INTRODUCTION TO MATLAB

Ross L. Spencer

Department of Physics and Astronomy

Brigham Young University

This is a tutorial to help you get started in Matlab. To find more details see the very helpful book *Mastering MATLAB 6* by Duane Hanselman and Bruce Littlefield. Examples of Matlab code in this pamphlet are in typewriter font `like this`. As you read through the sections below type and execute in Matlab all of the examples, either at the ≫ command line prompt or in a test program you make called `test.m`. Longer sections of code are flagged with the characters %begin and %end. All of the Matlab code between these two flags can be found in the file `tutorial.m` which you can find on the Physics 330 web page at www.physics.byu.edu.

This booklet can also be used as a reference manual because it is short, it has lots of examples, and it has a table of contents and an index. It is almost true that the basics of Matlab are in sections 1-9 while physics applications are in sections 9-17. Please tell me about mistakes and make suggestions to improve it (ross_spencer@byu.edu).

Contents

1 Running Matlab

1.1 Starting

Get on a department PC or buy Student Matlab for your own machine and start the program. *Note: the student version of Matlab is cheap, powerful, and even has part of Maple in it. You should buy it while you still have a student ID because after you graduate it's very expensive. Try www.mathworks.com.*

1.2 It's a Calculator

You can use Matlab as a calculator by typing commands at the $\gg$ prompt, like these. Try them out.

```
1+1
2*3
5/6
exp(-3)
atan2(-1,2)
```

And just like many hand calculators, `ans` in Matlab means the last result calculated (like % in Maple).

```
sin(5)
ans
```

Note that Matlab's trig functions are permanently set to radians mode. Note also that the way to enter numbers like 1.23×10^{15} is

```
1.23e15
```

And here's another useful thing. The up-arrow key $\uparrow$ will display previous commands. And when you back up to a previous command that you like, hit Enter and it will execute. Or you can edit it when you get to it (use $\leftarrow$, $\rightarrow$, and Del), then execute it. Try doing this now to re-execute the commands you have already typed.

1.3 Making Script Files

Most of the work you will do in Matlab will be stored in files called *scripts* containing Matlab commands to be executed over and over again. To make a script, first browse or type in the current directory window on the tool bar to put yourself in the directory where you want to store your Matlab scripts. Then open a new text file in the usual way by clicking on the empty document on the tool bar, or open an old one to be edited. A script can be filled with a sequence of Matlab commands that will be executed from top to bottom just as if you had typed them on the command screen. These files have .m extensions (automatically added in Windows) and can be executed by typing their names (without the .m extension) in the command window. For example, if you create a script called `test.m` you can execute it in the command window by typing `test`. *Do not choose file names that start with numbers, like* `430lab1a.m`*.* When Matlab receives the start of a number on the command line it thinks a calculation is coming, and since `430lab1a` is not a valid calculation Matlab will give you an error.

During a session keep this file open so you can modify and debug it. And remember to save the changes you make (Ctrl-s is a quick way) or Matlab in the command window won't know that you have made changes to the script.

Document your code by including lines in it that begin with %, like this. (Note: don't try to execute these lines of code; they just illustrates how to put comments in.)

```
% This is a comment line
```

Or you can put comments at the end of a line of code like this:

```
f=1-exp(-g*t)    % compute the decay fraction
```

You may need to type lines into your script that are too long to see well. To make the code look better you can continue program lines onto successive lines by using the ... syntax: (Note: don't try to execute these lines of code; they just illustrate how to continue long lines.)

```
a=sin(x)*exp(-y)*...
  log(z)+sqrt(b);
```

Finally, nearly always begin your scripts with the `clear` command. This is like Maple's restart and makes sure that you don't have leftover junk active in Matlab that will interfere with your code.

1.4 Running Script Files

When you have a script written and saved, go to the window with the Matlab command prompt ≫ and type the name of your file without the `.m` extension, like this:

```
test
```

Your script will then run. (Or you can use the `debug` menu in the editor, or the "Save and Run" shortcut key, `F5`.)

1.5 Pause command

A `pause` command in a script causes execution to stop temporarily. To continue just hit Enter. You can also give it a time argument like this

```
pause(.2)
```

which will cause the script to pause for 0.2 seconds, then continue. And, of course, you can ask for a pause of any number or fractions of seconds. Note, however, that if you choose a really short pause, like 0.001 seconds, the pause will not be so quick. Its length will be determined instead by how fast the computer can run your script.

1.6 Online Help

If you need to find out about something in Matlab you can use `help` or `lookfor` at the ≫ prompt. There is also a wealth of information under `Help Desk` in the `Help` menu of Matlab's command window. For example maybe you are wondering about the `atan2` function mentioned in Sec. 1.2. Type

```
help atan2
```

at the ≫ prompt to find information about how this form of the inverse tangent function works. Also type

```
help bessel
```

to find out what Matlab's Bessel function routines are called. `Help` will only work if you know exactly how Matlab spells the topic you are looking for.

`Lookfor` is more general. Suppose you wanted to know how Matlab handles elliptic integrals. `help elliptic` is no help, but typing

```
lookfor elliptic
```

will tell you that you should use

```
help ellipke
```

to find what you want.

1.7   Making Matlab Be Quiet

Any line in a script that ends with a semicolon will execute without printing to the screen. Try, for example, these two lines of code in a script or at the ≫ command prompt.

```
a=sin(5);
b=cos(5)
```

Even though the variable `a` didn't print, it is loaded with $\sin(5)$, as you can see by typing this:

```
a
```

1.8   Debugging

When your script fails you will need to look at the data it is working with to see what went wrong. In Matlab this is easy because after you run a script all of the data in that script is available at the Matlab ≫ prompt. So if you need to know the value of `a` in your script just type

```
a
```

and its value will appear on the screen. You can also make plots of your data in the command window. For example, if you have arrays x and y in your script and you want to see what $y(x)$ looks like, just type `plot(x,y)` at the command prompt.

The following Matlab commands are also useful for debugging:

```
who        lists active variables
whos       lists active variables and their sizes
what       lists .m files available in the current directory
```

1.9 Arranging the Desktop

A former student, Lance Locey, went directly from this Introduction to Matlab to doing research with it and has found the following way of arranging a few of Matlab's windows on the desktop to be very helpful. (A visual representation of this layout appears on the next page.)

1. Make the command window wide and not very tall, stretching across the bottom of the desktop.

2. Open a script editing window (click on the open-a-file icon on the tool bar) and place on the left side, just above the command window.

3. Click on View on the tool bar, then on Workspace to open the workspace window, and place it at the upper right.
4. In the command window type `a=1` so that you have a variable in the workspace window, then double click on the yellow name icon for a in the workspace window to make the Array Editor appear. Then place this window just below the workspace window and above the command window, at the left side of the desktop (see the next page.)

With these windows in place you can simultaneously edit your script, watch it run in the command window, and view the values of variables in the array editor. You will find that this is very helpful later when our scripts become more complicated, because the key to debugging is to be able to see the sizes of the arrays that your script is using and the numerical values that they contain. In the next section you will work through a simple example of this procedure.

Workspace Window

Array Editor

Script Editing Window

Command Window

>>

1.10　Sample Script

To help you see what a script looks like and how to run and debug it, here is a simple one that asks for a small step-size h along the x-axis, then plots the function $f(x) = e^{-x/6} \cos x$ from $x = 0$ to $x = 20$. The script then prints the step-size h and tells you that it is finished. The syntax and the commands used in this script are unfamiliar to you now, but you will learn all about them shortly. Type this script into an M-file called `test.m` using Matlab's editor with the desktop arranged as shown on the previous page. Save it, then run it by typing

```
test
```

in the command window. Or, alternatively, press `F5` while the editing window is active and the script will be saved, then executed. Run the sample script below three times using these values of h: 1, 0.1, 0.01. As you run it look at the values of the variables h, x, and f in the workspace window at the upper right of the desktop, and also examine a few values in the array editor just below it so that you understand what the array editor does.

Sample script:

```
clear; % clear all variables from memory
close all; % close all figure windows

h=input(' Enter the step-size h - ') ;
x=0:h:20;  % build an array of points [0,h,2h,...,20]
f=exp(-x/6).*cos(x);  % build the array [f(0),f(h),...f(20)]

plot(x,f)

fprintf(' Plot completed, h = %g \n',h)
```

1.11　Breakpoints and Stepping

When a script doesn't work properly, you need to find out why and then fix it. It is very helpful in this debugging process to watch what the script does as it runs, and to help you do this, Matlab comes with two important features: breakpoints and stepping.

　　To see what a breakpoint does, put the cursor on the `x=0:h:20` line in the sample script above and either click on Breakpoints on the tool bar and select Set/Clear, or press F12. Now press F12 repeatedly and note that the little red dot at the beginning of the line toggles on and off, meaning that F12 is just an on-off switch set set a breakpoint. When the red dot is there it means that a breakpoint has been set, which means that when the script runs it will execute the instructions in the script until it reaches the breakpoint, and then it will stop. Make this happen by pressing F5 and watching the green arrow appear on the line with the breakpoint. Look at the workspace window and note that h has been given a value, but that x has not. This is because the breakpoint stops execution just before the line on which it is set.

　　Now you can click on the Debug icon on the tool bar to see what to do next, but the most common things to do are to either step through the code executing each line in turn (F10) while watching what happens to your variables in the workspace and array editor windows,

or to just continue after the breakpoint to the end (F5.) Take a minute now and use F10 to step through the script while watching what happens in the other windows.

When you write a script to solve some new problem, you should always step through it this way so that you are sure that it is doing what you designed it to do. You will have lots of chances to practice debugging this way as you work through the examples in this book.

2   Variables

Maple has over 100 different data types; Matlab has just two: the matrix and the string. In both languages variables are not declared, but are defined on the fly as it executes. *Note that names in Matlab are case sensitive, so watch your capitalization.* Also: please don't follow the ridiculous trend of making your code more readable by using long variable names with mixed lower- and upper-case letters and underscores sprinkled throughout. Newton's law of gravitation written in this style would be coded this way:

```
Force_of_1_on_2 = G*Mass_of_1*Mass_of_2/Distance_between_1_and_2^2
```

You are asking for an early end to your programming career via repetitive-stress injury if you code like this. Do it this way:

```
F=G*m1*m2/r12^2
```

2.1   Numerical Accuracy

All numbers in Matlab have 15 digits of accuracy. When you display numbers to the screen, like this

```
355/113
```

you may think Matlab only works to 5 significant figures. This is not true; it's just displaying five. If you want to see them all type

```
format long e
```

The four most useful formats to set are

```
format short    (the default)
format long
format long e
format short e
```

Note: e stands for exponential notation.

2.2 π

Matlab knows the number π.

```
pi
```

Try displaying π under the control of each of the three formats given in the previous section. Note: you can set `pi` to anything you want, like this, `pi=2`; but please don't.

2.3 Assigning Values to Variables

The assignment command in Matlab is simply the equal sign. For instance,

```
a=20
```

assigns 20 to the variable `a`, as Maple would do with $a := 20$.

2.4 Matrices

Matlab thinks the number 2 is a 1x1 matrix:

```
N=2
size(N)
```

The array

```
a=[1,2,3,4]
size(a)
```

is a 1x4 matrix (row vectors are built with commas); the array

```
b=[1;2;3;4]
size(b)
```

is a 4x1 matrix (column vectors are built with semicolons, or by putting rows on separate lines-see below.)

The matrix

```
c=[1,2,3;4,5,6;7,8,9]
size(c)
```

is a 3x3 matrix (row entries separated by commas, different rows separated by semicolons.) It should come as no surprise that the Mat in Matlab stands for matrix.

When matrices become large the , and ; way of entering them is awkward. A more visual way to type large matrices in a script is to use spaces in place of the commas and to press the Enter key at the end of each row in place of the semicolons, like this:

```
A = [ 1  2  3  4
      5  6  7  8
      9 10 11 12
     13 14 15 16]
```

This makes the matrix look so nice in a script that you probably ought to use it exclusively.

When you want to access the elements of a matrix you use the syntax `A(row,column)`. For example, to get the element of `A` in the 3rd row, 5th column you would use `A(3,5)`. And if you have a matrix or an array and you want to access the last element in a row or column, you can use Matlab's **end** command, like this:

```
c(end)
A(3,end)
```

2.5 Strings

A string is a set of characters, like this

```
s='This is a string'
```

And if you need an apostrophe in your string, repeat a single quote, like this:

```
t='Don''t worry'
```

And if you just want to access part of a string, like the first 7 characters of s (defined above) use

```
s(1:7)
```

Some Matlab commands require options to be selected or set by using strings. Make sure you enclose them in single quotes, as shown above. If you want to know more about how to handle strings type `help strings`.

3 Input, Calculating, and Output

3.1 Input

To have a script request and assign a value to the variable N from the keyboard use

```
N=input(' Enter a value for N - ')
```

If you enter a single number, like 2.7, then N will be a scalar variable. If you enter an array, like this: [1,2,3,4,5], then N will be an array. If you enter a matrix, like this: [1,2,3;4,5,6;7,8,9], then N will be a 3x3 matrix. And if you don't want the variable you have entered to echo on the screen, end the input command line with a semicolon.

You can also enter data from a file filled with rows and columns of numbers. Matlab reads the file as if it were a matrix with the first line of the file going into the first row of the matrix. If the file were called `data.fil` and looked like this

```
1 2 3
4 5 6
7 8 9
```

then the Matlab command

```
load data.fil
```

would produce a matrix called `data` filled with the contents of the file.

3.2 Calculating

Matlab only crunches numbers. It doesn't do any symbolic algebra, so it is much less capable than Maple, but because it doesn't have to do hard symbolic stuff, it can handle numbers much faster than Maple can. (Testimonial: "I, Scott Bergeson, do hereby certify that I wrote a data analysis code in Maple that took 25 minutes to run. When I converted the code to Matlab it took 15 seconds.") Here's a brief summary of what it can do with numbers, arrays, and matrices.

3.3   Add and Subtract

Matlab knows how to add and subtract numbers, arrays, and matrices. As long as A and B are two variables of the same size (e.g., both 2x3 matrices), then $A + B$ and $A - B$ will add and subtract them as matrices:

```
A=[1,2,3;4,5,6;7,8,9]
B=[3,2,1;6,4,5;8,7,9]
A+B
A-B
```

3.4   Multiplication

The usual multiplication sign * has special meaning in Matlab. Because everything in Matlab is a matrix, * means matrix multiply. So if A is a 3x3 matrix and B is another 3x3 matrix, then $A * B$ will be their 3x3 product. Similarly, if A is a 3x3 matrix and C is a 3x1 matrix (column vector) then $A * C$ will be a new 3x1 column vector. And if you want to raise A to a power by multiplying it by itself n times, you just use

```
A^n
```

For a language that thinks everything in the world is a matrix, this is perfectly natural. Try

```
A*B
A*[1;2;3]
A^3
```

But there are lots of times when we don't want to do matrix multiplication. Sometimes we want to take two big arrays of numbers and multiply their corresponding elements together, producing another big array of numbers. Because we do this so often (you will see many examples later on) Matlab has a special symbol for this kind of multiplication:

```
.*
```

For instance, the dot multiplication between the arrays [a,b,c] and [d,e,f] would be the array [a*d,b*e,c*f]. And since we might also want to divide two big arrays this way, or raise each element to a power, Matlab also allows the operations

```
./  and .^
```

For example, try

```
[1,2,3].*[3,2,1]
[1,2,3]./[3,2,1]
[1,2,3].^2
```

These "dot" operators are very useful in plotting functions and other kinds of signal processing. (You are probably confused about this dot business right now. Be patient. When we start plotting and doing real calculations this will all become clear.)

