbers can be generated as

```
def prime_numbers(n):
    primes=[]; candidate=2
    while len(primes)<n:
        candidate_is_prime=True
        for p in primes:
            if p*p>candidate: break
            if candidate%p==0 : candidate_is_prime=False; break
        if candidate_is_prime: primes.append(candidate)
        candidate+=1
    return primes
```

8.5.2 Additive recurrences (lattice rules)

For any irrational number α the sequence

$$s_n(\alpha) \doteq \text{frac}(n\alpha) \tag{8.13}$$

(where $\text{frac}(x)$ is the fractional part of x) has discrepancy approaching $1/N$. This sequence can be also defined as an additive recurrence relation,

$$s_{n+1}(\alpha) = \text{frac}(s_n + \alpha) \equiv (s_n + \alpha) \mod 1 . \tag{8.14}$$

In a d-dimensional space one chooses for each dimension a separate sequence with its own irrational number $\alpha_i|_{i=1,...,d}$ (one possibility is to take the square roots of prime numbers modulo 1). The corresponding quasi-random vector $\mathbf{x}$ is then given as

$$\mathbf{x}(n) = \{\text{frac}(n\alpha_1), \ldots, \text{frac}(n\alpha_d)\} . \tag{8.15}$$

This formula is sometimes also referred to as *lattice rule*.

Here is an implementation of this algorithm in Python

```
class Lattice:
    def __init__(self,dim):
        self.alphas=[math.sqrt(p) % 1 for p in prime_numbers(dim)]
    def x(self,n): return [n*alpha % 1 for alpha in self.alphas]
```

and an illustration of such sequence is shown on Figure 8.2 (center).

8.6 Implementations

Table 8.3: C implementation of the stratified sampling algorithm

```c
#include<math.h>
#include<stdlib.h>
#include<stdio.h>
#define FOR(k) for(int k=0;k<dim;k++)
#define RND (double)rand()/RAND_MAX

double strata(
    int dim, double f(int dim,double*x),
    double*a,double*b,
    double acc,double eps,
    int n_reuse,double mean_reuse)
{
int N=16*dim;
double V=1; FOR(k) V*=b[k]-a[k];
int n_left[dim], n_right[dim];
double x[dim], mean_left[dim], mean_right[dim],mean=0;
FOR(k){ mean_left[k]=0; mean_right[k]=0; n_left[k]=0; n_right[k]=0; }
for(int i=0;i<N;i++){
    FOR(k) x[k]=a[k]+RND*(b[k]-a[k]);
    double fx=f(dim,x);
    mean+=fx;
    FOR(k){
        if(x[k]>(a[k]+b[k])/2){ n_right[k]++; mean_right[k]+=fx; }
        else                  { n_left[k]++; mean_left[k]+=fx; }
        }
    }
mean/=N;
FOR(k){ mean_left[k]/=n_left[k]; mean_right[k]/=n_right[k]; }

int kdiv=0; double maxvar=0;
FOR(k){
    double var=fabs(mean_right[k]-mean_left[k]);
    if(var>maxvar){ maxvar=var; kdiv=k; }
    }

double integ=(mean*N+mean_reuse*n_reuse)/(N+n_reuse)*V;
double error=fabs(mean_reuse-mean)*V;
double toler=acc+fabs(integ)*eps;
if(error<toler)return integ;

double a2[dim],b2[dim]; FOR(k)a2[k]=a[k]; FOR(k)b2[k]=b[k];
a2[kdiv]=(a[kdiv]+b[kdiv])/2; b2[kdiv]=(a[kdiv]+b[kdiv])/2;
double integ_left=
    strata(dim,f,a,b2,acc/sqrt(2),eps,n_left[kdiv],mean_left[kdiv]);
double integ_right=
    strata(dim,f,a2,b,acc/sqrt(2),eps,n_right[kdiv],mean_right[kdiv]);
return integ_left+integ_right;
}
```

Chapter 9

Nonlinear equations

9.1 Introduction

Non-linear equations (or *root-finding*) is a problem of finding a set of n variables $\mathbf{x} \doteq \{x_1, \ldots, x_n\}$ which satisfy a system of n non-linear equations

$$f_i(x_1, ..., x_n) = 0 \Big|_{i=1,...,n} . \tag{9.1}$$

In vector notation the system is written as

$$\mathbf{f}(\mathbf{x}) = 0 , \tag{9.2}$$

where $\mathbf{f}(\mathbf{x}) \doteq \{f_1(x_1, \ldots, x_n), \ldots, f_n(x_1, \ldots, x_n)\}$.

In one-dimension, $n = 1$, it is generally possible to plot the function in the region of interest and see whether the graph crosses the x-axis. One can then be sure the root exists and even figure out its approximate position to start one's root-finding algorithm from. In multi-dimensions one generally does not know if the root exists at all, until it is found.

The root-finding algorithms generally proceed by iteration, starting from some approximate solution and making consecutive steps—hopefully in the direction of the suspected root—until some convergence criterion is satisfied. The procedure is generally not even guaranteed to converge unless starting from a point close enough to the sought root.

We shall only consider the multi-dimensional case here since i) the multi-dimensional root-finding is more difficult, and ii) the multi-dimensional routines can also be used in the one-dimensional case.

9.2   Newton's method

Newton's method (also reffered to as Newton-Raphson method, after Isaac Newton and Joseph Raphson) is a root-finding algorithm that uses the first term of the Taylor series of the functions f_i to linearise the system (9.1) in the vicinity of a suspected root. It is one of the oldest and best known methods and is a basis of a number of more refined methods.

Suppose that the point $\mathbf{x} = \{x_1, \ldots, x_n\}$ is close to the root. The Newton's algorithm tries to find the step $\Delta \mathbf{x}$ which would move the point towards the root, such that

$$f_i(\mathbf{x} + \Delta \mathbf{x}) = 0 \Big|_{i=1,\ldots,n} . \tag{9.3}$$

The first order Taylor expansion of (9.3) gives a system of linear equations,

$$f_i(\mathbf{x}) + \sum_{k=1}^{n} \frac{\partial f_i}{\partial x_k} \Delta x_k = 0 \Big|_{i=1,\ldots,n} , \tag{9.4}$$

or, in the matrix form,

$$J\Delta \mathbf{x} = -\mathbf{f}(\mathbf{x}), \tag{9.5}$$

where J is the matrix of partial derivatives,

$$J_{ik} \doteq \frac{\partial f_i}{\partial x_k} , \tag{9.6}$$

called the *Jacobian matrix*. In practice, if derivatives are not available analytically, one uses finite differences,

$$\frac{\partial f_i}{\partial x_k} \approx \frac{f_i(x_1, \ldots, x_k + \delta x_k, \ldots, x_n) - f_i(x_1, \ldots, x_k, \ldots, x_n)}{\delta x_k} , \tag{9.7}$$

where the step δx_k is usually (unless the user knows better) chosen as $\delta x_k = |x_k|\sqrt{\epsilon}$ where ϵ is machine precision. For double-precision numbers $\sqrt{\epsilon} = 2^{-26}$. As a rule of thumb one should always try to rescale one's problem such that the typical scale of one's variables is around unity.

The solution $\Delta \mathbf{x}$ to the linear system (9.5)—called the Newton's step—gives the approximate direction and the approximate step-size towards the solution.

The Newton's method converges quadratically — that is, the error is squared at each step — if $\mathbf{x}$ is sufficiently close to the solution. Otherwise the full Newton's step $\Delta \mathbf{x}$ might actually diverge from the solution. Therefore in practice a more conservative step, $\lambda \Delta \mathbf{x}$ with $\lambda < 1$, is usually taken. The strategy of finding the optimal λ is referred to as *line search*.

It is typically not worth the effort to find λ which minimizes $\|\mathbf{f}(\mathbf{x} + \lambda \Delta \mathbf{x})\|$ exactly, since $\Delta \mathbf{x}$ is only an approximate direction towards the root. Instead, an inexact but quick minimization strategy is usually used, called the *backtracking line search*, where one first attempts the full step, $\lambda = 1$, and then backtracks, $\lambda \leftarrow \lambda/2$, until the condition

$$\|\mathbf{f}(\mathbf{x} + \lambda \Delta \mathbf{x})\| < \left(1 - \frac{\lambda}{2}\right) \|\mathbf{f}(\mathbf{x})\| \tag{9.8}$$

is satisfied. If the condition is not satisfied for sufficiently small $\lambda_{\min}$ the step is taken with $\lambda_{\min}$ simply to step away from the difficult place and try again.

