

ANDY RIDGWELL

GEO111 – NUMERICAL SKILLS IN GEOSCIENCE

UNIVERSITY OF CALIFORNIA, RIVERSIDE / DEPARTMENT OF EARTH SCIENCES
2015/6

Copyright © 2016 Andy Ridgwell

<http://www.seao2.info/teaching.html>

Except where otherwise noted, content of this document is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 3.0 license (CC BY-NC-SA 3.0) (<http://creativecommons.org/licenses/by-nc-sa/3.0/>)

First printing, October 2016

Contents

1	<i>Elements of ... Computers and software</i>	19
2	<i>Elements of ... MATLAB and data visualization</i>	21
2.1	<i>Using the MATLAB software</i>	22
2.1.1	<i>Starting MATLAB</i>	22
2.1.2	<i>The command line</i>	22
2.1.3	<i>MATLAB GUI</i>	22
2.1.4	<i>Help(!)</i>	23
2.2	<i>Basic concepts</i>	24
2.2.1	<i>Variables</i>	24
2.2.2	<i>Numerical expressions</i>	26
2.2.3	<i>Relational and logical operators</i>	27
2.2.4	<i>Functions (built-in)</i>	28
2.2.5	<i>Miscellaneous commands</i>	28
2.3	<i>Vectors and arrays #1</i>	30
2.3.1	<i>Creating vectors</i>	30
2.3.2	<i>Basic vector manipulation</i>	30
2.3.3	<i>Addressing elements in vectors</i>	31
2.4	<i>Basic graphing (aka. 'data visualization')</i>	33
2.4.1	<i>Plotting</i>	33
2.4.2	<i>Graph labelling</i>	34
2.4.3	<i>Sub-plots</i>	34
2.4.4	<i>Saving graphics and figures</i>	35

2.5	<i>Vectors and arrays #2</i>	36
2.5.1	<i>Creating matrices and arrays</i>	36
2.5.2	<i>Basic matrix manipulation</i>	37
2.5.3	<i>Some matrix math :(</i>	39
2.6	<i>Loading and saving data</i>	40
2.6.1	<i>Where am I?</i>	40
2.6.2	<i>Loading and importing data</i>	41
2.6.3	<i>Saving and exporting data</i>	41
2.6.4	<i>Loading and saving the workspace</i>	41
2.7	<i>Basic data processing</i>	42
2.8	<i>Yet more graphing</i>	46
2.8.1	<i>Modifying lines/symbols in plot</i>	46
2.8.2	<i>Plotting multiple data-sets</i>	46
2.8.3	<i>Scatter plots</i>	47
2.8.4	<i>Histograms</i>	49
2.8.5	<i>Simple 2D data and bitmap visualization</i>	50
3	<i>Elements of ... programming</i>	51
3.1	<i>Introduction to scripting (programming!) in MATLAB</i>	52
3.1.1	<i>Programming good practice</i>	53
3.1.2	<i>Debugging the bugs in buggy code</i>	55
3.2	<i>Functions</i>	58
3.3	<i>Conditionals '101'</i>	60
3.3.1	<i>if ...</i>	60
3.3.2	<i>switch ...</i>	64
3.4	<i>Loops '101'</i>	66
3.4.1	<i>for ...</i>	66
3.4.2	<i>Other loop configurations and usages</i>	69
3.4.3	<i>Fun(!) worked examples</i>	70
3.5	<i>Loops and conditionals ... together(!)</i>	74
3.5.1	<i>for ... and conditionals</i>	74
3.5.2	<i>while ...</i>	76

3.6	<i>Even more (and loopier) loops</i>	79
4	<i>Further ... Computers and software</i>	83
5	<i>Further ... MATLAB and data visualization</i>	85
5.1	<i>Further data input</i>	86
5.1.1	<i>Formatted text (ASCII) input</i>	86
5.1.2	<i>Importing ... Excel spreadsheets</i>	89
5.1.3	<i>Importing ... netCDF format data</i>	90
5.2	<i>Further data processing</i>	94
5.2.1	<i>Data interpolation</i>	95
5.3	<i>Further (spatial / (x,y,z)) plotting</i>	96
5.3.1	<i>Plotting maps</i>	103
6	<i>Further ... Programming</i>	105
6.1	<i>find!</i>	106
	<i>Bibliography</i>	113
	<i>Index</i>	115

List of Figures

1	Course schedule: weeks 1 through 5.	13
2	Course schedule: weeks 6 through 10.	14
2.1	Default output of plot.	33
2.2	A plot illustrating axis auto-scaling (maximum x and y values now slightly larger than 10 and 100, respectively).	34
2.3	A (only very slightly) improved plot.	34
2.4	Arrangement of subplots.	35
2.5	Spline fit to measured changes in CO ₂ concentration in Law Dome ice core, following <i>Etheridge et al.</i> [1996].	41
2.6	proxy reconstructed past variability in atmospheric CO ₂ .	42
2.7	Proxy reconstructed past variability in atmospheric CO ₂ (sorted data).	43
2.8	Proxy reconstructed past variability in atmospheric CO ₂ (sorted data).	46
2.9	Proxy reconstructed past variability in atmospheric CO ₂ (sorted data).	47
2.10	Proxy reconstructed past variability in atmospheric CO ₂ (scatter plot).	47
2.11	Proxy reconstructed past variability in atmospheric CO ₂ (scatter plot).	47
2.12	A 2D plot of some random gridded model data.	50
2.13	A 2D plot of some random gridded model data ... but with the underlying data matrix re-orientated before plotting.	50
3.1	Output from the (bug-fixed version of) <code>plot_some_dull_stuff m-file</code> .	57
3.2	Extremely unappealing blocky plot of Earth surface temperature (who cares with month? – the graphics are too poor to matter ...).	73
3.3	Continental outline (of sorts).	79
3.4	Another continental outline (of sorts).	79
3.5	Another go at the continental outline!	81
5.1	Very basic imaging (image) of an array (2D) of data – here, global bathymetry.	97
5.2	Slightly improved very basic imaging (imagesc) of bathymetry data.	97
5.3	Example result of basic usage of the contour function.	99
5.4	Example usage of <code>contourf</code> , with the hot <i>colormap</i> (giving dark/brown colors as deep ocean, and light/white as high altitude).	99
5.5	Example usage of <code>contour</code> , contouring only the zero height isoline, and providing a label.	99

- 5.6 Usage of `contour` but with lon/lat values created by `meshgrid` function and passed in (and with the hot *colormap* (giving dark/brown colors as deep ocean, and light/white as high altitude). 101
- 5.7 Example contour plot including `meshgrid`-generated lon/lat values. Result of `contourf(lon, lat, temp7, 30)`, where the data file was `temp7.tsv`, with some embellishments. 103
- 6.1 Proxy reconstructed past variability in atmospheric CO₂ (scatter plot). 110
- 6.2 Proxy reconstructed past variability in atmospheric CO₂ (scatter plot). 111

List of Tables

About the Course

GEO111 will provide an introduction to computer programming and numerical modelling for Earth and Environmental Science problems. It will provide a chance to learn a computer programming language and all the elements that constitute it, including concepts in number bases and types, logical constructs, debugging, etc. The course will develop programming skills step-wise, applying them at each point to practical questions and outcomes, such as data processing and visualization. How complex environmental processes can be encapsulated and approximated, and numerical models thereby constructed, will be illustrated. Guided opportunities will be provided to build a 'DIY' climate model and in doing so, further develop programming and modelling skills at the same time as reinforcing basic concepts in climate dynamics through practice in addition to theory.

The cumulating objectives of the course are to:

1. develop an understanding of how computers and the internet work and hence foster a critical understanding of modern technology,
2. provide hands-on training in how computer programs are written and numerical models constructed, and
3. develop both general (transferable) as well as specific numerical and analytical skills applicable to the Earth and Environmental Sciences.

The associated learning goals are firstly; to provide, through hands-on practical exploration, factual knowledge and an understanding of:

- Number bases, how computers work plus computer programs and their basic building blocks. (Learning Outcome 2).
- Numerical models and the representation of time. Construction and application of a variety of models spanning box models of biogeochemical cycles and population dynamics, through 1D reaction-diffusion models of surface Earth processes, to 3D gridded global models. (Learning Outcomes 1 and 2).
- The Greenhouse effect and basic climate feedbacks. (Learning Outcome 1).
- Awareness of how webpages and the internet work. (Learning Outcomes 2 and 4).
- The use of numerical models in addressing scientific questions and testing hypotheses as well as the limitations of numerical models. (Learning Outcomes 2 and 4).

and provide transferable skills in

- Written communication and presentation. (Learning Outcome 3).
- Problem solving and logical analysis, fault-finding. (Learning Outcomes 4 and 5).
- Computer programming. (Learning Outcomes 2 and 4).

Course logistics

Format

The weekly format of the Class is: one 1-hour lecture, together with one 3-hour computer practical session, plus a 2-hour interactive lecture/discussion session of worked problems and examples. The computer practical class is the central element, and will consist of structured exercises leading step-by-step through the components of computer programming and numerical model construction, debugging, and testing, plus applications to common geosciences problems. The lecture starting each week will outline the basics and introduce the key concepts of the week. The purpose of the 2-hour lecture/discussion session ending the week is to ensure all the concepts are understood and misconceptions resolved and will be a mix of presentation and worked-through examples, plus questions and discussion.

Timetable

Assessment

The course will be assessed as follows:

- Midterm paper – 50%
- Finals paper – 50%

The mid-term paper will be a written exam, consisting of a mixture of multiple choice and short-answer format questions. Its purpose is to test basic knowledge of computers and programming, plus general concepts and basic commands in MATLAB. The testable content comprises the material covered in the lectures and lab sessions (in weeks #1-5). The exam will be 2 hours long. The mid-term paper will constitute 50% of the total assessment of the course.

The Finals paper will be based upon the development and application of a computer model written in MATLAB and will be presented in the form of a science paper describing the model, its evaluation, application to a specified science question, plus discussion of model caveats and suggestions for future improvements. The scope of the model exercise will be somewhat restricted with a short menu of possible choices, but with considerable flexibility in terms of exactly what is done and explored with it (i.e. there is some slight possibility of actually having fun!). This will constitute the remaining 50% of the total assessment of the course.

	Monday am (1)	Monday am (2)	Friday pm
WEEK	Lecture A 08:10-09:00 GEOL 1444	Computing lab 09:10-12:00 Watkins 2101	Lecture B 14:10-16:00 Watkins 2101
19 / 23rd Sept.			Introduction to the course [Chapter 0] Format and content of the course. Office hours. Overview of course assessment.
26 / 30th Sept.	Computing '101' #1 [Chapter 1] Basic constituents and functioning of computers. Computer operating systems, programs, and software. Compiled and interpreted languages. Introduction to MATLAB.	Lab-based exercises [Chapter 2] MATLAB basics, including variables and matrixes, data I/O. Data processing in MATLAB. Basic statistics. Basic plotting and data visualization.	Worked examples & Q&A
3 / 7th Oct.	Fundamentals of computer programming [Chapter 3] Bases, logic and logic gates. Loops and conditionals. Subroutines and functions.	Lab-based exercises [Chapter 3] Sub-programs and functions. Code de-bugging practice.	Worked examples & Q&A
10 / 14th Oct.	Algorithms and problem-solving [Chapter 6] Algorithms and numerical techniques. Search and sort algorithms. Programming best practice and debugging.	Lab-based exercises [Chapter 6]	Worked examples & Q&A
17 / 21st Oct.	Data visualization [Chapter 5]	Lab-based exercises [Chapter 5] 2-D plotting and interpolation. Re-gridding. Data binning and histograms. Accessing and visualizing netCDF format data.	Worked examples & Q&A
24 / 28th Oct.	Computing '101' #2 [Chapter 4] How the internet 'works'. Computer networks. Web pages and basic html. Program GUIs.	Lab-based exercises [Chapter 4] Building a MATLAB GUI.	MIDTERM EXAM

Figure 1: Course schedule: weeks 1 through 5.

	Monday am (1)	Monday am (2)	Friday pm
WEEK	Lecture A 08:10-09:00 GEOL 1444	Computing lab 09:10-12:00 Watkins 2101	Lecture B 14:10-16:00 Watkins 2101
31 / 4th Nov.	Introduction to numerical modelling [Chapter 7] Time-stepping and integration techniques. Numerical stability and accuracy. Model code structure.	0-D (box) modelling [Chapter 7] Example: simple global radiocarbon box (reservoir) model.	Worked examples & Q&A
7 / 11th Nov.	Global climate modelling #1 [Chapter 8] Fundamental climate system processes and their representation in models. Surface energy budget and greenhouse gases. Heat capacity. Atmospheric transport.	Lab-based exercises [Chapter 8] Example: 0-D climate (energy balance) model.	VETERANS DAY
14 / 18th Nov.	Global climate modelling #2 [Chapter 8] Evaluation of numerical models. Model-data assessment. Numerical models in the literature, the work of the IPCC, and model 'inter-comparisons'.	Lab-based exercises [Chapter 8] Example: 1-D climate (energy balance) model.	Worked examples & Q&A
21 / 25th Nov.	Biogeochemical modelling [Chapter 9] Derivation and application of empirical relationships. Geochemical and biogeochemical models. Modelling ocean biogeochemical cycles. 1-D reaction-transport models.	Lab-based exercises [Chapter 9] Example: 1-D gas diffusion and consumption in soils.	THANKSGIVING
28 / 2nd Dec.	Earth system modelling [Chapter 10] Modelling feedbacks in the Earth system.	Lab-based exercises [Chapter 10] Example: 'Daisy World'.	Worked examples & Q&A

Figure 2: Course schedule: weeks 6 through 10.

Office Hours

There are no specific Office Hours, but rather an open invitation to drop by¹ (excluding Thursdays) and/or email² questions. Part of the purpose of the lab session on Fridays is to provide an opportunity for further clarification of the course material and to go through worked examples.

¹ My office is in the Geology building, room 464 (basement floor).

² andy@seao2.org

Course text

There is no one (or even two between them) commercial (published) course texts that covers both basic computer programming and numerical modelling at a suitable level, and certainly not in the context of MATLAB. Hence the reason for this *e*-book – to provide a 1-stop shop for a range of information and practical tutorials in useful and commonly used data manipulation and visualization, numerical techniques, and programming methodologies.

In conjunction with the course text (this document), a recommended (but not required) course textbook is *Matlab (Third Edition): A Practical Introduction to Programming and Problem Solving*³, which provides a good general introduction to MATLAB and covers similar material to some of the course.

³ Stormy Attaway. *Matlab (Third Edition): A Practical Introduction to Programming and Problem Solving*. Butterworth-Heinemann, 2013

For additional reading (on both MATLAB and numerical modelling), try:

- *The Climate Modelling Primer (4th Edition)*, by Kendal McGuffie and Ann Henderson-Sellers. Wiley-Blackwell (2014). ISBN: 978-1-119-94336-5.
- *Introduction to MATLAB (3rd Edition)*, by Delores M. Etter. Prentice Hall (2014). ISBN: 978-0133770018.
- *Mathematical Modelling of Earth's Dynamical Systems, A Primer*, by Rudy Slingerland and Lee Kump. Princeton (2011). ISBN: 978-0-691-14514-3.

How to use this Textbook

Elements of ... Computers and software

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetur id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

Nam dui ligula, fringilla a, euismod sodales, sollicitudin vel, wisi. Morbi auctor lorem non justo. Nam lacus libero, pretium at, lobortis vitae, ultricies et, tellus. Donec aliquet, tortor sed accumsan bibendum, erat ligula aliquet magna, vitae ornare odio metus a mi. Morbi ac orci et nisl hendrerit mollis. Suspendisse ut massa. Cras nec ante. Pellentesque a nulla. Cum sociis natoque penatibus et magnis dis parturient montes, nascetur ridiculus mus. Aliquam tincidunt urna. Nulla ullamcorper vestibulum turpis. Pellentesque cursus luctus mauris.

Elements of ... MATLAB and data visualization

HELLO NEWBIES! This first lab's porpoise is to start to get you familiar with what MATLAB is all about and understand how to import and manipulate (array) data in this software environment and do some basic plotted (aka 'data visualization'). If you are clever, you might find menu items or buttons to click that will do the same thing as typing in boring commands at the command line. In fact, you would have to be pretty dumb not to notice all that brightly colored eye-candy in the GUI (Graphical User Interface – i.e., menus, buttons, and stuff) at the top of the screen. However, you will get to grips with programming much quicker if you stick with the instructions and do almost everything that is asked of you using the command line (rather than doing stuff via the GUI), at least to start with. You'll just have to trust me for now ... We'll start with the very basics and things that you could easily do in Excel instead, and build up.

GRAPHICS is one of the important strengths of **MATLAB**. Although other software packages and scripting languages exist that perhaps have the edge on **MATLAB** in terms of visually appealing plots and graphs, **MATLAB** is worlds apart from e.g. **Excel**.

2.1 Using the MATLAB software

2.1.1 Starting MATLAB

To start with: find the MATLAB icon on the desktop; run the program. You should see a number of sub-windows arranged within the main MATLAB window, hopefully including at the very least, the *Command Window*¹. Depending on whether you have used MATLAB before and it has remembered your settings, windows may also include: *Command History*, *Workspace*, *Current Folder*. If instead you see; 'Tetris', 'Grand Theft Auto: San Andreas', and 'World Championship Pool', then you have the wrong software running and are going to find learning MATLAB rather hard. However, there is big \$\$\$ to be made in on-line gaming tournaments these days. Would you really rather be a geologist and spend the rest of your days hitting rocks with a hammer? If so, read on ...

¹ Conveniently labelled Command Window
– you cannot possibly fail to identify it
...

2.1.2 The command line

When MATLAB initially starts up, the *Command Window* should display the following text:

```
Academic License
»
```

or in older versions of the software:

```
To get started, select MATLAB Help or Demos from the Help
menu.
»
```

but in either case, with a vertical blinking line (cursor) following the double 'greater than' symbols².

If you are unfamiliar with using command-line driven software ... Don't Panic! Nothing bad can happen, regardless of what you do. Well, almost. It is possible to accidentally clear MATLAB's memory of the results of calculations and data processing and close plots and graphs before you have saved them, but MATLAB remembers all the commands you type, so in theory it is perfectly possible to quickly reproduce anything lost. (Later on we will be placing the sequence of commands into a file (that is saved) and so ultimately, MATLAB should turn out to be mostly fool-proof.)

² Note that in nerd-speak the » is called the command 'prompt' and is prompting you to type some input (Commands, swear words, etc.). See – the computer is just sat there waiting for you to command it to go do something (stupid?). If one does not appear at the bottom of whatever is in the *Command Window* is means that MATLAB is busy doing something extremely important. Or perhaps, MATLAB may have completely died. Either way, it will not accept any new/further commands until it is done calculating/dying.

2.1.3 MATLAB GUI

There are lots of fancy looking icons and pretty colors and you could spent all day staring at them and not getting any work done. Or

learn good programming practice. Which is why we mostly will ignore the eye-candy and little (if any) guidance will be given as to the functionality of the GUI. Look at this as a lesson for the user (to read the Help, textbook, on-line documentation, or simple go Google for an answer³).

³ i.e. Internet fishing

2.1.4 *Help(!)*

Press F1 or click on the question mark icon on the tool-bar, to bring up the indexed and searchable MATLAB documentation.⁴

⁴ It is also possible to obtain context-specific help, e.g. on a specific (built-in) *function*, which we'll see in due course.

2.2 Basic concepts

2.2.1 Variables

A *variable* is, in a sense, a pointer to a location in computer memory where a piece of information is stored⁵. A variable is associated a name to make things rather more easy and convenient. The name can be anything you like in MATLAB, as long as it does not contain numbers or special characters. So actually, you are only allowed sequences of letters (otherwise known as ‘words’). But you can create a variable name based on 2 (or more) words, separated by an underscore (_). Valid variable names would include:

```
A
B
cat
derpyhooves
this_is_boring_stuff
BIG
big6
```

Variables are entirely useless unless they have some information assigned to them. In fact, you can type in any of the variable names above (at the command line) and MATLAB will deny it knows what you are talking about⁷.

So far so useless – you need to *assign* something to it. Which brings us to quite ‘what’ and ‘how’. First of, you need to know that variables can have the following *types*:

- **Integer** – An integer number is a counting number, i.e. 1, 2, 3, ... and including zero and negative integers. MATLAB has different representations for integer numbers, depending on how large a number you need to represent (and how much memory it will need to allocate to storing it). This is something of a throw-back to the days when computers only had $1/10000000^{th}$ of the memory of your iPhone and were slower than a lemon.
- **Real (floating point)**⁸ – A real number can have a non-integer component, e.g. 1.5 or $6.022140857 \times 10^{23}$. Real numbers also come in different precisions in MATLAB (also to do with memory allocation and speed), determining not just the number of decimal places that can be represented, but also the maximum size.
- **String (character)** – One or more characters, but now allowing spaces (unlike in the case of naming variables).
- **Logical** – true or false.
- **etc** – No, not a real type, but to note that MATLAB defines and recognises a whole bunch of other types, including **Complex**

⁵ In the bad old days, this pointer was the actual address in memory and might have looked something like f04da105.

⁶ Note that MATLAB distinguishes between lower and UPPER case letters in a variable (i.e. BIG and big would represent two different and distinct variables).

⁷ Technically, MATLAB reports: Undefined function or variable which tells you it is neither a function name (more on this later), nor is defined as having any information associated with it.

⁸ The distinction (sort of) is that *floating point* is a specific representation of a **real** number.

(MATLAB can handle *complex numbers*) and **Object** (we will also not worry about *objects*, which can incorporate a combination of types. At least, not yet ...).

The first thing to learn is to ideally, do not attempt to mix up (combine) variables of different types. MATLAB is very forgiving when it comes to combining an *integer* and a *real* number in the same calculation, but in other programming languages, this should be avoided. However, even in MATLAB, *strings* and *reals* (or *integers*) are very different things. When necessary, different variable types can be converted between (see **Variable Type Conversion** Box).