3.5 Complex Arithmetic

Matlab works as easily with complex numbers as with real ones. The variable i is the usual imaginary number $i = \sqrt{-1}$, unless you are so foolish as to assign it some other value, like this:

```
i=3
```

If you do this you no longer have access to imaginary numbers, so don't ever do it. If you accidentally do it the command `clear i` *will restore it to its imaginary luster.* By using i you can do complex arithmetic, like this

```
z1=1+2i
% or you can multiply by i, like this
z1=1+2*i

z2=2-3i
% add and subtract
z1+z2
z1-z2
% multiply and divide
z1*z2
z1/z2
```

And like everything else in Matlab, complex numbers work as elements of arrays and matrices as well.

When working with complex numbers we quite often want to pick off the real part or the imaginary part of the number, find its complex conjugate, or find its magnitude. Or perhaps we need to know the angle between the real axis and the complex number in the complex plane. Matlab knows how do all of these

```
z=3+4i
real(z)
imag(z)
conj(z)
abs(z)
angle(z)
```

Matlab also knows how to evaluate many of the functions discussed in the next section with a complex argument. Perhaps you recall Euler's famous formula $e^{ix} = \cos x + i \sin x$? Matlab knows it too.

```
exp(i*pi/4)
```

Matlab knows how to handle complex arguments for all of the trig, exponential, and hyperbolic functions. It can do Bessel functions of complex arguments too.

3.6 Mathematical Functions

Matlab knows all of the standard functions found on scientific calculators and even many of the special functions that Maple can do. I will give you a list below of a bunch of them,

but first you need to know an important feature that they all share. They work just like the "dot" operators discussed in the previous section. This means, for example, that it makes sense to take the sine of an array: the answer is just an array of sine values, e.g.,

```
sin([pi/4,pi/2,pi])= [0.7071 1.0000 0.0000]
```

OK, here's the list of function names that Matlab knows about. You can use online help to find details about how to use them. Notice that the natural log function $\ln x$ is the Matlab function `log(x)`.

cos(x)	sin(x)	tan(x)	sec(x)	csc(x)	cot(x)
acos(x)	asin(x)	atan(x)	atan2(y,x)		
exp(x)	log(x)	[log(x) is ln(x)]	log10(x)	log2(x)	sqrt(x)
cosh(x)	sinh(x)	tanh(x)	sech(x)	csch(x)	coth(x)
acosh(x)	asinh(x)	atanh(x)			
sign(x)	airy(n,x)	besselh(n,x)	besseli(n,x)	besselj(n,x)	besselk(n,x)
bessely(n,x)	beta(x,y)	betainc(x,y,z)	betaln(x,y)	ellipj(x,m)	ellipke(x)
erf(x)	erfc(x)	erfcx(x)	erfinv(x)	gamma(x)	gammainc(x,a)
gammaln (x)	expint(x)	legendre(n,x)	factorial(x)		

3.7   Housekeeping Functions

Here are some functions that don't really do math but are useful in programming.

abs(x)	the absolute value of a number (real or complex)
clc	clears the command window; useful for beautifying printed output
ceil(x)	the nearest integer to x looking toward $+\infty$
clear	clears all assigned variables
close all	closes all figure windows
close 3	closes figure window 3
fix(x)	the nearest integer to x looking toward zero
fliplr(A)	flip a matrix A, left for right
flipud(A)	flip a matrix A, up for down
floor(x)	the nearest integer to x looking toward $-\infty$
length(a)	the number of elements in a vector
mod(x,y)	the integer remainder of x/y; see online help if x or y are negative
rem(x,y)	the integer remainder of x/y; see online help if x or y are negative
rot90(A)	rotate a matrix A by 90°
round(x)	the nearest integer to x
sign(x)	the sign of x and returns 0 if x=0
size(c)	the dimensions of a matrix

Try `floor([1.5,2.7,-1.5])` to see that these functions operate on arrays and not just on single numbers.

3.8   Output

Hold it: so far you may have executed most of the example Matlab commands in the command window. From now on it will prepare you better for the more difficult material coming

16

up if you have both a command window and an M-file window open. Put the examples to be run in the M-file (call it `junk.m`), then execute the examples from the command window by typing

```
junk
```

OK, let's learn about output. To display printed results you can use the `fprintf` command. For full information type

```
help fprintf
```

but to get you started, here are some examples. Try them so you know what each one produces. (Here's a hint: the stuff inside the single quotes is a string which will print on the screen; % is where the number you are printing goes; and the stuff after % is a format code. A g means use "whatever" format; if the number is really big or really small, use scientific notation, otherwise just throw 6 significant figures on the screen in a format that looks good. The format 6.2f means use 2 decimal places and fixed-point display with 6 spaces for the number. An e means scientific notation, with the number of decimal places controlled like this: 1.10e.)

```
fprintf('  N =%g \n',500)

fprintf('  x =%1.12g \n',pi)

fprintf('  x =%1.10e \n',pi)

fprintf('  x =%6.2f \n',pi)

fprintf(' x =%12.8f  y =%12.8f \n',5,exp(5))
```

Note: \n is the command for a new line. If you want all the details on this stuff, look in a C-manual or Chapter 10 of *Mastering Matlab 6*.

This command will also write output to a file. Here is an example from online help that writes a file filled with a table of values of x and $\exp(x)$ from 0 to 1. Note that when using Windows that a different line-feed character must be used with \r\n replacing \n (see the `fprintf` below.)

(Note: in the example below, don't put the comments in `junk.m`; just type in the code lines.)

```
%********************************************************
% build an array of x-values from 0 to 1 in steps of 0.1
% (using the colon command discussed in the next section
% Its syntax is x=xstart:dx:xend.)
%********************************************************

x=0:.1:1;

% check the length of the x array

N=length(x)
```

```
% build a 2xN matrix with a row of x and a row of exp(x)

y=[x;exp(x)];

%*********************************************************
% open a file to write the matrix in - fid is a unit number
% to be used later when we write the file and close it
% there is nothing special about the name fid - fxd works too,
% any variable is OK
%*********************************************************

fid=fopen('file1.txt','w')

%*********************************************************
% write the matrix to the file - note that it will be in the
% current Matlab directory.  Type cd to see where you are.
% Matlab writes the file as two columns instead of two rows.
%*********************************************************

fprintf(fid,'%6.2f  %12.8f \r\n',y)

% close the file

st=fclose(fid);
```

After you try this, open `file1.txt`, look at the way the numbers appear in it, and compare them to the format commands `%6.2f` and `%12.8f` to learn what they mean. Write the file again using the `%g` format for both columns and look in the file again.

4   Arrays and x-y Plotting

4.1   Colon (:) Command

Simple plots of y vs. x are done with Matlab's plot command and arrays. These arrays are easily built with the colon (:) command. To build an array x of x-values starting at $x = 0$, ending at $x = 10$, and having a step size of $h = .01$ type this:

```
clear;close all;  % close the figure windows
x=0:.01:10;
```

Note that the semicolon at the end of this line is crucial, unless you want to see 1001 numbers scroll down your screen. If you do make this mistake and the screen print is going to take forever, ctrl-c will rescue you.

An array like this that starts at the beginning of an interval and finishes at the end of it is called a cell-edge grid. A cell-center grid is one that has N subintervals, but the data points are at the centers of the intervals, like this

```
dx=.01;
x=.5*dx:dx:10-0.5*dx;
```

Both kinds of grids are used in computational physics. (Note: Matlab's `linspace` command also makes cell-edge grids. Check it out with `help linspace`.)

And if you leave the middle number out of this colon construction, like this

```
t=0:20;
```

then Matlab assumes a step size of 1. You should use the colon command whenever possible because it is a pre-compiled Matlab command. Tests show that using : is about 20 times faster than using a loop that you write yourself (discussed in Chapter 9). To make a corresponding array of y values according to the function $y(x) = \sin(5x)$ simply type this

```
y=sin(5*x);
```

Both of these arrays are the same length, as you can check with the `length` command (Note that commands separated with commas just execute one after the other, like this:)

```
length(x),length(y)
```

4.2 xy Plots, Labels, and Titles

To plot y vs. x, just type this

```
close all; % (don't clear--you will lose the data you want to plot)
plot(x,y,'r-');
```

The 'r-' option string tells the plot command to plot the curve in red connecting the dots with a continuous line. Other colors are also possible, and instead of connecting the dots you can plot symbols at the points with various line styles between the points. To see what the possibilities are type `help plot`.

And what if you want to plot either the first or second half of the x and y arrays? The `colon` and `end` commands can help:

```
nhalf=ceil(length(x)/2);
plot(x(1:nhalf),y(1:nhalf),'b-')
plot(x(nhalf:end),y(nhalf:end),'b-')
```

To label the x and y axes, do this after the `plot` command: (Note that Greek letters and other symbols are available through LaTex format–see *Greek Letters, Subscripts, and Superscripts* in Section 4.8.)

```
xlabel('\theta')
ylabel('F(\theta)')
```

And to put a title on you can do this:

```
title('F(\theta)=sin(5 \theta)')
```

You can even build labels and titles that contain numbers you have generated; simply use Matlab's `sprintf` command, which works just like `fprintf` except that it writes into a string variable instead of to the screen. You can then use this string variable as the argument of the commands `xlabel`, `ylabel`, and `title`, like this:

```
s=sprintf('F(\\theta)=sin(%i \\theta)',5)
title(s)
```

Note that to force LaTex symbols to come through correctly when using `sprintf` you have to use two backslashes instead of one.

4.3   Overlaying Plots

Often you will want to overlay two plots on the same set of axes. There are two ways you can do this.

```
%***********************************************************
% Here's the first way -- just ask for multiple plots on the
% same plot line
%***********************************************************

close
y2=cos(x);  % load y2 with cos(x), the second function

% plot both

plot(x,y,'r-',x,y2,'b-')

%***********************************************************
% Here's the second way -- after the first plot tell Matlab
% to hold the plot so you can put a second one with it
%***********************************************************

close all;
plot(x,y,'r-')
hold on
plot(x,y2,'b-')
```

You can now call as many plots as you want and they will all go on the same figure. To release it use the command

```
hold off
```

4.4   xyz Plots: Curves in 3-D Space

Matlab will draw three-dimensional curves in space with the `plot3` command. Here is how you would do a spiral on the surface of a sphere using spherical coordinates.

```
clear;close all;
dphi=pi/100; % set the spacing in azimuthal angle

N=30; % set the number of azimuthal trips
phi=0:dphi:N*2*pi;

theta=phi/N/2; % go from north to south once

r=1;  % sphere of radius 1

% convert spherical to Cartesian
x=r*sin(theta).*cos(phi);
y=r*sin(theta).*sin(phi);
z=r*cos(theta);

% plot the spiral
plot3(x,y,z,'b-')
axis equal
```

4.5   Logarithmic Plots

To make log and semi-log plots use the commands `semilogx`, `semilogy`, and `loglog`. They work like this.

```
x=0:.1:8;
y=exp(x);

semilogx(x,y);
title('Semilogx')
pause
semilogy(x,y);
title('Semilogy')
pause
loglog(x,y);
title('Loglog')
```

4.6   Generating Multiple Plots

You may want to put one graph in figure window 1, a second plot in figure window 2, etc.. To do so, put the Matlab command `figure` before each plot command, like this

```
x=0:.01:20;
f1=sin(x);
f2=cos(x)./(1+x.^2);
```

```
figure
plot(x,f1)
figure
plot(x,f2)
```

And once you have generated multiple plots, you can bring each to the foreground on your screen either by clicking on them and moving them around, or by using the command `figure(1)` to pull up figure 1, `figure(2)` to pull up figure 2, etc.. This might be a useful thing to use in a script. See online help for more details.

4.7   Controlling the Axes

You have probably noticed that Matlab chooses the axes to fit the functions that you are plotting. You can override this choice by specifying your own axes, like this.

```
close all;
x=.01:.01:20;
y=cos(x)./x;
plot(x,y)
axis([0 25 -5 5])
```

Or, if you want to specify just the x-range or the y-range, you can use `xlim`:

```
plot(x,y)
xlim([ 0 25])
```

or `ylim`:

```
plot(x,y)
ylim([-5 5])
```

And if you want equally scaled axes, so that plots of circles are perfectly round instead of elliptical, use

```
axis equal
```

4.8   Greek Letters, Subscripts, and Superscripts

When you put labels and titles on your plots you can print Greek letters, subscripts, and superscripts by using the LaTex syntax. (See a book on LaTex for details.) To print Greek letters just type their names preceded by a backslash, like this:

\alpha	\beta	\gamma	\delta	\epsilon	\phi
\theta	\kappa	\lambda	\mu	\nu	\pi
\rho	\sigma	\tau	\xi	\zeta	

You can also print capital Greek letters, like this \Gamma.

To put a subscript on a character use the underscore character on the keyboard: θ_1 is coded by typing `\theta_1`. And if the subscript is more than one character long do this: `\theta_{12}` (makes θ_{12}). Superscripts work the same way only using the $\wedge$ character: use `\theta\wedge{10}` to print θ^{10}.

To write on your plot you can just click A on the figure tool bar, then click in the figure window where you want the text to appear and type the text. You can use LaTex Greek here too. Or, you can use Matlab's `text` command in the format:

```
text(10,.5,'Hi');
```

which will place the text "Hi" at position $x = 10$ and $y = 0.5$ on your plot.

4.9 Changing Line Widths, Fonts, Etc..

If you want to change the look of anything on your plot, like the font style or size of text, the width of the lines, the font style and size of the axis labels, etc.., just left click on the thing you want to change until it is highlighted, then right click on it and select Properties.

This will take care of simple plots, but if you want to make publication quality figures you will have to work harder. See the section at the end of this booklet titled Plots for Publication by Tom Jenkins.

5 Surface, Contour, and Vector Field Plots

5.1 Meshgrid and Ndgrid

Matlab will also display functions of the type $f(x, y)$, either by making a contour plot (like a topographic map) or by displaying the function as height above the xy plane like a perspective drawing. Start by defining arrays x and y that span the region that you want to plot, then create the function $f(x, y)$ over the plane, and finally either use `contour`, `surf`, or `mesh`. The example that follows is one long piece of Matlab code. Put it in the file `test.m`, put pauses in appropriate places, and slowly execute your way through it. like this

```
%begin 5.1

%********************************************************
% Define the arrays x and y
% Warning: don't make the step size too small or you will
% kill the system
%********************************************************

clear;close all;
x=-1:.1:1;y=0:.1:1.5;

%********************************************************
% Use meshgrid to convert these 1-d arrays into 2-d matrices of
% x and y values over the plane
%********************************************************

[X,Y]=meshgrid(x,y);

%********************************************************
% Get f(x,y) by using f(X,Y).  Note the use of .* with X and Y
% rather than with x and y
%********************************************************
```

```
f=(2-cos(pi*X)).*exp(Y);

%***********************************************************
% Note that this function is uphill in y between x=-1 and x=1
% and has a valley at x=0
%***********************************************************

surf(X,Y,f);
%end 5.1
```

Make sure your Workspace window is open, then click on X and Y to view them with the Array Editor. Look at them until you understand how meshgrid has turned your one-dimensional arrays x and y into their two-dimensional versions X and Y.

