

Chapter 2

Getting Started with Programming

In this text we integrate the use programming techniques and tools in our study of physics. You will therefore need to know a few programming basics in order to profit from this approach. However, if you do not have a relevant background in introductory scientific programming, do not despair. Experience shows that you can learn to program through your first physics course—many students have done this successfully and with good results. In order to prepare you for the main text, this chapter provides an introduction to programming.

2.1 A Matlab Calculator

When you start matlab, you get a window where you can type commands to be executed. Click on the Command window and type:

```
>> 9*4  
ans = 36
```

Notice the difference between the text *you* type, which is preceded by `>>`, and the results generated by the program, which are typeset without indentation.

Matlab can be used as an advanced calculator by typing expressions on the command line:

```
>> 3*2^3+4  
ans = 28
```

Standard operators are plus (+), minus (-), multiplication (*), division (/), and power (^). Powers of ten are input using *e*:

```
>> 4.5e4  
ans = 45000  
>> 2.5e-10  
ans = 2.5000e-010
```

which also shows how matlab displays numbers.

Matlab has most mathematical functions and constants built in, such as `pi`, `cos`, `sin`, and `exp`:

```
>> 4*pi
ans = 12.5664
```

Matlab uses radians for the trigonometric functions:

```
>> sin(pi/6)
ans = 0.5000
```

You can find a list of useful syntax, functions and expressions in the summary.

Matlab becomes more useful when you have a formula you want to use. For example, you may want to use the formula:

$$T_F = \frac{9}{5}T_C + 32, \quad (2.1)$$

to find the temperature, T_F , in Fahrenheit, given the temperature T_C in centigrade. We may type this formula directly into matlab:

```
>> TF = 9/5*TC + 32
???: Undefined function or variable 'TC'.
```

Ooops. That did not work, because matlab does not yet know the value of T_C . We give T_C a value and retype the formula:

```
>> TC = 40
TC = 40
>> TF = 9/5*TC + 32
TF = 104
```

Instead of retyping the formula, you can use the up arrow to find your previously typed commands and execute them again. We have now defined a variable, `TC`. You can see the value given to `TC` by typing:

```
>> TC
TC = 40
```

Notice that we assigned a value to the variable `TF` through a calculation. We have not introduced a function for `TF`. What does this mean? It means that if you change the value of `TC`, the value of `TF` will not change automatically unless you retype the formula for `TF`. You can check this by assigning `TC` a new value, and then ask matlab for the value of `TF`:

```
>> TC = 50
TC = 50
>> TF
TF = 104
```

This is an important aspect of a programming language such as matlab: a variable does not change value unless you assign a new value to it!

2.2 Scripts and Functions

However, we do not want to type in the whole formula each time we want to calculate a new value for T_F . Instead we can make a *script*, a group of several statements, or a *function*, similar to an internal function such as `sin`.

Scripts

We can group several statements into a *script*, which we can reuse. You do this by opening a new file in the File menu in matlab: File → New → Blank M-file. This opens a new window with an editor. Here you can now type (or copy) the commands we already used:

```
TC = 40;  
TF = 9/5*TC + 32
```

Notice that we added a semicolon `;` behind `TC = 40`. This stops matlab from printing out the value of `TC` after it is assigned. Now, we need to save the script. In the editor window you do: File → Save. You must give the script a name and choose where to place it. This will generate a file with an extension `.m`—we call such a file an m-file, because it shows that the file contains a matlab script/program. You run the program from the editor window by typing the `F5` key, or by pressing the green arrow on the top of the editing window. As a result the commands in the script are executed as if they were typed into the matlab window, and the resulting output is shown in the matlab window:

```
TF =    104
```

You can now change the value of `TC` in the script and rerun the script to redo the calculation for another temperature. Alternatively, you can run the script by typing the name of the script (the name you gave it when you saved it) on the command line. For example, if you called the script `convert.m`, you can run the script by typing:

```
>> convert  
TF =    104
```

Notice that we wrote the script so that the temperature `TC` is assigned inside the script. This means that if you change the value of `TC` on the command line, for example by typing:

```
>> TC = 45;
```

and then run the script—the script will not use this new value of `TC`, but instead use the value from inside the script.

Writing scripts to solve simple problems will be our standard operating procedure throughout this text. This is an efficient way to develop a simple program, change the parameters (such as changing `TC`), and rerun the program with new parameters. While this is practical for developing short programs and solving simple problems,

it is not good programming practice. In general, we encourage the development of good programming practices, but in this text we will prioritize making the code as simple as possible.

Functions

From a programming perspective, it is better to introduce a *function* to calculate the temperature. A user defined function acts just like a predefined mathematical function such as `sin` or `exp`. We define a function by opening a new m-file: Push `File` → `New` → `Blank M-file`. We define a function by typing the following into the editor:

```
function TF=convertF(TC)
TF = 9/5*TC + 32;
return
```

and save it with the name `convertF.m`.