The second and perhaps rather more important thing, is how to assign a value to a variable (and in fact, create the variable in the first place). Programming languages such as FORTRAN require you to define the variable beforehand and assign it a type. MATLAB allows you to define and assign a value to a variable all at the same time, and it will kindly work out the correct type based on the value you assign to it. You assign a value using the assignment operator `=`⁹. For example:

```
A = 10
```

will assign the value 10 to the variable A. If you type this at the command line, MATLAB will kindly repeat what you have just told it and report the value of A back to you:

```
A =
10
```

Note that you do not need to add a space before and/or after the assignment operator (`=`). This is something of a personal programming and aesthetics preference, i.e. whether to pad things out with spaces or not. (Chose what you feel happiest with and later on, whatever leads to the fewest programming mistakes ...)

MATLAB will also report in the *Workspace* window, the name and value, type (called *Class*), etc of all your current variables (just one currently?). Actually, it is not all quite so simple. If you take a look at the *Class* of the variable A in the display window – it is listed as *double* (a *real* number) rather than an *integer*. So by default, if MATLAB does not know what you really want, it defines A as a double precision real number¹⁰.

The next complication comes when assigning a string (a sequence of characters) to a variable. For example, try:

```
B = apple
```

and MATLAB is far from happy. As it turns out, a sequence of characters can also refer to a *function*¹¹ in MATLAB, and this is what

Variable Type Conversion

MATLAB provides a variety of *functions* (see later) for converting between different *types* of *variables*. The most commonly-used/useful ones are as follows:

1. converting from a number to a *string* (s)
 - `s = num2str(N)`, where N is any number type variable
 - `s = int2str(I)`, where I is an integer
2. converting from a *string* (s) to a number
 - `x = str2num(s)`, where N is (generally) a double precision (*real*) number

Case #1 (`num2str`) is generally the most useful, e.g. in adding specific captions to plots (with caption text based on the value of a numerical variable) – examples are given later.

⁹ This is NOT 'equals' in MATLAB. We will see the *equality operator* shortly. `=` assigns the value or variable on its right the variable on the left.

¹⁰ If you genuinely wanted an integer, there are ways to do this, such as using a type conversion function from *real* to *integer* (see above).

¹¹ You will see *functions* shortly. For now – note that they are 'special' (reserved) words that perform some action and hence cannot also be used for a variable name.

MATLAB looks for (i.e. a match to apple in the list or variable (and function) names). To delineate apple as a string, you need to encase it in (single¹²) quotation marks:

```
B = 'apple'
```

Just as MATLAB creates new variables on the fly, you can re-assigned values to an existing variable, even if this means changing the *type*, e.g.

```
A = 'banana'
```

has now replaced the real number 10 with the character string **banana** in variable A. This is reflected in the updated variable list details given in the *Workspace* window (and a *Class* now listed as char).

Finally, it is possible to suppress output to the *Command Window* when making *assignments* – simply an a semi-colon (;) to the end of the *assignment* statement, i.e.

```
C = 'banana';
```

now does not results in anything being echoed to the command line (but the *Workspace* is still updated to reflect this variable assignment). If you wish to see the contents of the variable, you can either just type its name at the command line, or view its value as listed in the *Workspace* window.

2.2.2 Numerical expressions

You can do normal maths in MATLAB. Or at least, something that looks at least a little intuitive. (In fact, I often use MATLAB as a calculator.) The primary/common numerical expressions are:

- **exponentiation** — ^ — raises one number of variable to the power of a second, e.g. a^b , a to the power b, which is written in MATLAB as a^b .
- **multiplication** — × — e.g. $a \times b$, written in MATLAB as $a*b$.
- **division** — / — (written as you would expect).¹³
- **addition** — + — (guess).
- **subtraction** — - — again, obvious/intuitive.

The order in which the numerical operators are written down is important and MATLAB will execute them in a specific order (operators higher up the list, executed first), i.e. first ^, then *, /, and last +, -. There is also 'negation', when you change the sign of a variable, and which is executed immediately after exponentiation. The assignment operator (=)¹⁴ comes last. If you are unclear about the order numerical operators are carried out, then place parentheses () around the component of the calculation you wish to be carried

¹² Double "" quotation marks will not work.

¹³ Entertainingly, it turns out that if you write the reverse, backslash character (\) in the equation, you divide the over way (i.e. denominator divided by numerator).

¹⁴ This is **NOT** 'equals to', as you'll see shortly.

out first to enforce a particular order (this can also help in making an equation easier to read and ultimately, easier to debug code). For example, consider:

```
A = 3;
B = 6;
C = 2;
D = C*(A/B+1)
E = C*A/(B+1)
F = C*A/B+1
G = A*C/B+1
```

Try these out (and make up your own combinations) and confirm that the answers are what you would expecty them to be.

2.2.3 *Relational and logical operators*

We will see more of *relational and logical operators* later when we start to get into some proper coding. For now, you only need to know that a relational operator is one of:

- **greater than** — MATLAB symbol `>`
- **less than** — MATLAB symbol `<`
- **greater than or equal to** — MATLAB symbol `>=`
- **less than or equal to** — MATLAB symbol `<=`
- **equality** — MATLAB symbol `==`
- **inequality** — MATLAB symbol `~=`

and test the relationship between 2 variables. Note in particular, that the equality symbol (that tests the equivalence between two variables) is represented by TWO = characters (`==`), and remember that a single `=` character is the assignment operator.

In everyday language, the answer to any one of these relational tests would be a 'yes' or a 'no'. But in MATLAB (and other computer languages), the answer is given as the binary (logical) equivalent where 'yes' is represented by 1 and 'no' by 0. You can also use `true` (1) and `false` (0), e.g. `A = true` returns:

```
A =
    1
```

Finally, the *logical operators* (again, more on this later) are:

- **or** — symbol `||`
- **and** — symbol `&&`
- **not** — symbol `~`

For now – be familiar with how numerical expressions are written in MATLAB (you'll need to be using these from the outset), and keep in mind the existence of *relational and logical operators*.

2.2.4 Functions (built-in)

MATLAB provides numerous built-in **functions**¹⁵. These functions have specific names assigned to them, so care needs to be taken not to give a variable the same name as a function to avoid getting confused further down the road. Giving an exhaustive list (and brief description) is outside the scope of this document¹⁶. Common functions will be progressively introduced as this text progresses. Note that in addition to the on-line Help documentation, information on how to use a function and example uses is provided by typing `help` and then the function name (separated by a space) at the command line.

MATLAB also provides several built-in mathematical constants (saving having to define a variable with the appropriate number). These are simply variables that have been already defined and assigned values, but which you cannot change (hence the term 'constant'). For instance, the value of π , is assigned to a built-in variable with the name **pi**. You can access (display) its value by typing its name at the command line:

```
» pi
ans =
    3.1416
```

In this example, the use of the *function* is rather trivial – you need to tell the `pi` absolutely nothing, and it spits back the same thing (the value of π) each and every time. In most other *functions*, you will find that you have to pass some information, and the return value will depend on the input. (This will all become apparent in due course ...)

2.2.5 Miscellaneous commands

Related to what you have seen so far and will see soon, useful miscellaneous commands include:

- **clear** — Removes all variables from the workspace.
- **clear all** — (Removes all information from the workspace.)
- **close** — Closes the current figure window.
- **clear all** — (Closes all figure windows.)
- **exit** — Exits MATLAB and hence enables additional drinking time in the bar.

Note that a useful trick – if you want to re-use a previously used command but don't want to type it in all over again, or want to issue a command very similar to a previously-used one – is to hit the UP arrow key until the command you want appears. This can also be edited (navigate with LEFT and RIGHT arrow keys, and use Delete

¹⁵ We will be constructing our own later, at which point it should become apparent that there is nothing particularly special about them.

¹⁶ A full list of functions can be found in the MATLAB Help Documentation under *functions*.

and Backspace to get rid of characters) if needs be. Hit Enter to make it all happen.

2.3 Vectors and arrays #1

So far, your variables have all be what are known as *scalars* – i.e. single numbers (or strings). One of the most powerful things about MATLAB is its ability to represent vectors (1D columns or rows of numbers or strings) and arrays – 2D and higher dimensional regular grids of numbers or strings. (*matrix*¹⁷ is the name commonly given to a 2-D array.)

2.3.1 Creating vectors

Vectors are 1-D arrangements of numbers (or strings). You can enter them into MATLAB as a list of space-separated value, encased in (square) brackets, [], e.g.

```
B = [0.5 1.0 1.5 2.0 2.5]
```

or with the value comma-separated:

```
B = [0.5, 1.0, 1.5, 2.0, 2.5]
```

Either way, you end up with a vector on its side as a single row of numbers which in math-speak would look like:

$$B = \begin{pmatrix} 0.5 & 1.0 & 1.5 & 2.0 & 2.5 \end{pmatrix}$$

You can also create the equivalent, upright orientated vector (as a single column of numbers) by separating the elements by a semi-colon:

```
C = [0.5; 1.0; 1.5; 2.0; 2.5]
```

which gives the maths-speak representation:

$$C = \begin{pmatrix} 0.5 \\ 1.0 \\ 1.5 \\ 2.0 \\ 2.5 \end{pmatrix}$$

2.3.2 Basic vector manipulation

There are several basic and very useful ways of manipulating vectors (and as we'll see later – matrices). To start with, you might want to determine the orientation and length of a vector. There are several different ways to go about this, which in order of grown-up-ness are:

1. Display the contents of the vector in the command window by typing its name at the command line. Obviously, this will quickly become useless for very large vectors¹⁸.

¹⁷ Not to be confused with the film starting Keanu Reeves.

The **colon operator** can be used to much more rapidly create vectors (as long as the elements form a simple sequence in value) as compared to typing in the list of values explicitly. There are two variants to the syntax:

```
A = j:k
```

and

```
A = j:i:k
```

In the first example, j and k are the minimum and maximum values in the sequence of numbers in the vector. MATLAB completes the sequence by assuming that the values monotonically increase and that the elements are separated by one (1.0) in value. e.g.

```
>> A = 0:3
```

```
A =
```

```
0 1 2 3
```

Note that MATLAB is not inclined to let you directly create a vector of elements that decrease in value (you'll need to flip this puppy about to re-order it if that is what you want – see later).

In the second example, i is the increment MATLAB will use to complete the sequence from j to k. In the example in the text, you could have created the array B by typing:

```
>> B = 0.5:0.5:2.5
```

```
B =
```

```
0.5000 1.0000 1.5000
```

```
2.0000 2.5000
```

(More commonly, you might place the colon operator and its min/(increment)/max values inside a pair of brackets, i.e. A = [0:3]. so that it is unambiguous that you are creating an array

¹⁸ Try creating a vector from 1 to 100,000 and then displaying it ...

2. Refer to the Workspace window, although this also ends up a total Fail for long vectors.
3. Use the `length` or `size` function (see Box).

If you find that you want a different orientation (row vs. column) of the a vector, the vector can be flipped around (converting row-to-column and column-to-row) using the *transpose operator* (`.'`), e.g.:

```
D = B'
```

will turn the vector B into one (assigned to the variable D) with the same orientation as C.

You can also re-order the values in a vector (hence addressing the restriction in using the colon operator to create a vector that the values must be monotonically increasing rather than decreasing). Depending on the orientation of the vector, you can use either the `flipud` (for column vectors), or `fliplr` (for row vectors), functions to re-order the elements.

2.3.3 Addressing elements in vectors

Values can be extracted from a vector by specifying the index (technically, this should be an integer, but MATLAB is pretty forgiving and you can get away with using a real number when specifying an index) of the element required (counting along, left-to-right, or top-to-bottom, depending on the vector orientation), e.g.

```
» B(5)
ans =
    2.5000
```

or:

```
» C(3)
ans =
    1.5000
```

(In this text, I will refer to accessing a particular element (or elements) of a vector (or array) via its index as addressing. Unless I forget, then I might say something else. You'll have to keep on your toes – don't expect consistency here!)

There is a MATLAB function `end` (see Box) that enables you to easily address (accessing via its index) the very last value in a vector (in MATLAB, the index of the first position is always 1).

For addressing more than one element of a vector at a time, you can use the colon operator (see Box).

As well as reading out an existing value of a vector, you can also replace an existing value by assigning the new value to the appropriate index position. e.g. to replace the first element with a value of 0.0:

`length`

You can determine the length of a vector A with ...

```
length(A)
```

returning its integer length, and which could in turn be assigned to a variable, e.g. `B = length(A)`. (Technically, `length` returns the largest dimension of an array.)

`size` (use #1)

Returns both dimensions, even though for a vector, one of them always has a value of 1. This does allow you to determine its orientation though, as for the example of `A = [1:10]`:

```
» size(A)
ans =
     1 10
```

(1 row and 10 columns). For `A = A'`:

```
» size(A)
ans =
    10 1
```

(10 rows and 1 column).

`flipud`, `fliplr`

These two functions allow you to re-order a vector. Their use is simple:

```
» B = flipud(A)
```

will invert the order of elements of a column vector, and:

```
» B = fliplr(A)
```

will invert the order of elements of a row vector. Simple! Lesson over.

The **transpose operator**, in MATLAB-speak, "returns the nonconjugate transpose of A". Who knows what that means. In slightly more everyday (i.e. down here on Earth) language, it: "interchanges the row and column index for each element". Or sort of, just interchanges the rows and columns. The operation can be written:

```
» B = A.'
```

or

```
» B = transpose(A)
```

In practice, you can get away with being lazy (and in fact this is how it was in the old days, and just write):

```
» B = A'
```

(but get into the habit of using the formally correct, **Mathworks** official and UN-approved, syntax of `.'`).

```
B(1) = 0.0
```

(Here, you are saying that you would like to assign the value of 0.5 to the element in the vector given by the index 1. The previous content of the array at index position 1 is simply over-written.)

You can access more than a single element of a vector at a time, by means of the **colon operator**, `:` to define a min, max range of indices. For example:

```
>> B(2:4)
ans =
1.0000
1.5000
2.0000
```

To select all elements:

```
>> B(:)
ans =
0.5000
1.0000
1.5000
2.0000
2.5000
```

end

Represents the largest index in a vector when addressing it, or in MATLAB-speak: "end can ... serve as the last index in an indexing expression".

2.4 Basic graphing (aka. 'data visualization')

So far ... I suspect this is heavy-going and there is a lot to try and remember, such as command names, although knowing just that certain commands exist, is enough to start with and MATLAB Help can be used later to find out the exact name (and usage syntax). All this, and we have not even gotten on to matrices (2-D arrays) yet ... So, we'll take a diversion to look at some basic plotting techniques that will make sense now that you can create vectors of numbers to plot (and later, important some 'real' data). Unless you have forgotten how to create vectors already ... :(

2.4.1 Plotting

The command `figure` creates a figure window, which is where MATLAB displays its graphical output ... but on its own, without anything in it ... useless. So, let's put something in it, with the simplest possible graphical way of displaying data called `plot`. But first – create yourself a dummy dataset to plot. You are going to need to create yourself a pair of vectors – these can have any values (all numbers though) in them that you like, but perhaps aim for 1 vector with values counting up from 1 to 10 – this will form your *x*-axis, and the 2nd column ... whatever you like.¹⁹

As always, refer to the MATLAB help text on `plot` before using it (also refer to Box). The key information that will get you started is at the very top:

`PLOT(X,Y)` plots vector *Y* versus vector *X*.

So, you need to pass it your *x*-axis data vector (by its variable name), followed by your *y*-axis data vector (by its variable name) – comma separated. Do this, and depending on just what or how random your *y*-axis data was, you should end up with something like Figure 2.1 in a window captioned "Figure 1".²⁰

This ... is easily the least professional plot ever. And one that breaks all the most basic rules of scientific presentation, such as an absence of any labelling axes. There is also no title, although here in the course text I have added a figure caption in the document so I can sort of get away with it. But this is the default state of the basic plot function and you'll just have to deal with it (i.e. add a series of commands to add missing elements of the plot).

Note that by default, MATLAB also scales both axes to reasonably closely match the range of values. In the example here, the default min and max axes limits in fact turn out to be the min and max values in the *x* and *y*-axis data because the data is composed of

¹⁹ Looking ahead – you could create a *y*-axis vector formed of the squares of the numbers in the *x*-axis vector:

```
» Y = X.^2
```

(The `.^` bit says to square the value of each and every element in the vector.)

`plot`

The MATLAB function `plot` ... plots. More specifically, it plots pairs of (*x*,*y*) data and by default, does not plot the points explicitly but joins the (*x*,*y*) locations up by straight line segments. MATLAB calls these a '2D line plot', although there are plotting options that allow you only to display the individual (*x*,*y*) points (making it like the scatter function, which we'll see later).

Its most basic usage is:

```
plot(X,Y)
```

where *X* and *Y* are vectors – of the same length (important), but not necessarily of the same orientation (i.e. if one was a row vector and one a column vector, MATLAB would work it out, although it is perhaps best to avoid such a situation arising).

There are many options that go with this function, some of which we'll see and use later. You can also input matrixes as *X* and *Y* apparently. But I have absolutely no clue as to what might happen. I suspect that the plot will end up looking like a bad acid trip.

²⁰ If you cannot see the figure window ... check that the window is not hidden behind the main MATLAB program window!




Figure 2.1: Default output of `plot`.

relatively simply/whole numbers. If however the maximum y value was vary slightly larger, you'd see that MATLAB would adjust the maximum y -axis limit to the next convenient value so as to preserve a relatively simple series of labelled tick marks in the axis scale. In fact, why not try that – replace your maximum data value, with a value that is very slightly larger (an example is given in Figure 2.2).

²¹ Then re-plot and note how it has changed (if at all – it will depend somewhat on what data you invented in the first place).

2.4.2 Graph labelling

You have two options for editing the figure and e.g. adding axis labels. Firstly, you can use the GUI and the series of menu items and icons at the top of the Figure window to manipulate the figure. I suspect you'll prefer this ... but it is not very flexible, or rather, it requires your input each and every time you want to make changes or additions to a figure. The second possibility is to issue a series of commands at the command line. (The advantage with the latter we'll see later when we introduce `m-files`.) For now, I'll illustrate a few basic commands:

1. The first, obvious thing to do is to add axis labels. The commands are simple – `xlabel` and `ylabel`. They each take a string as an input, which is the text you would like to appear on the axis. If you change your mind, simply re-issue the command with the text you would like instead.
2. The command for title, perhaps unsurprisingly, is `title`. Again, pass the text you would like to appear as a string (in inverted commas `"`), or pass a the name of variable that contains a string (no `'` then needed).
3. You might want to specify the axis limits. The command is `axis` and it takes a vector of 4 values as its input – in order: minimum x , maximum x , minimum y , and maximum y value. e.g. `axis([0 10 -100 100])` would specify an x -axis running from 0 to 10, and a y -axis from -100 to 100.

Information as to how to use all of these commands can be found via MATLAB help. But a typical sequence, that gives rise to the improved plot shown in Figure 2.3, is given in the margin.

2.4.3 Sub-plots

You can also have more than one plot in a single Figure window. As an example, create some sine waves using the `sin` function (see help) over the range $0 < x < 2\pi$, e.g.:

²¹ If you have created a dummy dataset in which the value in the last row is the largest, replacing it is simple – remember the use of end in addressing an element in an array. If your dataset does not monotonically increase and the largest value falls somewhere in the middle ... you could cheat²¹ and open the array in the variable editor and discover which row it occurs on.




Figure 2.2: A plot illustrating axis auto-scaling (maximum x and y values now slightly larger than 10 and 100, respectively).




Figure 2.3: A (only very slightly) improved plot.

Example of adding axis labels and a plot title ...

```
» xlabel ...
('Whatever values');
» ylabel ...
('Whatever values
squared');
» title ...
('A plot of some ...
values vs. their ...
squares');
```

```

» x = 0:0.1:2*pi;
» y = sin(x);
» y2 = sin(2*x);

```

(Note how in the first line, the colon operator is used to create an x vector from 0 to 2π , in steps of 0.1. The second and third lines calculate the sine of all the x values, and sine of 2 times the x values, respectively, and assign the results to a pair of new vectors, y and $y2$.)

To place several different plots on the same figure uses the **subplot** command²². The subplot command is used as: `subplot(m,n,p)` where m is the number of rows of plots you want to have in your figure, n is the number of columns of plots in your figure, and p is the index of the plot you wish to create (see: Figure 2.4).

The basic code then goes something like:

```

» figure(1);
» subplot(2,2,1);
» plot(x,y);
» subplot(2,2,2);
» plot(x,y2);
» subplot(2,2,3);
» plot(x,-y);
» subplot(2,2,4);
» plot(x,-y2);

```

In this case, the 3rd and 4th subplots simply display the inverse of the curves in the subplots above.

2.4.4 Saving graphics and figures

You might just want to save the figure. (Why create it in the first place in fact if you are just going to throw it away ... ?) Again, you can do this via the GUI or at the command line²³. From the GUI, you have the option to save the figure in a way that can be loaded later and re-edited – this is the `.fig` format option. Or you can save (export) in a variety of common graphics formats (although once saved in this format, the graphics can only be edited later using a graphics package).

You can also close figure windows (see Box). No seriously. They are not forever. ;)

²² » help subplot




Figure 2.4: Arrangement of subplots.

²³ To export a graphic at the command line, use the **print** function. To cut a long story short (see: `help print`), to print to a postscript file:

```
print('-dpsc2', FILENAME)
```

where `FILENAME` is the filename as a string or a variable containing a string.

To close the current (active) Figure window, the command is:

```
» close
```

To close all currently open Figure windows:

```
» close all
```

2.5 Vectors and arrays #2

A matrix is another special case of an array – this time 2-D (rather than 1-D in the case of a vector). MATLAB totally hearts them.

2.5.1 Creating matrices and arrays

You can enter matrices (2-D arrays) into MATLAB in several different ways:

1. Enter an explicit list of elements. To enter the elements of a matrix, there are only a few basic conventions:
 - Separate the elements of a row with blanks or commas.
 - Use a semicolon, ; , to indicate the end of each row.
 - Surround the entire list of elements with brackets, [].
2. Load matrices from external data files.
3. Generate matrices using built-in functions.

As AN EXAMPLE, type in the following at the command prompt:

```
A = [15 7 11 6; 13 1 6 10; 21 17 5 3; 5 15 20 9]
```

MATLAB then displays the matrix you just entered²⁴:

```
A =
15 7 11 6
13 1 6 10
21 17 5 3
5 15 20 9
```

²⁴ Remember that you can add an ; to the end would prevent the assignment being displayed.