It is sometimes helpful when doing physics problems in two dimensions to have the matrix $f(i, j)$ correspond to $f(x, y)$, so that the index i represents x and the index j represents y. Unfortunately, meshgrid does this backwards, as you can see by checking the lengths of x and y, and then the size of f in the example above:

```
length(x)
length(y)
size(f)
```

But Matlab has another command called ndgrid which is similar to meshgrid but does the conversion to two dimensions the other way round. The plots look the same, but $f(i, j)$ is now in the $f(x, y)$ order:

```
[X,Y]=ndgrid(x,y);
f=(2-cos(pi*X)).*exp(Y);
surf(X,Y,f);
length(x)
length(y)
size(f)
```

5.2   Contour Plots and Surface Plots

```
%begin 5.2

%***********************************************************
% make a contour plot by asking Matlab to evenly space N contours
% between the minimum of f(x,y) and the maximum f(x,y) (the default)
%***********************************************************

N=40;
contour(X,Y,f,N);
title('Contour Plot')
xlabel('x')
ylabel('y')
pause
```

```
%*********************************************************
% You can also tell Matlab which contour levels you want to plot.
% To do so load an array (V in this case) with the values of the
% levels you want to see.  In this example they will start at
% the minimum of f, end at the maximum of f, and there will
% be 21 contours.
%
% You can even print contour labels on the plot, which is
% a big help, by assigning the plot to a variable name
% and using the clabel command, as shown below.  Only
% every other contour is labeled in this example
%*********************************************************

close all;
top=max(max(f));  % find the max and min of f
bottom=min(min(f));
dv=(top-bottom)/20; % interval for 21 equally spaced contours
V=bottom:dv:top;
cs=contour(X,Y,f,V);
clabel(cs,V(1:2:21)) % give clabel the name of the plot and
                     % an array of the contours to label
title('Contour Plot')
xlabel('x')
ylabel('y')

pause

%*********************************************************
% Now make a surface plot of the function with the viewing
% point rotated by AZ degrees from the x-axis and
% elevated by EL degrees above the xy plane
%*********************************************************

close all;
surf(X,Y,f); % or you can use mesh(X,Y,f) to make a wire plot
AZ=30;EL=45;
view(AZ,EL);
title('Surface Plot')
xlabel('x')
ylabel('y')
pause

%*********************************************************
% If you want to manually change the viewing angle of
% a surface plot, click on the circular arrow icon
% on the figure window, then click and move the
% pointer on the graph.  Try it until you get the
```

```
% hang of it.
%
% Here's a piece of code that lets you fly around the
% surface plot by continually changing the viewing angles
% and using the pause command; I think you'll be impressed
%**********************************************************

close all;
surf(X,Y,f);
title('Surface Plot')
xlabel('x')
ylabel('y')
EL=45;
for m=1:100
   AZ=30+m/100*360;
   view(AZ,EL);
   pause(.1);  % pause units are in seconds
end

pause

%**********************************************************
% This same trick will let you make animations of
% both xy and surface plots.  To make this surface
% oscillate up and down like a manta ray you could
% do this.
%**********************************************************

dt=.1;

for m=1:100
   t=m*dt;
   g=f*cos(t);

   surf(X,Y,g);
   AZ=30;EL=45;
   view(AZ,EL);
   title('Surface Plot')
   xlabel('x')
   ylabel('y')
   axis([-1 1 -1 1 min(min(f)) max(max(f))])
   pause(.1)

end

%end 5.2
```

Note that you can find all kinds of cool stuff about surface and contour plotting by typing

```
help graph3d
```

and then checking out these commands by using help on each one. Another good source is *Mastering Matlab 6*, Chapters 26-32.

5.3 Evaluating Fourier Series

Matlab will make graphical displays of infinite series of the kind discussed in Physics 318 and Physics 441 quickly and easily. Consider this solution for the electrostatic potential in a long tube of rectangular cross-section bounded by long metal strips held at different potentials. The tube has width $2b$ in x and width a in y. The two strips at $x = -b$ and $x = +b$ are held at potential V_0 while the strips at $y = 0$ and $y = a$ are grounded. (See *Introduction to Electrodynamics, Third Edition* by David Griffiths, pages 132-134.) The electrostatic potential is given by

$$V(x, y) = \frac{4V_0}{\pi} \sum_{n=0}^{\infty} \frac{1}{(2n+1)} \frac{\cosh\left[(2n+1)\pi x/a\right]}{\cosh\left[(2n+1)\pi b/a\right]} \sin\left[(2n+1)\pi y/a\right]$$

Here is a piece of Matlab code that evaluates this function on an xy grid and displays it as a surface plot.

```
%begin 5.3

clear;close all;

% set some constants

a=2;b=1;Vo=1;

% build the x and y grids

Nx=80;Ny=40;
dx=2*b/Nx;dy=a/Ny;
x=-b:dx:b;y=0:dy:a;

% build the 2-d grids for plotting

[X,Y]=meshgrid(x,y);

% set the number of terms to keep
% and do the sum

Nterms=20;

% zero V out so we can add into it

V=zeros(Ny+1,Nx+1);
```

```
% add the terms of the sum into V

for m=0:Nterms
   V=V+cosh((2*m+1)*pi*X/a)/cosh((2*m+1)*pi*b/a).*sin((2*m+1)*pi*Y/a)/(2*m+1);
end

% put on the multiplicative constant

V=4*Vo/pi*V;

% surface plot the result

surf(X,Y,V)
xlabel('x');
ylabel('y');
zlabel('V(x,y)')

%end 5.3
```

5.4   Vector Field Plots

Matlab will plot vector fields for you with arrows. This is a good way to visualize flows, electric fields, magnetic fields, etc.. The command that makes these plots is `quiver` and the code below illustrates its use in displaying the electric field of a line charge and the magnetic field of a long wire. Note that the vector field components must be computed in Cartesian geometry.

```
%begin 5.4

clear;close

x=-5.25:.5:5.25;y=x; % define the x and y grids (avoid (0,0))
[X,Y]=meshgrid(x,y);

% Electric field of a long charged wire

Ex=X./(X.^2+Y.^2);
Ey=Y./(X.^2+Y.^2);

% make the field arrow plot

quiver(X,Y,Ex,Ey)
title('E of a long charged wire')

axis equal  % make the x and y axes be equally scaled
pause

% Magnetic field of a long current-carrying wire
```

```
Bx=-Y./(  X.^2+Y.^2);
By=X./(X.^2+Y.^2);

% make the field arrow plot

quiver(X,Y,Bx,By)
axis equal
title('B of a long current-carrying wire')
pause

%*********************************************************
% The big magnitude difference across the region makes most arrows too small
% to see.  This can be fixed by plotting unit vectors instead (losing all
% magnitude information
%*********************************************************

B=sqrt(Bx.^2+By.^2);
Ux=Bx./B;Uy=By./B;
quiver(X,Y,Ux,Uy);
axis equal
title('B(wire): unit vectors')
pause

%*********************************************************
% Or, you can still see qualitative size information
% without such a big variation in arrow size by
% having the arrow length be logarithmic.  If s is
% the desired ratio between the longest arrow and
% the shortest one, this code will make the  appropriate
% field plot.
%*********************************************************

Bmin=min(min(B));
Bmax=max(max(B));
s=2;  % choose an arrow length ratio
k=(Bmax/Bmin)^(1/(s-1));

logsize=log(k*B/Bmin);
Lx=Ux.*logsize;
Ly=Uy.*logsize;
quiver(X,Y,Lx,Ly);
axis equal
title('B(wire): logarithmic arrows')

%end 5.4
```

There may be too much detail to really see what's going on in some field plots. You can work around this problem by clicking on the zoom icon on the tool bar and then using the mouse to define the region you want to look at. Clicking on the zoom-out icon, then clicking on the figure will take you back where you came from.

6 Vector Products, Dot and Cross

Matlab will do dot and cross products for you with the commands `dot` and `cross`, like this:

```
a=[1,2,3];b=[3,2,1];
dot(a,b)
cross(a,b)
```

(Cross products only work for three-dimensional vectors but dot products can be used with vectors of any length.)

7 Linear Algebra

Almost anything you learned about in your linear algebra class Matlab has a command to do. Here is a brief summary of the most useful ones for physics.

7.1 Solve a Linear System

Matlab will solve the matrix equation $\mathbf{Ax} = \mathbf{b}$, where $\mathbf{A}$ is a square matrix, where $\mathbf{b}$ is a known column vector, and where $\mathbf{x}$ is an unknown column vector. For instance, the system of equations

$$x + z = 4 \tag{1}$$
$$-x + y + z = 4 \tag{2}$$
$$x - y + z = 2 \ , \tag{3}$$

which is solved by $(x, y, z) = (1, 2, 3)$l, is handled in Matlab by defining a matrix A corresponding to the coefficients on the left side of this equation and a column vector b corresponding to the coefficients on the right (see the example below.) By use of the simple symbol \ (sort of "backwards divide") Matlab will use Gaussian elimination to solve this system of equations, like this:

```
% here are A and b corresponding to the equations above

A=[ 1, 0,1
   -1, 1,1
    1,-1,1 ];
b=[4
   4
   2];
```

```
% now solve for x and see if we obtain [1;2;3], like we should

x=A\b
```

7.2   Max and Min

The commands `max` and `min` return the maximum and minimum values of an array. And with a slight change of syntax they will also return the indices in the array at which the maximum and minimum occur.

```
x=0:.01:5;
y=x.*exp(-x.^2);

% take a look at the function so we know what it looks like

plot(x,y)

% find the max and min

ymin=min(y)
ymax=max(y)

% find the max and min along with the array indices imax and imin
% where they occur

[ymin,imin]=min(y)
[ymax,imax]=max(y)
```

7.3   Matrix Inverse

The `inv` command will compute the inverse of a square matrix. For instance, using the matrix

```
A=[1,0,-1;-1,1,1;1,-1,1]
```

we find

```
% load C with the inverse of A

C=inv(A)

% verify by matrix multiplication that A*C is the identity matrix

A*C
```

7.4   Transpose and Hermitian Conjugate

To find the transpose of the matrix A just use a single quote with a period, like this

```
A.'
```

To find the Hermitian conjugate of the matrix A (transpose of A with all elements replaced with their complex conjugates) type

```
A'
```

(notice that there isn't a period). If your matrices are real, then there is no difference between these two commands and you might as well just use `A'`. Notice that if `a` is a row vector then `a'` is a column vector. You will use the transpose operator to switch between row and column vectors a lot in Matlab, like this

```
[1,2,3]
[1,2,3]'
[4;5;6]
[4;5;6]'
```

7.5 Special Matrices

Matlab will let you load several special matrices. The most useful of these are given here.

```
% eye:
% load I with the 4x4 identity matrix (the programmer who invented this
% syntax must have been drunk)

I=eye(4,4)

% zeros:
% load Z with a 5x5 matrix full of zeros

Z=zeros(5,5)

% ones:
% load X with a 3x3 matrix full of ones

X=ones(3,3)

% rand:
% load Y with a 4x6 matrix full of random numbers between 0 and 1
% The random numbers are uniformly distributed on [0,1]

Y=rand(4,6)

% And to load a single random number just use

r=rand

% randn:
% load Y with a 4x6 matrix full of random numbers with a Gaussian
% distribution with zero mean and a variance of 1
```

```
Y=randn(4,6)
```

7.6 Determinant

Find the determinant of a square matrix this way

```
det(A)
```

7.7 Norm of Vector (Magnitude)

Matlab will compute the magnitude of a vector `a` (the square root of the sum of the squares of its components) with the `norm` command

```
a=[1,2,3]
norm(a)
```

7.8 Sum the Elements

For arrays the command `sum` adds up the elements of the array:

```
% calculate the sum of the squares of the reciprocals of the
% integers from 1 to 10,000

n=1:10000;
sum(1./n.^2)

% compare this answer with the sum to infinity, which is pi^2/6

ans-pi^2/6
```

For matrices the `sum` command produces a row vector which is made up of the sum of the columns of the matrix.

```
A=[1,2,3;4,5,6;7,8,9]
sum(A)
```

7.9 Selecting Rows and Columns

Sometimes you will want to select a row or a column of a matrix and load it into an array. This is done with Matlab's all-purpose colon (:) command.

To load a column vector b with the contents of the third column of the matrix A use:

```
b=A(:,3)
```

Recall that the first index of a matrix is the row index, so this command tells Matlab to select all of the rows of A in column 3.

To load a row vector c with the contents of the second row of the matrix A use:

```
c=A(2,:)
```

You can also select just part of row or column like this:

```
c=A(2,1:2)
```

which takes only the first two elements of the second row.

7.10   Eigenvalues and Eigenvectors

To build a column vector containing the eigenvalues of the matrix A in the previous section use

```
E=eig(A)
```

To build a matrix V whose columns are the eigenvectors of the matrix A and another matrix D whose diagonal elements are the eigenvalues corresponding to the eigenvectors in V use

```
[V,D]=eig(A)

% to select the 3rd eigenvector and load it into
% a column vector use

v3=V(:,3)  % i.e., select all of the rows (:) in column 3
```

7.11   Fancy Stuff

Matlab also knows how to do singular value decomposition, QR factorization, LU factorization, and conversion to reduced row-echelon form. And the commands `rcond` and `cond` will give you the condition number of a matrix. To learn about these ideas, consult a textbook on linear algebra. To learn how they are used in Matlab use the commands;

```
help svd
help QR
help LU
help rref
help rcond
help cond
```

8   Polynomials

Polynomials are used so commonly in computation that Matlab has special commands to deal with them. The polynomial $x^4 + 2x^3 - 13x^2 - 14x + 24$ is represented in Matlab by the array `[1,2,-13,-14,24]`, i.e., by the coefficients of the polynomial starting with the highest power and ending with the constant term. If any power is missing from the polynomial its coefficient must appear in the array as a zero. Here are some of the things Matlab can do with polynomials. Try each piece of code in Matlab and see what it does.

8.1   Roots of a Polynomial

```
% find the roots of a polynomial

p=[1,2,-13,-14,24];
r=roots(p)
```

8.2 Find the polynomial from the roots

If you know that the roots of a polynomial are 1, 2, and 3, then you can find the polynomial in Matlab's array form this way

```
r=[1,2,3];
p=poly(r)
```

8.3 Multiply Polynomials

The command `conv` multiplies two polynomial coefficient arrays and returns the coefficient array of their product.

```
a=[1,0,1];
b=[1,0,-1];
c=conv(a,b)
```

Stare at this result and make sure that it is correct.

8.4 Divide Polynomials

Remember synthetic division? Matlab can do it with the command `deconv`, giving you the quotient and the remainder.

```
% a=x^2+x+1 and b=x+1

a=[1,1,1];b=[1,1];

% now divide b into a finding the quotient and remainder

[q,r]=deconv(a,b)
```

After you do this Matlab will give you `q=[1,0]` and `r=[0,0,1]`, which means that $q = x + 0 = x$ and $r = 0x^2 + 0x + 1 = 1$ so

$$\frac{x^2 + x + 1}{x + 1} = x + \frac{1}{x + 1}$$

8.5 First Derivative

Matlab can take a polynomial array and return the polynomial array of its derivative:

```
a=[1,1,1,1]
ap=polyder(a)
```

8.6 Evaluate a Polynomial

If you have an array of x-values and you want to evaluate a polynomial at each one, do this:

```
% define the polynomial

a=[1,2,-13,-14,24];
```

```
% load the x-values

x=-5:.01:5;

% evaluate the polynomial

y=polyval(a,x);

% plot it

plot(x,y)
```

8.7   Fitting Data to a Polynomial

If you have some data in the form of arrays (x,y), Matlab will do a least-squares fit of a polynomial of any order you choose to this data. In this example we will let the data be the sine function between 0 and π and we will fit a polynomial of order 4 to it. Then we will plot the two functions on the same frame to see if the fit is any good. Before going on to the next section, try fitting a polynomial of order 60 to the data to see why you need to be careful when you do fits like this.

```
x=linspace(0,pi,50);

% make a sine function with 1% random error on it

f=sin(x)+.01*rand(1,length(x));

% fit to the data

p=polyfit(x,f,4);

% evaluate the fit

g=polyval(p,x);

% plot fit and data together

plot(x,f,'r*',x,g,'b-')
```

9   Loops and Logic

To use Matlab to solve many physics problems you have to know how to write loops and how to use logic.

9.1 Loops

A loop is a way of repeatedly executing a section of code. It is so important to know how to write them that several common examples of how they are used will be given here. The two kinds of loops we will use are the `for` loop and the `while` loop. We will look at `for` loops first, then study `while loops` a bit later in the logic section.