What do these statements do? We define a function by the command `function`. First, we write what the function should return. Here, the function returns the variable `TF`. We introduce the name of the function, `convertF`, and the arguments that we need to input when we use the function—here the only argument is the temperature in centigrades, `TC`. Inside the function we must calculate the value of `TF`, because this is the value the function is supposed to calculate.

We call our new function by typing:

```
>> convertF(45)
ans = 113
```

Notice that matlab requires each such user-defined function to be in a separate file, and that the file must be in the search path for matlab. This means that the file `convertF.m` must be in the current directory or in the standard matlab directory for this to work. I suggest that you always save the functions you need in the same directory as you save the scripts you are currently working on, and that you make new directories for each problem.

A particular feature of functions is that the internal calculations and variables used inside the function are lost as soon as the function is finished. For example, matlab may use several calculation steps if we call the `sin` function, but this is hidden from us. Outside the function, we only see the result of the function. To illustrate this, we could break our short function into several steps:

```
function TF=convertF2(TC)
ratio = 9/5;
constant = 32;
TF = factor*TC+ constant;
return
```

Here, we have introduced two internal variables, `ratio` and `constant`, that are forgotten as soon as the function is finished. For example, if we type:

Table 2.1 Measurement of masses m_i as a function of volumes V_i

i	1	2	3	4	5	6
m_i (kg)	1	2	4	6	9	11
V_i	0.13l	0.26l	0.50l	0.77l	1.15l	1.36l

```
>> convertF2(45)
ans = 113
>> ratio
???' Undefined function or variable 'ratio'.
```

we see that matlab does not know the value of `ratio` after the function has done its work.

Functions are powerful and necessary tools of more advanced programming techniques and will be gradually introduced throughout the text. But initially we try to make the programs as simple as possible, and we will therefore use simple scripting as our main tool.

2.3 Plotting Data-Sets

Matlab not only works as a numerical calculator, it also has advanced data visualization capabilities. For example, as part of a laboratory exercise you may have measured the volume and mass of a set of steel spheres. You number the measurements using the index, i , and record masses m_i and volumes V_i in Table 2.1, where we have used that 1 litre = 1l = 1 dm³.

Such a sequence of numbers are stored in an *array* (or a *vector*) in matlab. We define the sequence of masses and volumes in matlab using

```
>> m = [1 2 4 6 9 11];
>> V = [0.13 0.26 0.50 0.77 1.15 1.36];
```

We can find an individual mass value by:


```
>> m(1)
ans = 1
>> m(4)
ans = 6
```

There are now 6 values for the masses, numbered $m(1)$ to $m(6)$. Notice that matlab starts enumeration at 1, so that 1 is the first element and 6 is the last element. We call the array m a 6×1 array or a vector of length 6. The volumes are stored the same way:

```
>> V(1)
ans = 0.1300
>> V(4)
ans = 0.7700
```

The enumeration of the two arrays is identical: element $m(4)$ of the masses corresponds to element $V(4)$ of the volumes.

Fig. 2.1 The plot of V as a function of m (left) and with the plot of $\sin(x)$ as a function of x for 10 points (solid line) and for 1000 points (dotted line) (right)



The relation between m and V is illustrated by plotting V as a function of m . This is done by the `plot` command:

```
>> plot(m,V, 'o')
```

where the string `'o'` ensures that a small circle is plotted at each data-point. The `plot` command makes a “scatter” plot—it contains a point for each of the data-points $m(i)$, $V(i)$ in the two arrays. The two arrays must therefore be the same length—they must have the same number of elements. We annotate the axes by:

```
>> xlabel('m (kg)')
>> ylabel('V (l)')
```

where the `xlabel` refers to the first array m in `plot(m, V, 'o')`, and the `ylabel` refers to the second array—the V array. The resulting plot is shown in Fig. 2.1.

Where did the units (kg and liters) go when we defined the mass m and the volume V ? We cannot use units when we introduce digital representations of the numbers. We can only input numbers into `matlab`, and we have to keep track of the units. This is why we specified the units along the axes in the `xlabel` and `ylabel` commands.

2.4 Plotting a Function

`Matlab` cannot plot a function such as $\sin(x)$ directly. We must first generate two sequences of numbers, one sequence for the x 'es and one sequence for the corresponding values of $\sin(x)$, and then plot the two sequences against each other. While this may sound complicated, `matlab` has functions that ensure that you can almost directly write the mathematical expression into `matlab`.

Loops

We want to make a sequence of x 'es, such as 0.0, 0.1, 0.2, 0.3, ... etc, and then for each x_i we want to calculate the corresponding value for $\sin(x_i)$ in Table 2.2, where we generate x_i from 0.0 to 10.0 in steps of 0.1.