Here is a typical algrorithm for the Newton's method with backtracking line search and condition (9.8),

```
repeat
    calculate the Jacobian matrix J
    solve JΔx = −f(x) for Δx
    λ ← 1
    while ‖f(x + λΔx)‖ > (1 − λ/2) ‖f(x)‖ and λ ≥ 1/128 do λ ← λ/2
    x ← x + λΔx
until converged (e.g. ‖f(x)‖ < tolerance)
```

A somewhat more refined backtracking linesearch is based on an approximate minimization of the function

$$\phi(\lambda) \doteq \frac{1}{2}\|\mathbf{f}(\mathbf{x} + \lambda \Delta \mathbf{x})\|^2 \tag{9.9}$$

using interpolation. The values $\phi(0) = \frac{1}{2}\|\mathbf{f}(\mathbf{x})\|^2$ and $\phi'(0) = -\|\mathbf{f}(\mathbf{x})\|^2$ are already known[1]. If the previous step with certain λ_{trial} was rejected, we also have $\phi(\lambda_{\text{trial}})$. These three quantities allow to build a quadratic approximation,

$$\phi(\lambda) \approx \phi(0) + \phi'(0)\lambda + c\lambda^2 , \tag{9.10}$$

where

$$c = \frac{\phi(\lambda_{\text{trial}}) - \phi(0) - \phi'(0)\lambda_{\text{trial}}}{\lambda_{\text{trial}}^2} . \tag{9.11}$$

The minimum of this approximation (determined by the condition $\phi'(\lambda) = 0$),

$$\lambda_{\text{next}} = -\frac{\phi'(0)}{2c} , \tag{9.12}$$

---

[1]
$$\frac{\partial \phi}{\partial \lambda}|_{\lambda=0} = \sum_i f_i(\mathbf{x} + \lambda \Delta \mathbf{x})\frac{\partial f_i(\mathbf{x} + \lambda \Delta \mathbf{x})}{\partial \lambda}|_{\lambda=0} = \sum_i f_i \sum_j \frac{\partial f_i(\mathbf{x} + \lambda \Delta \mathbf{x})}{\partial x_j}\Delta x_j|_{\lambda=0} = -\sum_i f_i(\mathbf{x})^2$$

becomes the next trial step-size.

The procedure is repeated recursively until either condition (9.8) is satisfied or the step becomes too small (in which case it is taken unconditionally in order to simply get away from the difficult place).

9.3   Quasi-Newton methods

The Newton's method requires calculation of the Jacobian matrix at every iteration. This is generally an expensive operation. *Quasi-Newton* methods avoid calculation of the Jacobian matrix at the new point $\mathbf{x} + \lambda\Delta\mathbf{x}$, instead trying to use certain approximations, typically rank-1 updates.

9.3.1   Broyden's rank-1 updates

Broyden's algorithms [7] estimate the Jacobian $J + \Delta J$ at the point $\mathbf{x} + \Delta\mathbf{x}$ using the finite-difference approximation,

$$(J + \Delta J)\Delta\mathbf{x} = \Delta\mathbf{f} , \tag{9.13}$$

where $\Delta\mathbf{f} \doteq \mathbf{f}(\mathbf{x} + \Delta\mathbf{x}) - \mathbf{f}(\mathbf{x})$ and J is the Jacobian at the point $\mathbf{x}$.

Equivalently, one can apply the same approximation for the update $B + \Delta B$ of the inverse Jacobian matrix,

$$\Delta\mathbf{x} = (B + \Delta B)\Delta\mathbf{f} , \tag{9.14}$$

The matrix equation (9.14) is under-determined in more than one dimension as it contains only n equations to determine n^2 matrix elements of ΔB. Broyden suggested to choose ΔB (or ΔJ) as a rank-1 update that is linear $\Delta\mathbf{x}$. A rank-1 update has exactly n free parameters which can be determined form the (inverse secant) equation (9.14).

For example, one can choose the update in the form

$$\Delta B = \mathbf{c}\Delta\mathbf{x}^\mathsf{T} , \tag{9.15}$$

where $\mathbf{c}$ is an unknown vector. Inserting this ansatz into (9.14) and solving for $\mathbf{c}$ gives the update

$$\Delta B = \frac{\Delta\mathbf{x} - B\Delta\mathbf{f}}{\Delta\mathbf{x}^\mathsf{T}\Delta\mathbf{f}}\Delta\mathbf{x}^\mathsf{T} . \tag{9.16}$$

Here is a list of several rank-1 updates of the inverse Jacobian matrix,

1. "Good Broyden's method",

$$\Delta B = \mathbf{c}\Delta\mathbf{x}^\mathsf{T}B \;\Rightarrow\; \Delta B = \frac{\Delta\mathbf{x} - B\Delta\mathbf{f}}{\Delta\mathbf{x}^\mathsf{T}B\Delta\mathbf{f}}\Delta\mathbf{x}^\mathsf{T}B . \tag{9.17}$$

2. "Bad Broyden's method",

$$\Delta B = \mathbf{c}\Delta\mathbf{f}^\mathsf{T} \;\Rightarrow\; \Delta B = \frac{\Delta\mathbf{x} - B\Delta\mathbf{f}}{\Delta\mathbf{f}^\mathsf{T}\Delta\mathbf{f}}\Delta\mathbf{f}^\mathsf{T} \;. \tag{9.18}$$

3. Yet another method,

$$\Delta B = \mathbf{c}\Delta\mathbf{x}^\mathsf{T} \;\Rightarrow\; \Delta B = \frac{\Delta\mathbf{x} - B\Delta\mathbf{f}}{\Delta\mathbf{x}^\mathsf{T}\Delta\mathbf{f}}\Delta\mathbf{x}^\mathsf{T} \;. \tag{9.19}$$

In practice if one wanders too far from the point where J was first calculated the accuracy of the updates may decrease significantly. In such case one might need to recalculate J anew. For example, two successive steps with $\lambda_{\min}$ might be interpreted as a sign of accuracy loss in J and subsequently trigger its recalculation.

Again, if the denominator in the update formula becomes too small one has to discard the update as it is obviously wrong. If this happens two steps in a row, one has to recalculate one's B-matrix.

Here is a typical quasi-Newton algorithm,

```
calculate the inverse Jacobian matrix B = J⁻¹
repeat
    Δx = −Bf(x)
    λ = 1
    while ‖f(x + λΔx)‖ > (1 − λ/2) ‖f(x)‖ and λ ≥ 1/128 do λ = λ/2
    x = x + λΔx
    if λ ≥ 1/128 update B = B + ΔB else recalculate B = J⁻¹
until converged (e.g. ‖f(x)‖ < tolerance)
```

The matrix B is updated if the linesearch succeeds, that is, the step-parameter λ is not too small; in case the step-parameter becomes small the step is accepted unconditionally and the B-matrix is recalculated.

Chapter 10

Minimization

10.1 Introduction

Minimization (maximization) is the problem of finding the minimum (maximum) of a given—generally non-linear—real valued function $\phi(\mathbf{x})$ of an n-dimensional argument $\mathbf{x} \doteq \{x_1, \ldots, x_n\}$. The function is often called the *objective function* or the *cost function*.

Minimization is a simpler case of a more general poblem—*optimization*—which includes finding the best available values of the objective function within a given domain and/or subject to given constrains.

Minimization is not unrelated to root-finding: at the minimum all partial derivatives of the objective function vanish,

$$\frac{\partial \phi}{\partial x_i} = 0 \bigg|_{i=1\ldots n} , \tag{10.1}$$

and one can alternatively solve this system of (non-linear) equations.

10.2 Local minimization

Local minimization refers to a group of algorithms that move from one candidate solution to another candidate solution by applying local changes and moving "downhill" until a solution deemed optimal is found (or the alotted time is elapsed).

10.2.1   Newton's method

Newton's method is based on the quadratic approximation of the objective function $\phi(\mathbf{x})$ in the vicinity of the suspected minimum,

$$\phi(\mathbf{x} + \Delta\mathbf{x}) \approx \phi(\mathbf{x}) + \nabla\phi(\mathbf{x})^{\mathsf{T}}\Delta\mathbf{x} + \frac{1}{2}\Delta\mathbf{x}^{\mathsf{T}}H(\mathbf{x})\Delta\mathbf{x} \,, \tag{10.2}$$

where the vector $\nabla\phi(\mathbf{x})$ is the gradient of the objective function at the point $\mathbf{x}$,

$$\nabla\phi(\mathbf{x}) \doteq \left\{ \frac{\partial\phi(\mathbf{x})}{\partial x_i} \right\}_{i=1\ldots n} , \tag{10.3}$$

and $H(\mathbf{x})$ is the *Hessian matrix* – a square matrix of second-order partial derivatives of the objective function at the point $\mathbf{x}$,

$$H(\mathbf{x}) \doteq \left\{ \frac{\partial^2\phi(\mathbf{x})}{\partial x_i \partial x_j} \right\}_{i,j \in 1\ldots n} . \tag{10.4}$$

The minimum of the quadratic form (10.2), as function of $\Delta\mathbf{x}$, is found at the point where its gradient with respect to $\Delta\mathbf{x}$ vanishes,

$$\nabla\phi(\mathbf{x}) + H(\mathbf{x})\Delta\mathbf{x} = 0 \,. \tag{10.5}$$

This gives an approximate step towards the minimum, called the *Newton's step*,

$$\Delta\mathbf{x} = -H(\mathbf{x})^{-1}\nabla\phi(\mathbf{x}) \,. \tag{10.6}$$

The original Newton's method is simply the iteration,

$$\mathbf{x}_{k+1} = \mathbf{x}_k - H(\mathbf{x}_k)^{-1}\nabla\phi(\mathbf{x}_k) \,, \tag{10.7}$$

where at each iteration the full Newton's step is taken and the Hessian matrix is recalculated. In practice, instead of calculating H^{-1} one rather solves the linear equation (10.5).

Usually the Newton's method is modified to take a smaller step $\mathbf{s}$,

$$\mathbf{s} = \lambda\Delta\mathbf{x}, \tag{10.8}$$

with $0 < \lambda < 1$. The factor λ can be found by a backtracking algoritm similar to that in the Newton's method for root-finding. One starts with $\lambda = 1$ and than backtracks, $\lambda \leftarrow \lambda/2$, until the *Armijo condition*,

$$\phi(\mathbf{x} + \mathbf{s}) < \phi(\mathbf{x}) + \alpha\mathbf{s}^{\mathsf{T}}\nabla\phi(\mathbf{x}) \,, \tag{10.9}$$

is satisfied (or the minimal λ (say, 1/1024) is reached, in which case the step is taken unconditionally). The parameter α can be chosen as small as 10^{-4}.