The `for` loop looks like this:

```
for n=1:N
.
.
.
end
```

which tells Matlab to start n at 1, then increment it by 1 over and over until it counts up to N, executing the code between `for` and `end` for each new value of n. Here are a few examples of how the `for` loop can be used.

9.1.1 Summing a series with a `for` loop

Let's do the sum

$$\sum_{n=1}^{N} \frac{1}{n^2}$$

with N chosen to be a large integer.

```
s=0; % set a variable to zero so that 1/n^2 can be repeatedly added to it
N=10000;  % set the upper limit of the sum

for n=1:N % start of the loop

% add 1/n^2 to s each time, then put the answer back into s

   s=s+1/n^2;
end % end of the loop

fprintf(' Sum = %g \n',s)  % print the answer
```

You may notice that summing with a loop takes a lot longer than the colon (:) operator way of doing it:

```
% calculate the sum of the squares of the reciprocals of the
% integers from 1 to 10,000

n=1:10000;
sum(1./n.^2)
```

Try both the loop way and this `sum` command way and see which is faster. To slow things down enough that you can see the difference change 10,000 to 100,000. (On my Unix work-station the : way is 21 times faster than the loop, so use : whenever you can.) To do timing checks use the `cputime` command. Look it up in online help.

9.1.2 Products with a `for` loop

Let's calculate $N! = 1 \cdot 2 \cdot 3 \cdot ...(N-1) \cdot N$ using a for loop that starts at $n = 1$ and ends at $n = N$, doing the proper multiply at each step of the way.

```
P=1;  % set the first term in the product
N=20;  % set the upper limit of the product

for n=2:N % start the loop at n=2 because we already loaded n=1
  P=P*n; % multiply by n each time and put the answer back into P
end % end of the loop

fprintf(' N!  = %g \n',P)  % print the answer
```

Now use Matlab's `factorial` command to check that you got the right answer:

```
factorial(20)
```

You should be aware that the `factorial` command is a bit limited in that it won't act on an array of numbers in the way that cos, sin, exp etc. do. A better factorial command to use is the gamma function $\Gamma(x)$ which extends the factorial function to all complex values. It is related to the factorial function by $\Gamma(N+1) = N!$, and is called in Matlab using the command `gamma(x)`. so you could also check the answer to your factorial loop this way:

```
gamma(21)
```

9.1.3 Recursion relations with `for` loops

Suppose that we were solving a differential equation by substituting into it a power series of the form

$$f(x) = \sum_{n=1}^{\infty} a_n x^n$$

and that we had discovered that the coefficients a_n satisfied the recursion relation

$$a_1 = 1 \quad ; \quad a_{n+1} = \frac{2n-1}{2n+1} a_n$$

To use these coefficients we need to load them into an array a so that $a(1) = a_1, a(2) = a_2, ...$ Here's how we could do this using a `for` loop to load $a(1)...a(20)$.

```
a(1)=1;  % put the first element into the array
N=19;  % the first one is loaded, so let's load 19 more

for n=1:N %  start the loop
   a(n+1)=(2*n-1)/(2*n+1)*a(n); % the recursion relation
end

disp(a) %  display the resulting array of values
```

Note that the recursion relation was translated into Matlab code just as it appeared in the formula: `a(n+1)=(2*n-1)/(2*n+1)*a(n)`. The counting in the loop was then adjusted to fit by starting at $n = 1$, which loaded $a(1+1) = a(2)$, then $a(3)$, etc., then ended at $n = 19$,

which loads $a(19 + 1) = a(20)$. Always make the code you write fit the mathematics as closely as possible, then adjust the other coding to fit. This will make your code easier to read and you will make fewer mistakes.

9.2 Logic

Often we only want to do something when some condition is satisfied, so we need logic commands. The simplest logic command is the `if` command, which works like this. (Several examples are given; try them all.)

```
clear;
a=1;b=3;

% If the number a is positive set c to 1; if a is 0 or negative,
% set c to zero

if a>0
   c=1
else
   c=0
end

% if either a or b is non-negative, add them to obtain c;
% otherwise multiply a and b to obtain c

if a>=0 | b>=0  % either non-negative
   c=a+b
else
   c=a*b  % otherwise multiply them to obtain c
end
```

You can build about any logical condition you want if you just know the basic logic elements. Here they are

Equal	==
Less than	<
Greater than	>
Less than or equal	<=
Greater than or equal	>=
Not equal	~=
And	&
Or	|
Not	~

There is also a useful logic command that controls loops: `while`. Suppose you don't know how many times you are going to have to loop to get a job done, but instead want to quit looping when some condition is met. For instance, suppose you want to add the reciprocals of squared integers until the term you just added is less than 1e-10. Then you would change the loop in the $\sum 1/n^2$ example to look like this

```
clear
term=1  % load the first term in the sum, 1/1^2=1
s=term;  % load s with this first term

% start of the loop - set a counter n to one

n=1;

while term > 1e-10  % loop until term drops below 1e-10
   n=n+1; % add 1 to n so that it will count: 2,3,4,5,...
   term=1/n^2; % calculate the next term to add
  s=s+term;  % add 1/n^2 to s until the condition is met
end  % end of the loop

fprintf(' Sum = %g \n',s)
```

This loop will continue to execute until `term`<1e-10. Note that unlike the `for` loop, here you have to do your own counting, being careful about what value n starts at and when it is incremented ($n = n + 1$). It is also important to make sure that the variable you are testing (`term` in this case) is loaded before the loop starts with a value that allows the test to take place and for the loop to run (`term` must pass the `while` test.)

Sometimes `while` is awkward to use because you would rather just loop lots of times checking some condition and then break out of the loop when it is satisfied. The `break` command is designed to do this. When `break` is executed in a loop the script jumps to just after the `end` at the bottom of the loop. Here is our sum loop rewritten with `break`

```
clear
s=0; % initialize the sum variable

% start of the loop

for m=1:1000000
   term=1/n^2;

% add 1/n^2 to s until the condition is met
   s=s+term;
   if term < 1e-10
      break
   end
% end of the loop
end
```

9.3 Secant Method

Here is a real example that you will sometimes use to solve difficult equations of the form $f(x) = 0$ in Matlab. (Maple's `fsolve` command and Matlab's `fzero` command (which you will learn to use shortly) do this for you automatically; here you will see how they do it.)

The idea is to find two guesses x_1 and x_2 that are near a solution of this equation. (You find about where these two guesses ought to be by plotting the function and seeing about

Figure 1: The sequence of approximate points in the secant method.

where the solution is. It's OK to choose them close to each other, like $x_1 = .99$ and $x_2 = .98$.) Once you have these two guesses find the function values that go with them: $f_1 = f(x_1)$ and $f_2 = f(x_2)$ and compute the slope $m = (f_2 - f_1)/(x_2 - x_1)$ of the line connecting the points. (You can follow what is happening here by looking at Fig. 1.) Then fit a straight line through these two points and solve the resulting straight line equation $y - f_2 = m(x - x_2)$ for the value of x that makes $y = 0$, i.e., solve the line equation to find

$$x_3 = x_2 - \frac{f_2}{m} = x_2 - \frac{f_2(x_2 - x_1)}{f_2 - f_1}$$

as shown in Fig. 1. This will be a better approximation to the solution than either of your two initial guesses, but it still won't be perfect, so you have to do it again using x_2 and the new value of x_3 as the two new points. This will give you x_4 in the figure. You can draw your own line and see that the value of x_5 obtained from the line between (x_3, f_3) and (x_4, f_4) is going to be pretty good. And then you do it again, and again, and again, until your approximate solution is good enough.

Here's what the code looks like that solves the equation $\exp(-x) - x = 0$

```
%begin 9.3

clear;close all;

%**********************************************************
% Define the function as an in line function (See Section 13
% for more details.)
%**********************************************************

func=inline('exp(-x)-x','x');

% First plot the function

x=0:.01:2;
```

```
f=func(x);
plot(x,f,'r-',x,0*x,'b-')

%************************************************************
% (Note that the second plot is just a blue x-axis (y=0)
% 0*x is just a quick way to load an array of zeros the
% same size as x)

% From the plot it looks like the solution is near x=.6

% Secant method to solve the equation exp(-x)-x = 0

% Use an initial guess of x1=0.6
%************************************************************

x1=0.6;

% find f(x1)

f1=func(x1);

% find a nearby second guess

x2=0.99*x1;

% set chk, the error, to 1 so it won't trigger the while
% before the loop gets started

chk=1;

% start the loop

while chk>1e-8

% find f(x2)

   f2=func(x2);

% find the new x from the straight line approximation and print it

   xnew = x2 - f2*(x2-x1)/(f2-f1)

% find chk the error by seeing how closely f(x)=0 is approximated

   chk=abs(f2);

% load the old x2 and f2 into x1 and f1; then put the new x into x2
```

```
  x1=x2;f1=f2;x2=xnew;

% end of loop
end

%end 9.3
```

(Note: this is similar to Newton's method, also called Newton-Raphson, but when a finite-difference approximation to the derivative is used it is usually called the secant method.)

9.4   Using Matlab's Fzero

Matlab has its own zero-finder which probably is similar to the secant method described above. To use it you must make a special M-file called a *function*, which we will discuss in more detail in Sec. 13. Here I will just give you a sample file so that you can see how **fzero** works. This function file (called **fz.m** here) evaluates the function $f(x)$. You just need to build it, tell Matlab what its name is using the @-sign syntax illustrated below, and also give Matlab an initial guess for the value of x that satisfies $f(x) = 0$.

In the section of code below you will find the Matlab function **fz(x)** and the line of code invoking **fzero** that finds the root. The example illustrated here is $f(x) = \exp(-x) - x = 0$.

Here is the function M-file **fz.m** used by **fzero** in this example:

```
%beginfunction fz.m

function f=fz(x)

% evaluate the function fz(x) whose
% roots are being sought

f=exp(-x)-x;

%end fz.m
```

Here is the Matlab code that uses **fzero** and **fz** to do the solve:

```
%**********************************************************
%  Here is the matlab code that uses fz.m to find
%  a zero of f(x)=0 near the guess x=.7
%  Note that the @ sign is used to tell Matlab that
%  the name of an M-file is being passed into fzero
%**********************************************************

x=fzero(@fz,.7)
```

Figure 2: The centered derivative approximation works best.

10 Derivatives and Integrals

Matlab won't give you formulas for the derivatives and integrals of functions like Maple will. But if you have closely spaced arrays filled with values of x and $f(x)$ Matlab will quickly give you numerical approximations to the derivative $f'(x)$ and the integral $\int_a^b f(x)dx$.

10.1 Derivatives

In your first calculus class the following formula for the derivative was given:

$$\frac{df}{dx} = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h}$$

To do a derivative numerically we use the following slightly different, but numerically more accurate, form:

$$\frac{df}{dx} = \lim_{h \to 0} \frac{f(x+h) - f(x-h)}{2h}$$

It's more accurate because it's centered about the value of x where we want the derivative to be evaluated.

To see the importance of centering, consider Fig. 2. In this figure we are trying to find the slope of the tangent line at $x = 0.4$. The usual calculus-book formula uses the data points at $x = 0.4$ and $x = 0.5$, giving tangent line a. It should be obvious that using the "centered" pair of points $x = 0.3$ and $x = 0.5$ to obtain tangent line b is a much better approximation.

As an example of what a good job centering does, try differentiating $sinx$ this way:

```
dfdx=(sin(1+1e-5)-sin(1-1e-5))/2e-5

% take the ratio between the numerical derivative and the
% exact answer cos(1) to see how well this does

format long e
dfdx/cos(1)
```

You can also take the second derivative numerically using this formula

$$\frac{d^2 f}{dx^2} = \lim_{h \to 0} \frac{f(x+h) - 2f(x) + f(x-h)}{h^2}$$

For example,

```
d2fdx2=(sin(1+1e-4)-2*sin(1)+sin(1-1e-4))/1e-8
```

```
% take the ratio between the numerical derivative and the
% exact answer -sin(1) to see how well this does
```

```
format long e
d2fdx2/(-sin(1))
```

You may be wondering how to choose the step size h. This is a little complicated; take a course on numerical analysis and you can see how it's done. But until you do, here's a rough rule of thumb. To approximate the first derivative of $f(x)$ use $h = 10^{-5}$; to approximate the second derivative use $h = 10^{-4}$.

If you want to differentiate a function defined by arrays x and f, then the step size is already determined; you just have to live with the accuracy obtained by using $h = \Delta x$, where Δx is the spacing between points in the x array. *Notice that the data must be evenly spaced for the example I am going to give you to work.*

The idea is to approximate the derivative at $x = x_j$ in the array by using the function values f_{j+1} and f_{j-1} like this

$$f'(x_j) \approx \frac{f_{j+1} - f_{j-1}}{2h}$$

This works fine for an N element array at all points from x_2 to x_{N-1}, but it doesn't work at the endpoints because you can't reach beyond the ends of the array to find the needed values of f. So we use this formula for x_2 through x_{N-1}, then use linear extrapolation to find the derivatives at the endpoints, like this

```
%begin 10.1
```

```
dx=1/1000;
x=0:dx:4;
N=length(x);
f=sin(x);
```

```
% Do the derivative at the interior points all at once using
% the colon command
```

```
dfdx(2:N-1)=(f(3:N)-f(1:N-2))/(2*dx);
```

```
% linearly extrapolate to the end points (see the next section)
```

```
dfdx(1)=2*dfdx(2)-dfdx(3);
dfdx(N)=2*dfdx(N-1)-dfdx(N-2);
```

```
% now plot both the approximate derivative and the exact
```
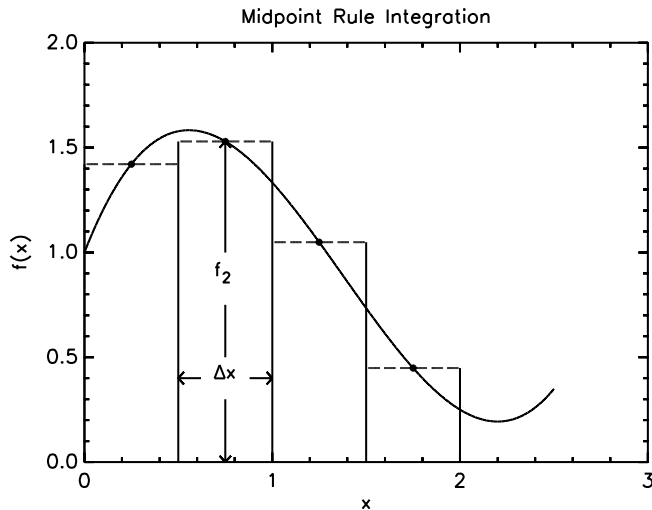
Figure 3: The midpoint rule works OK if the function is nearly straight across each interval.

```
% derivative cos(x) to see how well we did

plot(x,dfdx,'r-',x,cos(x),'b-')

%end 10.1
```

The second derivative would be done the same way. Matlab has its own routines for doing derivatives; look in online help for `diff` and `gradient`.