Table 2.2 Sequence of i , x_i and $\sin(x_i)$

i	1	2	3	4	5	...	n
x_i	0.0	0.1	0.2	0.3	0.4	...	10.0
$\sin(x_i)$	$\sin(0.0)$	$\sin(0.1)$	$\sin(0.2)$	$\sin(0.3)$	$\sin(0.4)$...	$\sin(10.0)$

How do we generate such an array in matlab? First, we have to generate the array.¹ How many elements do we need? Going from 0.0 to 10.0 in steps of 0.1 we need:

$$n = \frac{10.0 - 0.0}{0.1} + 1, \quad (2.2)$$

steps, where we have added one in order to include the last step (otherwise we would stop at 9.9 instead of at 10.0. We define an array of this length by:

```
>> n = ceil((10.0-0.0)/0.1)+1;
>> x = zeros(n,1);
```

Here, the function `ceil()` rounds up. The function `zeros(n,1)` generates and returns an array of size n by 1 which is filled with zeros. Now we need to fill the array:

```
>> x(1) = 0.0;
>> x(2) = 0.1;
>> x(3) = 0.2;
>> x(4) = 0.3;
>> x(5) = 0.4;
...

```

Fortunately, there is a more efficient way of doing this—by using a `for`-loop. A `for`-loop allows us to loop through a list of values 1, 2, 3, ..., n for the variable i , and then execute a set of commands at each step—exactly what we need. We can replace the long list of `x(1) = 0.0;` etc by the loop:

```
>> for i = 1:n
    x(i) = (i-1)*0.1;
end
```

If you type this in, nothing will execute until you type `end` Also note that the semicolon `;` is more important here, otherwise you will get 101 results written in your matlab window...

You can check the generated values of `x` by:

```
>> x
x =
Columns 1 through 7
    0.0000    0.1000    0.2000    0.3000    0.4000    0.5000    0.6000
    ....

```

¹In matlab it is not necessary to define the size of the array before it is filled. We could just fill it as we go along, but this is not good coding practice, it will lead to very slow codes for large arrays, and may cause surprising errors in your programs. We will therefore always predefine the size of arrays.

Notice how we specify the range of the loop, by specifying a sequence of numbers `1:n`. Typing `1:n` at the command prompt gives you exactly the list of values for `i`.

Now, let us put this into a small script. And let us also calculate the value for the function $\sin(x)$:

```
x0 = 0.0; x1 = 10.0; dx = 0.1;
n = ceil((x1-x0)/dx) + 1;
x = zeros(n,1);
y = zeros(n,1);
for i = 1:n
    x(i) = x0 + (i-1)*dx;
    y(i) = sin(x(i));
end plot(x,y);
```

Which both generates and plots the function $\sin(x)$. The variables `x0`, `x1`, and `dx` provide the start, stop and step of the x -values used. You can now change them and rerun the script to generate plots for other ranges or with other resolutions.

Notice that `y(i) = sin(x(i))` must appear inside the loop—that is before the `end`. Otherwise it would only be executed once, using the value `i` had at the end of the loop. Putting commands outside a loop that should be inside a loop is a common mistake—sometimes also done by experienced programmers.

The while-loop

The `for`-loop is probably the loop-structure you will use the most, but there are also other tools for making a loop, such as the `while`-loop. In the `while`-loop the commands inside the loop are executed until the expression in the `while` command is true. It does not automatically update a counter either. For example, we can change the script to calculate $\sin(x)$ to use a `while`-loop instead by the following changes:

```
...
y = zeros(n,1);
i = 0;
while i<=n
    i = i + 1;
    x(i) = x0 + (i-1)*dx;
    y(i) = sin(x(i));
end
...
```

Notice that we must assign `i=0` before the loop and `i=i+1` inside the loop, since we now need to update the counter “manually” inside the loop. We also introduce an “expression” `i<=n` which may be false (having the value 0) or true (having the value 1). The loop continues until the expression becomes false. Notice that a common source of error is to generate `while` loops that continue forever, for example because you have forgotten to update the counter inside the loop. You will notice this because your program never ends: Matlab will show “Busy” in the lower left corner of its window and never stop or plot your results.