Once you have entered the matrix, it is automatically remembered in the MATLAB workspace. You can refer to it simply as A.

Now go find the array you have just created in the *Workspace* window. Double-click on its name icon and see what goodies appear on the screen. This is a fancy array editor which looks a bit like one of those dreadful spreadsheet things. You can see that this might be handy to edit, view, and keep track of at least moderate quantities of data. This is a useful facility to have. However, we are going to concentrate on the command-line operation of MATLAB in the Lab because that will give you far more power and flexibility in applying numerical techniques to problem solving, and will form the basis of scripting (computer programming by another name) that we will see in a few lectures time. Close down this nice toy to leave just the original windows.

Elements in the matrix can be addressed using the syntax:

```
A(i,j)
```

where i is the row number, and j is the column number. It is very very easy to keep forgetting in which order the rows and columns are indexed., but I'll tell you here and now before I forget:

rows, columns

(You can always create a test matrix and access a specific element to check if in doubt!) In the example above:

```
» A(1,3)
ans =
    11
```

(i.e. the value of the element in the 1st row, 3rd column, is 11).

In general, the same functions and operators that applied to vectors and you saw earlier, also apply to matrixes (or specific dimensions of matrices). (See Boxes.)

Finally – a fundamental way of accessing data that you need to learn and be familiar with, is to employ the colon operator to select specific columns (or rows) of data. You'll find that this skill ends up inherent to many of your attempts to process and graph data. For instance, if your (x,y) data to plot ended up in MATLAB workspace in matrix form (it very commonly does) rather than as 2 sperate vectors (as you had when you first plotted anything), you will need to select separately the x (e.g. 1st column) data, and the y (2nd column) data, and pass these to the `plot` function. For the example of matrix A above, all the first column data can be selected by typing `A(:,1)`²⁵, which says all the rows ($:$) in the first column. Similarly, all the 2nd column data alone can be selected by `A(:,2)`. (You'll practice this endlessly later on and hopefully get it!)

2.5.2 Basic matrix manipulation

You can treat vectors and matrices (or parts of vectors and matrices), mathematically, as you would treat single values (i.e. *scalars*) but unlike a scalar, the transformation is applied to all specified elements of the array. This applies for all the basic numerical operators²⁶. For example, for vector B in the earlier example,

```
» 2*B
ans =
    0  2  3  4  5
```

and

```
» B-1.5
ans =
-1.5000 -0.5000  0  0.5000  1.0000
```

Similarly as for vectors, you can access more than a single element of a matrix by means of the colon operator, `:`. For example:

`A(:,1)` – selects the 1st column
`A(3,:)` – selects the 3rd row
`A(2:3,2:3)` – selects the 2×2 matrix of values lying in the centre of A , while `A(1:2,:)` selects the top half (first 2 rows) of the matrix.

²⁵ Remembering the HUGE hint above in 100 pt font as to the order of rows and columns ...

You can also determine the shape of your array using the `size` function. For a 2D array (matrix), when you pass it the name of your array, it returns the number of rows followed by the number of columns (in that order).

²⁶ Technically ... or at least to be consistent with other operations, you might write multiplication as `.*` rather than just plain old `*`. The preceding dot tells MATLAB not to treat this as matrix multiplication but to carry out the operation on each element in turn. In this case, it is the same thing (and both notations work the same), but later, is not. (This will make more sense when you get to see it in action, later.)

QUESTION: Multiply all the elements of A by the number 17. Assign the answer to a 3rd array (C). What is the value of the element C(2,3)? How would you ask for the 4th row, 2nd column element of the array C, and what is its value?

QUESTION: What is the sum of the 4th column of C ? (Sure – you also do it by using a calculator, but you will not always have such a small data-set as here. Perhaps you'll get a much larger data-set in the assessed exercise ;) So, practice doing it properly.)

QUESTION: What is the sum of the 2nd row of C? `sum` gives returns the sums of each column, and so on its own;

```
» C
C =
255 119 187 102
221 17 102 170
357 289 85 51
85 255 340 153
» sum(C)
ans =
918 680 714 476
```

gives you a row vector consisting of the sums of the individual columns of the matrix C above.

This is where the `transpose` function (`'`) comes in handy (see earlier). In this case, it flips a (2D) matrix around its leading diagonal (columns become rows, and rows, columns)²⁷.

```
» C'
ans =
255 221 357 85
119 17 289 255
187 102 85 340
102 170 51 153
```

(transposing the matrix turns the rows into columns)

```
» sum(C')
ans =
663 510 782 833
```

Now you get a row vector consisting of the sums of the individual columns of the matrix C, but since you have transposed the matrix C first, these four values are actually equal to the row sums.

Finally, you could transpose the answer:

```
» sum(C')'
ans =
663
510
782
833
```

The *function* `sum` ... sums things. The MATLAB Help documentation (`help sum`) says:

'If A is a vector, `sum(A)` returns the sum of the elements.'
'If A is a matrix, `sum(A)` treats the columns of A as vectors, returning a row vector of the sums of each column.'

²⁷ This is almost true. Technically the function you want is `.'`, as `'` will change the sign of any imaginary components. For real numbers, they are the same.

In addition to `transpose`, other useful array manipulation functions include:

`flipud` – flips the matrix in the up/down direction
`fliplr` – flips the matrix in the left/right direction
`rotate` – rotates the matrix
(As always, refer to the `help` on specific functions.)

now with a row vector gives you a format that looks like the row sums of the original matrix C .²⁸

2.5.3 *Some matrix math :(*

We will not concern ourselves with multiplying vectors and matrices together ... just yet ...

²⁸ Note how you can combine multiple functions in the same statement to create `sum(C')'`. However, to start with, it is much safer to do each step separately and hence be sure what you are doing.

2.6 Loading and saving data

There are a number of different ways to load/import data into the MATLAB *Workspace*. Rather than try and tediously list and describe the commands and syntax and blah blah, we'll be going through a couple of (hopefully) slightly-less tedious data-based examples as we progress through the course text. In this way, if nothing else, you might accidentally learn some 'science' even if nothing much about MATLAB ...

2.6.1 Where am I?

Before anything – you need to know where you are. If your file to load is not in the directory MATLAB is using, it will not find it. And if you save something and have no idea where it is being saved ... that can hardly go well.

By default, MATLAB looks to a file directory located within its installation directory (\$MATLAB/data). So, where the load command requires a filename to be passed, you will need to enter either the full location of the file; i.e., starting with the drive letter (e.g. as per displayed in the **Windows Filemanger** address bar, or the relative path to where the file is located (e.g. if there is a subdirectory called data, you will pass data/sediment_core_d180.txt²⁹). Alternatively, you can change the MATLAB directory that you are working in. (This works similar to UNIX/LINUX for those of you who are familiar with navigating your way around these operating systems.) You can make the download directory the default directory for working from by typing:

```
» cd DIRECTORY_PATH
```

where DIRECTORY_PATH is the path to the directory in which the data file has been saved³⁰, remembering that DIRECTORY_PATH is a string (i.e. enclosed in "). Or ... you can add a 'search path' (addpath) so that MATLAB knows where to look. (Note that both these alternative possibilities can be implemented from the GUI.)

There is also, of course, the GUI – from the File menu the option Import Data... will run the data import Wizard – note that you might have to select All Files (*.*) from the file type option box in order to find the file. I'll leave you to work the rest out for yourselves ... Maybe try importing the data into MATLAB this way once you have done it successfully using the load function at the command line.

²⁹ Remember that this is a *string* type.

³⁰ You can view the files that are present in the directory that you are working in by typing (more LINUX-speak): `ls`.

Load

Loads variable from a file into the workspace. The syntax is:

```
» load(filename)
```

where filename is the name of the file (remember: as a string, it needs to be enclosed in quotation marks). The file might be plain text (ASCII) or a MATLAB workspace file (see below), in which case it should have the file extension .mat. To force MATLAB to treat the file input as ASCII or a MATLAB workspace file, pass a second parameter (separated from the filename by a comma) – '-ascii' for ascii, and '-mat' for a MATLAB workspace file.

Note that in loading an ASCII data file, any line starting with a % is ignored. Also note that the data must be in a column format with no missing data.

For an ASCII file, the name of the variable created to hold the data being imported is automatically generated. So in the example of the data file being called 'twilight.txt', the variable generated will be called twilight. You can instead choose to assign the imported data to a variable name of your choice, by e.g.:

```
» sparkle =  
load('twilight.txt');
```


2.6.2 Loading and importing data

The simplest way (other than via the MATLAB GUI and the beautiful green **Import Data** icon) is to use the `load` function (see Box)³¹.

As a brief exercise and practice using `load` – first download the data file `etheridge_etal_1996.txt` from the course webpage of `www.seao2.org`. You might start by viewing the contents of the file by opening it in any text viewer. This is always a good place to start as it enables you to see what you are getting yourself in to (i.e. format of the file, any potential formatting issues, approximate size and complexity of the dataset, etc). Then import the data into the MATLAB workspace using the `load` command. Try simply typing the name of the variable that was automatically created (or the one you chose, if you assigned the imported data to a specific variable name – see Box) to provide a crude view of the data. Then double click on the name of the variable in the MATLAB Workspace window. This should open up a spreadsheet-like window in which the data can be viewed, sorted, and even edited. Finally, plot the data and remember to label it appropriately³². You should end up with something like Figure 2.5.

³¹ There is also a much more flexible way of loading text-based data using the function `textscan`, but that also requires files to be explicitly opened and closed using `fprintf`. We'll see a little of this later.)

³² FYI: the x-axis data is year, and the y-axis data is the mixing ratio of CO₂ in air in units of ppm.

2.6.3 Saving and exporting data

Arrays of numbers can be saved in a plain text (ASCII format) by means of the `save` function in a simple reverse of the use of `load` (see Box).

2.6.4 Loading and saving the workspace

The entire workspace (including all variables and their values, or just the values in a single variable if you wish) can be saved to a file and then later re-opened. The file format is specific to the MATLAB program and the file-name extension by default is `.mat`. You might find this very helpful to use in long lab exercise or large modelling projects, particularly if you do not come back to work at the exact same computer each time or wish to use continue the same piece of work on a laptop elsewhere.

save

Saves variables from the workspace to a file. There are two main forms (syntaxes) of the command:

```
» save(filename)
```

which saves the entire workspace to a `.mat` file (with the filename given by the string `filename` (in quotation marks), and:

```
» ...
   save(filename,A,'-ascii')
```

saves the data in the variable `A` (which must be given as a string, i.e. also enclosed in quotation marks) in plain text (ASCII) format.




Figure 2.5: Spline fit to measured changes in CO₂ concentration in Law Dome ice core, following *Etheridge et al.* [1996].

2.7 Basic data processing

As an example to kick-off some data-processing tricks, load in the Phanerozoic CO₂ dataset: `paleo_CO2_data.txt`. You can just import it into MATLAB using the `load` function. However, there is a complication here – unlike the ice core CO₂ dataset, you now have 4 columns in the array³³. The first column is age (Ma), the second the mean CO₂ value, and the 3d and 4th are the low and high, respectively, uncertainty limits. Remembering (I hope!) how to reference specific columns of data in a matrix³⁴ – plot the mean paleo CO₂ value as a function of age (in Ma). If you closed the previous Figure window (see earlier), it is not essential to explicitly open one (using the `Figure` command) – when you use the `plot` command, if there is no open Figure window, **MATLAB** will kindly open one for you. How thoughtful. The result should be something like 2.6. O dear ...

So ... that was not so successful. What is happening in the default behaviour of `plot`, is that the location defined by each subsequent row of data is being joined to the previous one with a line. This was fine for the ice-core CO₂ example dataset because time progressed monotonically in the first column, e.g. the data was ordered as a function of time. If you view the paleo CO₂ data, this is not the case. (In fact, in the original, full version of the data, ordering is by proxy type first, and then study citation, and only then age ...).

Your options are then:

1. You could import the data into **Excel**, then re-order (sort) it, then export it, then re-load it ...
2. You could sort it in **MATLAB** using the GUI variable view window. But lets not cheat for now.
3. You could sort it in **MATLAB** at the command line. How? Well, a reasonable gamble, which actually turns out to be a total win, is to try:

```
» help sort
```

Actually ... not quite. Reading the help text carefully (and you can always try it out and see what exactly it does if you are not sure), `sort` will sort all columns independently of each other, whereas we want the first column sorted and the remaining columns linked to this order. Under see also, MATLAB lists `sortrows` as a possibility. The help text on this looks a little more promising. It is still slightly opaque, so the best thing to do is to try it (and view the results)! This looks rather better. The resulting of plotting this is 2.7. (This is a good illustration of a guess of a function that was

³³ Remember that you can diagnose its size with `size` (or refer to the Workspace window)

³⁴ HINT: the colon operator (see earlier).




Figure 2.6: proxy reconstructed past variability in atmospheric CO₂.

not quite what was needed, but following up on the help suggestions leads to a more appropriate function.) At least now the curve is reminiscent of past changes in global temperature and the geological Wilson cycle, with high values in the Cretaceous and Jurassic and then lower again in the Carboniferous (roughly matching the progression of ice and hot house (and then back to recent ice ages) climates).

As a second example and to get you familiar with some additional very basic data processing, we are going to transform a sediment core $\delta^{18}\text{O}$ time-series into an estimated history of glacial-interglacial changes in sea-level. The scientific backstory is ...

Throughout the late Neogene³⁵, sea level has risen and fallen as continental ice sheets have waned and waxed. The main cause of sea-level change has been variation in the total volume of continental ice and resulting change in the fraction of the Earth surface H_2O contained in the ocean. Today more than 97% of the Earth surface H_2O is in the ocean and less than 2% is stored as ice in continental glaciers, with groundwater making up the bulk of the remainder. Of the total continental ice (ice above sea level), 80% is contained in the east Antarctic ice sheet, 10% in the west Antarctic ice sheet, and the final 10% in the Greenland ice sheet. (If all present-day continental ice were to melt, sea level would rise by 70 m.) During the last glacial maximum (LGM), sea level was about 125 m lower than present, equivalent to 3% more surface H_2O stored as continental ice. Because of its relationship to continental ice volume, an accurate late Neogene sea-level curve has been a long-term goal of scientists interested in ice-age cycles and their causes.

Glacial ice has a lower $^{18}\text{O}/^{16}\text{O}$ isotopic ratio than mean seawater³⁶. When ice volume is high, seawater has relatively high $^{18}\text{O}/^{16}\text{O}$ ratio. When ice volume is low, seawater has relatively low $^{18}\text{O}/^{16}\text{O}$ ratio. If the average $^{18}\text{O}/^{16}\text{O}$ ratio of glacial ice is constant with time, then changes in the average $^{18}\text{O}/^{16}\text{O}$ ratio of seawater will linearly approximate changes in the total volume of ice and by inference, sea-level. We (at least, I am) are interested in all this because knowing how ice volume and sea-level changed over the glacial-interglacial cycles has all sorts of important implications for understanding how climate change (e.g. via ice sheet albedo) and global carbon cycling and atmospheric CO_2 (e.g. via changes in the area of exposed continental shelves and carbon stored in soils and above-ground vegetation).

To start with we need to reconstruct past changes in the oxygen isotopic composition of the ocean. Handily, the $^{18}\text{O}/^{16}\text{O}$ ratio of foraminiferal calcite isolated from marine sediments is primarily a function of the $^{18}\text{O}/^{16}\text{O}$ ratio of the water together with the tempera-




Figure 2.7: Proxy reconstructed past variability in atmospheric CO_2 (sorted data).

³⁵ 23.03 millions years ago (end of the Oligocene) to present is the Neogene Period in Earth history.

³⁶ Basically – as moisture derived from the tropical ocean (and land) surface moves to high latitudes, condensation occurs and some of the moisture is lost as rain. In condensating water vapor, ^{18}O is preferentially incorporated into the liquid phase, meaning that the remaining water vapour has lower $^{18}\text{O}/^{16}\text{O}$. Eventually, the residual water vapour might fall as snow on an ice sheet. Hence why ice sheets at the LGM will have a lower $^{18}\text{O}/^{16}\text{O}$ than mean seawater.

ture of the water³⁷. By measuring the $^{18}\text{O}/^{16}\text{O}$ value of calcite down-core we are sampling $^{18}\text{O}/^{16}\text{O}$ with a progressively older age. In this way we can reconstruct how ocean $^{18}\text{O}/^{16}\text{O}$ has changed over time. These measurements are reported in units of parts per thousand (‰) and written as $\delta^{18}\text{O}$.

How to turn (scale) changes in $\delta^{18}\text{O}$ into sea-level change? Evidence from dated coral reef terraces suggest that sea-level was around 117 m lower at the peak of the last glacial (ca. 19 ka). We could then assume that the change in $\delta^{18}\text{O}$ from modern (preindustrial) to LGM equates to 117 m sea-level change, and hence create a continuous past sea-level curve from all the $\delta^{18}\text{O}$ data by applying a simple scaling factor³⁸. So:

- You first need the foraminiferal calcite $\delta^{18}\text{O}$ data. (Unless you want to go drill a long cylinder of mud from 3000 m down in the Atlantic Ocean, pick out all the microscopic foraminifera of a single species from samples of mud that you have carefully washed, blah blah blah ...) So, from the course web page; download the file `sediment_core_d180.txt` and save it locally.
- Load this file into MATLAB.
- If you have successfully loaded in the data-file, you should see a named icon for the array appear in the Workspace window. Try viewing the file in the two different ways:

1. At the Command line (`>>`), type in the array name. Because of the length of the data-file we imported, the contents of the array should have whizzed past you on the screen in a highly inconvenient fashion. You can use the scroll bar on the right of the Command Window window to move up and view the data that you can't see (the younger age $\delta^{18}\text{O}$ numbers). Note that as MATLAB imports data into an array from a file, it names the array it creates following that of the filename, but without the extension (the `'.txt'` bit).
2. Double-click on the array's icon in the Workspace window. Marvel at the fancy spreadsheet-like things that appear. Note that you can edit the data, add and delete rows and columns, and all sorts of stuff in this window, just like you can in Excel. Amuse yourself by scrolling down to the end of the data-set in the Array Editor and adding a new piece of data on line 784; age (column 1); 783 (ka); sea-level (column 2); 0.0 (m).

At the command line, list the contents of the array again to view the change you have at the end of the data-set. Use the **up arrow** to bring up the command you want rather than typing it in again.

³⁷ We we will not concern ourselves with temperature corrections here (in any case, it turns out that the temperature effect has the same sign as and is closely related to the ice volume effect) but instead assume that foraminiferal calcite $\delta^{18}\text{O}$ only reflect changes in (global) ice volume and sea-level.

³⁸ Conceptually, this is no different from saying that the difference between the freezing and boiling point of pure water (at 1 atm pressure) on the Celsius scale – 100°C, maps onto the equivalent interval on the Fahrenheit scale – 180°F (212–32 °F), and hence providing a means of converting a record of past changes in Fahrenheit, into degrees Celsius (and *vice versa*).

Now delete this new row. Note that it is easy to get confused with which row number you need to address – although the data starts from year 0, MATLAB always counts the index (the sequential integer counting of the row or column number) of a location from 1 (one). (So age 10 ka is on line 11, and age 200 on line 201, etc.)

- So far everything has been in $\delta^{18}\text{O}$ units and time as kyr. As a warm-up – try converting the units of time to years by multiplying the first column of the data array by 1000.0 and assigning it back into itself (this is not as weird and nonsensical as it sounds).

To estimate past changes in sea-level we need to scale the $\delta^{18}\text{O}$ values to reflect the equivalent changes in sea-level rather than changes in isotopic composition. We know that sea-level is 0 m (relative to modern) at 0 years ago and -117 m at 19,000 years ago. Try the following (you are going to have to **think**, but maybe also use the HINT in the margin):

Scale the $\delta^{18}\text{O}$ so that it represents changes in sealevel, relative to modern (0 m)³⁹.

- Plot it (changes in sea-level compared to modern, vs. time). And **nicely**.

Reminder: for a $n \times m$ array data, the first row is:

`data(1, :).`

The last row is:

`data(end, :).`

To find out the number of rows is:

» `length(data).`

The total size, in rows \times columns, can be found by:

» `size(data)`

(and also by referring to the **Value** column in the **Workspace** window)

³⁹ HINT – first determine the difference in $\delta^{18}\text{O}$ between time zero and 19 ka. This gives you the range of $\delta^{18}\text{O}$ that maps onto a sea-level change of 117 m. You also might transform the $\delta^{18}\text{O}$ data such that it has a value of zero at 0 ka (but retains the original amplitude of variability).

2.8 Yet more graphing

This section covers how to create slightly fancier plots in MATLAB and combines this with some more data loading practice.

2.8.1 Modifying lines/symbols in plot

The first deviant activity you can engage in with `plot`, it to graph the data without the line joining the points. Scrolling a little the way down » `help plot`, it turns out that there are a number of options for color, line style, and marker symbol that you list together as a single parameter, straight after the parameters for x and y vectors. By default, MATLAB plots a solid line in blue with no marker points. Obviously, we could forego the sorting and plot a sane graphic (hopefully) by plotting just points and having no line between them. Hell, you could even be radical and use a different color ... Or, you could specify a symbol and no line. The choice of colors is your oyster, as they (almost don't) say. e.g. Figure 2.8.

2.8.2 Plotting multiple data-sets

So far, so good. But so boring, although simple marker-only and joined-by-line plots have their place. For a start, the original data-set included an estimate of the uncertainty in the CO₂ reconstructions in the form of the min and max plausible value for each 'central' (best guess?) estimate. **Excel** can make plots incorporating errors, including non-symmetric errors, relatively easily. What about in **MATLAB**? Actually, I have absolutely no idea. (This would make such a good exercise for the reader, as they (do) say.)

Personally, I might have been tempted to draw vertical bars alongside the data (most likely). Or plotted in different symbols, the min and max values as points. Or plotted min and max lines as a bounding envelope. All of these require some further little trick in MATLAB, which involves the command `hold`. This is nice and simple and can be on, or off.

» `hold on` – will enable you to add additional elements to a graphic,

» `hold off` – returns to the default in which a new graphic replaces the current on in a Figure window.