10.2 Definite Integrals

There are many ways to do definite integrals numerically, and the more accurate these methods are the more complicated they become. But for everyday use the midpoint method usually works just fine, and it's very easy to code. The idea of the midpoint method is to approximate the integral $\int_a^b f(x)dx$ by subdividing the interval $[a, b]$ into N subintervals of width $h = (b - a)/N$ and then evaluating $f(x)$ at the center of each subinterval. We replace $f(x)dx$ by $f(x_j)h$ and sum over all the subintervals to obtain an approximate integral. This method is shown in Fig. 3. Notice that this method should work pretty well over subintervals like $[1.0, 1.5]$ where $f(x)$ is nearly straight, but probably is lousy over subintervals like $[0.5, 1.0]$ where the function curves.

```
%begin 10.2

close all;
N=1000;
a=0;b=5;
dx=(b-a)/N;
x=.5*dx:dx:b-.5*dx;   % build an x array of centered values
f=cos(x);   % load the function

% do the approximate integral
```
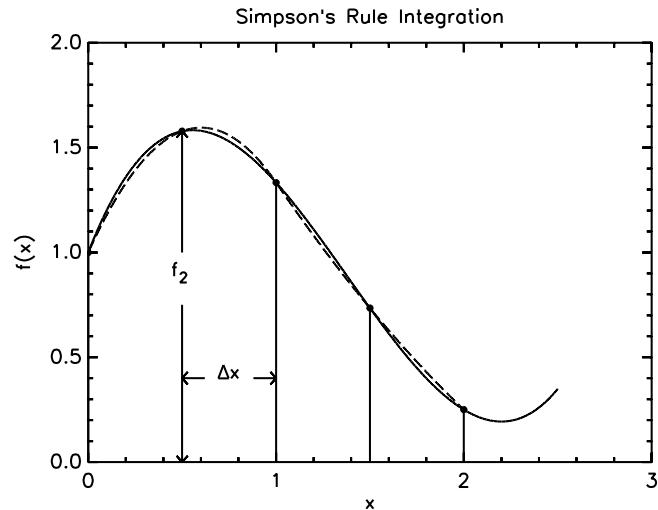
Figure 4: Fitting parabolas (Simpson's Rule) works better.

```
s=sum(f)*dx

% compare with the exact answer, which is sin(5)

s-sin(5)

%end 10.2
```

If you need to do a definite integral in Matlab, this is an easy way to do it. And to see how accurate your answer is, do it with 1000 points, then 2000, then 4000, etc., and watch which decimal points are changing as you go to higher accuracy. (A function that does this for you is given in the section titled *Make Your Own Functions.*)

And if you need to find the indefinite integral, in the section below titled *Make Your Own Functions* there is a piece of code that will take arrays of $(x, f(x))$ and calculate the indefinite integral function $\int_a^x f(x')dx'$.

10.3   Matlab Integrators

Matlab also has its own routines for doing definite and indefinite integrals using both data arrays and M-files; look in online help for the commands `trapz` (works on arrays, not functions), `cumtrapz`, `quad`, `quadl`, and `dblquad`, or in *Mastering Matlab 6*, Chapter 23. As an example, the Matlab routine `quad` approximates the function with parabolas (as shown in Fig. 10.3) instead of the rectangles of the midpoint method. This parabolic method is called Simpson's Rule. As you can see from Fig. 10.3, parabolas do a much better job, making `quad` a standard Matlab choice for integration. As an example of what these routines can do, here is the way you would use `quad` to integrate $\cos x e^{-x}$ from 0 to 2:

Note that these integration routines need function M-files, as `fzero` did. These will be discussed more fully in Chapter 13, so for now just use `fint.m`, given below, as a template for doing integrals with `quad`, `quadl`, etc..

```
%beginfunction fint.m
%***********************************************************
% define the function to be integrated in fint.m

function f=fint(x)

%***********************************************************
% Warning: Matlab will do this integral with arrays of x,
% so use .*, ./, .^, etc. in this function.  If you forget
% to use the .-form you will encounter the error:
%
%         Inner matrix dimensions must agree.

f=cos(x).*exp(-x);

%end fint.m
%***********************************************************


%***********************************************************
%begin 10.3a

% once fint.m is stored in your current directory
% you can use the following commands to integrate.

% simple integral, medium accuracy
quad(@fint,0,2)

% integrate with specified relative accuracy
quad(@fint,0,2,1e-8)

% integrate with specified relative accuracy
% using quadl (notice that quadl is faster--
% always try it first)

quadl(@fint,0,2,1e-8)

%***********************************************************
%end 10.3a
```

Or you can use an `inline function` like this to avoid making another file:

```
%begin 10.3b

f=inline('exp(-x).*cos(x)','x')

quadl(f,0,2,1e-8)
```

```
%end 10.3b
```

And if parabolas are good, why not use cubics, quartics, etc. to do an even better job? Well, you can, and Matlab does. The `quadl` integration command used in the example above uses higher order polynomials to do the integration and is the best Matlab integrator to use.

Matlab also has a command `dblquad` that does double integrals. Here's how to use it.

```
%beginfunction f2int.m

% First define the integrand as a function of x and y

function f=f2int(x,y)
   f=cos(x*y);

%end f2int.m

%**********************************************************
% Now here's how to obtain the double integral over
% the xy rectangle (0,2)X(0,2).  It runs lots
% faster if you use the 'quadl' option, as shown below
%**********************************************************

dblquad(@f2int,0,2,0,2,1e-10,'quadl')
```

11   Interpolation and Extrapolation

Since Matlab only represents functions as arrays of values a common problem that comes up is finding function values at points not in the arrays. Finding function values between data points in the array is called *interpolation*; finding function values beyond the endpoints of the array is called *extrapolation*. A common way to do both is to use nearby function values to define a polynomial approximation to the function that is pretty good over a small region. Both linear and quadratic function approximations will be discussed here.

11.1   Linear Interpolation and Extrapolation

A linear approximation can be obtained with just two data points, say (x_1, y_1) and (x_2, y_2). You learned a long time ago that two points define a line and that the two-point formula for a line is

$$y = y_1 + \frac{(y_2 - y_1)}{(x_2 - x_1)}(x - x_1)$$

This formula can be used between any two data points to linearly interpolate. For example, if x in this formula is half-way between x_1 and x_2 at $x = (x_1 + x_2)/2$ then it is easy to show that linear interpolation gives the obvious result $y = (y_1 + y_2)/2$.
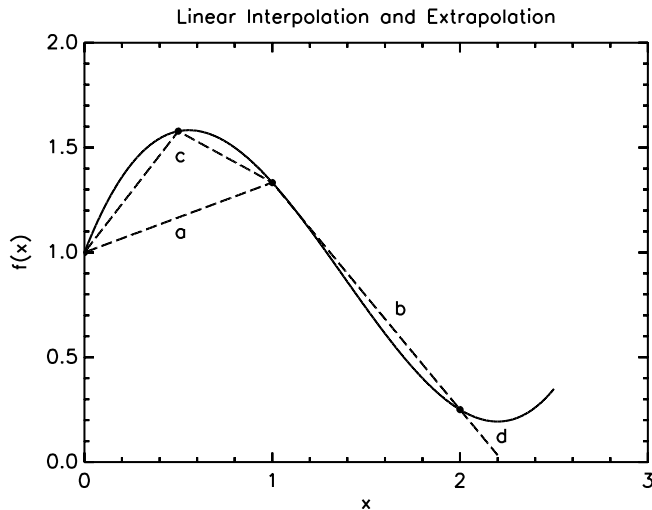
Figure 5: Linear interpolation only works well over intervals where the function is straight.

But you must be careful when using this method that your points are close enough together to give good values. In Fig. 11, for instance, the linear approximation to the curved function represented by the dashed line "a" is pretty poor because the points $x = 0$ and $x = 1$ on which this line is based are just too far apart. Adding a point in between at $x = 0.5$ gets us the two-segment approximation "c" which is quite a bit better. Notice also that line "b" is a pretty good approximation because the function doesn't curve much.

This linear formula can also be used to extrapolate. A common way extrapolation is often used is to find just one more function value beyond the end of a set of function pairs equally spaced in x. If the last two function values in the array are f_{N-1} and f_N, it is easy to show that the formula above gives the simple rule

$$f_{N+1} = 2f_N - f_{N-1}$$

which was used in the Matlab code in Sec. 10.1

You must be careful here as well: segment "d" in Fig. 11 is the linear extrapolation of segment "b", but because the function starts to curve again "d" is a lousy approximation unless x is quite close to $x = 2$.

11.2   Quadratic Interpolation and Extrapolation

Quadratic interpolation and extrapolation are more accurate than linear because the quadratic polynomial $ax^2 + bx + c$ can more easily fit curved functions than the linear polynomial $ax + b$. Consider Fig. 11.2. It shows two quadratic fits to the curved function. The one marked "a" just uses the points $x = 0, 1, 2$ and is not very accurate because these points are too far apart. But the approximation using $x = 0, 0.5, 1$, marked "b", is really quite good, much better than a two-segment linear fit using the same three points would be.

Unfortunately, the formulas for quadratic fitting are more difficult to derive. (The Lagrange interpolation formulas, which you can find in most elementary numerical analysis books, give these formulas.) But for equally spaced data in x, Taylor's theorem, coupled with the approximations to the first and second derivatives discussed in the section on numerical derivatives, make it easy to derive and use quadratic interpolation and extrapolation.
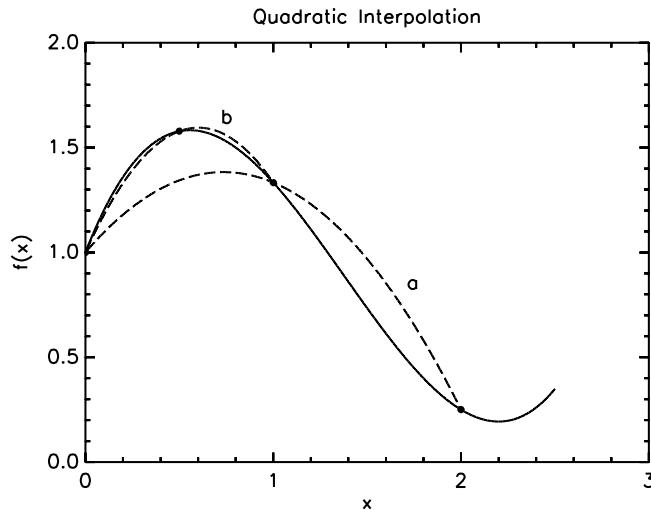
Figure 6: Quadratic interpolation follows the curves better if the curvature doesn't change sign.

I want to do it this way because it uses the approximate derivative formulas we used in Sec. 10.1 and illustrates a technique which is widely used in numerical analysis.

You may recall Taylor's theorem that an approximation to the function $f(x)$ near the point $x = a$ is given by

$$f(x) \approx f(a) + f'(a)(x - a) + \frac{1}{2}f''(a)(x - a)^2 + ...$$

Let's use this approximation (ignoring all terms beyond the quadratic term in $(x - a)$) near a point (x_n, f_n) in an array of equally spaced x values. The grid spacing in x is h. An approximation to Taylor's theorem that uses numerical derivatives in this array is then given by

$$f(x) \approx f_n + \frac{f_{n+1} - f_{n-1}}{2h}(x - x_n) + \frac{f_{n-1} - 2f_n + f_{n+1}}{2h^2}(x - x_n)^2$$

This formula is very useful for getting function values that aren't in the array. For instance, it is easy to use this formula to obtain the interpolation approximation to $f(x_n + h/2)$

$$f_{n+1/2} = -\frac{1}{8}f_{n-1} + \frac{3}{4}f_n + \frac{3}{8}f_{n+1}$$

and also to find the quadratic extrapolation rule

$$f_{N+1} = 3f_N - 3f_{N-1} + f_{N-2}$$

11.3 Interpolating With `polyfit` and `polyval`

You can also use Matlab's polynomial commands to build an interpolating polynomial. Here is an example of how to use them to find a 5th-order polynomial fit to a crude representation of the sine function.

```
%begin 11.3
```

```
% make the crude data set with dx too big for
% good accuracy

dx=pi/5;

x=0:dx:2*pi;

y=sin(x);

% make a 5th order polynomial fit to this data

p=polyfit(x,y,5);

% make a fine x-grid

xi=0:dx/20:2*pi;

% evaluate the fitting polynomial on the fine grid

yi=polyval(p,xi);

% and display the fit, the data, and the exact sine function

plot(x,y,'b*',xi,yi,'r-',xi,sin(xi),'c-')
pause

% display the difference between the polynomial fit and
% the exact sine function

plot(xi,yi-sin(xi),'b-')

%end 11.3
```

11.4   Matlab Commands Interp1 and Interp2

Matlab has its own interpolation routine `interp1` which does the things discussed in this section automatically. Suppose you have a set of data points $\{x, y\}$ and you have a different set of x-values $\{x_i\}$ for which you want to find the corresponding $\{y_i\}$ values by interpolating in the $\{x, y\}$ data set. You simply use any one of these three forms of the `interp1` command:

```
 yi=interp1(x,y,xi,'linear')
 yi=interp1(x,y,xi,'cubic')
 yi=interp1(x,y,xi,'spline')
```

    Here is an example of how each of these three types of interpolation works on a crude data set representing the sine function.

```
%begin 11.4a
```

```
% make the crude data set with dx too big for
% good accuracy

dx=pi/5;
x=0:dx:2*pi;
y=sin(x);

% make a fine x-grid

xi=0:dx/20:2*pi;

% interpolate on the coarse grid to
% obtain the fine yi values

% linear interpolation

yi=interp1(x,y,xi,'linear');

% plot the data and the interpolation

plot(x,y,'b-',xi,yi,'r-')
pause

% cubic interpolation

yi=interp1(x,y,xi,'cubic');

% plot the data and the interpolation

plot(x,y,'b-',xi,yi,'r-')
pause

% spline interpolation

yi=interp1(x,y,xi,'spline');

% plot the data and the interpolation

plot(x,y,'b-',xi,yi,'r-')

%end 11.4a
```

Matlab also knows how to do 2-dimensional interpolation on a data set of $\{x, y, z\}$ to find approximate values of $z(x, y)$ at points $\{x_i, y_i\}$ which don't lie on the data points $\{x, y\}$. You would use `zi = interp2(x,y,z,xi,yi,'linear')`, `zi = interp2(x,y,z,xi,yi,'cubic')`, or `zi = interp2(x,y,z,xi,yi,'spline')`. This will work fine for 1-dimensional data pairs $\{x_i, y_i\}$, but you might want to do this interpolation

for a whole bunch of points over a 2-d plane, then make a surface plot of the interpolated function $z(x, y)$. Here's some code to do this and compare these three interpolation methods (linear, cubic, and spline).

```
%begin 11.4b

x=-3:.4:3;
y=x;

% build the full 2-d grid for the crude x and y data
% and make a surface plot

[X,Y]=meshgrid(x,y);
Z=cos((X.^2+Y.^2)/2);
surf(X,Y,Z);
pause

%**********************************************************
% now make a finer 2-d grid, interpolate linearly to
% build a finer z(x,y) and surface plot it.

% Note that because the interpolation is linear the mesh is finer
% but the crude corners are still there
%**********************************************************

xf=-3:.1:3;
yf=xf;
[XF,YF]=meshgrid(xf,yf);
ZF=interp2(X,Y,Z,XF,YF,'linear');
surf(XF,YF,ZF);
pause

%**********************************************************
% Now use cubic interpolation to round the corners.  Note that there is
% still trouble near the edge because these points only have data on one
% side to work with, so interpolation doesn't work very well
%**********************************************************

ZF=interp2(X,Y,Z,XF,YF,'cubic');
surf(XF,YF,ZF);
pause

%**********************************************************
% Now use spline interpolation to also round the corners and see how
% it is different from cubic.  You should notice that it looks better,
% especially near the edges.  Spline interpolation is often the
% best.
%**********************************************************
```

```
ZF=interp2(X,Y,Z,XF,YF,'spline');
surf(XF,YF,ZF);
pause

%end 11.4b
```

For more detail see *Mastering Matlab 6*, Chapter 19.

12   FFT (Fast Fourier Transform)

12.1   Fourier Analysis

Suppose that you went to a Junior High band concert with a tape recorder and made a recording of Mary Had a Little Lamb. Your ear told you that were a whole lot of different frequencies all piled on top of each other, but perhaps you would like to know exactly what they were. You could display the signal on an oscilloscope, but all you would see is a bunch of wiggles. What you really want is the spectrum: a plot of sound amplitude vs. frequency. If this is what you want, Matlab can give it to you. It works like this.