You may wonder what the point of the `while`-loop is, since it looks more cumbersome than the `for`-loop. We will use the `while` loop when we want to continue a calculation for an unknown number of steps. For example, you may want to find

the motion of a falling ball up to the point where it hits the ground. However, you may not know beforehand how many steps are needed before it hits the ground. If the position of the ball is $x(t) = 1000 - 4.9t^2$, we can calculate the position as a function of time in steps of dt as long as x is positive using the following program:

```
t0 = 0.0; t1 = 10.0; dt = 0.01;
n = ceil((t1-t0)/dt) + 1;
t = zeros(n,1);
y = zeros(n,1);
i = 1;
t(1) = t0;
y(1) = 100.0-4.9*t(1)^2;
while y(i)>0.0
    i = i + 1;
    t(i) = t0 + (i-1)*dt;
    y(i) = 100.0-4.9*t(i)^2;
end
```

This script needs some explanation. First, we notice we update both now is t and y and not only x as before. In addition, we see that we calculate the value of $y(1)$ before the loop starts, otherwise the loop would never start. This is also a common mistake. Therefore, ensure that you understand why and what is done before the `while`-loop starts in this script.

Vectorization

While loops are generally powerful and useful methods, there is a much simpler way to generate sequences of numbers and plot functions in matlab—and the method also allows your code to follow the mathematical formulas more closely. This method is called *vectorization*.

We can make a sequence of x 'es in several ways using functions that are built into matlab instead of using a loop. For example, the function `linspace` generates a sequence of equidistant numbers from 0.0 to 10.0:

```
>> x = linspace(0,10,10)
x =
    Columns 1 through 10
    0 1.111 2.222 3.333 4.444 5.555 6.666 7.777 8.888 10.000
```

In this case we generated 10 numbers, but you are free to choose your own resolution. An alternative to specifying the number of points you want, as we do with `linspace`, is to specify the step size. The expression `(0.0:0.3:10.0)` returns an array starting at the value 0 and ending at 10.0 in steps of 0.3²:

```
>> x = (0:0.3:10.0)
ans =
    Columns 1 through 7
    0 0.3000 0.6000 0.9000 1.2000 1.5000 1.8000
    ....
```

²Notice the small difference between the two methods: Using `linspace` ensures that the first and the last numbers are included in the list, but when you use `(0:0.3:10.0)` the last number is 9.9 and not 10.0!

Ok—so that was simply an easier way of generating the array x . Why all the fuzz? Because of a powerful feature of matlab called vectorization: We can apply the function $\sin(x)$ to the whole array x . Matlab will then apply the function to each of the elements in x and return a new array with the same number of elements as x . The three lines:

```
>> x = linspace(0,10,10);
>> y = sin(x);
>> plot(x,y);
```

are equivalent to the program:

```
x0 = 0.0;
x1 = 10.0;
dx = 1.0;
n = ceil((x1-x0)/dx) + 1;
x = zeros(n,1);
y = zeros(n,1);
for i = 1:n
    x(i) = x0 + (i-1)*dx;
    y(i) = sin(x(i));
end plot(x,y);
```

Notice how simple the vectorized code is—it is almost identical to the mathematical formula. We only have to define the range of x -values before we call the $\sin(x)$ function. Beautiful and powerful.

The program above generates the solid line in Fig. 2.1. However, this plot has too sharp corners, because we have too few data-points. Let us generate 1000 x -values and plot $\sin(x)$ with this resolution in the same plot:

```
>> x = linspace(0,10,1000);
>> y = sin(x);
>> hold on;
>> plot(x,y,':');
>> hold off;
```

The result is shown in Fig. 2.1 with a dotted line. Here, we have used a few more tricks. We use the command `hold on` to ensure that matlab does not generate a new plot, which would remove the previous one, but instead plots the data in the same plot as we have already used. Typing `hold off` stops this behavior—otherwise all subsequent plots will be part of the same plot. We have also used the string `':'` to tell matlab that we want a dotted line. You can find more plotting methods in the summary at the end of the chapter.

The vectorization technique is very general, and usually allows you to translate a mathematical formula to matlab almost formula by formula. For example, we plot the function

$$f(x) = x^2 e^{-ax} \sin(\pi x), \quad (2.3)$$

from $x = 0$ to $x = 10$ by typing (when $a = 1$):

```
>> x = linspace(0,10,1000);
>> a = 1.0;
>> f = x.^2.*exp(-a*x).*sin(pi*x);
>> plot(x,f)
```

Again, we have introduced a few more tricks of the trade. Notice that we do not type multiply `*`, but we use the special notation `.*`, which means *element-wise* multiplication. If you only type `*` matlab will assume that you want to do a matrix multiplication (as you may know from a course on linear algebra)—which is something else. Here, we want to perform an element-wise multiplication, that is, we want to find $f(x_i)$ from:

$$f(x_i) = x_i^2 e^{-ax_i} \sin(\pi x_i), \quad (2.4)$$

where the operation is done for each of the elements, x_i , in the array `x`. If you do not use element-wise multiplication you will typically get a strange error message from matlab.³ You should therefore use element-wise multiplication by default. To illustrate what happens during element-wise multiplication type:

```
>> x = [1 2 3 4];
>> y = x*x
??? Error using ==> mtimes
Inner matrix dimensions must agree.
```

However, if you use element-wise multiplication, the expected result is found:

```
>> x = [1 2 3 4];
>> y = x.*x
y =      1      4      9     16
```

As soon as you have learned to transcribe mathematical expressions from the mathematical notation to matlab you are ready to calculate and plot any function.