AS AN EXAMPLE – set » `hold on`, and then plot the minimum and maximum CO₂ values (columns #3 and #4) in different symbols and different colors, on top of your existing plot. If you want to then label what different lines or sets of points are, you can add a legend with the `legend` command. For instance you have managed to successfully

The main (i.e. not an exhaustive list) data display options for the `plot` function are:

- (1) point style
`.` – point, `o` – circle, `x` – x-mark,
`+` – plus, `*` – star, `s` – square, `d` –
diamond, `v` – triangle (down).
- (2) line style
`-` – solid, `:` – dotted, `--` – dashed, and
when specifying a point style, not
specifying a line style results in no
line.
- (3) color
`b` – blue, `g` – green, `r` – red, `y` –
yellow, `k` – black, `w` – white.




Figure 2.8: Proxy reconstructed past variability in atmospheric CO₂ (sorted data).

plot the mean CO₂ values as discrete black circles, and the minimum and maximum uncertainty limits as blue and red lines, respectively, you could call:

```
» legend('Mean CO2', 'Lower uncertainty limit', 'Upper uncertainty limit');
```

and it should end up looking like Figure 2.9.

2.8.3 Scatter plots

We'll stay with the Phanerozoic proxy (CO₂) data, but put a different (graphical) spin on it.

Consider ... **scatter**. In fact, don't just consider it, help on it. The simplest possible usage is, apparently:

```
SCATTER(X,Y) draws the markers in the default size and color.
```

(where X and Y are vectors). This almost could not be more straightforward. Make yourself an X and Y vector out of the loaded-in dataset (or if you are feeling brave, you can pass in directly the appropriate parts of the dataset array), close the existing Figure window⁴⁰, and scatter-plot the (mean) CO₂ data.

Perhaps a little disappointingly, the default (Figure 2.10) (plus added labels) looks a little like one of the plots before. However, scatter can plot color-filled symbols, but more powerfully, can scale the fill color to a 3rd data value (vector), in a sort of pseudo 3D *x-y-z* plot. For instance, it will be duplicating information that is already presented (*y*-axis), but you could color-code the points, by the *y*-value, i.e. the atmospheric CO₂ value. e.g.

```
SCATTER(data(:,1),data(:,2),20,data(:,2))
```

draws the markers with an (area) size of 20 (points), in different colors. Coloring just the outlines of the circles is perhaps not ideal (difficult to see all of the color differences), so the circles can be filled in instead (and you could make them a little larger too):

```
SCATTER(data(:,1),data(:,2),40,data(:,2),'filled')
```

resulting in Figure 2.11.

ONE FINAL EXAMPLE in this section to introduce some new plotting functions, but also to quickly go back over some basic array manipulation and processing. The data we will be analysing is a series of seismic readings from the USGS. The quake data are extracted between -5 and 20 lat, and between 90 and 105 lon, starting Dec 26, 2004 and ending June 30, 2005. The data file can be found on the




Figure 2.9: Proxy reconstructed past variability in atmospheric CO₂ (sorted data).

⁴⁰ See earlier.




Figure 2.10: Proxy reconstructed past variability in atmospheric CO₂ (scatter plot).




Figure 2.11: Proxy reconstructed past variability in atmospheric CO₂ (scatter plot).

magnitude estimates - the area of dense aftershocks often delineates the part of the fault that ruptured, and scaling laws relate rupture length to magnitude.

Create a figure with multiple panels, showing:

- In the top LH corner plot the day 0-91 quakes, and color-code (or size-code) the markers for their magnitude.
- In the top RH plot the day 92 onwards quakes, and color-code (or size-code) the markers for their magnitude.
- In the bottom LH corner plot day 0-91 quakes, and color-code (or size-code) the markers for their depth.
- In the bottom RH plot the day 92 onwards quakes, and color-code (or size-code) the markers for their depth.

2.8.4 Histograms

We could also visually analyse the data as a histogram. Type `help hist` in the Command Window for a description of the `hist` function. The histogram must be supplied with a vector defining the 'bins' in which to sum the data. Here is your chance to use the colon operator again. O happy day.

1. To plot the frequency distribution of quakes as a function of their magnitude we need to create a series of bins to define the different magnitude ranges. How about bins with boundaries at magnitude; 1.0, 2.0, 3.0, ... 10.0. One complication is that the values in the vector `M` define the middle of the bins in the `hist` function and not the boundaries. The mid-points of this will be; 1.5, 2.5, 3.5, ... 9.5, and this is the vector you need to create and assign to a vector `M` (i.e., a vector array starting at 1.5, ending at 9.5, and with increments of 1.0).
2. Having created `M`, plot the histogram of quake frequency vs. quake magnitude by issuing:

```
» hist(data_USGS(:,5),M);
```

Question: what is the most frequent magnitude range of 'quake?

3. Now plot the histogram of quake frequency against time (i.e., day number) up to day number 186. You will have to assign a new vector of values to `M`, one that starts at 0.5 and ends at 185.5. Omori's Law says that the number of aftershocks per day should decrease following a power law – does this look to be the case (approximately)? (One problem is that the small earthquakes are missing which makes it appear not to work so well!)

4. Try this again (i.e., frequency of quakes vs. time), but investigate the effect of changing the bin size – try making the bins about 1 month (30 days) in duration. Note that now M must start at 15.0 (the mid-point of the first monthly bin). Sometimes changing the bin size can help if the data is noisy, but sometimes you lose important information. Which was better do you think – can you still see a power-law decay in quake frequency following each major event with the data in monthly bins? If you want, experiment with other bin sizes to see how the data comes out. There is not always a ‘right’ answer in plotting data and sometimes you just have to experiment a little to see what looks good.

Don’t forget that all the plots you make should be appropriately labelled ... Save them as a `fig` file if you think you might want to edit them again, and/or export as an image.

2.8.5 Simple 2D data and bitmap visualization

There are 2 different simple MATLAB commands for visualizing a 2D dataset (i.e. a matrix) as a bitmap image (and via a 3rd command, viewing various bitmap photo and image format files too). As something (2D data) to play with – load in the matrix `model_grid.txt`.

First off – as before, view the data in the array viewer, just to get a feel for what you are dealing with here (although you are unlikely to be much wiser after doing so). So go ahead and employ the `pcolor` function in its simplest possible usage (see Box). You can see (Figure 2.12) that it is ... something. Maybe a little like the continents, but up-side-down at the very least. What to do?

Well, it is a good job that you remember how to re-orientate arrays, right?⁴² If you guess right first time (three different basic transformations of a matrix were described), you get Figure 2.13.

Next try something very similar. but using the `image` function. Now the model grid is the correct way around! I have absolutely no idea why and why it is reading the matrix dimensions differently from `pcolor`. I am sure you could Google and find out. But you would have to actually care first.

What is the point of this? So you have the ability to simply visualise a gridded dataset. Later, we’ll be doing it properly and it gets rather more involved when you have to create matrixes to describe the grid dimensions (e.g. lon and lat) for yourself.

As your very last exercise – find an image on the internet that amuses you, download it, load it into MATLAB (using `imread`), visualize it using `image`, and then ... well, that depends on how amusing it is. Maybe try plotting something on top of it (using `hold on`) or simply go home.




Figure 2.12: A 2D plot of some random gridded model data.




Figure 2.13: A 2D plot of some random gridded model data ... but with the underlying data matrix re-orientated before plotting.

⁴² You don’t? See earlier in the Chapter ...

`pcolor`

MATLAB claims that `pcolor(C)` plots; "a rectangular array of cells with colors determined by `C`. Actually, I believe MATLAB on this. So if you have a matrix, MATLAB will plot a regular arrays of cells, with each cell representing one of the elements in the matrix, and will color that cell according to the value. (`pcolor` will by default, autoscale how the color scale maps onto the data in the matrix such that both extreme ends of the color scale are used.)

`image`

You can import an image, such as in `.jpg`, `.tiff`, or `.png` format, using `imread` – simply pass it the name of an image file (as a string, this variable name needs to be encased in inverted commas) and assign the results to a variable name of your choice. Then plot (using `image`) that variable.

3

Elements of ... programming

NERD. This is what you are now going to become. And lose all your social skills. And sit at home all day in front of your computer. Which has become your only friend.

You will achieve this higher state of Being by starting to learn to write and use *scripts* and *functions* (aka `m-files`) in **MATLAB**. Actually, at this point you are now writing computer programs (of a sort) rather than endlessly typing stuff at the command line in the forlorn hope that something useful might occur. You will also be doing a great deal of code debugging ...

3.1 Introduction to scripting (programming!) in MATLAB

Commands in MATLAB can become very lengthy, and you typically end up with multiple lines of code to get anything even remotely useful done. And as you have noticed, it can take a lot of time to enter in all these lines. When when you log off and go home ... it is all gone. ¹ ... If only there was some way of storing all these commands in such a way that they could be worked on and run again with the press of a button (as a wild guess, how about **F5**?), without having to enter them all in, all over again from scratch ...

Your wish is granted. In **MATLAB**, it is possible to store all of your commands in a single text file, and then request that they are all executed (sequentially) at one go. **MATLAB** gives this text file a fancy name (because it is a very fancy piece of software, after all) – a *script*², otherwise known as an **m-file**. To create a new **m-file**; from the File menu, select **Script** (a common type of **m-file**)³. You will see a text editor (more fancy-ness) appear in front of your very eyes, containing your requested (but currently empty) **m-file**. Save the **m-file** to your directory of choice. Alternatively, simply create a new (blank) text file and saving it with the extension **.m**, rather than e.g. **.txt**, creates you a (script) **m-file**. From an **m-file**, you can issue all the **MATLAB** commands you previously would have entered individually, line-by-tedious-line, at the command line. Furthermore, having created and saved a **MATLAB** script, it can be executed again and as many times as you like.

You can execute an **m-file** by typing its name into the **Command window** (omitting the **.m** file extension). Ensure that **MATLAB** is operating in the same directory as the directory that you have saved your **m-file**. You can also run the *script* (**m-file**) by hitting the big bright green Run icon button at the top of the **m-file** editor⁴. The short-cut for running it is to whack your paw down on the Function Key **F5**.

OK – you are now ready for your very first program ... inevitably ... this has to be to print 'Hello World' to the screen. No, really. (Google it.) Create a new **m-file**, calling it e.g. `hello_world.m`. You need to use the function `disp` (see Box or type » `help disp`) as always, for function syntax and usage), which will print to the screen, either any text you specify (in inverted commas), or the value of a variable (which could also contain character information). For now, simply pass the text directly. Your program needs just a single line in the **m-file**:

```
disp('hello, world')
```

Save the file (to your working directory). Run it at the command line by typing its name (omitting the **.m** extension). Your first program

¹ **MATLAB** remembers all the commands used in previous session (although this may not be the case of shared, lab computers) and lists them in the **Command History window**. You can recover and re-execute a previous command in this list by double-clicking it. You can also re-run more than one line at a time by selecting multiple lines and pressing **F9** (or Evaluate Selection from the (R-mouse button in Windows) context menu).

m-file

... is nothing more than a simple text file, in which a series of one or more **MATLAB** commands are written and which via the **.m** extension, **MATLAB** interprets as a program file (*script* or *function*) that can be edited and executed (or rather, the list of commands inside, can be executed in sequential order).

Assume a similar convention to that for *variables* in the naming of **m-files**.

² The conception of a *function*, will be introduced later.

³ In order version of **MATLAB**: **File/New** menu, and select: **Blank M-file**.

⁴ In older versions of **MATLAB** – select: **Debug/Run** from the 'debug' menu of the **Editor window**.

is a success! (Surely you could not screw up a single line program ... ?⁵) You could extend this to a mighty 2-line program by defining the string as a variable and displaying the contents of the variable, i.e.,

```
message = 'hello, world';
disp(message)
```

For further practice – pick one of any of the previous exercises in which multiple lines of code were required, place them into a new **m-file** (either by re-typing them in or copying them out of the **Command History window**), save the file (to the same directory that you are working from), and run it by typing its name at the command line (omitting the **.m** extension).

3.1.1 Programming good practice

A few tips about good practice in (**MATLAB**) programming before we go on (and on and on and on):

- Choose helpful variable names so that it is clear what each variable represents. Avoid **excessively** short names, except for simple index and counting variables. At the other extreme – excessively long names, which might be wonderfully descriptive, can lead to even simple calculation stretching over multiple lines of code (which can make it more difficult to see what is going on in the code overall).
 - Use comments within your m-file to add explanation and commentary on your program. Anything after a % on the same line is considered a comment⁶, and is ignored by **MATLAB**.
 - Structure the code nicely. You can break the code up into sections, e.g. by adding a blank line. You might also start each section with a label summarizing that it is going to do (via the addition of a *comment*).
 - To start with – program in as a simple step-by-step way as possible. Breaking a complex calculation into several lines of simpler calculations is much easier to debug and work out what you were doing later, particularly if comments are also added. For all practical purposes – at this level, everything will run just as fast whether as a complex calculation on one line, or simple bite-sized calculation spread over 4 lines with comment in between.
 - Always save your changes before running your program (or you may unknowingly be running the previous version).
 - If using the script to do some plotting, sometimes (but not always) it is convenient to add at the top of the m-file,
- ```
close all;
```

This command close all currently open figures, plots, images, etc.

#### disp

... displays something (the contents of a variable) to the screen. Actually, its effect is basically identical to leaving off the semi-colon (;) from the end of a line. In the example of:

```
disp(X)
```

where the contents of X is a string, you get the text displayed.

Note that the difference between using `disp` and simply typing the variable name:

```
disp(X)
```

is ... well, find out for yourself!

#### Creating help text in an m-file

**MATLAB** allows you to create a 'help' section in the **m-file** – text that is outputted to the screen if you type `help` on that particular *script* (or *function*). The text is defined by a block of comment lines at the very top of the script file (or after the function definition in the case of a function). The last sequential comment line is taken to be the end of the help section. Note that the help section can be a minimum of one single line. A typical basic format is:

1. Name of (in capitals), and very brief summary, of the script (/function).
2. List and description of the different forms of use (if there are one or more optional parameters) including definition of the input parameters.
3. Examples.
4. A See also section listing similar or related scripts or functions.

<sup>6</sup> Your % comment can start on a new line, or follow on from the end of a line of code, whichever is more helpful.

An illustration (and a far from perfect illustration) of a short script exhibiting at least a few examples of good practice, is:

```
function [dum_temp] = ch4_ebm_basic(dum_S0)
% 0D case of EBM - analytical solution
% function takes one parameter - the solar constant (units of
% W m-2) [NB. modern value: 1370.0]
% define constants
const_0C = 273.15; % (units: K)
const_sigma = 5.67E-8; % Stefan-Boltzmann constant (units: W
% m-2 K-1)
% define model parameters
par_emiss = 0.62; % (non-dimensional)
par_albedo = 0.3; % mean albedo
% solve for surface temperature
% equilibrium equation:
% (1.0-par_albedo)*(par_S0/4.0) = par_emiss*const_sigma*loc_temp^4.0
% then re-arranged to:
loc_temp = ...
((1.0-par_albedo)*(dum_S0/4.0)/par_emiss/const_sigma)^0.25;
% convert temperature units (Kelvin to Celsius) and set value
% of return variable
dum_temp = loc_temp - const_0C;
end
```

which also illustrates one possibility for variable naming convention ('constants' (variables which never change in value) start with a const\_ and parameters (variables whose values might be changed) with par\_, temporary ('local') variables with loc\_ and variables passed into and out of the function: dum\_). Note use of the semi-colon at the end of every line to prevent (here unwanted) printing of results to the screen. In the file, you can create as much 'ASCII art' as you like if it helps to make the code clearer, e.g. adding separator comment lines ...

```
% -----
```

... or highlighting certain section headers, e.g.

```
% *** PLOTTING SECTION ***
```

If it (a line) starts with a percentage symbol, then **MATLAB** ignores it and you can type whatever you like after it (on the same line).

Your Hello World program might look like the following once it has had a little tune-up (although in this example this is pretty much over-kill):

```
% program to print 'Hello World' to the screen
% *** START ***
```

```
% first - define the text to display and assign it to the
variable message
message = 'hello, world';
% second - display the contents of variable message
disp(message)
% *** END ***
```

Finally, and related to the next subsection – code in stages, testing the (partial) code at each step. Do not try and write all the code in one go and only try it out at the end<sup>7</sup>.

<sup>7</sup> Because it will not work 99 times out of 100 ...

### 3.1.2 Debugging the bugs in buggy code

What programming is mostly about is not writing new code so much as debugging<sup>8</sup> what you have already written. Key then is to reduce the incidence of bugs occurring in the first place, and when they do occur, firstly to have code that lends itself to debugging and secondly, knowing how to go about the debugging. The first two facets are at least partly addressed through good programming practice (see earlier)<sup>9</sup>.

<sup>8</sup> The art of fault-finding in computer code.

Here's an example to try out to start to see what might be involved in debugging, loosely based on a previous plotting example – go create a new **m-file** called: **plot\_some\_dull\_stuff.m**<sup>10</sup>. Then add the following lines to the file:

<sup>9</sup> And by the discipline of software engineering, which is way out of scope of this course.

```
% my dull plotting program
% first, initialize variables and close existing figure
windows
close all;
x = -2*pi:0.1:2*pi;
y1 = sin(x);
y2 = cos[x];
% open a figure window and plot a sine graph
figure;
plot(x,y1,'r');
% add a cosine graph
hold on;
plot(x,y2,k);
```

<sup>10</sup> Remember – you are advised to name your m-files as something vaguely descriptive of what the script actually does (and you do not have to go with this choice, although it might turn out to be perfectly descriptive ;) (i.e. you do not have to call it this!)

and then run it (refer to above for how).

Pretty dull stuff eh? Wait – maybe you didn't get a figure appearing on the screen with a pair of sines and cosines on. Has **MATLAB** given you an error? If you typed in the above 'correctly', you should see:

```
Error: File: plot_some_dull_stuff.m Line: 6 Column: 9
Unbalanced or unexpected parenthesis or bracket.
```

Actually ... if this were your program, you should have paid attention to earlier and not have written it all at once before testing it! But

at least **MATLAB** is giving you some sort of feedback. The actual error reported might not always mean that much to you but the line number at which the problem occurred is gold-dust. The line of code it does not like is line 6<sup>11</sup>, which is:

```
y2 = cos[x];
```

Maybe the mistake is already obvious? If it is – go fix it and re-run the program. If not, maybe test out the line more simply, passing in a value directly to the function `cos` and not bother assigning the result to a different variable, e.g.

```
» cos[0.0]
```

to which you get told:

```
» cos[0.0]
cos[0.0]
↑
Error: Unbalanced or unexpected parenthesis or bracket.
```

Now you have reduced the use of the `cos` command to its simplest, whilst retaining the usage in your program that seemed to cause an issue. Hopefully, now the error is apparent. If still not, check out help on the `cos` function, or search `cos` in the **MATLAB** help (from the question mark icon in the toolbar).

*Is it important to recognise that (1) bugs will not always be flagged by MATLAB with a line number, and you can have valid code but nonsensical results, and (2) the mistake is often made earlier in the code than when MATLAB flags up a problem line.*

Other strategies for helping debug include:

1. Checking the what the values of the variables were at the point at which the program derped – the current (and the point of program crash) variable values are listed in the **Workspace window**.
2. Changing the relevant variable value(s) (here `x`) and re-typing the problem line to see if it makes a difference<sup>12</sup>.
3. Commenting out (%) lines of code temporarily, or adding in additional (temporary) lines of code, and re-running. Where coding in bite-sized chunks is an advantage in this respect, is that if a program stops working after you have added a new section of code, you can go comment out the new code (never normally just delete it all), check that the original section of code still works, and then line-by-line, un-comment the new code until the problem line is found.
4. You can also put your program on hold just before the problem line and explore the state of the variables at that point (see Box),

<sup>11</sup> Note that although **MATLAB** ignores comment lines (in the context of executing code), it does count them when telling you which line of the program code an error occurs at.

<sup>12</sup> This is sort of similar to the example given of simply testing a specific value directly.



although in this particular example of a bug, MATLAB does not allow this, presumably because it feels that the mistake is simple and can be easily fixed.

Once you have fixed this, re-run the program. Ha ha – it still does not work. (It is far from unusual to have multiple mistakes in the same piece of code, hence why writing the code in chunks and testing each time is helpful.) Now we have a problem on line 12:

```
Undefined function or variable 'k'.
```

```
Error in tmp2 (line 12)
plot(x,y2,k);?
```

Now **MATLAB** does not like function or variable 'k' because it cannot find that it has ever been defined. Is k meant to be a function or variable? Look up `help plot` to remind yourself of the correct syntax if the problem is not immediately obvious.

Once you have fixed the second bug; saved, and re-run the script, you should see Figure 3.1.

#### Debugging – breakpoints

Breakpoints are indicators in the code that tell MATLAB to pause at that point. This allows for in-depth testing of variable values and lines of code without having to exit the program.

To add a breakpoint in the code – click in the (grey) margin of the code editor on the problem line or before, and MATLAB adds a red circle to indicate a 'breakpoint' has been set. The presence of a breakpoint tells MATLAB to pause at that line.

To unset a breakpoint, click on the red circle or you can clear one or more from the drop-down **Breakpoints** menu in the toolbar.




Figure 3.1: Output from the (bug-fixed version of) `plot_some_dull_stuff` m-file.

### 3.2 Functions

*Functions* in **MATLAB**, are really just fancy scripts. Again – just plain old lines of code in a text file that is given a **.m** extension (making it an **m-file**). The big difference from a *script* in MATLAB is that a *function* can take variables as input and/or return an output (in contrast, a script takes no input and returns no outputs, other than plots or data files that might be saved).

A *function* is defined (and differentiated from a *script*) by a special line at the very start<sup>13</sup> of the m-file (see Box).