10.2.2 Numerical calculation of gradient and Hessian matrix

The *forward* finite-difference approximation to the derivative $f'(x)$ of a function $f(x)$ is given as

$$f'(x) \approx \frac{f(x + \delta x) - f(x)}{\delta x} , \qquad (10.10)$$

where δx is usually chosen as $(1 + |x|)\sqrt{\epsilon}$ where ϵ is the machine precision (for "double" numbers $\sqrt{\epsilon} = 2^{-26}$). This gives the following approximation for the gradient,

$$\nabla \phi(x)_i = \frac{\partial \phi(x)}{\partial x_i} = \frac{\phi(x_1, \ldots, x_i + \delta x_i, \ldots, x_n) - \phi(x_1, \ldots, x_n)}{\delta x_i} , \qquad (10.11)$$

where $\delta x_i = (1 + |x_i|)\sqrt{\epsilon}$.

Similarly, the Hessian matrix is given as second forward finite difference,

$$\frac{\partial^2 \phi}{\partial x_i \partial x_j} = \frac{\partial}{\partial x_i} \nabla \phi_j = \frac{\nabla \phi(x_1, \ldots, x_i + \delta x_i, \ldots, x_n)_j - \nabla \phi(x_1, \ldots, x_n)_j}{\delta x_i} \qquad (10.12)$$

where the step-size δx_i however must be larger than in the first finite difference, for example, $\delta x_i = (1 + |x_i|)\epsilon^{1/4}$.

The *central* (possibly better) finite difference formula for the first derivative is given as,

$$f'(x) \approx \frac{f(x + \delta x) - f(x - \delta x)}{2\delta x} . \qquad (10.13)$$

This gives the following approximation for the gradient,

$$\nabla \phi(x)_i = \frac{\partial \phi(x)}{\partial x_i} = \frac{\phi(x_1, \ldots, x_i + \delta x_i, \ldots, x_n) - \phi(x_1, \ldots, x_i - \delta x_i, \ldots, x_n)}{2\delta x_i} , \qquad (10.14)$$

and the following approximation for the Hessian matrix,

$$H_{jk} = \frac{\partial^2 \phi}{\partial x_j \partial x_k} \approx$$
$$\frac{\phi(\mathbf{x} + \delta \mathbf{x}_k + \delta \mathbf{x}_j) - \phi(\mathbf{x} + \delta \mathbf{x}_k - \delta \mathbf{x}_j) - \phi(\mathbf{x} - \delta \mathbf{x}_k + \delta \mathbf{x}_j) + \phi(\mathbf{x} - \delta \mathbf{x}_k - \delta \mathbf{x}_j)}{4\delta x_k \delta x_j} . \qquad (10.15)$$

where $\delta \mathbf{x}_k$ is a vector in the direction k with the length $\delta x_k = (1 + |x_k|)\sqrt{\epsilon}$ where ϵ is the *machine epsilon*. In practice the calculation of the gradient should reuse the ϕ-values from the calculation of the Hessian matrix.

10.2.3   Quasi-Newton methods

Quasi-Newton methods are variations of the Newton's method which attempt to avoid recalculation of the Hessian matrix at each iteration, trying instead certain updates based on the analysis of the gradient vectors. The update δH is usually chosen to satisfy the condition

$$\nabla\phi(\mathbf{x} + \mathbf{s}) = \nabla\phi(\mathbf{x}) + (\mathrm{H} + \delta\mathrm{H})\mathbf{s} , \tag{10.16}$$

called *secant equation*, which is the Taylor expansion of the gradient.

The secant equation is under-determined in more than one dimension as it consists of only n equations for the n^2 unknown elements of the update δH. Various quasi-Newton methods use different choices for the form of the solution of the secant equation.

In practice one typically uses the inverse Hessian matrix (often—but not always—denoted as B) and applies the updates directly to the inverse matrix thus avoiding the need to solve the linear equation (10.5) at each iteration.

For the inverse Hessian matrix the secant equation (10.16) reads

$$(\mathrm{B} + \delta\mathrm{B})\mathbf{y} = \mathbf{s} , \tag{10.17}$$

or, in short,

$$\delta\mathrm{B}\mathbf{y} = \mathbf{u} , \tag{10.18}$$

where $\mathrm{B} \doteq \mathrm{H}^{-1}$, $\mathbf{y} \doteq \nabla\phi(\mathbf{x} + \mathbf{s}) - \nabla\phi(\mathbf{x})$, and $\mathbf{u} \doteq \mathbf{s} - \mathrm{B}\mathbf{y}$.

One usually starts with the identity matrix as the zeroth approximation for the inverse Hessian matrix and then applies the updates.

If the minimal λ (say, 1/1024) is reached during the bactracking line-search—which might be a signal of lost precision in the approximate (inverse) Hessian matrix—it is advisable to reset the current inverse Hessian matrix to identity matrix.

Table 10.2.3 lists one possible algorithm of the quasi-newton method with updates.

Broyden's update

The Broyden's update is chosen in the form

$$\delta\mathrm{B} = \mathbf{c}\mathbf{s}^{\mathsf{T}} . \tag{10.19}$$

where the vector $\mathbf{c}$ is found from the condition (10.18),

$$\mathbf{c} = \frac{\mathbf{u}}{\mathbf{s}^{\mathsf{T}}\mathbf{y}} . \tag{10.20}$$

Sometimes the dot-product $\mathbf{s}^{\mathsf{T}}\mathbf{y}$ becomes very small or even zero which results in serious numerical difficulties. One can avoid this by only performing update if the condition $|\mathbf{s}^{\mathsf{T}}\mathbf{y}| > \epsilon||\mathbf{s}||||\mathbf{y}||$ is satisfied where ϵ is a small number, say 10^{-6}.

Table 10.1: Quasi-newton minimisation algorithm with updates.

```
set the inverse Hessian matrix to unity, B = 1
repeat until converged (e.g. ‖∇φ‖ < tolerance) :
    calculate the Newton's step Δx = −B∇φ
    if not ∇φ · Δx < 0 :
        reset B = 1 and Δx = −∇φ
    do linesearch starting with λ = 1 :
        if φ(x + λΔx) < φ(x) or λ < λ_min :
            accept the step x = x + λΔx
            update B = B + δB
            break linesearch
        backtrack λ/ = 2
```

Symmetric Broyden's (aka BFGS) update

The Broyden's update (10.19) is not symmetric (while the Hessian matrix should be) which is an obvious drawback. Therefore a beter approximation might be the symmetric Broyden's update,

$$\delta B = \mathbf{a}\mathbf{s}^\mathsf{T} + \mathbf{s}\mathbf{a}^\mathsf{T} . \tag{10.21}$$

The vector $\mathbf{a}$ is again found from the condition (10.18),

$$\mathbf{a} = \frac{\mathbf{u} - \gamma \mathbf{s}}{\mathbf{s}^\mathsf{T}\mathbf{y}} , \tag{10.22}$$

where $\gamma = (\mathbf{u}^T \mathbf{y})/(2\mathbf{s}^T \mathbf{y})$.

Again one only performs the update if $|\mathbf{s}^\mathsf{T}\mathbf{y}| > \epsilon ||\mathbf{s}|| ||\mathbf{y}||$.

This update is generally known as Broyden–Fletcher–Goldfarb–Shanno (BFGS) update.

SR1 update

The symmetric-rank-1 update (SR1) in chosen in the form

$$\delta B = \mathbf{v}\mathbf{v}^\mathsf{T} , \tag{10.23}$$

where the vector $\mathbf{v}$ is again found from the condition (10.16), which gives

$$\delta B = \frac{\mathbf{u}\mathbf{u}^T}{\mathbf{u}^T\mathbf{y}} . \tag{10.24}$$

Again, one only performs the update if denominator is not too small, that is, $|\mathbf{u}^\mathsf{T}\mathbf{y}| > \epsilon ||\mathbf{u}|| ||\mathbf{y}||$.

Other popular updates

The wikipedia article "Quasi-Newton method" list several other popular updates.

10.2.4   Downhill simplex method

The *downhill simplex method* [18] (also called "Nelder-Mead" or "amoeba") is a com-
monnly used minimization algorithm where the minimum of a function in an n-dimensional
space is found by transforming a simplex—a polytope with $n+1$ vertexes—according
to the function values at the vertexes, moving it downhill until it converges towards
the minimum.

The advantages of the downhill simplex method is its stability and the lack of use
of derivatives. However, the convergence is realtively slow as compared to Newton's
methods.

In order to introduce the algorithm we need the following definitions:

- Simplex: a figure (polytope) represented by $n+1$ points, called vertexes, $\{\mathbf{p}_1, \ldots, \mathbf{p}_{n+1}\}$
  (where each point $\mathbf{p}_k$ is an n-dimensional vector).

- Highest point: the vertex, $\mathbf{p}_{\mathrm{hi}}$, with the highest value of the function: $\phi(\mathbf{p}_{\mathrm{hi}}) = \max_k \phi(\mathbf{p}_k)$.

- Lowest point: the vertex, $\mathbf{p}_{\mathrm{lo}}$, with the lowest value of the function: $\phi(\mathbf{p}_{\mathrm{lo}}) = \min_k \phi(\mathbf{p}_k)$.