The technique of vectorization is a powerful and efficient technique. Matlab is usually very fast at calculating vectorized commands. And we can write very elegant programs using such techniques, ensuring that the matlab code follows the mathematical formulation closely, which makes the code easy to understand.

2.5 Random Numbers

Sometimes you need randomness to enter your physical simulation. For example, you may want to model the motion of a tiny dust of grain in the air bouncing about due to random hits by the air molecules, so called Brownian motion. We model this by having the grain move a *random* distance during a given time interval. How do we create random numbers on the computer? Unfortunately, we cannot generate really random numbers, but most programs have decent pseudo-random number generators, which generate a sequence of numbers that appear to be random. In matlab, we can simulate the throw of a dice using

```
>> randi(6)
ans =      3
```

where `randi(n)` generates a random integer between 1 and n , where each outcome has the same probability. If you type the command several times, you will get a new

³But sometimes you may by accident succeed in doing a matrix multiplication, and you may get rather surprising errors. This is a type of error that may be difficult to track down.

answer each time. Matlab includes several functions that return random numbers: It can generate random real numbers between 0 and 1 using the `rand` function and normal-distributed numbers (with average 0 and standard derivation 1) using the function `randn`.

2.6 Conditions

Now, if we return to discuss the motion of the grain of dust, we want to model its motion according to a simple rule: I throw a dice. If I get between 1 and 3, the grain moves a step forward, otherwise it moves a step backward. How can we handle such conditions? We need a set of conditional statements, so that we can perform a given set of commands when a particular condition is fulfilled. We need an `if`-statement:

```
if (expr)
    <statement a1>
    <statement a2>
    ..
else
    <statement b1>
    <statement b2>
    ..
end
```

Here the expression (`expr`) is an expression such as `randi(6)>3` which may be true or false. If the expression is true, statements `a1`, `a2`, ... are executed, otherwise the statements `b1`, `b2`, ... are executed.

Let us use this to find the motion of the grain. Every time we throw the dice, the grain moves a distance $dx = \pm 1$. If the grain is at position x_i at step i , the grain will be at a position

$$x_{i+1} = x_i + dx . \quad (2.5)$$

at step $i + 1$. We can use this rule and an `if`-statement to write the script to find the position at subsequent steps $i = 1, 2, 3, \dots$:

```
n = 1000;
x = zeros(n,1);
for i = 1:n-1
    if (randi(6)<=3)
        dx = -1;
    else
        dx = 1;
    end
    x(i+1) = x(i) + dx;
end
```

The resulting motion is shown in Fig. 2.2.

We will use `if`-statements throughout the text, often to enforce particular conditions on the motion. For example, we could add a level of complexity to the motion of the grain by requiring that the grain moved inside a narrow channel of width 10: The grain cannot move outside a region spanning from -5 to $+5$:




Fig. 2.2 Plot of the position $x(i)$ of a random walker (a bouncing grain of dust) as a function of the number of steps i done (*left*), and when the walker is constrained to the zone $-5 \leq x \leq +5$ (*right*)

```
n = 1000;
x = zeros(n,1);
for i = 1:n-1
    if (randi(6)<=3)
        dx = -1;
    else
        dx = +1;
    end
    x(i+1) = x(i) + dx;
    if (x(i)>5)
        x(i) = 5;
    end
    if (x(i)<-5)
        x(i) = -5;
    end
end plot(x), xlabel('i'), ylabel('x(i)')
```

The resulting motion x_i as a function of i is shown in Fig. 2.2.

For the interested reader, we include a particularly compact formulation of the random walk, which you may have fun trying to understand.

```
>> x = cumsum(2*(randi(6,1000,1)<=3)-1);
>> plot(x)
```

You can research this expression by using the `help`-function in matlab:

```
>> help cumsum
```

2.7 Reading Real Data

When you work with physics you need to handle real data: NASA publishes data for the motion of most stellar objects; your mobile phone has an accelerometer; and a GPS that measures thousands of data-points in a few seconds. You do not want to type in these numbers by hand. Therefore you need to be able to read files containing

numbers. For example, the motion of a sprinter running 100m is given in the file 100m.d.⁴ The file looks like this if you open it in a text editor (such as emacs):

```
0.0000000e+000 -2.1155775e-001
1.0000000e-002 -1.7485406e-001
2.0000000e-002 -1.3798607e-001
3.0000000e-002 -1.0095306e-001
4.0000000e-002 -6.3754256e-002
5.0000000e-002 -2.6388915e-002
...
```