12.2   Matlab's FFT

First you need the data in digital form, so you have to run your recording through an analog-to-digital converter to build a file of numerical signal values. Matlab will read such data files in several ways, but the simplest to use with a time series would just be a text file with two columns, one for time t and one for the signal $f(t)$. Or perhaps it could just be the signal in one column and you know what the time step τ is because you know the sampling rate.

If the 2-column file were called `signal.dat` and if it had two columns Matlab would read it this way:

```
load signal.dat;
% the data is now stored in the variable signal as an Nx2 matrix
% unpack it into t and f(t) arrays

t=signal(:,1);
tau=t(2)-t(1);
f=signal(:,2);
```

Or if the file only contains one column of data and you know the interval between the data points then Matlab would read it this way

```
load signal.dat;
f=signal;
N=length(f);
tau=.001;  % tau is the time interval between the points in f
t=0:tau:(N-1)*tau;
```

Now you have a series of data points (t_j, f_j) equally spaced in time by time interval τ and there are N of them covering a total time from $t = 0$ to $t = T = (N-1)\tau$. The Matlab

function `fft` will convert the array $f(t)$ into an array $g(\omega)$ of amplitude vs. frequency, like this

```
g=fft(f);
```

If you look at the array g you will find that it is full of complex numbers. It is complex because it stores phase information as well as amplitude information, just like the Fourier transform of mathematical analysis. For comparison, here is the formula for the analytic Fourier transform followed by the set of sums (the sum is done for each value of discrete frequency index k) that Matlab does with its `fft` command:

$$\text{Fourier transform}: \quad \mathcal{F}(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} f(t)e^{-i\omega t}dt$$

$$\text{FFT}, \quad k = 0 : N-1 \quad : \quad F_{k+1} = \sum_{j=0}^{N-1} f_{j+1}e^{-i2\pi jk/N}$$

Often we don't want the phase information, so we work instead with the so-called *power spectrum*, gotten this way:

```
P=abs(g).^2
```

Now we would like to plot $P(\omega)$ vs. angular frequency ω so we can see the values of the frequencies in our signal. The spectrum $P(\omega)$ is ready to go, but to see what the frequencies are, carefully compare the analytic and FFT formulas above, using $j = t_j/\tau$. We find that the frequency interval from one point in P to the next is $\Delta f = 1/(N\tau)$ or $\Delta\omega = 2\pi/(N\tau)$; the lowest frequency is 0; and the highest frequency in the array is $f_{max} = (N-1)/(N\tau)$ or $\omega_{max} = 2\pi(N-1)/(N\tau)$. Hence, the frequency array f (in cycles per second) and the ω array (in radians per second) would be built this way.

```
N=length(f);

% build f (Hertz)

df=1/(N*tau);
f=0:df:1/tau-df;

% build w (radians/sec)

dw=2*pi/(N*tau);
w=0:dw:2*pi/tau-dw;
```

An important thing to notice from the definition of the frequency step size $\Delta\omega$ is that if you want to distinguish between two frequencies ω_1 and ω_2 in a spectrum, then you must have $\Delta\omega << |\omega_2 - \omega_1|$. You can refine this resolution in frequency by choosing a large value for the length of the time series, t_{final}.

Another important thing to notice is that the maximum frequency you can detect is $\omega_{max} \approx 2\pi/\tau$, so if you want to see high frequencies you need a small time step, which

usually means lots of data points. (As you will see shortly the maximimum frequency you can reliably detect is only half of this value due to aliasing.)

This puts the peaks at the right places, but you may be wondering what the huge vertical scale means. The numbers are large for two reasons. (1) The FFT is not a Fourier transform, but is instead a related sum (see above). To convert g, the sum from FFT, to something closer to a Fourier integral you have to multiply it by the time step τ and divide by 2π, as indicated in the definition of the Fourier transform above. This will reduce the size of the FFT because $\tau/(2\pi)$ is typically a small number. (2) The Fourier transform just tends to be large anyway because instead of amplitude, it is amplitude density (amplitude per unit frequency). So if the signal is confined to a tiny range in ω, its density will be huge. Consider, for instance, the Fourier transform of $\cos\omega_0 t$:

$$\mathcal{F}(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \cos\omega_0 t e^{-i\omega t} dt = \frac{1}{2} \left(\delta(\omega + \omega_0) + \delta(\omega - \omega_0) \right)$$

Yes, those are delta-functions, and they are infinite (but they have finite area.) So when Matlab does the FFT on periodic data, the result is a bunch of approximate delta functions, which have large amplitudes. Because of this problem, most people just ignore the units on the vertical axis of the FFT and concentrate on the relative heights of the peaks.

You will find that all data sets are not created equal when `fft` is applied to them. Sometimes it will give you the answer really fast and sometimes it will be slow. You can make it run at optimum speed if you always give it data sets with 64, 128, 1024, 16384, etc. (powers of 2) data points in them. (Use `help fft` to see how to give `fft` a second argument so powers of 2 are always used.)

Here is an example to show how this process works.

```
%begin 12.2
%***********************************************************
% build a time series made up of 5 different frequencies
% that lasts for about 2000 seconds,
% then use fft to display the spectrum
%***********************************************************

N=16384;tau=2000/N;
t=0:tau:(N-1)*tau;

% Make a signal consisting of angular frequencies
% w=1, 3, 3.5, 4, and 6

f=sin(t)+.5*sin(3*t)+.4*sin(3.5*t)+.7*sin(4*t)+.2*sin(6*t);
plot(t,f)
pause

% the time plot is very busy and not very helpful

% now take the fft and display the power spectrum

g=fft(f);
P=abs(g).^2;
```

```
dw=2*pi/(N*tau);
w=0:dw:2*pi/tau-dw;
plot(w,P)
xlabel('\omega')
ylabel('P(\omega)')
title('Power Spectrum')
pause


%**********************************************************
% Notice that the right side of this plot is a mirror
% image of the left side.  This is called aliasing
% and it makes the upper half of the spectrum be
% completely bogus--we should ignore it.
% we do this by using the axis command to only look
% at half of it
%**********************************************************

axis([0 max(w)/2,0 max(P)]);

%end 12.2
```

Notice that the ω's where there are peaks in the spectrum are right where they should be: [1,3,3.5,4,6]. Also notice that the peak heights are qualitatively in the right relation.

Finally, you need to be aware of the limitations of this technique. As you can see from the discussion above, the lowest frequency you can detect with the FFT is $\Delta\omega = 2\pi/T$. So if you want see low frequency features in the spectrum you need data over a long time T. As you can also see, the highest frequency that you can see without aliasing trouble is $\omega_{max} = \pi/\tau$ (this important limiting frequency is called the *Nyquist frequency*), so to see high frequencies you need small time steps between data points. And to see both high and low frequencies you just need lots and lots of points, maybe even more than your computer memory will hold. So you need to design your data-taking carefully to capture the features you want to see without breaking your computer.

For more details see online help or *Mastering Matlab 6*, Chapter 21.

13   Make Your Own Functions: Inline and M-files

As you use Matlab to solve problems you will probably want to build your own Matlab functions. You can do this either by putting simple expressions into your code by using Matlab's `inline` command, or by defining function files with .m extensions called M-files.

13.1   Inline Functions

Matlab will let you define expressions inside a script for use as functions *within that script only.* For instance, if you wanted to use repeated references to the function

$$f(x, y) = \frac{\sin{(xy)}}{x^2 + y^2}$$

you would use the following syntax (to make both a line plot in x with $y = 2$ and to make a surface plot)

```
clear;close all;
f=inline('sin(x.*y)./(x.^2+y.^2)','x','y');

x=-8:.1:8;y=x;

plot(x,f(x,2))

[X,Y]=meshgrid(x,y);
surf(X,Y,f(X,Y))
```

The expression that defines the function is in the first string argument to `inline` and the other string entries tell `inline` which argument goes in which input slot when you use the function.

13.2   M-file Functions

These are subprograms stored in text files with .m extensions and are called M-files. A function is different than a script in that the input parameters it needs are passed to it with argument lists like Matlab commands (think about `sin(x)` or `plot(x,y,'r-')`). *Note, however, that anything inside Matlab functions is invisible in the command window.* So to debug a function you need to use print and plot commands in the function file. Or you can make it a stand-alone script by commenting out the function line so that its variables are available at the Matlab command level.

You can also pass information in and out of functions by using Matlab's `global` command to declare certain variables to be visible in all Matlab routines in which the global command appears. For instance, if you put the command

```
global a b c;
```

both in a script that calls `derivs.m` (see below) and in `derivs.m` itself, then if you give a, b, and c values in the main script, they will also have these values inside `derivs.m`. This construction will be especially useful when we use Matlab's differential equation solving routines in Chapter 16.

Rather than give you a syntax lecture I will just give you three useful functions as examples, with comments about how they work.

13.3 Derivative Function `derivs.m`

The first is a function called `derivs.m` which takes as input a function array y representing the function $y(x)$, and dx the x-spacing between function points in the array. It returns `yp` and `ypp`, numerical approximations to the first and second derivatives, as discussed in the section on numerical differentiation. First I will give you the script, then I will show you how to use it.

```
%beginfunction derivs.m

function [yp,ypp]=derivs(y,dx)

%***********************************************************
% This function numerically differentiates the array y which
% represents the function y(x) for x-data points equally spaced
% dx apart.  The first and second derivatives are returned as
% the arrays yp and ypp which are the same length as the input
% array y.  Either linear or quadratic extrapolation is used
% to load the derivatives at the endpoints.  The user decides
% which to use by commenting out the undesired formula below.
%***********************************************************

% load the first and second derivative arrays at the interior points

N=length(y);
yp(2:N-1)=(y(3:N)-y(1:N-2))/(2*dx);
ypp(2:N-1)=(y(3:N)-2*y(2:N-1)+y(1:N-2))/(dx^2);

% now use either linear or quadratic extrapolation to load the
% derivatives at the endpoints

% linear
%yp(1)=2*yp(2)-yp(3);yp(N)=2*yp(N-1)-yp(N-2);
%ypp(1)=2*ypp(2)-ypp(3);ypp(N)=2*ypp(N-1)-ypp(N-2);

% quadratic
yp(1)=3*yp(2)-3*yp(3)+yp(4);yp(N)=3*yp(N-1)-3*yp(N-2)+yp(N-3);
ypp(1)=3*ypp(2)-3*ypp(3)+ypp(4);ypp(N)=3*ypp(N-1)-3*ypp(N-2)+ypp(N-3);

%end derivs.m
```

To use this function first build an array of function values:

```
%begin 13.3

x=0:.01:10;
y=cos(x);
```

```
% Then, since the function returns two arrays in the form
% [yp,ypp], you would use it this way:

[fp,fpp]=derivs(y,.01);

% look at the approximate derivatives

plot(x,fp,'r-',x,fpp,'b-')

%end 13.3
```

Note that I didn't have to call them [yp,ypp] when I used them outside the function in the main script. This is because all variables inside functions are local to these programs and Matlab doesn't even know about them in the command window.

Sorry, I lied–I need to bore you with some syntax because everybody gets confused about the first line in a function program. The syntax of the first line is this

```
function output=name(input)
```

The word `function` is required. `output` is the thing the function passes back to whomever called it and its name is local to the function. If it is a single variable name then the code in the function needs to assign that variable name a value, an array of values, or a matrix of values. If the function returns more than one such result then the names of these results are put in square brackets, as in `derivs.m`. The function returns these results to the assigned variable(s), as in the derivs example:

```
[fp,fpp]=derivs(y,dx);
```

The keyword `name` in the function command line above should be the name of the `.m` file that contains the function code. (You can use other names, but you might drive yourself nuts if you do.) `input` is the argument list of the function. When the function program is called the arguments passed in and the arguments used in the function must match in number and type.

13.4 Definite Integral Function `defint.m`

This function uses the midpoint rule to integrate a function over a chosen interval using a chosen number of integration points. The function to be integrated must be coded in the sub function contained at the end of the function file `defint.m` (Note: Matlab's `quad` and especially `quadl` do the same thing, only better. This is just a simple example to show you how to program.)

```
%beginfunction defint.m

function s=defint(a,b,N)
```

```
%**********************************************************
% this function uses the midpoint rule on N subintervals
% to calculate the integral from a to b of the function
% defined in the sub function at the bottom of this
% function

% load dx and build the midpoint rule x-values
%**********************************************************

dx=(b-a)/N;
x=a+.5*dx:dx:b-.5*dx;

%**********************************************************
% use the function f(x) defined in the sub function below
% to obtain the midpoint approximation to the integral and assign
% the result to s
%**********************************************************

s=sum(f(x))*dx;

% here's the sub function

function y=f(x)

% define the function f(x) and assign it to y

y=cos(x);

%end defint.m
```

To use it, first edit the file `defint.m` so that the sub function at the bottom of the file contains the function you want to integrate. Then give `defint.m` the integration limits (say $a = 0$ and $b = 1$) and the number of points $N = 1000$ to use in the midpoint rule like this

```
defint(0,1,1000)
```

```
% or
```

```
s=defint(0,1,1000);
```

In the first case the approximate integral prints on the screen; in the second it doesn't print but is assigned to `s`.

13.5   Indefinite Integral Function `indefint.m`

This function takes an array of function values in `y` and an x-spacing `dx` and returns an approximate indefinite integral function $g(x) = \int_a^x y(x')dx'$. The function values must start

at $x = a$ and be defined at the edges of the subintervals of size h rather than at the centers as in the midpoint method. Because of this we have to use the trapezoid rule instead of the midpoint rule. The trapezoid rule says to use as the height of the rectangle on the interval of width h the average of the function values on the edges of the interval:

$$\int_x^{x+h} y(x')dx' \approx \frac{y(x) + y(x + h)}{2} h$$

Note that this function does exactly the same thing as Matlab's function `cumtrapz`.

```
%beginfunction indefint.m

function g=indefint(y,dx)


%**********************************************************
% returns the indefinite integral of the function
% represented by the array y.  y(1) is assumed to
% be y(a), the function value at the lower limit of the
% integration.  The function values are assumed to be
% values at the edges of the subintervals rather than
% the midpoint values.  Hence, we have to use the
% trapezoid rule instead of the midpoint rule:
%
% integral(y(x)) from x to x+dx is (y(x)+y(x+dx))/2*dx

% The answer is returned as an array of values defined
% at the same points as y
%**********************************************************

% the first value of g(x) is zero because at this first value
% x is at the lower limit so the integral is zero

g(1)=0;

N=length(y);

% step across each subinterval and use the trapezoid area
% rule to find each successive addition to the indefinite
% integral

for n=2:N

   % Trapezoid rule
   g(n)=g(n-1)+(y(n-1)+y(n))*.5*dx;

end

%end indefint.m
```

Take a minute now and use `derivs.m` and `indefint.m` to make overlaid plots of the function $f(x) = \cos x e^{-x}$, its first and second derivatives, and its indefinite integral $F(x) = \int_0^x f(x')dx'$ on the interval $x = [0, 3]$.

14 Fitting Functions to Data

A common problem that physicists often encounter is to find the best fit of a function (with a few variable parameters) to a set of data points. If you are fitting to a polynomial, then the easiest way to do this is with `polyfit`. But if you want to fit to a sine, a cosine, an exponential, a log, etc., then there is no simple command available. Pay attention to this section; it is useful.

Matlab has a very nice multidimensional minimizer routine called `fminsearch` that will do such general fits if you give it a half-decent initial guess for the fitting parameters. To see how to use `fminsearch` generally you can use online help, but here I will just show you how to use it for function fitting.