- Centroid: the center of gravity of all points, except for the highest: $\mathbf{p}_{\mathrm{ce}} = \frac{1}{n} \sum_{(k \neq \mathrm{hi})} \mathbf{p}_k$

The simplex is moved downhill by a combination of the following elementary oper-
ations:

1. Reflection: the highest point is reflected against the centroid, $\mathbf{p}_{\mathrm{hi}} \to \mathbf{p}_{\mathrm{re}} = \mathbf{p}_{\mathrm{ce}} + (\mathbf{p}_{\mathrm{ce}} - \mathbf{p}_{\mathrm{hi}})$.

2. Expansion: the highest point reflects and then doubles its distance from the
   centroid, $\mathbf{p}_{\mathrm{hi}} \to \mathbf{p}_{\mathrm{ex}} = \mathbf{p}_{\mathrm{ce}} + 2(\mathbf{p}_{\mathrm{ce}} - \mathbf{p}_{\mathrm{hi}})$.

3. Contraction: the highest point halves its distance from the centroid, $\mathbf{p}_{\mathrm{hi}} \to \mathbf{p}_{\mathrm{co}} = \mathbf{p}_{\mathrm{ce}} + \frac{1}{2}(\mathbf{p}_{\mathrm{hi}} - \mathbf{p}_{\mathrm{ce}})$.

4. Reduction: all points, except for the lowest, move towards the lowest points
   halving the distance. $\mathbf{p}_{k \neq \mathrm{lo}} \to \frac{1}{2}(\mathbf{p}_k + \mathbf{p}_{\mathrm{lo}})$.

Table 10.2 shows one possible algorithm for the downhill simplex algorithm.

Table 10.2: Downhill simplex (Nelder-Mead) algorithm

```
REPEAT :
  find highest , lowest , and centroid points of the simplex
  try reflection
  IF  φ(reflected) < φ(lowest) :
    try expansion
    IF  φ(expanded) < φ(reflected) :
      accept expansion
    ELSE :
      accept reflection
  ELSE :
    IF  φ(reflected) < φ(highest) :
      accept reflection
    ELSE :
      try contraction
      IF  φ(contracted) < φ(highest) :
        accept contraction
      ELSE :
        do reduction
UNTIL converged (e.g. size (simplex)<tolerance )
```

10.2.5   Gauss-Newton algorithm

The Gauss-Newton algorithm is designed to minimize an objective function $\phi(\mathbf{c})$ that is given as a sum of squares of several (non-linear) functions $r_i(\mathbf{c})$,

$$\phi(\mathbf{c}) = \sum_{i=1}^{n} r_i^2(\mathbf{c}) \, , \tag{10.25}$$

where $\{c_{k=1\ldots m}\}$ is the set of parameters of the objective function. In particular, the algorithm can be used to solve a non-linear least squares curve fitting problem where the function to minimize is given as

$$\chi^2(\mathbf{c}) = \sum_{i=1}^{n} \left(\frac{f(\mathbf{c}, x_i) - y_i}{\delta y_i} \right)^2 \, , \tag{10.26}$$

where $\{x_i, y_i \pm \delta y_i\}$ is the set of data to fit and $f(\mathbf{c}, x)$ is the fitting function that depends on a set of parameters $\mathbf{c}$.

The algorithm can also be used to find an approximate solution to an overdetermined (if $n > m$) system of non-linear equations

$$\mathbf{r}(\mathbf{c}) = 0 \, . \tag{10.27}$$

Just like the Newton's method the algorithm relies on the Taylor expansion of the objective function in the vicinity of the suspected minimum,

$$\phi(\mathbf{c} + \Delta\mathbf{c}) \approx \phi(\mathbf{c}) + \mathbf{g}^\mathsf{T}\Delta\mathbf{c} + \frac{1}{2}\Delta\mathbf{c}^\mathsf{T}\mathrm{H}\mathbf{c} , \tag{10.28}$$

where the (size-m) gradient $\mathbf{g}$ is given as

$$g_k = 2\sum_{i=1}^{n} r_i \frac{\partial r_i}{\partial c_k} \tag{10.29}$$

and the (square $m \times m$) Hessian matrix H is given as

$$\mathrm{H}_{jk} = 2\sum_{i=1}^{n} \left(\frac{\partial r_i}{\partial c_j}\frac{\partial r_i}{\partial c_k} + r_i\frac{\partial^2 r_i}{\partial c_j\partial c_k} \right) . \tag{10.30}$$

Now, in the Gauss-Newton method one ignores the second-derivative term in (10.30) which results in the following approximation for the Hessian matrix,

$$\mathrm{H}_{jk} \approx 2\mathrm{J}^\mathsf{T}\mathrm{J} , \tag{10.31}$$

where J is the (tall $n \times m$) Jacobian matrix of the $\{r_i\}$ functions,

$$\mathrm{J}_{ik} = \frac{\partial r_i}{\partial c_k} . \tag{10.32}$$

The approximation (10.31) may be valid in two cases,

1. The functions r_i are small in the vicinity of the minimum;

2. The functions r_i are only slightly non-linear such that the second derivatives are small in magnitude.

Using the Jacobian matrix the gradient of ϕ can be written as

$$\mathbf{g} = 2\mathrm{J}^\mathsf{T}\mathbf{r} . \tag{10.33}$$

From here the algorithm proceeds as in the usual Newton's method: one finds the Newton's step,
$$\Delta\mathbf{c} = -\mathrm{H}^{-1}\mathbf{g} \approx -\left(\mathrm{J}^\mathsf{T}\mathrm{J}\right)^{-1}\left(\mathrm{J}^\mathsf{T}\mathbf{r}\right) , \tag{10.34}$$
and then does the backtracking line-search.

Note that $\left(\mathrm{J}^\mathsf{T}\mathrm{J}\right)^{-1}\mathrm{J}^\mathsf{T}$ is the left pseudo-inverse of the matrix J. Therefore the Newton's step of the Gauss-Newton method for the objective function $\sum_{i=1}^{n} r_i^2$ is equivalent to the step of the root-finding Newton's method for the system of equations $\mathbf{r}(\mathbf{c}) = 0$.

10.3 Uncertainties of nonlinear least squares fit parameters

The *non-linear least squares* fit is the process of fitting a curve (a mathematical function, $f(\mathbf{c}, x)$, where $\mathbf{c}$ is the set of fitting parameters) to a set of data points with uncertainties, $\{x_i, y_i \pm \delta y_i\}$. However, unlike the ordinary least squares fit, where the fitting parameters enter linearly, in the non-linear least squares fit the parameters enter essentially non-linearly.

The fit is achieved by minimizing the sum of squares (hence the name) of the deviations of the curve from the data (called χ^2 in physics),

$$\chi^2(\mathbf{c}) = \sum_{i=1}^{n} \left(\frac{f(\mathbf{c}, x_i) - y_i}{\delta y_i} \right)^2 \equiv \sum_{i=1}^{n} r_i^2(\mathbf{c}) , \qquad (10.35)$$

where

$$r_i(\mathbf{c}) \doteq \frac{f(\mathbf{c}, x_i) - y_i}{\delta y_i} \qquad (10.36)$$

are the (weighted) residuals.

The χ^2 can be minimized in the space of the fitting parameters either using any of the general minimization algorithms or using the Gauss-Newton algortithm which is specifically designed for an objective function in the form of the sum of squares of some residuals.

The uncertainties of the fitting parameters can be estimated by i) Taylor expansion of χ^2 around the minimum; ii) linearizing the problem; iii) calculating the uncertainties using the same technique as for the ordinary least squares fit. In other words, we apply the Newton's method to find the solution of the minimization problem and then determine the uncertainties of the fitting parameters from the last Newton's step (the one that brings us to the minimum).

10.3.1 Linearization of nonlinear problem at minimum

The Newton's step $\Delta\mathbf{c}$ toward the minimum is found from the second order Taylor expansion of χ^2,

$$\chi^2(\mathbf{c} + \Delta\mathbf{c}) \approx \chi^2(\mathbf{c}) + \mathbf{g}^\mathsf{T}\Delta\mathbf{c} + \frac{1}{2}\Delta\mathbf{c}^\mathsf{T}\mathrm{H}\Delta\mathbf{c} , \qquad (10.37)$$

where the gradient $\mathbf{g}$ is given as

$$g_k = \frac{\partial\chi^2}{\partial c_k} = \sum_{i=1}^{n} 2r_i \frac{\partial r_i}{\partial c_k} . \qquad (10.38)$$

and the Hessian matrix H is given as

$$H_{jk} = \frac{\partial^2 \chi^2}{\partial c_j \partial c_k} = \sum_{i=1}^{n} \left(2 \frac{\partial r_i}{\partial c_j} \frac{\partial r_i}{\partial c_k} + 2 r_i \frac{\partial^2 r_i}{\partial c_j \partial c_k} \right) . \tag{10.39}$$

Close to the minimum the term with the second derivative can be (hopefully) neglected (since at the minimum $f(\mathbf{c}, x_i) \approx y_i$).