A total of 972 lines of data. The first column gives the time, measured in seconds, and the second column gives the position of the runner, measured in meters. Fortunately, it is very simple to read such as file into matlab. It is done by a single command:

```
>> load run100m.d
```

This generates a variable `run100m` with 972 rows and two columns, as seen by the command `whos`:

```
>> whos run100m
Name           Size           Bytes   Class   Attributes
run100m        971x2           15536   double
```

We split the data into two arrays `t` and `x` by:

```
>> t = run100m(:,1);
>> x = run100m(:,2);
```

and plot the data using

```
>> plot(t,x)
```

If you experience a problem where matlab cannot find the file, getting an error message like:

```
>> load run100m.d
??? Error using ==> load
Unable to read file run100m.d: No such file or directory.
```

It means that the file `run100m.d` is not in your current working directory. By default your current working directory is shown to the left in your matlab window—if not `selectWindow` → `2 Current Directory` from the menu. Here, you should see the file `run100m.d` listed. If not, you need to change the current working directory to where you placed the file when you downloaded it—and try again.

2.7.1 Example: Plot of Function and Derivative

Problem: Plot the function

$$f(x) = e^{-x^2}, \quad (2.6)$$

and its derivative by using the formula:

⁴<http://folk.uio.no/malthe/mechbook/100m.d>




Fig. 2.3 Plot of $f(x)$ as a function of x and its derivative df/dx as a function of x calculated using a numerical method

$$f'(x) \simeq \frac{f(x+h) - f(x-h)}{2h}, \quad (2.7)$$

as an approximation for the derivative on the interval $-5 \leq x \leq 5$. You may use the value $h = 0.001$ for h .

Solution: The function can be plotted directly by a vectorized approach:

```
>> x = linspace(-5,5,1000);
>> f = exp(-x.^2);
>> plot(x,f);
```

In order to use the numerical approximation for the derivative, we need to perform the approximation for each of the x -values in the x -array. We access them by a `for`-loop through the 1000 elements in the x -array:

```
>> h = 0.001;
>> df = zeros(1000,1);
>> for i = 1:1000
...   df(i) = (exp(-(x(i)+h)^2)-exp(-(x(i)-h)^2))/(2*h); ...end
>> plot(x,df);
```

The resulting plot is shown in Fig. 2.3.

Even simple problems such as these are useful to implement as scripts saved in a file, since this makes debugging—the process of finding and removing errors in the script—simpler. If you make a small mistake, you have to retype all the commands when you operate on the command line, but if you use a script, you simply make a small change in the script, rerun, and that is it.

Summary

Using matlab as a calculator:

- Direct calculations are done on the command line


```
>> 10.0*sin(pi/3)+4.0^3
```
- Defining and reusing variables

```
>> a = 2.0; b = 4.5; c = a^2 + b^2
```

- Vectorized plotting of functions

```
>> x=linspace(0,10,0.01); y=exp(-x).*sin(x); plot(x,y);
```

- Vectorized operations are denoted by a leading `.` and are done *element-wise*.

Functions and scripts:

- A script is a sequence of executable commands stored in a separate `.m`-file

- All variables are available on from the command line afterwards
- The script is run by typing F5 in the editor
- Scripts allow rapid rerunning a program after changes in parameters

- A function must be stored in a separate `.m` file. It has the syntax

```
function y = myfunction(a,b,c)
v = a.*b.*c;
d = v.^2;
y = 2*d;
return
```

- The name of the function (`myfunction`) should be the name of the `.m`-file. Variables defined inside the function, such as `v` and `d` are not available outside the function

Plotting:

- You plot two arrays `t` and `x` versus each other by

```
plot(t,x,'-b'); xlabel('t (s)'); ylabel('x (m)');
```

- Plotting markers and symbols are listed in Table 2.3.
- Plotting several data-sets in the same plot:

```
plot(t1,x1,'-b',t2,x2,'-r')
```

or

```
plot(t1,x1,'-b');
hold on; plot(t2,x2,'-r'); hold off
```

- Combining several plots into one figure:

```
subplot(2,1,1); plot(t1,x1,'-b');
subplot(2,1,2); plot(t2,x2,'-r');
```

- Saving a figure to a file: either by using the save button from the figure window, or you can save a figure `myfigure.pdf` as a pdf from the command line by

```
saveas(gcf,'myfigure.pdf','pdf');
```

Table 2.3 Line markers, colors and plotting symbols

Colors	Lines	Symbols	Symbols				
b	Blue	-	Solid	.	Point	v	Triangle (down)
g	Green	:	Dotted	o	Circle	^	Triangle (up)
r	Red	- .	Dashdot	x	X-mark	<	Triangle (left)
c	Cyan	-	Dashed	+	Plus	>	Triangle (right)
m	Magenta	(None)	No line	*	Star	p	Pentagram
y	Yellow			d	Diamond	h	Hexagram
k	Black						
w	White						