Suppose we have a set of data points (x_j, y_j) (perhaps read in with Matlab's `load` command as discussed in the Fast Fourier Transform section) and a proposed fitting function of the form $y = f(x, a_1, a_2, a_3, ...)$. For example, we could try to fit to an exponential function with two adjustable parameters a_1 and a_2 as is done in the example in `leastsq.m` below:

$$f(x, a_1, a_2) = a_1 e^{a_2 x} \quad .$$

Or you could fit to a cubic polynomial in x^2 with four adjustable parameters a_1, a_2, a_3, a_4 with this f:

$$f(x, a_1, a_2, a_3, a_4) = a_1 + a_2 x^2 + a_3 x^4 + a_4 x^6$$

which is what `polyfit` does.

In any case, what we want to do is choose the parameters $(a_1, a_2, a_3, ...)$ in such a way that the sum of the squares of the differences between the function and the data is minimized, or in mathematical notation we want to minimize the quantity

$$S = \sum_{j=1}^{N} (f(x_j) - y_j)^2$$

The first thing you need to do is to make a Matlab M-file called `leastsq.m` which evaluates the least-squares sum you are trying to minimize. It needs access to your fitting function $f(x, a)$, which is stored in the Matlab M-file `funcfit.m`. Here are examples of these two files.

This function `leastsq.m` is the one that you give to `fminsearch`. It allows `fminsearch` to talk to your fitting function `funcfit.m` by calculating S from the (x, y) data and the a_n's; it looks like this:

```
%beginfunction leastsq.m

function s=leastsq(a,x,y)
```

```
%*********************************************************
% leastsq can be passed to fminsearch to do a
% non-linear least squares fit of the function
% funcfit(a,x) to the data set (x,y).
% funcfit.m is built by the user as described here

% a is a vector of variable parameters; x and y
% are the arrays of data points
%*********************************************************

% find s, the sum of the squares of the differences
% between the fitting function and the data

s=sum((y-funcfit(a,x)).^2);

%end leastsq.m


%beginfunction funcfit.m

function f=funcfit(a,x)

%*********************************************************
% this function evaluates the fitting
% function f(x,a1,a2,a3,...) to be fit to
% data.  It is called by leastsq.

% a is a vector of variable fitting parameters a1, a2, ...
% that are used in the function and x is a
% vector of x-values

% the function returns a vector f=f(x,a1,a2,...)
%*********************************************************

% sample exponential function with 2 variable
% parameters

f = a(1)*exp(a(2)*x);

%end funcfit
```

With these two functions built and sitting in your Matlab directory we are ready to do the fit. Here is a piece of code that allows you to enter an initial guess for the fitting parameters, plot the initial guess against the data, then tell `fminsearch` to do the least squares fit. The behavior of `fminsearch` can be controlled by setting options with Matlab's `optimset` command. In the code below this command is used to set the Matlab variable `TolX`, which tells `fminsearch` to keep refining the parameter search until the parameters are determined to a relative accuracy of `TolX`. Finally, it plots the best fit against the data. I suggest you

give it a name (perhaps `datafit.m`) and save it for future use. The data needs to be sitting in the file `data.fil` as two columns of (x, y) pairs, like this

```
0.0   1.10
0.2   1.20
0.4   1.52
0.6   1.84
0.8   2.20
1.0   2.70
```

Let's try to fit this data to the function $f(x) = a_1 \exp(a_2 x)$.

OK–here's the fitting code; I call it `datafit.m`

```
%begin 14

%***********************************************************
% Uses fminsearch to least squares fit
% a function defined in funcfit.m to
% data read in from data.fil
%***********************************************************

clear;close

% read the data file and load x and y

load data.fil;
x=data(:,1);y=data(:,2);

% set up for the plot of the fitting function

xmin=min(x);
xmax=max(x);
npts=1001;
dx=(xmax-xmin)/(npts-1);
xplot=xmin:dx:xmax;

% set ifit to 0 and don't continue on to the fit until
% the user sets it to 1

ifit=0;

while ifit==0

disp(' Enter an initial guess for the function ')
a=input('parameters [a1,a2,...] in vector form [...]- ')

% plot the data and the initial function guess

yplot=funcfit(a,xplot);
```

```
plot(x,y,'b*',xplot,yplot,'r-')
xlabel('x')
ylabel('y')
title(' Initial Guess and Data')

ifit=input(' Enter 0 to try again, 1 to try to fit - ')

end

%*********************************************************
% Do the fit with the option TolX set; fminsearch will adjust a
% until each of its elements is determined to within TolX.
% If you think fminsearch could do better than it did, reduce TolX.
%*********************************************************

option=optimset('TolX',1e-5);
a=fminsearch(@leastsq,a,option,x,y)

% plot the data and the final function fit

yplot=funcfit(a,xplot);

% Plot the final fit and the data

plot(x,y,'b*',xplot,yplot,'r-')
xlabel('x')
ylabel('y')
title(' Final Fit and Data')

%end 14
```

It's a little painful to have to make three files to get this job done, but I suggest you learn how to use **fminsearch** this way. It comes up all the time.

15   Systems of Nonlinear Equations

Matlab's `fminsearch` can also be used to solve systems of nonlinear equations. Consider the following pretty-impossible-looking set of three equations in three unknowns (x, y, z).

$$\sin(xy) + \exp(-xz) - 0.95908 = 0 \tag{4}$$

$$z\sqrt{x^2 + y^2} - 6.70820 = 0 \tag{5}$$

$$\tan(y/x) + \cos z + 3.17503 = 0 \tag{6}$$

The way to talk `fminsearch` into solving this set is to invent a scalar function $S(x, y, z)$ which consists of the sum of the squares of the three functions given above:

$$S(x, y, z) = [\sin(xy) + \exp(-xz) - 0.95908]^2 +$$

$$\left[z\sqrt{x^2 + y^2} - 6.70820\right]^2 + [\tan(y/x) + \cos z + 3.17503]^2$$

If you look at $S(x, y, z)$ for a bit you will see that (1) it is always positive, (2) its smallest possible value is zero, and (3) if you can somehow find values of (x, y, z) that make it zero, then you have solved the system of three equations.

Note, however, that `fminsearch` is a *minimizer*, not a zero finder. So it may find a local minimum of $S(x, y, z)$ which does not satisfy $S = 0$. If it fails in this way you need to (1) know about it and (2) make another initial guess so `fminsearch` can take another crack at the problem.

Here are two pieces of Matlab code that will use `fminsearch` to solve systems like this. The first one calls `fminsearch` and the second one, which is a Matlab function (`eqsystem.m`) contains the equations to be solved.

```
%begin 15

%************************************************************
% Uses fminsearch to look for solutions to the
% nonlinear system of equations defined in the
% file eqsystem.m
%************************************************************

clear;

itry=0;

while itry ==0

disp(' Enter an initial guess for the solution')
x=input('of the system of equations in the form [x1,x2,...] - ')

% evaluate the scalar function S(x1,x2,...) so the
% user can see how good the guess is

s=eqsystem(x);
fprintf(' S(x1,x2,...) = %g \n',s)

itry=input(' Enter 0 to try again, 1 to try to solve - ')

end

x=fminsearch(@eqsystem,x)
```

```
s=eqsystem(x);
fprintf(' Final value of S(x1,x2,...) = %g \n',s)
disp(' Make sure it is close to zero')

%end 15


%beginfunction eqsystem.m

function s=eqsystem(xn)

% nonlinear system of equations routine for
% use with zerofinder and fminsearch

x=xn(1);y=xn(2);z=xn(3);

Eq1 = sin(x*y)+exp(-x*z)-0.95908;
Eq2 = z*sqrt(x^2+y^2) -6.70820;
Eq3 = tan(y/x)+cos(z)+3.17503;

s=Eq1^2+Eq2^2+Eq3^2;

%end eqsystem.m
```

When you solve this system and you are asked for an initial guess, try something near (0.8,1.8,2.5).

16   Ordinary Differential Equations

The standard way to write a differential equation, or a system of differential equations, in numerical work is as a first order system, like this:

$$\frac{d\mathbf{x}}{dt} = \mathbf{F}(\mathbf{x})$$

where $\mathbf{x}$ is a vector of unknown functions of the parameter t (often time in physics problems) and where $\mathbf{F}(\mathbf{x})$ is a vector-valued function of a vector argument. This makes perfect sense to a mathematician, but physicists usually need examples. Here are a couple.

16.1   Decay of a Radioactive Sample

If there are N atoms of an unstable element with an exponential decay rate of γ then the differential equation describing how N decreases in time is

$$\frac{dN}{dt} = -\gamma N$$

which is just a single first order differential equation.

16.2 Simple Harmonic Oscillator

The equation of motion of a mass m bouncing in a weightless environment on a spring with spring constant k is

$$\frac{d^2x}{dt^2} = -\omega_0^2 x \quad \text{where} \quad \omega_0 = \sqrt{\frac{k}{m}}$$

This is a second order differential equation rather than a first order system, so we need to change its form. This is done simply by using position $x(t)$ and velocity $v(t) = dx/dt$ as two unknown functions of time. The first order set consists of the definition of $v(t)$ in terms of $x(t)$ and the second order differential equation with d^2x/dt^2 replaced by dv/dt:

$$\frac{dx}{dt} = v \tag{7}$$

$$\frac{dv}{dt} = -\omega_0^2 x \tag{8}$$

It is always possible to use this trick of defining new functions in terms of derivatives of other functions to convert a high order differential equation to a first order set.

So let's assume that we have a first order set. How can we solve it? Maple will often give you an analytic expression for the solution; Matlab can only give you an array which is a numerical approximation. But Matlab will find this approximation very quickly and you have control over how it does it, so it is often more efficient to use Matlab for hard differential equations than to use Maple.

I am going to show you two methods. The first is simple, intuitive, and inaccurate. The second is a little more complicated, not terribly intuitive, but pretty accurate. There are many ways to numerically solve differential equations and the two I will show you are rather crude; Matlab has its own solvers (discussed in Sec. 16.5) which are better, but you will learn a bit about the ideas they are based on by studying these two methods.

16.3 Euler's Method

The first method is called Euler's Method (say "Oiler's Method"), and even though it's pretty bad, it is the basis for many better methods. Here's the idea.

First, quit thinking about time as a continuously flowing quantity. Instead we will seek the solution at specific times t_n separated by small time steps τ. Hence, instead of $\mathbf{x}(t)$ we will try to find $\mathbf{x}_n = \mathbf{x}(t_n)$. The hope is that as we make τ smaller and smaller we will come closer and closer to the true solution.

Since we are going to use discrete times and since the initial conditions tell us where to start, what we need is a rule that tells us how to advance from $\mathbf{x}_n$ to $\mathbf{x}_{n+1}$. To find this rule let's approximate the differential equation $d\mathbf{x}/dt = \mathbf{F}$ this way

$$\frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{\tau} = \mathbf{F}(\mathbf{x}_n, t_n)$$

In doing this we are assuming that our solution is represented as an array of values of both t and $\mathbf{x}$, which is the best that Matlab can do. If we already know $\mathbf{x}_n$, the solution at the present time t_n, then the equation above can give us $\mathbf{x}$ one time step into the future at time $t_{n+1} = t_n + \tau$:

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \mathbf{F}(\mathbf{x}_n, t_n)\tau$$

This is a little abstract, so let's use it to approximately solve the harmonic oscillator equation. For this case Matlab would use for **x** the vector [x,v] and for **F** the vector [v,-w∧2*x]. (Stare at the harmonic oscillator equation given above as a first order system until you can see that this is true.) Here's a script that uses this method to solve the harmonic oscillator equation.

```
%begin 16.3

% Use Euler's method to solve the harmonic oscillator equation
clear;close all;

% set the angular frequency
w=1;

% decide how long to follow the motion, 10 periods in this case
tfinal=2*pi/w*10;

% choose the number of time steps to take
N=input(' Enter the number of time steps to take - ')

% t=zeros(1,N+1);x=zeros(1,N+1);v=zeros(1,N+1); % uncomment this line to make
                                               % the code run faster

% calculate the time step
tau=tfinal/N;

% initialize the time array
t(1)=0;

% set the initial values of position and velocity
x(1)=1;v(1)=0;

% Do Euler's method for N time steps
for n=1:N
  t(n+1)=n*tau;
  x(n+1)=x(n) + v(n)*tau;
  v(n+1)=v(n) - w^2*x(n)*tau;
end

% plot the result and compare it with the exact solution
% which is x(t)=cos(w*t)
plot(t,=,'r-',t,cos(w*t),'b-')

%end 16.3
```

When you copy this code into a file and run it you will see that even if you take 1000 steps, the solution is not very good. No matter how small `tau` is, if you run long enough Euler's method will blow up.

Also note that if you try to run this script for many steps ($N = 50,000$, for instance) it
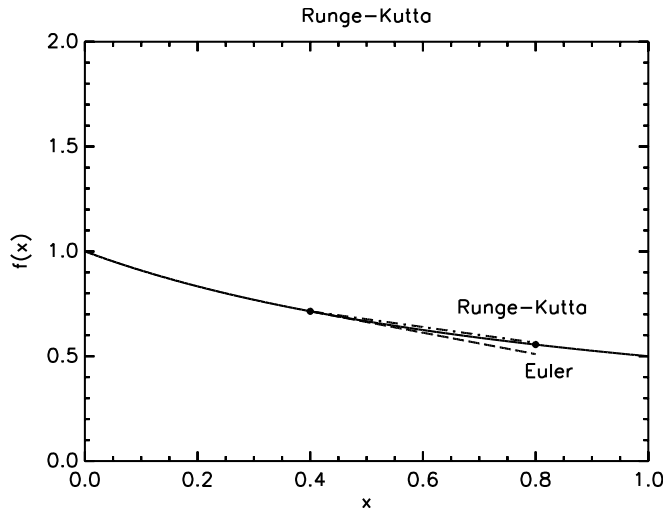
Figure 7: Runge-Kutta anticipates curving and beats Euler.

runs slow. The reason is that as you keep making the t, x, and v arrays longer and longer in the loop Matlab has to keep making new space for them in memory. But if you define them ahead of time to be big enough (see the commented lines just after the line $N = input...$ in the code above), the arrays are defined to be big enough before you start the loop and no time will be wasted increasing the array sizes. Run this script again with this line of code uncommented and watch how fast it runs, even if you choose $N = 500,000$.

16.4 Second-order Runge-Kutta

Here is a method which is still quite simple, but works a lot better than Euler. But it is, in fact, just a modification of Euler. If you go back and look at how we approximated $d\mathbf{x}/dt$ in Euler's method you can see one thing that's wrong with it: the derivative is not centered in time:

$$\frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{\tau} = \mathbf{F}(\mathbf{x}_n, t_n)$$

The left side of this equation is a good approximation to the derivative halfway between t_n and t_{n+1}, but the right side is evaluated at t_n. This mismatch is one reason why Euler is so bad.

Runge-Kutta attempts to solve this centering problem by what looks like a cheat: (1) Do an Euler step, but only by $\tau/2$ so that we have an approximation to [x,v] at $t_{n+1/2}$. These half-step predictions will be called [xhalf,vhalf]. (2) Then evaluate the function $\mathbf{F}(\mathbf{x}, t)$ at these predicted values to center the derivative

$$\frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{\tau} = \mathbf{F}(\mathbf{x}_{n+1/2}, t_{n+1/2})$$

This is the simplest example of a *predictor-corrector* method and it works lots better than Euler, as you can see by running the code given below. (Step (1) above is the predictor; step (2) is the corrector.)

You can see this difference between Runge-Kutta and Euler in Fig. 7, where the upward curve of the solution makes Euler miss below, while Runge-Kutta's half-way-out correction to the slope allows it to do a much better job.