Introducing the (tall $n \times m$) Jacobian matrix J of the residuals,

$$J_{ij} = \frac{\partial r_i}{\partial c_j} = \frac{1}{\delta y_i} \frac{\partial f(\mathbf{c}, x_i)}{\partial c_j} , \tag{10.40}$$

one can rewrite the gradient as

$$\mathbf{g} = 2 J^\mathsf{T} \mathbf{r} \tag{10.41}$$

and the Hessian matrix as

$$H = 2 J^\mathsf{T} J . \tag{10.42}$$

The corresponding Newton's step is determined by the equation

$$H \Delta \mathbf{c} = -\mathbf{g} , \tag{10.43}$$

or

$$J^\mathsf{T} J \Delta \mathbf{c} = -J^\mathsf{T} \mathbf{r} , \tag{10.44}$$

which gives the Newton's step to the minimum as

$$\Delta \mathbf{c} = -J^{-1} \mathbf{r} \tag{10.45}$$

where

$$J^{-1} = (J^\mathsf{T} J)^{-1} J^\mathsf{T} \tag{10.46}$$

is the pseudo-inverse of the matrix J.

10.3.2   Uncertainties of the fit parameters

Equation (10.45) defines $\Delta c_{k=1...m}$ as function of $y_{i=1...n}$. The question is, if y_i are determined with uncertainties δy_i, what are the uncertainties of Δc_k?

The answer is given by the *propagation of uncertainty* rule which says that the (co)variances $\delta c_k \delta c_j$ are given as

$$\delta c_k \delta c_j = \sum_i \frac{\partial \Delta c_k}{\partial y_i} \frac{\partial \Delta c_j}{\partial y_i} \delta y_i \delta y_i = \sum_i (J^{-1})_{ki} (J^{-1})_{ji} . \tag{10.47}$$

In matrix notation the covariance matrix $\Sigma_{kj} = \delta c_k \delta c_j$ is given as

$$\Sigma = J^{-1}J^{-\mathsf{T}} = (J^{\mathsf{T}}J)^{-1} \ . \tag{10.48}$$

The uncertainties of the fitting parameters are then given as the square roots of the diagonal elements of the covariance matrix,

$$\delta c_k = \sqrt{\Sigma_{kk}} \ . \tag{10.49}$$

Notice that within the approximation (10.42) (that should work well at the minimum) the covariance matrix is given via the inverse of the Hessian matrix, H^{-1}, at the minimum,

$$\Sigma = (J^{\mathsf{T}}J)^{-1} = 2H^{-1} \ , \tag{10.50}$$

which is the canonical texbook result. It can also be obtained from the Taylor expansion of the variation of χ^2 with respect to fit parameters at the minimum, where the gradient is zero,

$$\delta\chi^2 = \frac{1}{2}\delta c^{\mathsf{T}}H\delta c = \text{trace}\left(\frac{1}{2}\Sigma H\right) \ . \tag{10.51}$$

The uncertainties of fit parameters are determined by a unit variation of χ^2 per degree of freedom (that is, per fit parameter). That is, the matrix inside the trace operator must be the unit $m \times m$ matrix. This gives

$$\Sigma = 2H^{-1} \ . \tag{10.52}$$

Bibliography

[1] J.A. Nelder and R. Mead, "A simplex method for function minimization", Computer Journal 7 (1965) 308

Chapter 11

Global optimization

11.1 Introduction

Global optimization generally refers to the problem of locating a good approximation to the *global* minimum of a given objective function in a given search space that typically contains many local minima and is generally large enough to prohibit exhaustive enumeration.

When only a small sub-space of the search space can be realistically sampled within the allotted time, the stochastic methods — which use some form of randomness — usually come to the fore. In the following several popular stochastic global minimization algorithms are shortly described.

11.2 Randomized local minimizers

For a differentiable optimization problem a good local minimizer would typically converge to the nearest local miminum relatively fast. Therefore a class of "quick-and-dirty" global minimizers can be constructed by simply adding some elements of random sampling to local minimizers.

11.2.1 Local minimization from several random start-points

One strategy is to start the local miminizer several times from different (quasi)random starting points within the given search space (recording the best solution). The procedure is repeated until the alloted time is exhausted.

11.2.2   Local minimization from best random sample

Another strategy is to use the allotted time to (quasi)randomly sample the given search space and then run the local minimizer from the best sampled point.

11.3   Simulated annealing

Simulated annealing is a stochastic meta-heuristic algorithm for global minimization. The name and inspiration come from annealing—heating up and cooling slowly—in material science. The slow cooling allows a piece of material to reach a state with "lowest energy".

The objective function in the space of states is interpreted as some sort of potential energy and the points in the search space are interpreted as states of a certain physical system. The system attempts to make transitions from its current state to some randomly sampled nearest states with the goal to eventually reach the state with minimal energy – the global minimum.

The system is attached to a thermal reservoir with certain temperature T. Each time the energy of the system is measured the reservoir supplies it with a random amount of thermal energy sampled from the Boltzmann distribution,

$$P(E) = Te^{-E/T} \ . \tag{11.1}$$

If the temperature equals zero the system can only make transitions to the neighboring states with lower potential energy. In this case the algorithm turns merely into a local minimizer with random sampling.

If temperature is finite the system is able to climb up the ridges of the potential energy—about as high as the current temperature—and thus escape from local minima and hopefully eventually reach the global minimum.

One typically starts the simulation with some finite temperature on the order of the height of the typical hills of the potential energy surface, letting the system to wander almost unhindered around the landscape with a good chance to locate if not the best then at least a good enough minimum. The temperature is then slowly reduced following some annealing schedule which may be supplied by the user but must end with $T = 0$ towards the end of the allotted time budget.

Table 11.1 lists one possible variant of the algorithm. Here the function `neigbour` is system-dependent and should return a randomly chosen "neighbour" of the given state. For a continuous function, where the state is the position of the current approximation to the minimum, the neghbour could be a random position within radius R from the current position. The step-radius can be gradually reduced to zero toward the end of the simulation, like

$$R(t) = R_0 \cdot (1 - \frac{t}{t_a}) \ , \tag{11.2}$$

Table 11.1: Simulated annealing algorithm

```
state ← start_state
T ← start_temperature
energy ← E(state)
REPEAT :
  new_state ← neighbour(state)
  new_energy ← E(new_state)
  IF new_energy < energy :
    state ← new_state
    energy ← new_energy
  ELSE :
    do with probability exp(− (new_energy−energy)/T) :
      state ← new_state
      energy ← new_energy
  reduce_temperature_according_to_schedule(T)
UNTIL terminated
```

where R_0 is the initial radius, t is the running time, and t_a is the allotted time.

The temperature can be also reduced linearly,

$$T(t) = T_0 \cdot (1 - \frac{t}{t_a}) \, , \qquad (11.3)$$

where T_0 is the initial temperature.

11.4 Quantum annealing

Quantum annealing is a general global minimization algorithm which—like simulated annealing—also allows the search path to escape from local minima. However instead of the thermal jumps over the potential barriers quantum annealing allows the system to tunnel through the barriers.

In its simplest incarnation the quantum annealing algorithm allows the system to attempt transitions not only to the nearest states but also to distant states within certain "tunneling distance" from the current state. The transition is accepted only if it reduces the potential energy of the system.

At the beginning of the minimization procedure the tunneling distance is large—on the order of the size of the region where the global minimum is suspected to be located—allowing the system to explore the region. The tunneling distance is then slowly reduced according to a schedule such that by the end of the allotted time the tunneling distance reduces to zero at which point the system hopefully is in the state with minimal energy.

Table 11.2: Quantum annealing algorithm

```
state ← start_state
energy ← E(state)
R ← start_radius
REPEAT :
   new_state ← random_neighbour_within_radius(state,R)
   new_energy ← E(new_state)
   IF new_energy < energy :
      state ← new_state
      energy ← new_energy
   reduce_radius_according_to_schedule(R)
UNTIL terminated
```

11.5   Evolutionary algorithms

Unlike annealing algorithms, which follow the motion of only one point in the search space, the evolutionary algorithms typically follow a set of points called a *population* of individuals. Somewhat like the downhill simplex method which follows the motion of a set of points – the simplex.

The population evolves toward more fit individuals where fitness is understood in the sense of minimizing the objective function. The parameters of the individuals (for example, the coordinates of the points in the parameter space of the objective function) are called genes.

The algorithm proceeds iteratively in discrete steps where the population in each iteration is called a generation. In each generation the fitness of each individual—typically, the value of the objective function—is evaluated and the new generation is generated stochastically from the gene pool of the current generation through certain operations (like crossovers and mutations) such that the genes of more fit individuals have a better chance of propagating into the next generation.

Each new individual in the next generation can be produced from a pair of "parent" individuals of the current generation, as inspired by biology, but more than two "parents" can be used as well. The parents for a new individual are selected from the individuals of the current generation through a fitness based stochastic process where fitter individuals are more likely to be selected.

Generation of "children" continues until the population of the new generation reaches the appropriate size after which the iteration repeats itself.

The algorithm is terminated when the fitness level of the population is deemed sufficient or when the allocated budget is exhausted.

11.5.1 Particle Swarm Optimization (PSO)

One example of evolutionary optimization algorithms is the *particle swarm optimization* method [25] where a population of "particles" move (in discrete time-steps) through the parameter space of the objective function and sample the function at the particles' positions at each step. The movement of each particle is stochastic and is influenced by the particle's best position as well as the whole swarm's best position: at each time-step a particle gets a stochastic kick toward it's own best position and toward the swarm's best position. After a certain amount of steps the swarm is expected to converge to the best solution.

A particle number i carries three vector-parameters: its position $\mathbf{x}_i$ in the parameter space of the objective function; its best position $\mathbf{p}_i$ so far; and its velocity, $\mathbf{v}_i$. In addition the swarm as a whole remembers its global best position, $\mathbf{g}$.