Loops:

- **for-loops** run a counter sequentially through a list of values:

```
for i = 1:100
    x(i) = sin(i/100.0);
end
```

- **while-loops** run until a given expression is true

```
i = 0; while (i<100)
    i = i + 1;
    x(i) = sin(i/100.0);
end
```

Expressions:

- **if-statements** are used to run a sequence of commands given a particular expression is true:

```
if (x>10.0)
    y = 10.0;
else
    y = -10.0;
end
```

- Expressions return true (1) or false (0) and can be joined using logical operators such as *or* and *and* as shown in Table 2.4

Table 2.4 Expressions and operators used in matlab

Expression	Name	Example	Operator	Name	Example
==	Equal	(x==0.0)	&&	Logical AND	(x==0.0) && (y>0.0)
~=	Not-equal	(x~=0.0)		Logical OR	(x~=0.0) (y>0.0)
>=	Greater than or equal	(x>=0.0)			
<=	Less than or equal	(x<=0.0)			
>	Greater than	(x>0.0)			
<	Less than	(x<0.0)			

Exercises

2.1 Seconds.

(a) Write a script that calculates the number of seconds, s , given the number of hours, h , according to the formula $s = 3600h$.

(b) Use the script to find the number of seconds in 1.5, 12, and 24 h.

2.2 Spherical mass.

(a) Write a script that calculates the mass of a sphere given its radius r and mass density ρ according to the formula $m = (4\pi/3) \rho r^3$.

(b) Use the script to find the mass of a sphere of steel of radius $r = 1$ mm, $r = 1$ m, and $r = 10$ m.

2.3 Angle.

(a) Write a function that for a point (x, y) returns the angle θ from the x -axis using the formula $\theta = \arctan(y/x)$.

(b) Find the angles θ for the points $(1, 1)$, $(-1, 1)$, $(-1, -1)$, $(1, -1)$.

(c) How would you change the function to return values of θ in the range $[0, 2\pi]$?

2.4 Unit vector.

(a) Write a function that returns the two-dimensional unit vector, (u_x, u_y) , corresponding to an angle θ with the x -axis. You can use the formula $(u_x, u_y) = (\cos \theta, \sin \theta)$, where θ is given in radians.

(b) Find the unit vectors for $\theta = 0, \pi/6, \pi/3, \pi/2, 3\pi/2$.

(c) Rewrite the function to instead take the argument θ in degrees.

2.5 Plotting the normal distribution. The normal distribution, often called the Gaussian distribution, is given as:

$$P(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/(2\sigma^2)}, \quad (2.8)$$

where μ is the average and σ is the standard deviation.

(a) Make a function `normal(x, mu, sigma)` that returns the normal distribution value, $P(x, \mu, \sigma)$ as given by the formula.

(b) Use this function to plot the normal distribution for $-5 < x < 5$ for $\mu = 0$ and $\sigma = 1$.

(c) Plot the normal distribution for $-5 < x < 5$ for $\mu = 0$ and $\sigma = 2$ and for $\sigma = 0.5$ in the same plot.

(d) Plot the normal distribution for $-5 < x < 5$ for $\sigma = 1$ and $\mu = 0, 1, 2$ in three subplots above each other.

2.6 Plotting $1/x^n$. The function $f(x; n)$ is given as $f(x; n) = x^{-n}$.

(a) Make a function `fvalue(x, n)` which returns the value of $f(x; n)$.

(b) Use this function to plot $1/x$, $1/x^2$ and $1/x^3$ in the same plot for $-1 < x < 1$.

2.7 Plotting $\sin(x)/x^n$. The function $g(x; n)$ is given as:

$$g(x; n) = \frac{\sin(x)}{x^n}. \quad (2.9)$$

(a) Make a function `gvalue(x, n)` which returns the value of $g(x; n)$.

(b) Use this function to plot $\sin(x)/x$, $\sin(x)/x^2$ and $\sin(x)/x^3$ in the same plot for $-5 < x < 5$.

(c) Use the help function to find out how to place legends for each of the plots into the figure.

2.8 Logistic map. The iterative mapping $x(i+1) = r x(i) (1 - x(i))$ is called the logistic map.

(a) Make a function `logistic(x, r)` which returns the value of $x(i+1)$ given $x(i)$ and r as inputs.

(b) Write a script with a loop to calculate the first 100 steps of the logistic map starting from $x(1) = 0.5$. Store all the values in an array `x` with $n = 100$ elements and plot x as a function of the number of steps i :

(c) Explore the logistic map for $r = 1.0, 2.0, 3.0$ and 4.0 .