This is all not as weird as you might think. For example, you have already used the function `sin` – this takes a single input (angle in radians), and returns a single output (the sine of the angle). If you were to write your own function for `sin`, the file would start something like:

```
function [Y] = sin(X)
```

You can't, of course, go re-defining pre-defined MATLAB function names<sup>14</sup>. So how about if in your work, you found you frequently needed to use the square of the sine of a number. You could keep writing:

```
Y = (sin(X))^2
```

or, if you were a little more devious, you could create your own function for returning the square of the sine of a number. Your **m-file**, which here we'll call `sin2`, the contents of which would look like:

```
function [Y] = sin2(X)
Y = (sin(X))^2;
end
```

but of course with LOTS of comments to remind you what the function does etc. The new *function* is used pretty much as you would expect and have used previously, e.g.

```
» sin2(0.5)
```

will return the square of the sine of a value of 0.5 and dump the answer to the screen, and

```
» Y = sin2(0.5);
```

does the same but assigns the answer to the variable `Y` (and the semi-colon suppresses output to the screen).

Go create your own function now. Start by creating one that takes a single input and returns a value equal to the sine of the square of the value (rather than the square of the sine as above). When you are happy with this, create one with 2 inputs (see Box), that returns a

<sup>13</sup> Literally: line 1. Not even a comment line is allowed to appear before the *function* definition line.

#### Functions

The all-important fancy first line of a *function*, as defined in MATLAB help, looks like:

```
function [y1,...,yN] =
myfun(x1,...,xM)
```

Thanks MATLAB (this seems overly complex to say the least!)

OK – lets break this down. Lets assume that you call the **m-file** `calc_stuff`. The minimal definition of a function then looks like:

```
function [] = calc_stuff()
```

(The syntax is critical and the definition line must look like this.) Here we are saying – pass in not parameters and return no values either. So exactly like a normal script would work and you would execute the function `calc_stuff` by typing at the command line:

```
» calc_stuff()
```

(Maybe you can get away without the `()` bit.)

If you want to pass in a single parameter (here: `X`), then you define the function:

```
function [] =
calc_stuff(X)
```

(To pass in more than 1 variable, simply comma separated the variable names.)

To pass out a parameter (here: `Y`) (and no input):

```
function [Y] =
calc_stuff()
```

Lastly, at the end of the function, you include the line:

```
end
```

<sup>14</sup> Actually you can, but it is best not to.

value equal to the sine of the first input, divided by the cosine of the second input<sup>15</sup> (i.e.  $y = \frac{\sin(x_1)}{\cos(x_2)}$ ).

You have used other functions, perhaps without knowing it, and some of them return values, but because you have not attempted to assume the returned values to anything, you have not noticed. For example, `plot` and `scatter` are in fact functions, and return the ID of the plot graphic. We simply have not been asking for the returned value so far. As per **MATLAB** help:

```
H = SCATTER(...) returns handles to the scatter objects
created.
```

with the handle, `H`, being an identifier of the graphic which could prove to be useful if e.g. you would like to modify one of the properties of an existing graphic.

Finally, it is important to note that by default, any variables created within a *function* are TOP SECRET, and by that, I mean that they are not accessible to the main **MATLAB** workspace and do not appear listed in the **Workspace window**. To see that this is a non-Trump-able true fact, create the following function (basically, the first example but split into 2 steps):

```
function [Y] = sin2new(X)
tmp = sin(X);
Y = tmp^2;
end
```

Here, a variable `tmp` is created to hold the value of the partial calculation. It does not appear in the **Workspace window** when you use the function. The advantage of this is that you could create a second function that also created a temporary variable internally called `tmp` with both instances of `tmp` treated entirely separate and isolated by **MATLAB** (i.e. setting the value of one instance of `tmp` does not affect the value of the other). This also however does lead to some additional complications in debugging *functions* (see Box). Try setting a breakpoint at the start of the line where the square of `tmp` is calculated – note that `tmp` now appear in the **Workspace window**. Continue the function and when it terminates, note that `tmp` is now gone from the list.

<sup>15</sup> Mathematically, the answer is not valid for all possible values of the 2 inputs (why?), and later we'll learn how to pro-actively deal with such a situation.

#### Debugging – functions

Functions are a prime example of the importance of being able to pause code part the way through (e.g. by setting a breakpoint) because when a function terminates, or crashes, you get to see none of the values of any variables created within the function, unless they have been returned as output (and assuming here that the code did not crash and managed to get to the end). Setting a breakpoint allows you to interrogate the values of any internal variables.

### 3.3 Conditionals '101'

#### 3.3.1 if ...

One of the other main programming constructs is the conditional statement, in which the outcome (one or more statement(s)) is conditional on the 'truth' or otherwise of a given (i.e. it being true or false). This is embodied in **MATLAB** (and similarly in most languages) by the `if ... end` construct (see **Conditional Statements** Box).

In creating an `if ... end` construct, the statement tested for truth can be any one of:

1. A variable having a value of true (1) or false (0). e.g.

```
if happy
...
```

where happy is a variable.

2. A **MATLAB** function returning a true or false, e.g.

```
if isnan(A)
...
```

where variable A, may or may not be a NaN.

3. A relational operator (see earlier), i.e. one of e.g.:

```
>, <, <=, >=, ==, ~=, &&, ||
```

and applied to a pair of variables, one variable and one value, or two values, e.g.:

```
if A > B
...
```

where A and B are numbers.

---

AN INITIAL AND RATHER COMPUTER PROGRAMMING TEXTBOOK-LIKE EXAMPLE is as follows: designing a program (a **MATLAB** script saved as an **m-file**) that asks whether or not you like bananas, and if you answer 'yes', tells you 'Correct – they are a great fruit!'.

But before we worry about anything else (e.g. how to apply a conditional statement), you'll need to know about inputting information into a MATLAB program from the keyboard<sup>16</sup>. Amazingly, you can guess (I actually just did) the command for requesting input – it is `input` (for 'input' – a rare occasion when everything is logical and simple!) (see Box).

With this (how to get **MATLAB** to ask for input and then receive and do something with keyboard input) – firstly create a blank **m-file** and save with a 'suitable' filename. Maybe add a header comment to remind you what this script is going to do.

#### Conditional Statements

The principal *conditional* statement in **MATLAB** is: `if ... end`

The basic `if` structure is:

```
if EXPRESSION (IS TRUE)
 STATEMENT(S)
end
```

in which the code CODE is executed if EXPRESSION is evaluated as true. No code is executed otherwise (and STATEMENT is false).

A variant addition – **else** – which allows for an alternative block of code (OTHER STATEMENT(S)) to be executed if EXPRESSION is instead evaluated as false, is:

```
if EXPRESSION (IS TRUE)
 STATEMENT(S)
else
 OTHER STATEMENT(S)
end
```

Finally, there is 3rd variant including **elseif**:

```
if EXPRESSION (IS TRUE)
 STATEMENT(S)
elseif EXPRESSION (IS TRUE)
 OTHER STATEMENT(S)
else
 OTHER STATEMENT(S)
end
```

Now, assuming that the first **EXPRESSION** is not true, a second **EXPRESSION** is evaluated, and only if that second **EXPRESSION** is also not true, will the final possible **STATEMENT** be evaluated. (Here, this final variant is shown with an `else ...` included at the end, but this is not a formal requirement to include.)

<sup>16</sup> All programming languages have such a facility and many basic programs, at least in the Old Days prior to widespread GUIs, make use of keyboard input

Secondly, (and on the next line) – define the text (question) that you are going to ask and assign this string to the variable `my_question`. Then place the input command (on the next, now 3rd line) for string input, and assign the input string to the variable `my_answer`. You should have a program consisting of 3 lines – an initial comment line, a line defining the question and assigning this string to a handy variable (`my_question`), and a line taking the results of the input function, and assigning it to a second variable (`my_answer`).

Run the program thus far. You should see the question displayed, and when you type in an answer and hit **RETURN**, the program will end. Because your **m-file** is configured as a script and not a function (see earlier), you can see the variable `answer` in the variable list and can check its value – it should contain a string with the answer you gave to the question. Make sure it all works like this so far.<sup>17</sup>

OK – aside from the use of `input`, there is nothing new here. Yet. The purpose of the program is to give a reply that depends on the answer given. This is where we are going to utilize a *conditional statement* – depending on whether the answer is ‘yes’ or not, we are going to display a different message. This is a fundamental programming element – different code will execute depending on the value of a variable – here the ‘different code’ is a different message and the value of the variable is ‘yes’ or ‘no’ (or other answer).

You are going to add an ‘if ...’ statement to the code (starting on line 4) to test whether the answer, held in the variable `answer`, is equal to ‘yes’. In the language of **MATLAB** syntax (see Box), the *EXPRESSION* is whether the string contained in `my_answer` is ‘yes’. How do we ask **MATLAB** to compare the value of `my_answer` with ‘yes’? Once upon a time, long long ago, **MATLAB** was simple and helpful and you could write:

```
if (my_answer == 'yes')
 [MESSAGE]
end
```

where `[MESSAGE]` you will later replace by a message that you will display using the `disp` command that you saw before. (In this stupid example it might be: ‘Correct – they are a great fruit!’).

However ... life is no longer this simple. **MATLAB** is going to make us use the function `strcmp` (see Box). In using `strcmp` we might break things down into 2 steps – the first comparing the 2 strings (`my_answer` and ‘yes’) and returning to us a value of *true* or *false* that we will store in a new variable. In the second step, we’ll ask the conditional to act on the value of the variable. The code will now look like this:

```
comparison_result = strcmp(my_answer, 'yes');
```

#### input

There are two variants – one for inputting numerical information and one for inputting a string (test) (as 1 could be either the value one or a 1-character string ...).

For inputting a numerical value:

```
x = input(prompt)
```

will display the text in the string variable `prompt` and set the value of `x` when a number is entered and **RETURN** pressed.

For inputting a string:

```
str = input(prompt, 's')
```

will display the text in the string variable `prompt` and set the value of `str` when a string is entered and **RETURN** pressed. Note that the second parameter passed to the function `input('s')`, tells **MATLAB** that the input is a string rather than a number.

<sup>17</sup> HINT: When you type the answer, it appears on the screen immediately adjacent (and untidily) to the end of the question. You can make this look nice(r) by adding a space at the end of the question string you assigned to `prompt`, e.g. `prompt = 'Do you like bananas? ';`

**strcmp** For once, the **MATLAB** help explanation is relatively simple and straightforward:

```
tf = strcmp(s1,s2)
compares s1 and s2 and
returns 1 (true) if
the two are identical.
Otherwise, strcmp returns
0 (false).
```

Which is pretty well much how we expected asking: `s1 == s2` to pan out.

(In **MATLAB** help – `tf`, the variable name used in the example, is short for ‘true-false’.)

```

if comparison_result
 [MESSAGE]
end

```

Or, we could have made this more compact:

```

if strcmp(my_answer, 'yes')
 [MESSAGE]
end

```

Your code should now have the 3 lines from before (comment, define question, get input) followed by 4 lines of the conditional structure, comprising: the `strcmp` function, the `if ...`, use of `disp` to display a message, and last, `end`.

Re-run (after saving) the program and confirm that it works (asking whether you like bananas and if you answer 'yes', tells you 'Correct – they are a great fruit!'). If not – time to de-bug! Note that if you tested the code in two stages, any bug at this point is only in the conditional structure. Start by double-checking the syntax required for the `if ...` structure. You could also try commenting out the message line and re-running.

Next, you might display an alternative message if the answer is not 'yes'. Refer to **help** / the margin Box on `if ...` and note that you can extend the structure with an `elseif` which would be followed by a line displaying the alternative message (e.g. 'Then you need to get a life, apple-lover.')

<sup>18</sup>

You could extend this example further and tackle the situation of their being 3 possible answers – 'yes', 'no', and ... 'I don't know' (or any other answer). Now the basic structure becomes

```

if strcmp(my_answer, 'yes')
 [MESSAGE 1]
elseif strcmp(my_answer, 'no')
 [MESSAGE 2]
else
 [MESSAGE 3]
end

```

Here – we are now adding an `elseif ...` line (followed by its specific message) (and see Box/**help**). Maybe try this and test it fully – inputting a 'yes', a 'no', and some other answer, and confirming that you get the correct message displayed.

You could also turn this around, and test for any answer except 'no' (the `~` is making the test, not 'no'), i.e.

```

if ~strcmp(my_answer, 'no')
 [MESSAGE 1]
else
 [MESSAGE 3]
end

```

<sup>18</sup> And then the line with `end` after that – follow the prescribed structure \*exactly\*.

Now you are asking whether the answer is something other than 'no' (which might be 'yes', but not necessarily so) – in the logical construct – whether the (string) contents of answer are not equivalent to 'no'.

---

CONTINUING TO BEAT THIS SAME TIRED EXAMPLE TO DEATH ... what if some wise-crack answered 'YES' rather than 'yes'?<sup>19</sup> One could write:

```
if strcmp(my_answer, 'yes')
 [MESSAGE 1]
elseif strcmp(my_answer, 'YES')
 [MESSAGE 1]
end
```

This will work, but you might note that you have had to exactly duplicate the MESSAGE line. If instead of displaying a simple message, a complex calculation was carried out – all the lines of the code following the if ... would have to be exactly duplicated after the elseif ... . While it might seem trivial to simply copy-paste the required lines, this is<sup>20</sup> dangerous – if the first set of lines are ever changed (due to a bug-fix or simple further development of the code), the same changes MUST then be exactly duplicated in each and every instance, or the code will not longer work correctly. This is *very* easy to forget to do, particularly for extensive code or code that you have not looked at for ... years. Code duplication also makes the overall code unnecessarily long (and hence harder to look through).

Instead, we can nest statements containing relational operators. What does this mean? Well, in the example of the answer being 'yes' or 'YES', logically, what we want is:

- (1) the contents of answer is equivalent to 'yes'
- OR
- (2) the contents of answer is equivalent to 'YES'

In code, this is written:

```
strcmp(answer, 'yes') || strcmp(answer, 'YES')
```

Make sure you are happy with what this means (it is pretty well much exactly as it looks == logic).

So – go modify your code to allow for a 'YES' or a 'yes'. Hell, try allowing for a 'Y' or a 'y' as well.<sup>21</sup> (You could extend it to 'no' also but I think you get the point ...)

---

<sup>19</sup> This goes to the heart of all software testing – what if the user does something you were not expecting? Hence why all software undergoes extensive testing by user or people who did not test it. Sometimes there are pre-releases ('alpha' or 'beta' versions or simple 'pre-release') of software to all or specific parts of the user community, precisely to provide feedback, find bugs, and see whether they can break it ...

<sup>20</sup> Note quite in the same way that driving down a mountain highway with your eyes shut or hungry sharks are dangerous.

<sup>21</sup> Sort of for this reason and that there are many different ways of writing 'yes', software often requires you to answer 'yes' in a restricted number of ways – this restriction is made clear as part of the message that asks the question. Common is to restrict the answer to 'Y' or 'y'.

How many bananas could you eat in a day? I bet it is less than ten. We'll let the computer ask and if the answer is 10 or more, you (the computer) shouts: 'lier!'.<sup>22</sup>

The basic code is very similar to before. Create a new m-file, add a comment line, define your question ('How many bananas do you think you could you eat in a single day?') and then get **MATLAB** to ask it and pass back whatever is entered in at the command line. The only difference at this point – refer to the usage of `input` (see Box) – is that we want a number input rather than a string. You can call the variable into which you assign the result of `input`, the same as before, or to make it distinct, e.g. `n_bananas`, i.e.

```
n_bananas = input(my_question)
```

In the `if` statement, we now want to test whether the value of `n_bananas` is greater or equal to 10 (or equivalently, greater than 9), i.e.

```
if (my_answer >= 10)
 [MESSAGE 1]
else
 [MESSAGE 2]
end
```

or

```
if (my_answer > 9)
 [MESSAGE 1]
else
 [MESSAGE 2]
end
```

Write this code and get it going. Feel free to switch fruit / fruit consumption threshold, question/answers, or whatever.

### 3.3.2 *switch ...*

A less commonly used alternative to `if ...` is `switch ... case ...` and is helpful in the case of multiple possible correct answers and/or multiple different answers.

For instance, and back to the ... fruit ... you might want the same answer for multiple different kinds of fruit. Trying coding up the program that would give you 'A great fruit!' for any of 'banana', 'kiwi', 'apple', 'pineapple', and 'cucumber' (yes they are technically fruit – Google it). You will find either you have many lines of code and many duplicated lines of the same message, or a very long line after `if ...` with loads of `strcmp` and ORs (`||`). Using `switch ... case ...` the code instead might look like:

<sup>22</sup> This example is even more stupid than the last one. But no more stupid than in any computer programming textbook and it will at least demonstrate a subtly different usage of `if` ...



```

switch my_answer
 case {'banana', 'kiwi', 'apple', 'pineapple', and 'cucumber'}
 disp('A great fruit!')
 otherwise
 disp('yuck!')
end

```

where `my_answer` is the name of a fruit entered in, in response to input, e.g.

```
my_answer = input('What is your favourite fruit?', 's');
```

Note that for a list of multiple possible value, **MATLAB** requires the list after `case` to be encased in `{}`. For a single answer, it would just be:

```
case 'banana'
```

for a string, and for a number:

```
case 10
```

### Conditional Statements (2)

The other main *conditional* statement is: **switch ... case ... end**

The basic switch structure is:

```

switch VARIABLE
 case VALUE(s)
 STATEMENT(s)
 end

```

which deviates rather from how **MATLAB** describes it, but this makes more sense to me (and hopefully to you). Here, `VARIABLE` is a variable and it is compared with one or more `VALUE(s)`. If the value of `VARIABLE` matches that of the `VALUE(s)`, then `STATEMENT(s)` are executed.

A common variant adds a default set of `STATEMENT(s)` to be executed if the value of `VARIABLE` does not match any of the `VALUE(s)`, e.g.

```

switch VARIABLE
 case VALUE(s)
 STATEMENT(s)
 otherwise
 STATEMENT(s)
end

```

You can also have multiple case possibilities:

```

switch VARIABLE
 case VALUE(s)
 STATEMENT(s)
 case VALUE(s)
 STATEMENT(s)
 otherwise
 STATEMENT(s)
end

```

### 3.4 Loops '101'

The next main program construct that you are going to see is the *loop*. There are a number of different forms of this in **MATLAB** (see **Loops Box**) (and also in other programming languages), but the basic premise is the same – a designated block of code (one of more lines of code<sup>23</sup>), is repeated, until some condition is met. That condition might be something as simple as a count having been reached, e.g. the block of code is always executed  $n$  times, or the condition might be slightly more complex and involve a *conditional statement* (see later). Will explore a very basic loop though an example, almost as contrived as for conditionals :o)

#### 3.4.1 for ...

In this subsection we'll start with a very straight-forward and somewhat abstracted usage of `for ...`, which hopefully will get you in the mood for *loops*. Then we'll go through some slightly more problem-focused examples.

---

LOOPS GROUND ZERO. Basically – loops cycle through a series of numbers between specific limits, or if you like, 'count'. As the loop counts (cycles), it allows you to execute some code, so for each count (or cycle), the (same) block of code is executed. We'll worry about what you might 'do'<sup>24</sup> (i.e. the code fragment) in a loop, later.

Consider, or rather: create a new **m-file**<sup>25</sup> with the following loop:

```
for n=1:10
end
```

Save it. Run it. What did it do?

I bet you have absolutely no idea! It actually cycled around ten times, counting from  $n=1$  through  $n=10$ , but you would not know it as there was no code without the loop to do anything.<sup>26</sup>

There are 2 alternative crude debugging strategies you could take<sup>27</sup>:

1. Simply add a line within the loop with the name of the (counting) variable, e.g.

```
for n=1:10
 n
end
```

and it will spit out the value of  $n$  each time around the loop.

2. Print the value of  $n$  'properly'<sup>28</sup>, e.g.

#### Loops in MATLAB

**for**  
The basic `for ... end` structure is:

```
for n = VAL1:VAL2
 CODE
end
```

where VAL1 and VAL2 are the limits that  $n$  will count between (starting at VAL1 and ending at VAL2), meaning that STATEMENT(S) will be executed  $(VAL2-VAL1)+1$  times in total. STATEMENT(S) can be one or more lines of code, that will all be executed on each and every cycle of the loop.

The loop need not count in increments of one (1), the default, e.g.:

```
for n = VAL1:INC:VAL2
 CODE
end
```

counts with an increment of INC. It is also possible to count down (a negative value of INC).

#### while

The basic structure is similar to that for `for ... end`:

```
while STATEMENT (IS TRUE)
 CODE
end
```

while differs from if in that there are no alternative branches of code that can be executed. The while ... end loop cycles and CODE continued to be executed (for ever) until the STATEMENT is evaluated to be false.

<sup>23</sup> It is possible to for the block of code to be only a fragment of a single line and hence the entire loop plus code block, to be written on a single line.

<sup>24</sup> Note intentionally a joke. Actually, this is only funny if you know FORTRAN, and even then it is only marginally funny.

<sup>25</sup> Comment it!

<sup>26</sup> You get one clue – if you look in the variables **Workspace window**, you'll see there is a variable  $n$ , with a value of 10 – the last value it was assigned before the *loop* ended.

<sup>27</sup> Plus, you could add a breakpoint and view the value of  $n$  in the **Workspace window** each cycle around the loop.

<sup>28</sup> Although you can get away with just writing:

```
disp(n)
```

```

for n=1:10
 disp(str2num(n))
end

```

or you can tart this up even nicer by creating a string that provides more explicit information back to you, e.g.

```

for n=1:10
 my_string = ['The value of n is: ' str2num(n)]
 disp(my_string)
end

```

or if you are happy with more going on in a single line:

```

for n=1:10
 disp(['The value of n is: ' str2num(n)])
end

```

(but they work the same – check it).

LOOPS IN ACTION. So, consider the following (contrived) ‘problem’ – you want to be able to enter a series of numbers and return their sum (although equally one could perform and return all sorts of statistics).<sup>29</sup> The basic code is simple. Using the other (numerical input) form of input, for 2 numbers, it might look like (although in practice, your code is full of helpful comments, right?):

```

my_question = 'Please enter a number: ';
A = input(my_question);
my_question = 'Please enter a number: ';
B = input(my_question);
disp(['The sum of the numbers is: ' num2str(A+B)]);

```

The first 4 lines you should be A-OK with. Note that in line 5, 2 strings have been concatenated by enclosing ‘The sum of the numbers is: ’ and num2str(A+B) in a pair of brackets []. The string representing the number sum is itself created by adding A and B, and then converting the resulting number into a string using num2str (see earlier). As always – if you are happier breaking down the last line into its component parts, e.g.

```

answer = A+B;
answer_string = num2str(answer);
disp(answer_string);

```

then please do!