At each time-step Δt (usually equal unity) first the positions of the particles are updated,

$$\mathbf{x}_i = \mathbf{x}_i + \mathbf{v}_i \Delta t \ , \tag{11.4}$$

and the objective function is sampled at the new positions. Then the particles' best and the global best positions are updated. After that the velocities of the particles are stochastically updated according to the formula,

$$\mathbf{v}_i = w\mathbf{v}_i + u(\mathbf{p}_i - \mathbf{x}_i) + u(\mathbf{g} - \mathbf{x}_i) \ , \tag{11.5}$$

where u is a random number from a unit uniform distribution, and $w < 1$ us the damping parameter which ensures that the swarm gradually calms down (hopefully in the area of global minimum).

Table 11.3 lists one possible implementation of the algorithm.

11.5.2 Bare bones PSO (BBPSO)

A simpler variant of the PSO algorithm is the the so called "bare bones PSO" [19]. Here one dispences with the velocity of the particles and instead updates the positions of the particles using the following rule,

$$\mathbf{x}_i = G\left(\frac{\mathbf{p}_i + \mathbf{g}}{2}, \|\mathbf{p}_i - \mathbf{g}\|\right) \ , \tag{11.6}$$

where $\mathbf{x}_i$, $\mathbf{p}_i$ are the position and the best position of particle i, $\mathbf{g}$ is the global best position, and $G(\mathbf{x}, \sigma)$ is the Gaussian (normal) distribution with the mean $\mathbf{x}$ and standard deviation σ.

Table 11.3: Particle swarm optimization algorithm

```
Initialize  uniform  unit  random  number  generator  u;
Assume  the  time−step  is  equal  unity,  Δt = 1;
Initialize  particle  positions  randomly  within  given  rectangular
volume  V[a, b]  given  by  the  vectors  a  and  b,
   xᵢ = random  vector  within  V[a, b];
Initialize  particle  velocities  randomly,
   vᵢ = random  vector  within  V[a−b/2, b−a/2] 1/Δt ;
Initialize  local  best  positions,
   pᵢ = xᵢ;
Initialize  global  best  position,
   g = minᵢ(f(pᵢ));
REPEAT:
   Update  velocities  (the  damping  parameter  w ≈ 0.72),
      vᵢ = wvᵢ + U · (pᵢ − xᵢ) 1/Δt + U · (g − xᵢ) 1/Δt ;
   Update  positions,
      xᵢ = xᵢ + vᵢΔt;
   Update  local  bests,
      if  f(xᵢ) < f(pᵢ)  pᵢ = xᵢ;
   Update  global  best,
      if  f(xᵢ) < f(g)  g = xᵢ;
UNTIL  allotted  time  is  spent  or  converged
```

Generation of normally distributed sequences

Box-Muller transform   A normally distributed sequence of numbers, r, with zero mean and unit variance (called standard normal distribution) can be constructed using the *Box-Muller transform* from two sequences, u_1 and u_2, which are uniformly distributed on the unit interval $(0, 1]$. The transformation is given as

$$r_1 = \sqrt{-2\ln u_1}\cos(2\pi u_2) ,\tag{11.7}$$

$$r_2 = \sqrt{-2\ln u_1}\sin(2\pi u_2) .\tag{11.8}$$

The sequences r_1 and r_2 are independent random variables with a standard normal distribution.

A normal distribution R with a given mean m and variance σ can be constructed from the standard normal distribution r as

$$R = m + r\sigma .\tag{11.9}$$

Irwin-Hall distribution The size-n *Irwin-Hall distribution* is the distribution of the sum of n independent numbers, u_i, which are uniformly distributed on $(0, 1)$,

$$x = \sum_{i=1}^{n} u_i .$$
(11.10)

By the Central Limit Theorem as n increases the Irwin-Hall distribution $H_n(x)$ approaches the normal distribution $G_{\mu\sigma}(x)$ with the mean $\mu = n/2$ and the variance $\sigma^2 = n/12$,

$$\sqrt{\frac{n}{12}} H_n \left(x\sqrt{\frac{n}{12}} + \frac{n}{2} \right) \stackrel{n\to\infty}{\longrightarrow} G_{0,1}(x) .$$
(11.11)

This leads to a simple approximation where a standard normal distribution is given by the sum of 12 pseudorandom numbers on $(0, 1)$,

$$\sum_{i=1}^{12} u_i - 6 \approx G_{0,1}$$
(11.12)

Chapter 12

Artificial Neural Networks

12.1 Introduction

An *artificial neural network* is simply a mathematical function. Specifically, a vector function, $\mathbf{F}$, of a vector argument $\mathbf{x}$ (often called the *input signal*),

$$\mathbf{y} = \mathbf{F_p}(\mathbf{x}) \tag{12.1}$$

where the vector $\mathbf{y}$ is the return value (often called *response*), and the vector $\mathbf{p}$ is the set of internal parameters of the network. The parameters are tuned such that the network can perform some useful function like recognizing a pattern in a bitmap picture or approximating a solution to a differential equation. Tuning the parameters for a specific task is called *network training* or *learning* [27].

12.2 Applications

Neural networks have found numerous applications in many areas. In particlar they are used for pattern recognition. An example of pattern recognition is handwriting recognition. In this case the input to the network is the bitmap picture of a handwritten character: the vector $\mathbf{x}$ contains the RGB values of the bitmap's pixels. The response $\mathbf{y}$ of the network is the UTF code of the recognised character.

Another example is the traffic sign recognition network for vehicle control systems. In this case the input to the network is the bitmap picture from the vehicle's camera and the output is the code of the recognized traffic sign.

In physics neural networks are used as function approximants (for interpolation, regression, and numerical solution of differential equations) as well as in data processing and modelling of complex systems.

12.3   Graphical representation

Artificial neural networks are designed to resemble biological neural networks in the brains of animals, hence the name. Specifically, an artificial neural network is a collection of connected nodes called artificial neurons, typically represented graphically. A single artificial neuron is customarily pictured as

$$x_1 \qquad\qquad\qquad\qquad\qquad y_1 = w_1 r$$
$$r = f\left(\sum_i x_i\right)$$
$$x_2 \qquad\qquad\qquad\qquad\qquad y_2 = w_2 r$$
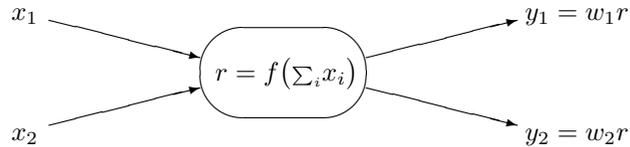
The neuron takes one or more input signals x_i, applies the neuron's *activation function*, f, to the sum of the input signals producing its response, r, and then sends one or more output signals, y_i, where the response is multiplied by weight-factors w_i.

Here is an example of a network that takes a 2-dimensional vector x as input and returns a 2-dimensional vector y as output,
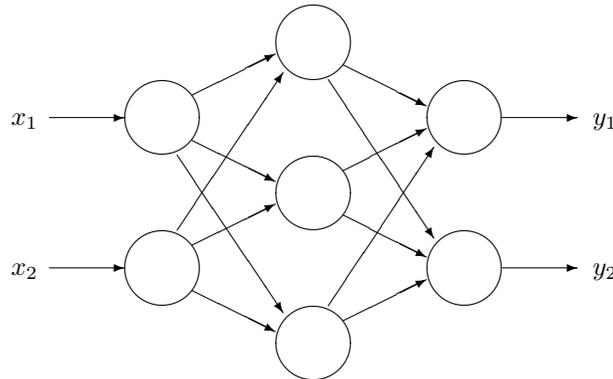
The left column of neurons (which receive the input signal $\mathbf{x}$) is called the *input layer*. The right column of neurons which send out the response $\mathbf{y}$ is called the *output layer*. The middle column of neurons is called the *hidden layer*. Each arrow connecting neurons carries its own weight-factor. Networks with one or more hidden layers are often called *deep neural networks*. Networks where the signal flow is always from left to right are called *feedforward* networks.

The activation function $f_i(x)$ of the i'th neuron is often chosen as

$$f_i(x) = f\left(\frac{x - a_i}{b_i}\right) \tag{12.2}$$

where f is the common activation function and where the shift a_i and scale b_i are the individual neuron's parameters. The total network parameter-vector $\mathbf{p}$ is then given as the collection of all weight-factors w_{ij}, shifts a_i, and scales b_i.

There exist many different types of networks with different numbers of neurons/layers and with different topological structures.

12.4 Training (learning)

Training is tuning the network's parameters $\mathbf{p}$ to better handle the given task. It typically involves minimization of certain *cost function* of network parameters, $C(\mathbf{p})$. The two major training paradigms are *supervised* and *unsupervised* learning.

Supervised learning uses a set of inputs paired with the desired outputs. Here the cost function is given by the difference between the network's output and the desired output. For example, in handwriting recognition the supervised learning can use a set of bitmaps with known handwritten characters with the cost function being the number of wrongly recognized characters.

Another example is the one-dimensional interpolation where the input data are pairs $\{x_i, y_i\}_{i=1\ldots n}$ (the table to interpolate) and the cost function is the average squared deviation,

$$C(\mathbf{p}) = \frac{1}{n} \sum_{i=1}^{n} \left(F_{\mathbf{p}}(x_i) - y_i\right)^2 \ . \tag{12.3}$$

In unsupervised learning the input data is given together with the cost function but without the correct output. For example, in solving a differential equation the cost function could be the average mismatch between the left- and right-hand sides of the differential equation.