72

```
%begin 16.4

% Runge-Kutta second order approximate solution to the harmonic oscillator
clear;close all;

% set the angular frequency
w=1;

% decide how long to follow the motion, 10 periods in this case
tfinal=2*pi/w*10;

% choose the number of time steps to take
N=input(' Enter the number of time steps to take - ')

% calculate the time step
tau=tfinal/N;

% initialize the time array
t(1)=0;

% set the initial values of position and velocity
x(1)=1;v(1)=0;

% Do Runge-Kutta for N time steps
for n=1:N
    t(n+1)=n*tau;

    % Predictor step .5*tau into the future
    xhalf=x(n) + v(n)*tau*.5;
    vhalf=v(n) - w^2*x(n)*tau*.5;

    % Corrector step
    x(n+1)=x(n) + vhalf*tau;
    v(n+1)=v(n) - w^2*xhalf*tau;

end

% plot the result and compare it with the exact solution
% x(t)=cos(w*t) and v(t)=-w*sin(w*t)

plot(t,x,'r-',t,cos(w*t),'b-')

%end 16.4
```

16.5 Matlab's Differential Equation Solvers

Matlab also has its own differential equation solvers and they are more accurate than the
simple methods discussed in Secs. 16.3 and 16.4. Here is a list borrowed from *Mastering*

Matlab 6 of these Matlab functions and what they do.

ode23:

An explicit, one-step Runge-Kutta low-order (2-3) solver. (Like the second-order Runge-Kutta method predictor-corrector discussed here in Sec. 16.4.) Suitable for problems that exhibit mild stiffness, problems where lower accuracy is acceptable, or problems where $\mathbf{F}(t, \mathbf{x})$ is not smooth (e.g. discontinuous).

ode45:

An explicit, one-step Runge-Kutta medium-order (4-5) solver. Suitable for non-stiff problems that require moderate accuracy. *This is typically the first solver to try on a new problem.*

ode113:

A multi-step Adams-Bashforth-Moulton PECE solver of varying order (1-13). Suitable for non-stiff problems that require moderate to high accuracy involving problems where $\mathbf{F}(t, \mathbf{x})$ is expensive to compute. Not suitable for problems where $\mathbf{F}(t, \mathbf{x})$ is discontinuous or has discontinuous lower-order derivatives.

ode23s:

An implicit, one-step modified Rosenbrock solver of order 2. Suitable for stiff problems where lower accuracy is acceptable, or where $\mathbf{F}(t, \mathbf{x})$ is discontinuous. *Stiff problems are those in which there are several different rates of change involved whose sizes differ by several orders of magnitude, or more.*

ode15s:

An implicit, multi-step solver of varying order (1-5). Suitable for stiff problems that require moderate accuracy. *This is typically the solver to try if* ode45 *fails or is too inefficient.*

Below you will find two sample scripts odetest and rhs which can use any of these Matlab solvers to solve and plot the solution of the harmonic oscillator equation. Note that they all work in this way:

(1) You define the right-hand side function for your set of first order differential equations in the M-file rhs.m.

(2) You choose the beginning and ending times to pass into the Matlab ode function.

(3) You choose the initial column vector y representing the initial conditions for your problem.

(4) You choose the ode solver control options by using Matlab's odeset function.

(5) You ask Matlab to give you a column of times t and a matrix of y-values by calling one of the ode solvers like this

```
[t,y]=ode45(@rhs,[tstart,tfinal],y,options);
```

(6) This command returns a column vector t of the discrete times between tstart and tfinal which ode45 chose to make the solution be as accurate as you asked it to be when you used odeset. You will also receive a matrix y with as many columns as you have unknowns in your set of ode's and with as many rows as you have times in t. And if you make the required accuracy smaller, you will receive more data points. *Important note: they*

will not be equally spaced in time. If you just want the solution at certain pre-set times $[t_n]$*, just replace the 2-element array* `[tstart,tfinal` *with an array of the times that you want:* `[t1,t2,t3,t4,...,tN]`*. For example, you could replace* `[tstart,tfinal]` *with the equally spaced array of times* `tstart:dt:tfinal`*.* If the position x is called `y(1)` in your set of equations then you can obtain an array containing these positions by extracting the first column, like this

```
x=y(:,1);
```

Once you have extracted the different components of your solution from `y`, i.e., x, v_x, z, v_z, etc., you can use Matlab's plotting and data analysis capabilities to slice and dice the data anyway you want.

Oh, and what if you really wanted the data to be equally spaced in time because, for example, you wanted to send the solution off to `fft` to find a spectrum? Just use `interp1` to interpolate the output of `ode45` onto an equally spaced time grid and away you go.

Here's the sample differential equation solving code.

odetest.m (main script):

```
%begin 16.5


%***********************************************************
% ordinary differential equation solver using
% Matlab's ode solvers and the M-file rhs.m to
% specify F(t,y)
%***********************************************************

clear;close all;

% declare the oscillator frequency to be global and set it
global w0;
w0=1;

% set the initial and final times
tstart=0;tfinal=200;

% set the initial conditions in the y0 column vector
y0=zeros(2,1);
y0(1)=.1; % initial position
y0(2)=0; % initial velocity

% set the solve options
options=odeset('RelTol',1e-8);

% do the solve
   [t,y]=ode45(@rhs,[tstart,tfinal],y0,options);

% unload the solution that comes back in y into x and v arrays
x=y(:,1);v=y(:,2);
```

```
%**********************************************************
% because Matlab's ode solvers don't use equally spaced
% time steps, and because you might want equal spacing,
% here's how you convert from Matlab's unequally-spaced (t,x,v)
% to equally spaced data (te,xe,ve)
%**********************************************************

N=length(t);
taue=(tfinal-tstart)/(N-1);
te=tstart + (0:taue:(N-1)*taue) ;
te=te';  % convert te to a column vector, to match t
xe=interp1(t,x,te,'spline');
ve=interp1(t,v,te,'spline');

%**********************************************************
% Note that you could have obtained equally-spaced points by
% telling ode45 to give you the solutions at times you specify.
% For instance, suppose you wanted 1024 points between t=0 and
% t=200.  You could build them like this (the code is commented):

% N=1024;
% taue=(tfinal-tstart)/(N-1);
% te=tstart + (0:taue:(N-1)*taue) ;
% [t,y]=ode45(@rhs,te,y,options);
% xe=y(:,1);ve=y(:,2);
%**********************************************************


% plot the position vs. time

plot(te,xe)
pause

% make a "phase-space" plot of v vs. x

plot(xe,ve)

%end 16.5



    rhs.m (Matlab function)

%beginfunction rhs.m

%**********************************************************
% right-hand side function for Matlab's ordinary
```

```
% differential equation solvers: simple harmonic
% oscillator example:

% x=y(1)         v=y(2)

%       dx                dv
%       --  = v    ;      -- = -w^2*x
%       dt                dt
%**********************************************************

function F=rhs(t,y)

% declare the frequency to be global so its value
% set in the main script can be used here
global w0;

% make the column vector F filled
% with zeros
F=zeros(length(y),1);

% build the elements of F

% the equation dx/dt=v means that F(1)=y(2)
F(1)=y(2);

%  the equation dv/dt=-w0^2*x means that F(2)=-w0^2*y(1)
F(2)=-w0^2*y(1);

%end rhs.m
```

16.6   Event Finding with Matlab's Differential Equation Solvers

Something you will want to do with differential equation solvers is to find times and variable values when certain events occur. For instance, suppose we are solving the simple harmonic oscillator and we want to know when the position of the oscillator goes through zero with positive velocity, as well as when the velocity is zero and decreasing. I have good news and bad news. The good news is that Matlab knows a way to do this. The bad news is that the way is a little involved. If you try to figure out how it works from online help you will be confused for a while, so I suggest that you use the example files given below. Read them carefully because I have put all of the explanations about how things work in the codes as comments. The main file is eventode.m and its right-hand side function is eventrhs.m. There is also an additional M-file to control the event-finding called events.m.

eventode.m (main script):

```
%begin 16.6
```

```
% eventode: example of event finding in Matlab's ode solvers

clear;close;

dt=.01;  % set the time step
y0=[0;1];  % put initial conditions in the [x;vx] column vector

% turn the eventfinder on by specifying the name of the M-file
% where the event information will be processed (events.m)
options=odeset('Events',@events,'RelTol',1e-6);

% call ode45 with event finding on and a parameter omega passed in
omega=1;

[t,y,te,ye,ie]=ode45(@eventrhs,[0,20],y0,options,omega);

%*********************************************************
% Here's what the output from the ode solver means:
% t: array of solution times
% y: solution vector, y(:,1) is x(t), y(:,2) is vx(t)
% te: array of event times
% ye: solution vector at the event times in te
% ie: index for the event which occurred, useful when you
%     have an array of events you are watching instead of
%     just a single type of event.  In this example ie=1
%     for the x=0 crossings, with x increasing, and ie=2
%     for the vx=0 crossings, with vx decreasing.

% separate the x=0 events from the vx=0 events
% by loading x1 and v1 with the x-positions and
% v-velocities when x=0 and by loading x2 and v2
% with the positions and velocities when v=0
%*********************************************************

n1=0;n2=0;
for m=1:length(ie)

    if ie(m)==1
       n1=n1+1;

% load event 1: x,v,t
       x1(n1)=ye(m,1);v1(n1)=ye(m,2);t1(n1)=te(m);
    end

% load event 2: x,v,t
    if ie(m)==2
```

```
        n2=n2+1;
        x2(n2)=ye(m,1);v2(n2)=ye(m,2);t2(n2)=te(m);
     end

end

% plot the harmonic oscillator position vs time
plot(t,y(:,1),'g-')
hold on

% plot the x=0 crossings with red asterisks and the v=0
% crossings with blue asterisks
plot(t1,x1,'r*')
plot(t2,x2,'b*')

%end 16.6
```

eventrhs.m (Matlab function to compute $[dy(1)/dt;dy(2)/dt]$

```
%beginfunction eventrhs.m
%--------------------------------------------------------------------

function rhs=eventrhs(t,y,omega)

% right-hand side for the simple harmonic oscillator
% make sure rhs is a column vector

rhs(1,1)=y(2);
rhs(2,1)=-omega^2*y(1);

%--------------------------------------------------------------------
%end eventrhs.m
```

events.m (Matlab function to control event finding by Matlab's ode solvers)

```
%beginfunction events.m
%--------------------------------------------------------------------

function [value,isterminal,direction] = events(t,y,omega)

%***********************************************************
% Locate the time and velocity when x=0 and x is increasing

% value array:  same dimension as the solution y.  An event is defined
%               by having some combination of the variables be zero.
%               Since value has the same size as y (2 in this case) we
%               can event find on two conditions.  Should be a column vector
```

```
%************************************************************

value(1,1) = y(1);  % load value(1) with the expression which,
                    % when it is zero, defines the event, y(1)=0 in this case.

value(2,1)=y(2);    % load value(2) with a second event condition, vx=0
                    % (y(2)=0) in this case.  If you don't want a second
                    % event just set value(2)=1 so it is never 0.

isterminal = [0 ; 0];  % this vector tells the integrator whether
                       % to stop or not when the event occurs.
                       % 1 means stop, 0 means keep going.  isterminal
                       % must have the same length as y (2 in this case).
                       % Should be a column vector

direction = [1 ; -1];  % direction modifier on the event:
                       % 1 means value=0 and is increasing;
                       % -1 means value=0 and is decreasing;
                       % 0 means value is zero and you don't care
                       % whether it is increasing or decreasing.
                       % direction must have the same length as y.
                       % should be a column vector

%------------------------------------------------------------------
%end events.m
```

17   Publication Quality Plots

With a bit of coaxing Matlab will make plots that are suitable for inclusion in a thesis or for publication in a journal. The material in this section is included courtesy of Tom Jenkins.

The default settings which Matlab uses to plot functions are usually OK for looking at plots on a computer screen, but they are generally pretty bad for generating graphics for articles to be published in a journal. Getting around this problem is the purpose of this section. First let's make a simple Matlab plot of the sine function.

```
clear;close all;
dx=.01;
x=0:dx:2*pi;
f=sin(x);

plot(x,f,'r-')
```

This figure has a couple of problems. The axes Matlab chose on which to plot the function aren't necessarily the ones we would pick, the line width in the plot is too small, making the plot look "spidery", and the lettering is too small. Fortunately, Matlab allows us to make better choices in these matters.

The way to control how Matlab handles these parameters, and almost any other parameters we might be interested in, is by using the `set` utility in Matlab. This allows us to specify how we want our plot to look, instead of letting Matlab use its own default settings. Here's an example.

```
%begin 17a

clear;close all;
dx=.01;
x=0:dx:2*pi;
f=sin(x);
plot(x,f,'r-','LineWidth',2.5)
aa=gca;
set(aa,'XLimMode','manual','YLimMode','manual',...
       'Xlim',[0 2*pi],'Ylim',[-1 1],'FontSize',20,'LineWidth',2.5,...
       'XTick',[0 .5 1 1.5 2 2.5 3 3.5 4 4.5 5 5.5 6],...
       'YTick',[-1 -.8 -.6 -.4 -.2 0 .2 .4 .6 .8 1])
xlabel('\theta')
ylabel('sin (\theta)')
title('sin (\theta) vs. \theta')

%end 17a
```

Execute this code and see what the plot looks like.

Although this structure is (of course) more complicated than the original plot, it does make a graph that's suitable for publication. Here is the same code, but in more detail (comments added).

```
%begin 17b

clear;close all;
% Defines the function we're going to plot
```

```
dx=.01;
x=0:dx:2*pi;
f=sin(x);

%***********************************************************
% Make the plot.  The LineWidth command tells Matlab how many 'points'
% wide to make each line.  The default is 0.5, so we've made the lines
% five times thicker than normal.
%***********************************************************

plot(x,f,'r-','LineWidth',2.5)

%***********************************************************
% Get the handle to the current axes.  Matlab uses 'handles' (a kind
% of label) to refer to each set of axes.  Matlab's default axes appear
% on the plot at this point, so we have to figure out the handle to
% this set of axes so we can change its properties.  The 'gca' command
% stands for Get Current Axes and returns a number (the label Matlab
% uses to refer to the axes); we store this number in the variable 'aa'.
% This way we can change the properties of 'aa' and Matlab will know
% what we're talking about.
%***********************************************************

aa=gca;

%***********************************************************
% Change the properties of these axes using the 'set' command.
% The syntax is
% set(aa,'PropertyName','PropertyValue').
% I have used multiple PropertyNames and PropertyValues in the same
% command, but this is the basic idea.  It tells Matlab to take the
% object with handle aa and set its PropertyName to PropertyValue.
% Note that PropertyNames must be enclosed in single forward quotes.

% 'XLimMode' and 'YLimMode' can be set to 'auto' or 'manual'.  Note
% that these settings are strings, so they need to be enclosed in
% single forward quotes as well.  If these PropertyNames are set
% to 'auto' Matlab sets its own limits for the left, right, top,
% and bottom edges of the graph.  If they are set to 'manual'
% we can use the 'XLim' and 'YLim' commands, followed by vectors
% telling where we want the edges of the graph to be located.

% 'FontSize' sets the size of the font to be used in the xlabel,
% ylabel, zlabel, and title commands for the plot.

% 'LineWidth' sets the width of the lines the axes (not the plot)
% are drawn with.  Its format is similar to the earlier 'LineWidth'
```

```
% command referring to the plot itself.

% 'XTick' and 'YTick' refer to the tick marks along the x and y
% axes respectively.  Their arguments are vectors; Matlab will
% put a tick mark at each vector element along that axis.  (Too many
% tick marks makes the plot illegible, so be careful).
%********************************************************

set(aa,'XLimMode','manual','YLimMode','manual',...
       'Xlim',[0 2*pi],'Ylim',[-1 1],'FontSize',20,'LineWidth',2.5,...
       'XTick',[0 .5 1 1.5 2 2.5 3 3.5 4 4.5 5 5.5 6],...
       'YTick',[-1 -.8 -.6 -.4 -.2 0 .2 .4 .6 .8 1])
xlabel('\theta')
ylabel('sin (\theta)')
title('sin (\theta) vs. \theta')

%end 17b
```

Here is the revised plot.

Index