So far so good. But what if you wanted 4 numbers summed ...

```

my_question = 'Please enter a number: ';
A = input(my_question);
my_question = 'Please enter a number: ';

```

<sup>29</sup> Obviously, one way to do this would be to enter the numbers into a file first, use the load function, and calculate the sum.

```

B = input(my_question);
my_question = 'Please enter a number: ';
C = input(my_question);
my_question = 'Please enter a number: ';
D = input(my_question);
disp(['The sum of the numbers is: ' num2str(A+B+C+D)]);

```

You can see whether this is going – firstly that you are duplicating more and more lines of code as the number of numbers increases. Secondly, and we'll come to that in a moment – what if the program does not know *a priori* how many numbers you want to sum?

You can see the code that is being repeated (here for input x):

```

my_question = 'Please enter a number: ';
x = input(my_question);

```

If you bothered to read the margin box earlier, you'd known that this is exactly what a *loop* can be used for. We therefore want something of the form:

```

for n = VAL1:VAL2
 my_question = 'Please enter a number: ';
 x = input(my_question);
end

```

The easy part is the configuration of the loop – in the previous example with 4 inputs, we would write:

```
for n = 1:4
```

and the loop with go around 4 times as the counter *n* counts from 1 (VAL1) to 4 (VAL2) in increments of 1 (the default behavior of the *colon operator*). Each time around the loop the block of (2 lines of) code is executed and a number is inputted. But what is still missing? Try it exactly like this and see if you can see what is going on, or rather, not going on. If you think it is not working as expected – try some debugging. See if you can come up with a solution once you see what the problem is. (Warning: the spoiler is in the margin.)

After having tried your own solutions, try out both of the given alternatives (assuming that one of them was not also your solution). Note that you are not given the complete code needed and some further debugging might be needed (but they do both work!).

Two things to be aware of in doing this:

1. If you set the maximum number of items quite high and then get bored and need to exit the program – press the key combination **Ctrl-C** and **MATLAB** will exit your program (but leave **MATLAB** running).
2. If you run the program a second time and use the vector approach, something very odd starts to happen to the reported sum.

It should be apparent if you tried it out, that the value of *x* at the very end of the program, is equal to the last value you entered. In other words, each time you go around the loop you are over-writing the previous entered value and end up with nothing to sum at the end. There are two (or more) possibilities to solve this:

1. You could keep a *running sum*. This would also avoid having to explicitly calculate a sum at the end, but you would not have saved the numbers as you went and no other stats would be possible.

You would do this by adding the inputted value to the existing value, i.e.

```
x = x + input(prompt);
```

where *x* is the running total. What this says is: take the current value of *x*, add the value if the user input, and place the total back into the variable *x*.

The only problem here ... is that MATLAB does not know what the very first value of *x* is – i.e. the value before the loop starts and that you then try and add `input(prompt)` to. The solution is to initialise the value of *x* before the loop starts, e.g.

```
x = 0;
```

2. Alternatively, you could add the newly inputted number to the end of an existing vector. In this way, you end up recording all the values that were inputted. e.g.

```
y = [y input(prompt)];
```

which says take the vector *y*, and add a further value (`input(prompt)`) to the end of it. At the end of the program (after the loop has terminated), you have to sum the contents of the vector *y*.

You can solve this (first try it out – running the program several times in a row to see what happens) either by initializing the vector `y`, just like you did for `x` in the 1st solution, i.e.

```
y = [];
```

(before the loop starts, of course), or you can clear the workspace using `» clear all` (clears *\*all\** variables), or clear just the problem variable (`y`) that will end up growing and growing and growing ... (`» clear y`).

### 3.4.2 Other loop configurations and usages

In the previous examples, the loop limits were fixed in the program itself – you’d have to edit the script code and save the file in order to be able to input and sum a different number of values. You could create a more flexible program by making the m-file a function rather than a script.<sup>30</sup> The idea here is to create a function that takes a single input. This input will be the maximum loop count. If the input variable was called `max_count`, then the loop structure would now look like:

```
for n = 1:max_count
 my_question = 'Please enter a number: ';
 x = input(my_question);
end
```

<sup>30</sup> There are other ways of adding flexibility to the loop count that we’ll see shortly.

Referring to the previous lessons on *functions* (as well as **help** if need be), create a function that when you call it, e.g. like:

```
» function_sum(5)
```

will request 5 inputs and display the sum.

Alternatively, your program (as a *script*), before the loop starts, could ask for the number of values to be entered, passing this to the variable `max_count`, with the loop then looking exactly like the above. In both cases you are substituting a fixed number (e.g. 4) for a variable that might contain any number. Equally, not only does the count not need to start at one, and the lower loop count limit could also be a variable (`min_count`?).

Finally, in addition to flexible loop count limits, the value of the increment in the count each time around the loop need not be one. For example:

```
for n = 10:10:100
 ...
end
```

is exactly equivalent in terms of the number of iterations carried out to

```
for n = 1:1:10
 ...
end
```

and which is the same as the default behavior of the colon operator:

```
for n = 1:10
 ...
end
```

The value of the loop counter *n* simply differs by a factor of 10 at every iteration between the top and bottom two versions.

### 3.4.3 *Fun(!) worked examples*

(Only one example to date. And not necessarily even fun.)

---

LOOPS, CAMERA, ACTION! (A more colorful example of loops in action.) What we are going to do is (load and) plot a sequence of monthly data-sets and put them together to create a movie (animated graphic) to illustrate the seasonality of temperature in global climate. You will hopefully thereby better appreciate the value of constructs such as *loops* in computer programming in saving you a whole bunch of effort and needless duplication of code. (Equally, you might not have wanted a movie as the end result, but simply a number of plots, all identical except in the specific array of data they were plotted from.)

First download all the monthly global surface temperature data-files on the course webpage (there are 12 files to download)<sup>31</sup>. Then you are going to want to plot them all ... which would get desperately tedious if you had to do this at the command line 12 times. Think how much more of your life you would be wasting if the data were weekly. Or monthly data for 1972 through 2003, some 372 separate data-files ... You would never have time to go get a coffee ever again(?)

Create a new *m*-file. Call it ... anything you like<sup>32</sup>. However, as well as appropriately naming your script file, add a *comment* on the first line of the file as a reminder to yourself of what it is going to do. Also, for now, it is helpful to include the command: `close all` (which closes all currently open figure windows) although this is far from essential.

To make an animation, we need to make a series of frames, with each one being a different monthly temperature plot (in sequence; Jan through Dec). The files are rather conveniently named: `temp1.tsv`, `temp2.tsv`, ... `temp12.tsv`<sup>33</sup>. We should start by loading this little lot in. For the first file we could write:

<sup>31</sup> In scripting, it is also possible to automate downloading files from the internet.

<sup>32</sup> `bob_the_builder.m` counts as 'anything you like', but that looks pretty lame and it certainly won't help you remember what the script does if you came back to it sometime in the future.

<sup>33</sup> Don't worry about the `.tsv` file extension – the file format is plain old text (ASCII) and could have instead been `.txt`.

```
temp = load('temp1.tsv');
```

or equally:

```
temp(:, :) = load('temp1.tsv');
```

and hence with a slight-of-hand, we could also write:

```
temp(:, :, 1) = load('temp1.tsv');
```

Can you see that these statements are identical? Run the script with one, then with the other, just to be sure. The last form is really useful, because we can now go on and write:

```
temp(:, :, 2) = load('temp2.tsv');
```

What you have done here is to load the January 2D (lon-lat) temperature distribution into the 1st 2D layer of the temp array, and then we have gone and created a second 2D layer on top of the first with the February data in it. Look at the **Workspace window** (or type `size(temp)`) – you now have a 3D ( $94 \times 192 \times 2$ ) array. Fancy! This is your first 3D array – there is nothing really conceptually different from the 2D arrays that you have already been using, we simply have a 3rd index for the third dimension (if it helps, you can think of a 3D array as being indexed by: row, column, layer).

You could go on and load in the March, April, etc data in a similar fashion, but you should be able to see a pattern forming here – each filename differs only in the number at the end of its name and this number corresponds not only to the number of the month, but will also correspond to the layer index of the 3D array that you will create. This is something that a loop could be used for while you go off for a coffee.

We first need to construct the loop framework. We'll call the month number counter variable, `month`. Create a loop (with nothing in it yet) with `month` going from 1 to 12.<sup>34</sup> Refer to the course text (this document!), and/or the MATLAB documentation, and/or the entirety of the internet, if necessary. The syntax (and examples) is described in full under » `help for`. Save the script (m-file) and run it<sup>35</sup>. What happens? Can you tell?

One way of following what is going on as **MATLAB** executes the commands within a script is to explicitly request that it tells you how it is getting on. You can use the function `disp` to help you follow what the program is doing (this is Old School debugging<sup>36</sup>). Within the loop, add the following line:

```
disp(month)
```

then save and re-run the script. Now you can see how the loop progresses. This sort of thing can be useful in helping to *debug* a program – it allows you to follow a program's progress, and if the program (or **MATLAB** script) crashes, then at least you will know at

<sup>34</sup> Don't forget to suitably comment what it is that the loop does with a line (or even 2, but don't write a whole essay) beginning with a %.

<sup>35</sup> Typing: the m-file filename without the extension.

<sup>36</sup> You can also add a **breakpoint** within the loop and thus can cycle through the loops one-by-one, thereby being able to check the status of the variables within the loop and how they change from iteration to iteration.

what loop count this happened at, even if you are not given any more useful information by **MATLAB**. Only when you are happy that you have constructed a loop that goes around and around 12 times with the variable month counting up from 1 to 12; comment out (%) the printing (`disp`) line<sup>37</sup> (unless you have grown rather attached to it) and move on.

We can construct filenames to load in by:

1. Forming a complete filename by concatenating sperate strings.

For example:

```
» filename = ['temp' '1' '.tsv']
```

will create the filename for the first dataset out of 3 components parts – a common elements of all the filenames ('temp'), the number of the month ('1'), and the file extension ('.tsv').

2. Converting a number value of a (count) variable to a string (the `num2str` function).

This is where the role of the loop counter (stored in the variable month) comes in. Each time around the loop, the value of variable month is the number of the month. All you have to do is to convert this value to a *string* and thereby automatically generate the correct month's filename each time (as per above).

Now add the following within the *loop* in your script;

```
filename = ['temp' num2str(month) '.tsv'];
```

and after it some debugging<sup>38</sup>:

```
disp(filename)
```

just to confirm that appropriate filenames are being generated. Save and run the script. Satisfy yourself that you know what it is doing. Can you see that you are now automatically generating all the 12 filenames in sequence? And this only takes 3 lines of code total (not including the debugging line), compared with 12 lines if you had to write down all the 12 file names long-hand.

*comment out* the `disp(filename)` line, and add a new line to load in each dataset from the new filename that is constructed each time the loop goes around.<sup>39</sup> Assign the new 2D data array to the temp array at the next layer number. Take a look at the **Workspace window** – note that you have an array (temp) that has size  $94 \times 192 \times 12$ . If temp is  $94 \times 192 \times 1$  then go back a page or so and go through the bit about loading data into a 3D array. You want to avoid over-writing the information that is already there, so the line; `temp = load(filename);` will not work (and you will only get a  $94 \times 92$  array after going 12 times around the loop). Why? (Again, look back a page-ish.)<sup>40</sup>

<sup>37</sup> Note that by commenting out a line rather than completely deleting it, if you want to print out the loop count in the future, all you have to do is to un-comment the line, rather than type in the command all over again. This can be really useful if your debug command is long, or particularly if you have a whole series of lines that are required to report the information you want to know.

<sup>38</sup> Or you can make use of a **breakpoint**.

<sup>39</sup> Remember that the load line goes inside the loop. (Why? Try writing it outside the loop (at the end) and see what happens if you like.)

<sup>40</sup> If you are still stuck, then stick up a paw.



At the end of (but still within) the loop (i.e., before the loop has completely finished), create a new figure window on one line, then plot (using `pcolor`) the monthly temperature data on the next line, and add the essential labelling stuff (lines after that). All within the loop still. This line should look something like:

```
pcolor(temp(:,:,month));
```

and should produce extremely exciting graphics as in Figure 3.2<sup>41</sup>. (Don't just type this line in blindly (maybe it doesn't 'work' anyway). Make sure that you understand what you are doing (otherwise why do GEO111 at all?).)

Save and run the script. Do you have 12 different temperature plots on the computer screen?<sup>42</sup> Note that this is where the `close all` command at the start of your script comes in useful. Because if you re-run the script, you won't then end up with 24 figure windows. And then 36 the time after that, and ... (There is actually no need to create a new figure window each time – comment out the command that creates a new figure window (`figure`). Save and re-run and note the difference.)

Finally ... look up **MATLAB** help on `getframe`. Then go back to your global temperature loading/plotting script and add the following line<sup>43</sup>:

```
M(month)=getframe;
```

Save and run. When MATLAB is all done, at the command line type in:

```
» movie(M,5,2)
```

and hopefully ... an animation of the progression of monthly surface air temperatures globally, should appear<sup>44</sup>.

If you want to play some more, just type `help movie` – there are controls for not only the number of times you loop through the complete animation, but also for the numbers of frames per second. But we will revisit this later – the 2D plotting you have done so far is \*very\* basic and there is no scale or sane x/y axes. Later we can also add the continental outlines that will help orient you and improve the quality of the graphical output.

Before you move – go look at your script – is it well commented? Would you be able to tell exactly what it does by the end of GEO111? What about next year? Are the *loop* contents indented? It is important that it is commented and laid out adequately.

<sup>41</sup> The 2D graphics will get \*much\* better later – one thing at a time!




Figure 3.2: Extremely unappealing blocky plot of Earth surface temperature (who cares with month? – the graphics are too poor to matter ...).

<sup>42</sup> If not, stick your paw up in the air for help ...

#### movie2avi

The function `movie2avi` converts an animation encoded in **MATLAB**'s `movie` format to an `avi` file, which is a common film format that can then be played in **Windows** (or other operating systems) without having to use **MATLAB** to display it. It is also a format that could e.g. be embedded in a **Powerpoint** presentation. A typical basic usage is:

```
» movie2avi(M, 'file.avi');
```

where `file.avi` is the output filename and `M` the input **MATLAB** movie name.

<sup>43</sup> Where to put the line? See the Example given in the help on this function. It is exactly what you are doing here.

<sup>44</sup> Note that the active Figure window may have disappeared behind some other windows so go rescue it to see what is happening.

### 3.5 Loops and conditionals ... together(!)

No surprise that you might combine both loops and conditionals in the same programming structure. In fact, this becomes very powerful and is an extremely common device in programming.

#### 3.5.1 for ... and conditionals

As an alternative to (or as well as) a fixed loop, or variable and (function) parameter passed controlled loop, we could specify a near infinite loop, but provide a get out of jail free. For example, within the loop, we could add a line that asks an additional question: 'Another input (y/n)?' We would test the answer and if no ('n'), exit the loop (and report the sum as before). This would look like:

```
my_question1 = 'Please enter a number: ';
my_question2 = 'Another input (y/n)? ';
for n = 1:1000000
 my_number = input(my_question1);
 my_string = input(my_question2, 's');
 if strcmp(my_string, 'n')
 break
 end
end
```

where 1000000 is simply chosen as a 'very large number' and one rather larger than the maximum number of numbers you could ever imagine entering<sup>45</sup>.

The key new command here is `break`. The way the code works (hopefully!) is that at the start of a new iteration of the loop, the 'another input' question is asked – if no further input is required, the loop exits via the `break` command. Otherwise (the else), the user is prompted for another input. Note that now we have loops and conditionals nested together, it helps even more to *indent* the code<sup>46</sup>. Also note that here – the two different questions (demands) outputted to the screen – 'Another input (y/n)?' and 'Please enter a number' – are pre-defined before the loop starts. These same lines could be placed within the loop, but re-defining the variable e.g. `my_question1` as 'Another input (y/n)?', each and every time, is redundant (i.e. it could instead simply be defined once at the start of the program). Also also note that in this code, the number entered in is assigned to the variable `my_number` rather than `n` as was used before – simply to help distinguish the number input from the string input (assigned to `my_string`).

It is up to you to 'do' (i.e. add or modify the code) something with the number entered in and stored in the variable `my_number`, as each

#### `break`

Simply – `break` terminates the execution of a `for` or `while` loop'. And from **help** a further clarification: 'Statements in the loop after the `break` statement do not execute.'

Slightly more complicated (but not much) in the case of nested loops – in this case, `break` exits only the loop in which it occurs.

#### Indenting code

Just do it (or let MATLAB do it). Even for a single loop or conditional, it is way easier to see what code is within the loop and what outside it, when the code inside starts several spaces in from the margin.

For nested loops and conditionals, it is even more important to keep (visual) track on what is going on.

Note that the indentation (or lack of) does not affect the execution of the code (unlike in e.g. Python).

<sup>45</sup> There is a better way of doing this, with the `while` construct, that we'll see shortly.

<sup>46</sup> **MATLAB** will do this for you if you click on the Indent icon. It will also indent the code as far as it reasonably can, as you type.

time around the loop, the previous value is over-written by the new input.

Currently, the program only exits upon entering 'n' to the question. Instead, we could have it exiting for any answer other than 'y':

```
my_question1 = 'Please enter a number: ';
my_question2 = 'Another input (y/n)? ';
for n = 1:1000000
 my_number = input(my_question1);
 my_string = input(my_question2, 's');
 if ~strcmp(my_string, 'y')
 break
 end
end
```

which compares `my_answer` and 'y', if this is not true (that they are the same), `break` is executed.

A MORE PRACTICAL EXAMPLE would be in saving a data file, to test for a filename already existing and if so, automatically modifying the new file name so as not to over-write the file.<sup>47</sup> The relevant function is `exist` and in the case of a test for a file, returns either 0 (the file does not exist in the MATLAB search path, although that does not rule out it existing somewhere else entirely), or 2 (the file exists).

Clearly(?), in the example of saving the movie file (using the `movie2avi` command), you might well want to test whether the filename that you have chosen already exists (i.e. the value returned by `exist` is 2). If so (i.e. the file exists), you need to modify the filename by means of a new concatenation, perhaps appending something like '\_NEW' to the end of the string<sup>48</sup>. If not, and the filename has not already been used, you can proceed as before – the equivalent of 'doing nothing'. Go ahead – try it (i.e. modify your code to avoid over-writing an existing filename).

You could start by defining a default filename in the code<sup>49</sup> that you will use if there is no clash with any existing file, e.g.

```
my_filename = 'GE0111_movie.avi'
```

Now test whether this filename already exists:

```
filename_check = exist(my_filename, 'file')
```

Finally, using an `if` statement, test whether the value of `filename_check` is equal to 2. If so, you are going to need to modify the filename string (`my_filename`). If not, you can let the *conditional* just end and proceed to saving. Modifying the filename is just as per for the example of loading global temperature distributions, e.g.

<sup>47</sup> Note that while in the m-file Editor, **MATLAB** asks you if you want to over-write an existing file, when saving a file directly from a program, no such dialogue box or warning is given.

<sup>48</sup> Recall that in using the `movie2avi` command, you pass a filename – simply modify the filename passed, in a similar way to in which you modified the filename for loading the temperature data.

#### **exist**

Tests for whether a specified variable, function, file, or directory exists, and in generally, which is these it is.

The general syntax and usage is:

```
exist('A')
```

to return what A is.

An extended syntax with a second passed parameter:

```
exist('A', 'file')
```

returns value of 2 is returned is A if a file, and for:

```
exist('A', 'dir')
```

returns a value of 7 is returned is A if a directory.

<sup>49</sup> Either near the very start of the program (neater), or just before you need to use the string (to save a file).

```
my_filename = ['NEW_' my_filename];
```

where here, we take the string contained in `my_filename`, we append a 'NEW\_' to the start<sup>50</sup>, and assign the new (longer) string back into the variable `my_filename`.

The file naming becomes a little awkward, so rather than the entire filename + extension, you might just store just the filename in the (`my_filename`) variable. i.e.

```
my_filename = 'GE0111_movie'
```

but the remembering when you test for the existence of a particular file, you must add the extension, i.e.

```
filename_check = exist([my_filename '.avi'],'file')
```

(here we create a new string [`my_filename '.avi'`] by concatenating `my_filename` with the extension `'.avi'`). If the filename exists, the new filename we generate can then be:

```
my_filename = [my_filename '_NEW'];
```

(adding the `'_NEW'` after, rather than before the existing filename string).

<sup>50</sup> Note that because the filename already has its `'.avi'` extension attached, you'll have to modify the start of the string.

### 3.5.2 *while ...*

We can re-frame the earlier example programs using the `while` construct rather than the `for` loop. But now ... you need to specify under what conditions the loop continues as the basic syntax (see earlier or **help**) is:

```
while STATEMENT (IS TRUE)
 CODE
end
```

Here – `STATEMENT (IS TRUE)` is the conditional. For instance and rather trivially, create the following as a program and run it<sup>51</sup>:

```
while true
 disp('sucker')
end
```

<sup>51</sup> You ... are going to need a **Ctrl-C** on this one ...

What has happened is that `true` is always ... `true`. Hence the condition is always met and the `while` loop loops forever. Conversely, `while false` would never loop, not even once. more interesting and useful is when the statement might change in value as the loop progresses.

Consider (and type up in a script):

```
n = 0;
while (n < 10)
 disp('sucker')
end
```

This also will loop for ever as  $n$  is initialized to 0 and hence the statement ( $n < 10$ ) is always true. But if we increment the value of  $n$  each time around the loop:

```
n = 0;
while (n < 10)
 disp('not a sucker')
 n = n + 1;
end
```

then the loop will execute exactly 10 times (just as per for  $n = 1:10$ ). You could also do this in reverse:

```
n = 10;
while (n > 0)
 disp('not a sucker')
 n = n - 1;
end
```

Now,  $n$  counts down from 10 and when it reaches a value of 0, it is no longer greater than zero and the statement ( $n > 0$ ) is false (and the loop terminates).