2.9 Euler's method. In mechanics, we often use Euler's method to determine the motion of an object given how the acceleration depends on the velocity and position of an object. For example, we may know that the acceleration $a(x, v)$ is given as $a(x, v) = -kx - cv$. If we know the position x and the velocity v at a time $t = 0$: $x(0) = x_0 = 0$ and $v(0) = v_0 = 1$, we can use Euler's method to find the position and velocity after a small timestep Δt :

$$v_1 = v(t_0 + \Delta t) = v(t_0) + a(v(t_0), x(t_0))\Delta t \quad (2.10)$$

$$x_1 = x(t_0 + \Delta t) = x(t_0) + v(t_0)\Delta t \quad (2.11)$$

$$v_2 = v(t_1 + \Delta t) = v(t_1) + a(v(t_1), x(t_1))\Delta t \quad (2.12)$$

$$x_2 = x(t_1 + \Delta t) = x(t_1) + v(t_1)\Delta t \quad (2.13)$$

and so on. We can therefore use this scheme to find the position $x(t)$ and the velocity $v(t)$ as function of time at the discrete values $t_i = i \Delta t$ in time.

- (a) Write a function `acceleration(v, x, k, C)` which returns the value of $a(x, v) = -kx - Cv$.
- (b) Write a script that calculates the first 100 values of $x(t_i)$ and $v(t_i)$ when $k = 10$, $C = 5$, and $\Delta t = 0.01$. Plot $x(t)$, $v(t)$, and $a(t)$ as functions of time.
- (c) What would you need to change to instead find $x(t)$ and $v(t)$ if the acceleration was given as $a(v, x) = k \sin(x) - Cv$?

2.10 Throwing two dice. You throw a pair of six-sided dice and sum the number from each of the dice: $Z = X_1 + X_2$, where Z is the sum of the results from dice 1, X_1 , and dice 2, X_2 . If we perform this experiment many times (N), we can find the average and standard deviation from standard estimators from statistics. The average, $\langle Z \rangle$, of Z is estimated from:

$$\langle Z \rangle = \frac{1}{N} \sum_{j=1}^N Z_j, \quad (2.14)$$

and the standard deviation, ΔZ , is estimated from:

$$\Delta Z = \frac{1}{N-1} \left(\sum_{j=1}^N (Z_j - \langle Z \rangle)^2 \right)^{1/2}. \quad (2.15)$$

- (a) Write a function that returns an array of N values for Z .
- (b) Write a function that returns an estimate of the average of an array z using the formula provided.
- (c) Write a function that returns an estimate of the standard deviation of an array z using the formula provided.
- (d) Find the average and standard deviation for $N = 100$ throws of two dice.

2.11 Reading data. The file `trajectory.dat`⁵ contains a list of numbers:

```
t0 x0 y0
t1 x1 y1
.. .. ..
tn xn yn
```

corresponding to the time $t(i)$ measured in seconds, and the x and y positions $x(i)$ and $y(i)$ measured in meters for the trajectory of a projectile.

- (a) Read the data file into the arrays 't, x, and y'.
- (b) Plot the x and y positions as function of time in two plots above each other.
- (c) Plot the (x, y) position of the object in a plot with x and y on the two axes.

⁵<http://folk.uio.no/malthe/mechbook/trajectory.dat>

2.12 Numerical integration of a data-set. The file `velocity.dat`⁶ contains a list of numbers:

```
t0 v0
t1 v1
.. .. ..
tn vn
```

corresponding to the time $t(i)$ measured in seconds, and the velocity $v(i)$ measured in meters per second for the trajectory of a projectile.

(a) Read the data file into the arrays `t`, and `v`.

(b) Plot $v(t)$ as function of time.

For a data-set $t(i)$, $v(i)$, you can estimate the function corresponding to the integral of $v(t)$ with respect to t at the times t_i using the iterative scheme:

$$y(t_1) \simeq y(t_0) + v(t_0) (t_1 - t_0) \quad (2.16)$$

$$y(t_2) \simeq y(t_1) + v(t_1) (t_2 - t_1) \quad (2.17)$$

$$\dots \simeq \dots \quad (2.18)$$

$$y(t_n) \simeq y(t_{n-1}) + v(t_{n-1}) (t_n - t_{n-1}) \quad (2.19)$$

where $v(t_i) = 'v(i)'$ and $t_i = 't(i)'$. You can assume that the motion starts at $y(t_0) = 0.0\text{m}$ at $t = t_0$.

(c) Write a script to calculate the time integral $y(t_i)$ of the dataset using this formula. Implement using a `for`-loop.

(d) Plot the position $y(t)$ and the derivative $v(t)$ as functions of time in two plots above each other.

⁶<http://folk.uio.no/malthe/mechbook/velocity.dat>