Chapter 13

Fast Fourier transform

Fast Fourier transform (FFT) is an efficient algorithm to compute the discrete Fourier transform (DFT). Computing DFT of a size-N vector in the naïve way, using the definition, takes $O(N^2)$ arithmetic operations, while an FFT can compute the same result in only $O(N \log N)$ operations. The difference in speed can be substantial, especially for large data sets. This improvement made many DFT-based algorithms practical. Since the inverse of a DFT is also a DFT, any FFT algorithm can be used for the inverse DFT as well.

The most well known FFT algorithms, like the Cooley-Tukey algorithm [9], depend upon the factorization of N. However, there are FFTs with $O(N \log N)$ complexity for all N, even for prime N.

13.1   Discrete Fourier Transform

For a set of complex numbers $\{y_n\}_{n=0,\dots,N-1}$ the DFT is defined as a set of complex numbers $\{Y_k\}_{k=0,\dots,N-1}$ given as

$$Y_k = \sum_{n=0}^{N-1} y_n e^{-2\pi i \frac{nk}{N}} \ . \tag{13.1}$$

The inverse DFT is given by

$$y_n = \frac{1}{N} \sum_{k=0}^{N-1} Y_k e^{+2\pi i \frac{nk}{N}} \ . \tag{13.2}$$

These transformations can be viewed as expansion of the vector $\mathbf{y} = \{y_n\}$ in terms

of the orthogonal basis of vectors $\mathbf{v}_k$,

$$\mathbf{v}_k = \{e^{\frac{2\pi i k}{N} n}\} \tag{13.3}$$

$$\mathbf{v}_k^\star \cdot \mathbf{v}_{k'} = \sum_{n=0}^{N-1} \left(e^{-2\pi i \frac{kn}{N}}\right) \left(e^{2\pi i \frac{k'n}{N}}\right) = \sum_{n=0}^{N-1} e^{2\pi i \frac{(k'-k)n}{N}} = N\delta_{kk'} \ . \tag{13.4}$$

The DFT represents the amplitude and phase of the different sinusoidal components in the input data y_n.

The DFT is widely used in different fields, like spectral analysis, data compression, solution of partial differential equations and others.

13.1.1   Applications

Data compression

Several lossy (that is, with certain loss of information) image and sound compression methods employ DFT as an approximation for the Fourier series. The signal is discretized and transformed, and then the Fourier coefficients of high/low frequencies, which are assumed to be unnoticeable, are discarded. The decompressor computes the inverse transform based on this reduced number of Fourier coefficients.

Noise filtering

DFT can be used to attempt to filter out noise from a noisy signal. First, the signal is Fourier transformed. Then either the Fourier components with small amplitude are removed (set to zero) or the Fourier components with high frequencies are removed (set to zero). The filtered signal is then given as the inverse DFT of the modified set of Fourier components. And example is shown on Figure 13.1.

Differential equations

Discrete Fourier transforms are often used to solve differential equations, where the DFT is used as an approximation for the Fourier series (which is recovered in the limit of infinite N). The advantage of this approach is that it expands the signal in complex exponentials $e^{i\omega t}$, which are eigenfunctions of differentiation: $\frac{d}{dt}e^{i\omega t} = i\omega e^{i\omega t}$. Thus in the Fourier representation (in the frequency domain) differentiation is simply multiplication by $i\omega$.

A linear differential equation with constant coefficients is transformed into an easily solvable algebraic equation. One then uses the inverse DFT to transform the result back into the ordinary representation (time domain). This approach belongs to the group of *spectral methods.*
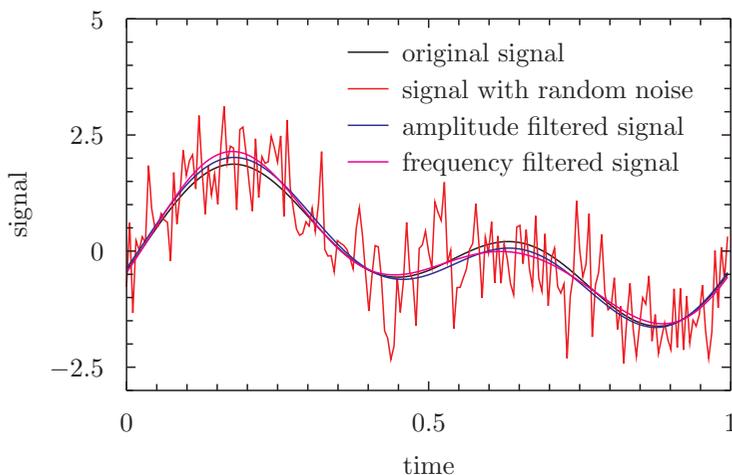
Figure 13.1: Noisy signal filtering using DFT

For example, consider the following first-order differential equation with a periodic boundary condition,

$$\frac{dy(t)}{dt} + 2y(t) = f(t) , \tag{13.5}$$

where $f(t)$ is a given periodic force-function with period T and where we try to find a solution $y(t)$ with the same period T.

First, we discretize the time-variable over one period of T and sample both $y(t)$ and $f(t)$ at N equally spaced points

$$t_n = n\Delta t , n = 0, \ldots, N - 1 \tag{13.6}$$

where $\Delta t = T/N$. That gives us the discrete sequences $\{y_n = y(t_n)\}$ and $\{f_n = f(t_n)\}$.

Now, we apply the DFT to both $\{y_n\}$ and $\{f_n\}$,

$$Y_k = \sum_{n=0}^{N-1} y_n e^{-2\pi i k n/N} , \quad F_k = \sum_{n=0}^{N-1} f_n e^{-2\pi i k n/N} . \tag{13.7}$$

If the DFT $\{Y_k\}$ is known, the function $y(t)$ is approximated as

$$y(t) \approx \sum_{k=0}^{N-1} Y_k e^{i\omega_k t} , \tag{13.8}$$

where $\omega_k = 2\pi k/T$. The derivative $dy(t)/dt$ in the same approximation is given as

$$\frac{dy(t)}{dt} \approx \sum_{k=0}^{N-1} i\omega_k Y_k e^{i\omega_k t} \, , \tag{13.9}$$

that is, the DFT of the derivative is given as

$$\mathrm{DFT}\left(\frac{dy(t)}{dt}\right) = \left\{i\frac{2\pi k}{T}Y_k\right\} \, . \tag{13.10}$$

Now we can take the DFT of the entire differential equation,

$$i\frac{2\pi k}{T}Y_k + 2Y_k = F_k \, . \tag{13.11}$$

This algebraic equation is can be solved for Y_k,

$$Y_k = \frac{F_k}{2 + i\frac{2\pi k}{T}} \, . \tag{13.12}$$

Finally, we obtain the (approximation of the) solution $y(t)$ as the inverse DFT of the obtained $\{Y_k\}$,

$$y_n = \frac{1}{N}\sum_{k=0}^{N-1} Y_k e^{2\pi i k n/N} \, . \tag{13.13}$$

Convolution and Deconvolution

FFT can be used to efficiently compute convolutions of two sequences. A convolution is the pairwise product of elements from two different sequences, such as in multiplying two polynomials or multiplying two long integers.

One example comes from data acquisition processes where the detector introduces certain (typically Gaussian) blurring to the sampled signal. A reconstruction of the original signal can be obtained by deconvoluting the acquired signal with the detector's blurring function. Let us say that we measure the original signal $\{y_n\}$ with a detector that slightly mixes the neighboring samples. The measured signal $\{s_n\}$ is than given as a linear superposition of several points, something like

$$s_n = \frac{1}{10}y_{n-1} + \frac{8}{10}y_n + \frac{1}{10}y_{n+1} \, . \tag{13.14}$$

Generally this is written as a *convolution* of two sequences,

$$s_n = \sum_m y_m \epsilon_{n-m} \, , \tag{13.15}$$

where ϵ_{n-m} represents the detector efficiency function (a narrow bell-shape function for a good detector, approaching the delta-function for the perfect detector). Formally, if we have a size-N sequence $\{y_n\}$, the convolution is defined as

$$s_n = \sum_{m=0}^{N} y_m \epsilon_{(n-m) \bmod N} \, . \tag{13.16}$$

Now the interesting thing is that the DFT of a convolution of two sequences is given by a product of DFT's of the sequences. Indeed,

$$
\begin{aligned}
\mathrm{DFT}(\{s_n\})_k \doteq S_k \;&=\; \sum_{n} e^{-2\pi i k n/N} \sum_{m=0}^{N} y_m \epsilon_{(n-m) \bmod N} \\
&=\; \sum_{k'} \sum_{n} e^{-2\pi i (k-k') n/N} \frac{1}{N} Y_{k'} E_{k'} \\
&=\; Y_k E_k \, .
\end{aligned}
\tag{13.17}
$$

Therefore the DFT of the original signal can be obtained by dividing the DFT of the measured signal by the DFT of the detector efficiency,

$$Y_k = \frac{S_k}{E_k} \, . \tag{13.18}$$

The original signal is then restored by inverse DFT of $\{S_k/E_k\}$. Note however that if the DFT of the detector efficiency $\{E_k\}$ has values close to zero for some frequencies, the division can significantly amplify noise present in the blurred signal $\{s_n\}$, leading to a noisy and unstable restoration. This is a common issue in deconvolution.