It is not always completely obvious whether even simple while loops like this execute 9 or 10 (or 11) times particularly when often you might come across while ( $n \geq 0$ ) that allows the loop to continue when  $n$  has reached a value of zero (but not below). So – spend a little while playing about with different while configurations and loop criteria.

Finally, note that the conditional statement in the while loop need not test for an integer being larger or smaller than some threshold. One could equally loop on the basis of a string equality/inequality. For example, taking the previous example using break could be re-coded with a while loop:

```
my_question1 = 'Please enter a number: ';
my_question2 = 'Another input (y/n)? ';
my_string = 'y';
while strcmp(my_string, 'y')
 my_number = input(my_question1);
 my_string = input(my_question2, 's');
end
```

and ends up a slightly shorter and more compact piece of code, omitting the need for a break or a nested structure. However, in this example, we do need to initialize the value of `my_string` (to 'y' – assuming that we want at least one number). Try it and then adjust it so that the loop proceeds as long as the answer is not 'n' (rather than as long as it is 'y')<sup>52</sup>. Note that as before – it is up to you to 'do' (i.e. add or modify the code) something with the number entered in an

<sup>52</sup> See earlier Example.

stored in the variable `my_number`, as each time around the loop, the previous value is over-written by the new input.

EXTENDING THE FILENAME CHECKING EXAMPLE ....<sup>53</sup> to fully integrate a loop and conditional. The problem with the previous code is that you checked for the existence only a default filename (and appended `'_NEW'` if a file already existed).

One (partial) solution would have been, rather than append a pre-defined string (`'_NEW'`) to the filename, would be to request that the user provide either a string to append, or a completely new filename. You have already see the input command in action, so you should be in a good position to code this modification up.<sup>54</sup>

A better solution (because even when asking for an alternative filename – what if that file exists too?) would be to keep checking for a filename clash and keep asking for a new filename, until a unique one is found. Who knows how many attempts this might take (to find an unused filename), so `while ...` would be a better choice of loop than `for ....` Because `exist` returns a 2 if the file already exists, a logical condition for `while` would be `while exist is returning 2`:

```
my_question = 'Please enter an alternative filename (without
the extension): ';
while (filename_check == 2)
 my_filename = input(my_question,'s');
 filename_check = exist([my_filename '.avi'],'file')
end
```

Within the loop, a new filename is requested and then check against the directory contents. What is missing is the initial value of `filename_check`. In a previous example, we simply set a value at the start. If we did that here, the first line of this code would look like:

```
filename_check = 2
```

In this case, we do not need a default filename as the user provides the very first filename that is tested. Alternatively, we could perform a single check before the *loop* starts:

```
my_question = 'Please enter an alternative filename ...
(without the extension): ';
my_filename = 'GE0111_movie';
filename_check = exist([my_filename '.avi'],'file')
while (filename_check == 2)
 my_filename = input(my_question,'s');
 filename_check = exist([my_filename '.avi'],'file')
end
```

<sup>53</sup> Which first time around did not actually combine loops and conditionals in the same structure. Rather, a loop came first in the program (loading in and plotting the temperature data), ended, and only then a conditional checking the filename.

<sup>54</sup> Effectively, all you have to do, if `exist` returns a 2 and the file already exists, is to ask for an alternative filename, and use the string entered in as the new filename (and don't forget to add the `'.avi'` extension to the end when saving)

### 3.6 Even more (and loopier) loops

[Further examples of increasingly extreme loopiness.]

**LOOPING THROUGH ARRAYS.** In plotting e.g. global temperature distributions, it would be nice to add on the continental outline. Currently and particularly with the very basic 2D plotting you have seen so far (`pcolor`) left to some extent guessing where the land and where the ocean is.

A pair of files are provided (from the website), comprising a series of pairs of lon-lat values that delineate the outline of the continents and all but the smallest of islands:

```
continental_outline_lat.dat
continental_outline_lon.dat
```

Download, and load these into the **MATLAB** workspace (in the 'usual way'). You should now have 2 vectors. Maybe view then in the **Variable Window** to get a better idea of what you are dealing with. Also keep an eye on the entries in the **Workspace Window** and perhaps the **Min** and **Max** values to give you an idea of the range (here: of longitude and latitude values). Try plotting these lon/lat locations. Use the scatter plotting function (which makes it all the easier as your data is in the form of 2 vectors already). You might need to reduce the size of the plotted points (refer to the earlier exercises, or `help`) and additionally, you might want to fill the points (up to you). Remember you can set the axis limits, which presumably should be 0 to 360 or -180 to 180, on the *x*-axis (longitude), and -90 to +90 on the *y*-axis (latitude). Font sizes of labels can also be increased if necessary. You might end up with something like Figure 3.3.

By plotting dots (points), the coastal outline at higher latitudes gets increasingly pixelated (why?). So, we might instead plot as lines between the lon-lat pairs. For this, you could simply use `plot`. Do this, and see if you get something like Figure 3.4..

Well ... interesting. If you think about it, as one continental outline is completed, the next lon-lat pair will be for the next continent or island. What `plot` does is to join up *\*all\** the adjacent points, which is why you get the straight lines criss-crossing the map with the start of each successive continent and island in the dataset joined to the end of the previous one.

The continental outline dataset is not actually that useless. There are additional files that specify which block of lon-lat pairs belong to a single shape (i.e. continent or island). Load in the 2 additional files:

```
continental_outline_start.dat
continental_outline_end.dat
```




Figure 3.3: Continental outline (of sorts).




Figure 3.4: Another continental outline (of sorts).

These vectors hold information regarding the start row and end row, of each shape. Again, view the contents of these vectors to get an idea of what is going on. For example, you'll see that the first entry is that the first shape starts on row 1 (`continental_outline_start.dat`), and ends on row 100 (`continental_outline_end.dat`). The 2nd shape starts on row 101, and ends on row 200. etc etc The simplest way too start dealing with all this, is to just plot the very first shape, defined by rows 1-100 of the lon and lat vectors. By now, you hopefully will be able to see that to plot rows 1-100 of lon and lat data, you are going to do:

```
plot(lon(1:100),lat(1:100));
```

(here I have named the arrays `lon` and `lat` for added convenience rather than the long-winded default file-name based versions (`continental_outline_lat`, `continental_outline_lon`)).

Well ... this is probably about as unexciting as it gets – a small piece of the Antarctic coastline. If you do a `hold on` and plot the next block (rows 101-200), you'll get the next chunk of coastline. (Try this and see.) You could keep going this – manually adding additional sections of the global continental outline. This could get tedious ... and it turns out that there are 283 different fragments to plot, all one after another. (This number comes from asking **MATLAB** the length of `continental_outline_start.dat` or `continental_outline_end.dat`.) This is, of course, why we need to get clever with a *loop* and automatically go through all 283 fragments, plotting them on on top of another in the same figure.

How? First you need to have the plot command in a more general form – you do not want to have to read the values out of the `continental_outline_start.dat` and `continental_outline_end.dat` files manually. Hopefully, it should be apparent that you can re-write the plot statement for the first fragment, as:

```
plot(lon(line_start:line_end),lat(line_start:line_end));
```

where for the first fragment, the values of `line_start` and `line_end` are given by `lstart(1)` and `lend(1)`, respectively (renaming the original vectors to shorten the variable name)<sup>55</sup>. Re-writing again:

```
plot(lon(lstart(1):lend(1)),lat(lstart(1):lend(1)));
```

Try this and check you still get the single piece of the Antarctic coastline.

Really, you should hopefully be making the mental leap to looking at (1) and thinking that it could be:  $(n)$ , where  $n$  is a loop counter which can go from 1 to 283<sup>56</sup> and hence loop through all the line fragments. Yes? For instance, setting  $n=1$ , and plot (with  $n$  replacing

#### Length

This function could almost not be simpler – just pass the name of a vector, and it returns its length (i.e. the number of rows, or columns, depending on the shape of the vector).

<sup>55</sup> You cannot use the obvious variable name `end` – why not?

<sup>56</sup> This number comes from a 5th file – `continental_outline.k.dat`, that numbers the continents/islands from 1 to 283. You don't need it, although downloading it, loading it, and determining the length of the vector gives you the loop limit and you would not have to go trusting me to write down 283 correctly without making a mistake ...



1 in the code fragment above) – you should again get that very first fragment. Try setting  $n=283$  and plot. Do you get the last fragment (what is it of<sup>57</sup>)?

So ... create yourself an `m`-file. Load in the lon-lat pairs as vectors (renaming then to something more manageable if you wish). Load in the vectors continuing the start and end information. Create a `do ... end` loop. Maybe print (`disp`) the loop count and run the program (after saving), just to check first that the loop is functioning correctly. Before the loop, create a Figure window. and set `hold on`. You now have a basic shall of a program – loading in the data, initializing a figure, and appropriate looping, but not yet actually doing anything within the loop.

In the *loop* all you need is the `plot` command, but with the start and end rows being a function of  $n$  (or whatever you call the loop counter). Set axis dimensions and label nicely (after the loop ends). Run it. Hopefully ... something like Figure 3.5 appears(?)

<sup>57</sup> An island at about 20N and -150E if you have done it correctly.

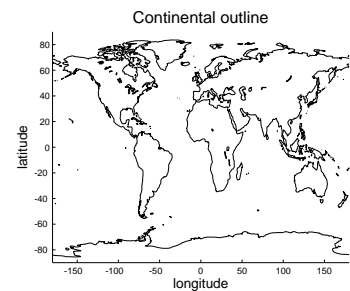


Figure 3.5: Another go at the continental outline!



## 4

### *Further ... Computers and software*

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetur id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

Nam dui ligula, fringilla a, euismod sodales, sollicitudin vel, wisi. Morbi auctor lorem non justo. Nam lacus libero, pretium at, lobortis vitae, ultricies et, tellus. Donec aliquet, tortor sed accumsan bibendum, erat ligula aliquet magna, vitae ornare odio metus a mi. Morbi ac orci et nisl hendrerit mollis. Suspendisse ut massa. Cras nec ante. Pellentesque a nulla. Cum sociis natoque penatibus et magnis dis parturient montes, nascetur ridiculus mus. Aliquam tincidunt urna. Nulla ullamcorper vestibulum turpis. Pellentesque cursus luctus mauris.



# 5

## *Further ... MATLAB and data visualization*

This chapter is something of a potpourri of MATLAB data and visualization methodologies and techniques, generally building on the basics covered in Chapter 2.

## 5.1 Further data input

Previously, you imported ASCII data into **MATLAB** using the `load` command<sup>1</sup>. You might not have realized it at the time, but the use of `load` requires that your data is in a fairly precise format. **MATLAB** says "ASCII files must contain a rectangular table of numbers, with an equal number of elements in each row. The file delimiter (the character between elements in each row) can be a blank, comma, semicolon, or tab character. The file can contain MATLAB comments (lines that begin with a percent sign, %)." Firstly, your data may not be in a simple format and often may contain both numerical values and string values. Secondly, your data may not even be in a text/ASCII format. For instance, your data maybe be in an Excel spreadsheet, or for spatial scientific data, an increasingly common format is called 'netCDF' (Network Common Data Form). In this section, we'll go through the basics and some examples of each.

### 5.1.1 Formatted text (ASCII) input

The general procedure that you need to follow to input formatted text data is as follows:

1. First, you need to 'open' the file – the command (function) for this is called `fopen` (see Box). You need to assign the results of this function to a variable for later use.  
What is going on and why this all differs so much from using `load`, where you only had to use a single command, is that you first have to open a connection to the file ... before you even read any of the contents in(!)<sup>2</sup>.
2. Secondly ... you can read the content in (finally!). The complications here include specifying the format of the data you are going to read in. You also need to tell **MATLAB** the ID of the file that you have opened (so it knows which one to read from). The function you are going to use to do this is called `textscan`.
3. Close the file using `fclose` (see Box). You are going to have to pass the ID of the open file again when you call this function (so **MATLAB** knows which file to close).
4. Lastly, you are going to have to deal with the special data structure that **MATLAB** has created for you ...

If you are interested (probably not) – the connection made to an open file is called a file *pipe*. Typically, you have multiple open file *pipes* at the same time in programs, and this is why obtaining and then specifying a unique ID for the *pipe* you wish to read or write through, is critical.

<sup>1</sup> Or maybe 'cheated' and used the **MATLAB** GUI ...

#### opening and closing files

**MATLAB** has a pair of commands for opening and closing files for read/write:

- `fopen` will open a file. It needs to be passed the name (and path if necessary) of the file (as a string), and will return an ID for the file (assign (save) this to a variable – you'll need it!).
- `fclose` ... will close the file. It requires the ID of the file (i.e. the variable name you assigned the result of calling `fopen` to) passed to it as a parameter.

#### textscan

According to (actually, paraphrased from) **MATLAB**:

```
C = textscan(ID,format)
```

"... reads data from an open text file into a cell array, *C*. The text file is indicated by the file identifier, *ID*. Use `fopen` to open the file and obtain the *ID* value. When you finish reading from a file, close the file by calling `fclose(ID)`."

The ID part should be straightforward (if not – follow through the Example).

The format bit is the complicated bit ... There is some help in a following Box and via the Example. Otherwise, there is a great deal of details and examples in **MATLAB help** – you could look at this as a sort of menu of possibilities, and given a particular file import problem, the best thing to do is simply scan through help, looking for something that matches (or is close to) your particular data problem (and/or ask Google).

<sup>2</sup> This is very common across all(?) programming languages.

AS AN INITIAL EXAMPLE to illustrate this alternative (and more flexible) means of importing of data, we are going to return to the paleo atmospheric CO<sub>2</sub> proxy dataset file – `paleo_CO2_data.txt`<sup>3</sup>. Assuming that you have already (previously) downloaded it, open it up in a text editor and view it – you should see 4 neatly (ish) aligned columns of numeric values ... and nothing else<sup>4</sup>.

OK – so having seen the format of the data in the ASCII file, you are going to work through the following steps<sup>5</sup>:

1. First 'open' the file – you will be using the function command **fopen**, and passing it the filename<sup>6</sup> (including the path to the file if necessary). So that you can easily refer to the file that you have opened later, assign the output of `fopen`<sup>7</sup> to a variable, e.g.
 

```
» fopen_id = fopen('paleo_CO2_data.txt');
```
2. Now ... this is where it gets a trickier – the function you are going to use now is called `textscan`. Refer to **help** on `textscan`, but as a useful minimum, you need to pass 3 pieces of information:
  - (a) The ID of the open file (you have assigned this to a handy variable (`fopen_id`) already.)
  - (b) The *format* of the file (see margin note). (This is where it gets much less fun, but hang in there!) You simply list, space-separated, and between a single set of quotation marks, one format option per element of data.

In this particular Example, there are 4 items of data (per row) – each of them is an integer or a floating point number<sup>8</sup>, depending on how you want to look at it. Assuming that the data is a floating point number, the *format* for the input of each number item, is `%f`.

The result of `textscan` is then assigned to a parameter, e.g.

```
my_data = textscan(fopen_id, '%f %f %f %f');
```

3. So far, so good! And you can now close the file:
 

```
» fclose(fopen_id);
```
4. Actually, it does get worse before the end of the tunnel ... what `textscan` actually returns, i.e. your read-in data, is placed into an odd structure call a *cell array*. It is not worth our while worrying about just what the heck this is, and if you view it in the **Variables** window (i.e. double click on the `cell array` name in the **Workspace** window), it does not display the simple table of 4 columns of data that maybe you were expecting. For now, we can transform this format into something that we are more familiar with using the `cell2mat` function, e.g.

<sup>3</sup> The version that you have used before – not to be confused with a version ending in `.dat` that we will look at shortly ...

<sup>4</sup> This 'nothing else' is important as it is the reason why you were previously able just to load the data.

<sup>5</sup> You can start off working at the command line if you wish, but ultimately, you are going to need to put everything into an **m-file**.

<sup>6</sup> For convenience, you could assign the filename (+ its path) to a (string) variable and then simply pass the variable name – remember, no ' ' needed for a variable naming containing a string (whereas ' ' is needed for the string itself).

<sup>7</sup> The output is a simple integer index, whose value is specific to the file that you have opened.

#### According to **MATLAB help**:

*"the format is a string of conversion specifiers enclosed in single quotation marks. The number of specifiers determines the number of cells in the cell array C."* Take this to mean that you need one format specifier, per column of data. The specifier will differ whether the data element is a number or character (and MATLAB will further enable you to create specific numerical types).

The format specifiers are all listed under `help textscan`. However, your Dummies Guide to `textscan` (and good for most common applications) is that the following options exist:

```
%d - (signed)integer
%f - floating point number
%s - string
```

**MATLAB** will automatically repeat the format for as many lines as there are of data. Alternatively you can specify precisely how many times you would like the format repeated (and hence data read in).

<sup>8</sup> At least, none of them are clearly strings, right?

```
my_data_array = cell2mat(my_data);
```

And now ... it is done, i.e. there exists a simple array, of 4 columns, the first being the age (Ma), the second being the CO<sub>2</sub> concentration value (units of ppm), and the 3rd and 4th; minimum and maximum error estimates in the proxy reconstructed value. :)

AS A FURTHER EXAMPLE, we are going to process a more complicated version of the paleo atmospheric CO<sub>2</sub> proxy dataset. The file is called `paleo_CO2_data.dat` and is available from the course webpage. An initial problem here is even opening up the file to view it – if you use standard **Windows** editors such as **Notepad** it fails to format it properly when displaying its contents<sup>9</sup>. The first lesson then in scientific computing then is to have access to a more powerful/flexible editor than default/built-in programs such as **Notepad**. One good (**Windows**) alternative is **Notepad++**<sup>10</sup>. So go open the file with this instead<sup>11</sup>. Note the format – there are a bunch of header lines and moreover, some of the columns are not numbers (but rather strings). So even if you were to edit out the headers with comments (%)<sup>12</sup>, you are still left with the problem of mis-matched columns. You could edit the file in **Excel** to remove the problematic columns ... but now this seems like a real waste of time to be editing data formats with one software package just to get it into a second! (Again, you could use the **MATLAB** GUI import functionality ... but it will be a healthy life experience for you to do it at the command line :o) )

OK – so having gotten an idea of the format of the ASCII data file, you are going to work again through the 4 steps:

1. First 'open' the file as before (fopen) and assigned the ID returned by the function to a variable `openfile_id2`.
2. Call `textscan`. However, we now want to pass 3 pieces of information (compared to 2 before):
  - (a) The ID of the open file.
  - (b) The *format* of the data.
  - (c) And now – a parameter, together with an (integer) value, to specify how many rows of the file, assumed to be the header information, to skip.

(Again – the result of `textscan` is then assigned to a variable which will represent a *cell array*.)

Lets do the easy bit first – to tell **MATLAB** to skip *n* lines of a file, you add the parameter 'HeaderLines' to the list of parameters passed to `textscan`, and then simply tell it how many lines to skip. In this Example, you'd add:

MATLAB claims that a **cell array** is "A cell array is a data type with indexed data containers called cells. Each cell can contain any type of data. Cell arrays commonly contain pieces of text, combinations of text and numbers from spreadsheets or text files, or numeric arrays of different sizes." I am sort of prepared to believe this.

Basically, in object-oriented speak, a cell array is an object, or rather, an array of objects. As MATLAB hints – the cells can contain \*anything\*. Your limitation previously is that an array had to be all floating point numbers, all integers, or all strings, and if strings, all the strings had to be the same size. For strings in particular, it is obvious that a more flexible format where a vector could contain both 'banana' and 'kiwi' is needed (try creating a 2-element vector with these 2 words and see what happens). You clearly might also want to link a number with a string (e.g. number of bananas) in the same array, rather than have to create 2 separate arrays.

#### cell2mat

Having created this weird format (cell array), now MATLAB has to give you a way of converting the data inside into something more usable. The function is `cell2mat`, which for a cell array C:

```
A = cell2mat(C);
```

will return the corresponding ('normal') array A.

Now this is only true if all the data in C is of the same type (e.g. all floating point numbers). If the data types are mixed or you only wish for a sub-set of the data to be extracted and converted, simply index the required part of the cell array (Examples on this later).

<sup>9</sup> If you use a **Mac** (or **linux**) however, all text editors should display the content just fine.

<sup>10</sup> Conveniently installed on the Watkins computer lab computers.

<sup>11</sup> Right-mouse-button-click over the file, then select **Open with** and then click on **Notepad++**.

<sup>12</sup> Recall that **MATLAB** ignore lines starting with a % and this includes loading in data lines using `load`.



```
my_data = textscan(openfile_id2, ... , 'HeaderLines', 3);
```

OK – now to dive back into the MATLAB syntax mire ... Lets just load in just the first 2 columns of data, and assume that they are both integers (and skipping the first 3 lines of the file as per above). We might guess that we could simply write:

```
my_data = textscan(openfile_id2, '%d %d', 'HeaderLines', 3);
```

Try it (including closing the file, and a call to `cell2mat`, as before). What has happened?

It seems that MATLAB translates your format ('%d,%d') into: 'read in a pair of integers, and keep automatically repeating this, until something else is encountered'. That something else is sequence of characters at the end of the first data line (line #4, because we skipped the first 3), that makes **MATLAB** think that it has finished (or rather, that it cannot reading in 2 pairs of integers any longer). This leaves you with 2 pairs of integers – i.e. a  $2 \times 2$  matrix (as you'll see if you look at `my_data_array`).

Here is a solution – we could omit all the information following the first 2 elements (something for Google to help with).<sup>13</sup>:

```
my_data = ...
 textscan(openfile_id2, '%d %d %*[^\\n]', 'HeaderLines', 3)
```

3. Now close the file:

```
fclose(openfile_id);
```

4. And now convert the results to something more human-readable:

```
my_data_array = cell2mat(my_data);
```

This should do it – a simple array, of 2 columns, the first being the age (Ma) and the second the CO<sub>2</sub> concentration value (units of ppm). :)

There must be some sort of important life lesson hidden here. Perhaps about only working with well-behaved data files, or using the GUI import functionality? But hopefully it does illustrate that messy files can be dealt with, without the need for laborious editing or processing in **Excel**.