13.2   Cooley-Tukey algorithm

In its simplest incarnation this algorithm re-expresses the size $N = 2M$ discrete Fourier transform $\{c_k\} = \mathrm{DFT}\left(\{x_n\}_{n=0,\ldots,N}\right)$ in terms of two DFTs of size M,

$$
\begin{aligned}
c_k \;&=\; \sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{nk}{N}} \\
&=\; \sum_{m=0}^{M-1} x_{2m} e^{-2\pi i \frac{mk}{M}} + e^{-2\pi i \frac{k}{N}} \sum_{m=0}^{M-1} x_{2m+1} e^{-2\pi i \frac{mk}{M}} \\
&=\; \begin{cases} c_k^{(\mathrm{even})} + e^{-2\pi i \frac{k}{N}} c_k^{(\mathrm{odd})} & , k < M \\ c_{k-M}^{(\mathrm{even})} - e^{-2\pi i \frac{k-M}{N}} c_{k-M}^{(\mathrm{odd})} & , k \geq M \end{cases} ,
\end{aligned}
\tag{13.19}
$$

Table 13.1: Naïve Python implementation of the Cooley-Tukey algorithm

```
import cmath

def dft (x, sign=−1) :
    N = len(x)
    w = cmath.exp(sign*2*cmath.pi*1j/N)
    return [sum( x[n]*w**(n*k) for n in range(N) ) for k in range(N)]

def fft (x, sign=−1) :
    N = len(x)
    if N % 2 == 0 :
        M = N//2
        ce = fft(x[0::2], sign)
        co = fft(x[1::2], sign)
        w = cmath.exp(sign*2*cmath.pi*1j/N)
        c = [ce[k]  +w**k*co[k]   for k in range(M)   ] \
           +[ce[k−M]+w**k*co[k−M] for k in range(M,N)]
        return c
    else : return dft(x,sign)

def ift (x) : return [z/len(x) for z in fft(x,+1)]
```

where $c^{(\text{even})}$ and $c^{(\text{odd})}$ are the DFTs of the even- and odd-numbered sub-sets of $\{x_n\}$. This re-expression of a size-N DFT as two size-$\frac{N}{2}$ DFTs is sometimes called the Danielson-Lanczos lemma. The exponents $e^{-2\pi i \frac{k}{N}}$ are called *twiddle factors*. The operation count by application of the lemma is reduced from the original N^2 down to $2(N/2)^2 + N/2 = N^2/2 + N/2 < N^2$.

For $N = 2^p$ Danielson-Lanczos lemma can be applied recursively until the data sets are reduced to one datum each. The number of operations is then reduced[1] to $O(N \ln N)$ compared to the original $O(N^2)$.

---

[1]Let $T(N)$ be the number of operations to calculate size-N DFT using Cooley-Tukey algorithm. Applying the Lanczos-Danielson lemma once gives

$$T(N) = 2T(N/2) + O(N) . \tag{13.20}$$

Applying it again and again for $N = 2^p$ gives

$$
\begin{aligned}
T(N) &= 2T(N/2) + O(N) \\
&= 4T(N/4) + 2O(N) \\
&= 8T(N/8) + 3O(N) \\
&= \ldots \\
&= NT(1) + pO(N) . \tag{13.21}
\end{aligned}
$$

Now, $T(1) = 0$ as it takes no operations. That gives $T(N) = O(N \log_2 N)$.

A naïve implementation of the Cooley-Tukey algorithm in Python is shown in Table 13.1. However this implementation makes a lot of unnecessary list operations. A more refined implementation that works on thesame array is shown in Table 13.2

The established library FFT routines, like FFTW and GSL, further reduce the operation count (by a constant factor) using advanced programming techniques like precomputing the twiddle factors, effective memory management and others.

13.3   Multidimensional DFT

For example, a two-dimensional set of data $x_{n_1 n_2}$, $n_1 = 1 \ldots N_1$, $n_2 = 1 \ldots N_2$ has the discrete Fourier transform

$$c_{k_1 k_2} = \sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} x_{n_1 n_2} e^{-2\pi i \frac{n_1 k_1}{N_1}} e^{-2\pi i \frac{n_2 k_2}{N_2}} \ . \tag{13.22}$$

Table 13.2: Csharp-implementation of the Cooley-Tukey algorithm

```
using static System.Math;
using static complex;
using static cmath;

public static partial class matlib{

public static void dfts
(int sign,int N,complex[] x,int ix,int stride,complex[] c,int ic){
    for(int k=0;k<N;k++){
        c[ic+k]=0;
        for(int n=0;n<N;n++)
            c[ic+k]+=x[ix+n*stride]*exp(sign*2*PI*I*n*k/N);
    }
}

public static void ffts
(int sign,int N,complex[] x,int ix,int stride,complex[] c,int ic){
    if(N==1) c[ic+0]=x[ix+0];
    else if(N%2==0){
        ffts(sign,N/2,x,ix+0       ,2*stride,c,ic+0   );
        ffts(sign,N/2,x,ix+stride,2*stride,c,ic+N/2);
        for(int k=0;k<N/2;k++){
            complex p=c[ic+k], q=exp(sign*2*PI*I*k/N)*c[ic+k+N/2];
            c[ic+k       ]=p+q;
            c[ic+k+N/2]=p−q;
            }
        }
    else dfts(sign,N,x,ix,stride,c,ic);
    }

public static complex[] fft(complex[] x){
    int N=x.Length;
    var c=new complex[N];
    ffts(−1,N,x,0,1,c,0);
    return c;
    }

public static complex[] ift(complex[] c){
    int N=c.Length;
    var x=new complex[N];
    ffts(+1,N,c,0,1,x,0);
    for(int i=0;i<N;i++)x[i]/=N;
    return x;
    }

}//class
```

Bibliography

[1] Hiroshi Akima. A new method of interpolation and smooth curve fitting based on local procedures. *Journal of Assoc. for Comp. Mach.*, 17(4):589–602, 1970.

[2] A.N.Krylov. On the numerical solution of equation by which are determined in technical problems the frequencies of small vibrations of material systems. *Izvestiia Akademii nauk SSSR (in Russian)*, 7(4):491–539, 1931.

[3] W.E. Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Quarterly of Applied Mathematics*, 9:17–29, 1951.

[4] R.M. Barnett. Review of particle properties. *Physical Review D*, 54:1, 1996.

[5] J.P. Berrut. Rational functions for guaranteed and experimentally well-conditioned global interpolation. *Comput. Math. Appl.*, 15(1):1–16, 1988.

[6] Przemyslaw Bogacki and Lawrence F. Shampine. A 3(2) pair of Runge–Kutta formulas. *Applied Mathematics Letters*, 2(4):321–325, 1989.

[7] C.G. Broyden. A class of methods for solving nonlinear simultaneous equations. *JSTOR*, 19(92):577–593, 1965.

[8] C.W. Clenshaw and A.R. Curtis. A method for numerical integration on an automatic computer. *Numerische Mathematik*, 2:197–205, 1960.

[9] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Math. Comput.*, 19(90):297–301, 1965.

[10] P.P.M. de Rijk. A one-sided jacobi algorithm for computing the singular value decomposition on a vector computer. *SIAM J. Sci. Stat. Comput.*, 10(2):359, 1989.

[11] M. Galassi et al. *GNU Scientific Library Reference Manual*. Network Theory Ltd, 3rd edition, 2009.

[12] Erwin Fehlberg. *Low-order classical Runge-Kutta formulas with step size control and their application to some heat transfer problems.* NASA Technical Report, 1969.

[13] Wallace Givens. Computation of plane unitary rotations transforming a general matrix to triangular form. *J. SIAM*, 6(1):26–50, 1958.

[14] Gene H. Golub. Some modified matrix eigenvalue problems. *SIAM Rev.*, 15(2):318–334, 1973.

[15] Gene H. Golub and John H. Welsch. Calculation of Gauss quadrature rules. *Mathematics of Computation*, 23(106):221–230, 1969.

[16] Pedro Gonnet. Increasing the reliability of adaptive quadrature using explicit interpolants. *ACM Trans. Math. Soft.*, 37(3):26:2–26:32, 2010.

[17] A.S. Householder. Unitary triangularization of a nonsymmetric matrix. *Journal of the ACM*, 5(4):339–342, 1958.

[18] J.A.Nelder and R.Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.

[19] James Kennedy. Bare bones particle swarms. *Proceedings of the 2003 IEEE Swarm Intelligence Symposium*, page 80–87, 2003.

[20] Aleksandr Semenovich Kronrod. *Nodes and weights of quadrature formulas. Sixteen-place tables.* Consultants Bureau, 1965.

[21] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, 45:255–282, 1950.

[22] Masatake Mori. Quadrature formulas obtained by variable transformation and the DE–rule. *Journal of Computational and Applied Mathematics*, 12&13:119–130, 1985.

[23] T. N. L. Patterson. The optimum addition of points to quadrature formulae. *Mathematics of Computation*, 22(104):847–856, 1968.

[24] Robert Piessens, Elise de Doncker-Kapenga, Christoph W. Überhuber, and David Kahaner. *QUADPACK: A subroutine package for automatic integration.* Springer-Verlag, 1983.

[25] R. Poli. Analysis of the publications on the applications of particle swarm optimisation. *Journal of Artificial Evolution and Applications*, 2008:1–10, 2008.

[26] B. Riemann. Über die darstellbarkeit einer function durch eine trigonometrische reihe. *Abhandlungen der Königlichen Gesellschaft der Wissenschaften zu Göttingen*, 13:87–132, 1868.

[27] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61(arXiv:1404.7828):85–117, 2015.

[28] G.R. Farrar W.H. Press. Recursive stratified sampling for multidimensional monte carlo integration. *Computers in Physics*, 4:190–195, 1990.