### 5.1.2 Importing ... Excel spreadsheets

If your data is contained in an Excel spreadsheet, and you want it in **MATLAB**, your options are:

1. Select some, or all, of the columns and rows in a specific worksheet, and either copy-paste this into a text file (but taking care that the worksheet column widths are formatted such that they

<sup>13</sup> This turns out to be specifying '%\*[^\\n]', which in effects sort of says: 'skip everything (all the fields) (%\*) up until the end of the line is found ([^\\n]).

are wider than the widest data element), or save in an ASCII format, with comma or tab delineations between columns. In either case, then load in the data using `load`, or if consisting of mixed numbers/text, go through the Hell that is `textscan` ....

2. Use **MATLAB** function `xlsread`.

So ... option #2 looks ... is looking the easiest ... :)

AS AN EXAMPLE, lets return to the paleo proxy CO<sub>2</sub> data again, but this time, as an Excel sheet. The data file you need is:

**paleo\_CO2\_data.xlsx**

(You may as well go load this into Excel just to take a look at the format and so subsequently, you'll know if you have imported it faithfully or not.)

From the help box on `xlsread`, it should be pretty apparent what you do. And in fact, I am going to leave you to work it out – try and import the age and CO<sub>2</sub> data (the numeric part of the data) from **paleo\_CO2\_data.xlsx**.

If you need to, you index a cell array, pretty well much like a normal array, except it has an alternative syntax. For a normal, numeric array `A`, you might write:

```
» A(4,3)
```

to reference the value in the 4th row, 3rd column. For a *cell array* `C`, to index the cell in the 4th row, 3rd column, you'd also write:

```
» C(4,3)
```

but you'd get a cell returned, not the value in the cell. If you want the value in the cell located at (4,3), you'd put the index in curly brackets:

```
» C{4,3}
```

and you'd get a value of 3000 returned in the example of `raw`.

### 5.1.3 Importing ... *netCDF* format data

Much of spatial, and particularly model-generated, scientific output, is in the form of *netCDF* files. This is a format designed as a common standard to facilitate sharing and transfer of spatial data, but in a way that enables e.g. a 'complete' description of dimensions and various types of meta-data to be incorporated along with the data. The format is platform independent and a variety of graphical viewers exist for viewing and interrogating the data. Most programming languages support the reading and writing of *netCDF* format data. **MATLAB** is no exception here.

#### **xlsread**

There are various uses (i.e. alternative allowed syntax) for `xlsread` for an Excel file with name `filename`. The 2 relevant and more useful ones look to be:

1. `num = xlsread(filename)`  
which will return the \*numeric\* data in the Excel file `filename` in the form of a matrix, `num`. Note that non-numeric (e.g. string) headers and/or columns, are ignored. Also note that `num` is a 'normal' numeric array and does not require any conversion.
2. `[num,txt,raw] = ...  
xlsread(filename)` will additionally return text data in a *cell array* `txt`, and \*everything\* in a cell array `raw`.

You can also specify a particular worksheet out of an Excel file to load in:

```
num = ...
xlsread(filename,sheet)
```

(and there are further refinements and options listed under **help**).

As per the previous subsection on data import, and indeed file read/write in programming languages in general – one opens a file and receives an ID for that file. The file can then be written to or read (including just interrogating its properties rather than necessarily extracting spatial data) using this ID. And of course, closed (using the ID). However, the *netCDF* standard is a little odd in how reading/writing is implemented and everything has to be done by determining the ID of a particular data variable or property of the file. As you'll see ...

The general approach for reading *netCDF* data is as follows:

1. Open the *netCDF* file by

```
ncid = netcdf.open(filename, 'nowrite');
```

where `filename` is the name of the *netCDF* file (which generally will end in `.nc`). `'nowrite'` simply tells **MATLAB** that this file is being open as read-only (this is the 'safe' option and prevents accidental deletion of over-writing of data).

2. This is the weird bit, as we cannot ask for the data we want automatically :o) Instead, given that we know<sup>14</sup> the name of the variable we want to access, we ask for its ID ...

```
varid = netcdf.inqVarID(ncid, NAME);
```

where `NAME` is the name of the variable (as a string), allowing us to then request the data:

```
data = netcdf.getVar(ncid, varid);
```

that says – assign the data represented by the variable `varid`, in the *netCDF* file with ID `ncid`, to the variable `data`.

So actually, not totally weird – you request the ID of the variable, then use that to get access to the data itself. The names of the **MATLAB** commands vaguely make sense in this respect – `inqVarID` for inquiring about the ID of a variable, and `getVar` for getting the variable (data) itself<sup>15</sup>.

3. Finally – close the file, by passing the ID variable into the function `netcdf.close`, i.e.

```
netcdf.close(ncid);
```

Note that you need to pass the ID of the *netCDF* file for each and every command (after `netcdf.open`) so **MATLAB** knows which *netCDF* object you are referring to.

<sup>14</sup> There are ways of listing the variables if not.

<sup>15</sup> It is beyond the scope of this course to worry about why in the case of *netCDF*, the function are all `netcdf.` something. Just to say, it involves objects and methods and is a common notation in object orientated languages (that nominally, **MATLAB** isn't).

---

FOR A *netCDF* EXAMPLE, we'll take the output of a low resolution Earth system model (**GENIE**). To start off, download the '2D marine sediment results' *netCDF* file – **fields\_sedgem\_2d.nc**. The data here

is relatively simple – a 2D distribution of bottom-water and surface sediment properties, saved at a single point in time. In other words, there are only 2 (spatial) dimensions to the data<sup>16</sup>.

OK – we'll start by opening the file (assuming that you have downloaded it!), remembering to assign its unique ID to some variable. Then, you'll want to get hold of (and assign to another variable), the ID of the variable we want to get hold of and plot – in this Example, it is called 'grid\_topo'. Having obtained the ID for this variable, you can then fetch it – assign it to a variable data. Then close the file.<sup>17</sup>

You should now have an array called data. It should be 36×36 in size. Why not plot it<sup>18</sup>. Can you guess what it might be? Is it in the correct orientation? (If not – correct it.)

Clearly what is missing are the x and y axis values, which you should have correctly deduced are longitude and latitude, respectively, with latitude presumably going from -90 to 90N, and longitude ... well, maybe it is not completely obvious exactly what the value of longitude is at the original.

A great strength of *netCDF* is the ability of this file format to also contain the grid (axis) details that the data is on. There are ways of finding out the names of the axis variables (dimensions), but for now, I'll give you them:

- 'lat' – is the latitude axis. (Technically, the axis values are the mid-points of the grid cells.)
- 'lon' – is the longitude axis.

The axes are held in the *netCDF* file as vectors and we can retrieve this (1D) data in a similar way to the 2D data:

```
varid = netcdf.inqVarID(ncid,'lat');
lat = netcdf.getVar(ncid,varid);
varid = netcdf.inqVarID(ncid,'lon');
lon = netcdf.getVar(ncid,varid);
```

in which we obtain the ID of the axis variable 'lat', then retrieve the axis data and assign it to a vector lat (and then likewise for longitude). Do this, and confirm that you get plausible vectors representing positions along a longitude and latitude axis.

The final task would then be to take the 2 axis vectors, and create a pair of matrices – one containing longitude values associated with the 2D data points, and one containing latitude values associated with the 2D data points. For this, you need to use *meshgrid*<sup>19</sup>. See if you can create the necessary lon/lat matrices and then plot the model topo data on its correct axes.<sup>20</sup>

The variable names of other data-sets that you might load and experiment with in terms of plotting function, color scale, and any

<sup>16</sup> Adding time would make it 3 dimensions (2 spatial + 1 of time). Adding height or depth in the ocean would also make it 3 (3 spatial). 3 spatial + time would make for a 4-dimensional dataset ...

<sup>17</sup> You should be able to do all of this without further hints – the sequence of commands and how they are used, is given in the introduction to this subsection.

<sup>18</sup> Your choice of 2D plotting function.

<sup>19</sup> See subsequent section.

<sup>20</sup> If you have flipped the data matrix around earlier when plotting, un-do this, or re-load the 2D data, or else the axes will no longer correspond to the data matrix orientation ...

other refinements that help visualise the data, include:

- 'ocn\_sal' – deep ocean salinity (units of per mil).
- 'ocn\_O2' – concentration of oxygen in bottom waters (units of  $\text{mol kg}^{-1}$ ).
- 'sed\_CaCO<sub>3</sub>' – % of calcium carbonate in surface sediments.

---

IN A RELATED *netCDF* EXAMPLE, we'll extend the problem to 3D – 2 spatial dimensions (longitude and latitude) and one of time. The file you need is called **fields\_biogem\_2d.nc**<sup>21</sup>. You are going to go through the same basic procedures of: opening the file, obtaining the variable ID, accessing the data using that ID, and closing the file. The name of the variable is called 'atm\_temp'. Create a script to do this all, calling the data array that you obtain by calling

```
netcdf.getVar(ncid,varid);
```

data3. How many dimensions does this array have? What are the lengths along each dimension? Can you guess which dimension of the 3 time is?

The name of the time axis variable is 'time', and you can access the times along this axis (i.e. the times at which the model saved a 2D spatial state) by:

```
varid = netcdf.inqVarID(ncid,'time');
times = netcdf.getVar(ncid,varid);
```

Ideally, you should be able, given the 3D array that you have obtained (from the data variable `atm_temp`), to specify and plot, the 1st model-projected surface air temperature distribution, as well as the last distribution. And given that the variables for latitude and longitude are also 'lat' and 'lon', you should be able to plot the temperature distribution with appropriate axes (and contoured).

You should also ... using `find`, be able to determine (and plot) the 2D data slice corresponding to the year (mid point) 1999.5.

Finally, test yourself and understanding to date, by creating an animation of how the surface air temperature in the model evolves over time.<sup>22</sup>

<sup>21</sup> The back-story is that this contains the 2D surface ocean and atmosphere fields from a model experiment in which the climate system was spun-up from rest and uniform values of everything, so as time progresses, the spatial patterns of the climate system start to evolve and stabilize.

<sup>22</sup> You have everything you need – the vector of times, and from this you can determine how many times there are and hence the number of iterations of a loop.

## 5.2 *Further data processing*

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetur id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

### 5.2.1 *Data interpolation*

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetur id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

### 5.3 Further (spatial / $(x,y,z)$ ) plotting

As you have seen earlier – the simplest possible way of taking a matrix of data values and plotting them spatially, as a function of  $(x,y)$  location, is the function `image`. In effect, this is treating your data as if it were an image – the data values being the ‘color’ of each pixel and the location in the matrix defining where in the image (row, column) the pixel is. The problem with this is that information regarding what is on the  $x$  and  $y$  axes is lost, be this distance, lat/lon, or some set of observed/experimental variables, or whatever. Instead, the points are evenly spaced on both axes. Moreover, the raw values are plotted and there is no possibility of interpolation/contouring or smoothing. One could regard scatter plotting as an improvement over this and a sort of  $x,y,z$  plotting, in as much as a 3rd dimension ( $z$  data value) can be represented through color and/or symbol shape and at times this can be quite effective. However, again, no interpolation/contouring or smoothing is possible with `scatter`.

For plotting true  $(x,y,z)$ /'3D' plots (i.e. data values in 2 spatial dimension), MATLAB provides a wide variety of more formal ways of plotting data spatially, including even the possibility of adding a 4th dimension representing the data value  $(x,y,z,zz)$  (see Box).

For a feel of what you should be able to learn to achieve using **MATLAB** – go to the following [webpage](#). In this data repository you can do things like re-plot with different longitude, latitude, and temperature ranges. Overlay the coastlines, and other useful things like that. You can also click through the different months of the year to get a feel for how the surface temperatures on Earth change with the seasons. Note that the graphic produced from this particular website is not particularly great, and you can all do better than this using **MATLAB** already. Presumably there are some lazy PhD students out there lacking the skills that you are (hopefully) learning. Perhaps they should take GEO111 (or maybe you are ...)?

---

AS AN EXAMPLE, load in the global topographic data file (**etopo1deg.dat**) from the course webpage. This is the height of the (solid) surface of the Earth relative to mean sealevel in meters, with the continents having a positive value and the ocean floor, negative. The data is conveniently on a  $1^\circ$  (longitude and latitude) grid. You could view the resulting elements of the 2D array in the Variable window if you like ... but at  $360 \times 180$  in size, there may not be much of use you can glean by visually inspecting the matrix<sup>23</sup>.

Try throwing the array into the `image` function see what happens (hopefully something like Figure 5.1). It had happened to come out

#### **$x,y,z$ PLOTTING**

**MATLAB** calls plots of a ( $z$ ) value as a function of both  $x$  and  $y$ , '3D'. Strictly, one could look at some of these functions as 2D, as `scatter` can plot a 3rd data ( $z$ ) value as different colors/shapes/sizes as a function of both  $x$  and  $y$  ... Anyway, the most commonly used/useful and fortunately simple, functions which create a 2D ( $x, y$ ) plot but with contours in the value of ( $z$ ), are:

1. **contour** – Plots a figure with the data contoured, with a range and increment between contours that is fully specifiable, color-coded or not, and labelled or not. Options are also provided for specifying how the contouring is done (and the data interpolated).
2. **contourf** – Similar to `contour`, except in between the (now simple black, by default) contours, a fill color is plotted and scaled to the data value.

For a genuine 3D plot, with surface height determined by the data in the 3rd dimension of the array, colors and/or contours in the data in the 4th array dimension, suitable functions include:

`surf`, `surfc`, `mesh`  
(but are not considered further here).

**imagesc** For a data array (matrix)  $A$ ,  
`imagesc(A)`

displays the data array as if a bitmap, but unlike `image` (see earlier), "uses the full range of colors in the colormap".

<sup>23</sup> More useful than are the summary details in the **Workspace window**, such as the apparent absence of NaNs and that the **Min** and **Max** Earth surface heights seem plausible.



displayed upsidedown<sup>24</sup>, then you'd need to flip the matrix upside-down using the command:

```
etopolddeg=flipud(etopolddeg);
```

and if the Earth instead appeared on its side<sup>25</sup>, you need to swap the rows and columns ( $x$  for  $y$  axis):

```
etopolddeg=etopolddeg';
```

It is not unusual for a first plotting attempt of spatial data to be incorrectly orientated and a little trial-and-error to get it straight is perfectly acceptable!

This is not exactly the prettiest of images. You can distinguish ocean (blue) from land (mostly brown, but other color pixels in places). Fortunately, MATLAB provides a variant of this plotting function, `imagesc`, that calculates the color scale to exactly span the min/max values in the data. Try it (and get something like Figure 5.2 hopefully).

The function `imagesc` also enables the range of data values the color range corresponds to, to be set. Refer to help on this function and see if you can plot just the above-sealevel, i.e. land surface heights, spanning zero (sealevel) to the maximum height<sup>26</sup>.

Which sort of in a round-about sort of way also brings us to how to set the color scale, which can be changed using the `colormap` command (see Box). Try out some different *colormaps* and re-plot the global topography data. What scales work well and what do not? Which scales help pick out details of e.g. ocean floor depth variation and which help pick out simple land-sea contrasts. Think about what one might want to highlight about global topography and what color scale might be best for this purpose?

---

STICKING WITH GLOBAL EARTH SURFACE TOPOGRAPHY, how else can we display the spatial data? For instance we might want to interpolate it, contour it, or simply get the longitude and latitude axes correct. Note that only by luck, because this particular dataset is 1 degree by 1 degree, the default axis scale in MATLAB when using `image` is approximately correct, although note that 'latitude' has been ordered in reverse and it goes from 1 to 180 rather than -90 to 90 ... We'll explicitly address this shortly.

To start with, you can simply use the `contour` function (see Box), passing only the matrix (of global topography values). Try this. Now you might want to think about flipping the matrix up-down, and/or left-right, as your plot should have come out looking like Figure 5.3.

Once you have fixed the orientation of the topography map, you might play about with the color scale (`colormap`) as before. You

<sup>24</sup> It doesn't in this particular case.

<sup>25</sup> Actually, in this example, it is OK in this respect too. Boring!




Figure 5.1: Very basic imaging (`image`) of an array (2D) of data – here, global bathymetry.




Figure 5.2: Slightly improved very basic imaging (`imagesc`) of bathymetry data.

<sup>26</sup> Don't forget the function `max`.

**colormap** MATLAB has a number of 'colormaps' built in – color scale that determine the colors that correspond to the data. The command to change the *colormap* from the default is:

```
» colormap NAME
```

where NAME is the name of the *colormap*. You can find a list of possible *colormaps* in `help` on `colormap` (in a table towards the bottom). But a brief summary is:

- **parula** – the current **MATLAB** default – chosen to provide a wide range of color and color intensity.
- **jet** – the old **MATLAB** default, but one which uses red and green in the same color, which should be avoided (why?).
- **hot**, **cool** – relatively simple color transitions but useful – hot is something like you'll see in publication figures.
- **pink** – another simple and at times useful transition and from dark (almost black) to white.

To return to the default *colormap*:

```
» colormap default
```

might also try the companion to `contour` – `contourf`. Re-orientating the matrix, switching to a different *colormap*, and plotting using `contourf`, might give you something like Figure 5.4.

OK, so a next refinement in plotting esp. maps and contour plots, is firstly to specify the range of the color scale, as we may not want the min-to-max range chosen by default by **MATLAB**, and the number of contours (e.g. in the topography example, they are pretty far apart and it is difficult to make out much detail). Both of these factors can be addressed simultaneously, by giving MATLAB a vector containing the value at which you want the contours drawn<sup>27</sup>.

Taking the global topography data – lets say you were interested only in low lying and shallow bathymetry, and wanted 20 contours intervals. Assuming a range in topographic height (relative to sealevel) of -1000 m to +1000 m, you should be able to deduce how to create the vector(?)<sup>28</sup>

Do this and check e.g. by opening up the vector in the **Variables window**. You should see the numbers from -1000 to 1000 in intervals of 100. Why, for instance, can you not simply write:

```
» v = [-1000:1000];
```

??? (Or rather: why might this not be a good idea ... ?)

Having created a specific vector of contours to plot, try it out. OK – so this is a little weird and maybe not so useful, but you get the point hopefully. So try plotting the following:

1. Just above sealevel topography, up to 10,000 m, in increments of 100 m.
2. Just the sealevel (coastline) contour ... trickier – create a vector with a value at zero, and a value either side – one very high and one very low. Use `contour` rather than `contourf`, although the latter produces a lovely land-sea mask!
3. Convert the data matrix of value in units of m, to ft, and plot the ocean floor (values equal to or below sealevel) in intervals of 1000 ft.
4. Finally – try some different color scales for the above. Think about which color scales best help illustrate the data, and whether `contour` or `contourf` is clearer. Also: how many contour intervals is 'best'? Your key is to make features clear, within the plot becoming cluttered or overly detailed.

The final refinement in contour plotting we'll look at is adding labels to the contours. The command to do this is `clabel` (for 'contour label') (see Box). Now, before anything, there is a slight complication. `clabel` needs to know details of the contours and graphics object with which to do anything with. For the purposes of this course,

<sup>27</sup> By default: MATLAB determines the minimum and maximum data values, and draws 10 equally spaced contours between these limits.

<sup>28</sup> If not, it is:

```
» v = [-1000:100:1000];
```

**contour** There are various uses of `contour`. The simplest is:

```
contour(Z)
```

where Z is a matrix. This ends up similar to `image` except with the data contoured rather than plotted as pixels (the 'similarity' here is that the x and y axis values are the number of the rows and columns of the data).

You can specify the values at which the contours are drawn, by passing a vector (v) of these values, e.g.

```
contour(X,v)
```

More involved and practical, is:

```
contour(X,Y,Z)
```

where X, Y, and Z, are all matrices of the \*same\* size (there is important). X and Y contain the x and y coordinate locations of y data values (contained in matrix Z). In the example of a map – X and Y contain the longitude and latitude values of the data values in Z.

Similarly, you can add a vector v containing the contours to be drawn, by:

```
contour(X,Y,Z,v)
```

you don't have to worry about the details of this, but simply need to know the following:

1. When you call `contour` (or `contourf`), 2 parameters are returned, which so far you have not cared about or even noticed. We now need them. SO when you call either plotting function, using the syntax:

```
[C,h] = contour(...)
```

which saves a matrix of data to `C`, and a ID (technically: graphics object 'handle') to `h`.

2. When you call `clabel`, pass these parameters back in, e.g.

```
clabel(C,h)
```

(in its most basic usage).

If you do this, in an earlier example of plotting just the zero height contour, and now using the most basic default usage of `clabel` (as above), you get, for good or for bad, Figure 5.5.

In the default usage of `clabel`, you'll get a label added on every contour that you plot. This ... can get kinda messy if you have lots and lots of contours plotted. You may well not need every single contour labelled, particularly if you also provide a color scale (see below). So you can also pass in a vector to tell MATLAB which contours to label. For example, if you have a contour interval vector:

```
v = [-1000:100:1000];
```

maybe you only want labels every 500m, so you'd use a vector:

```
w = [-1000:500:1000];
```

to specify the labelling intervals. The complete set of commands becomes:

```
>> v = [-1000:100:1000];
>> w = [-1000:500:1000];
>> [C,h] = contour(etopo1deg,v);
>> clabel(C,h,w);
```

Finally – missing from our color-coded plots so far, is a color scale to relate values to colors (although labelling the contours works as an OK substitute). The MATLAB command is simple:

```
>> colorbar
```

(and see Box for further usage). Try adding a *colorbar*, and in different places in the plot. Refer to the Box to try and add a caption to it ...




Figure 5.3: Example result of basic usage of the `contour` function.




Figure 5.4: Example usage of `contourf`, with the *hot colormap* (giving dark-/brown colors as deep ocean, and light/white as high altitude).




Figure 5.5: Example usage of `contour`, contouring only the zero height isoline, and providing a label.

#### **clabel**

```
>> clabel(C,h)
```

labels every contour plotted from  
`[C,h] = contour( ... );`  
 (or from `contourf`).

By prescribing and passing a vector `v` of contour intervals, you can label fewer/specific intervals rather than all of them (the default), e.g.

```
>> clabel(C,h,v)
```

IN THIS NEXT EXAMPLE, we'll address the issue with missing/incorrect lon/lat axis labels on the plots.

Each data point in the `etopo1deg` matrix should have one longitude value ( $x$ -axis) and one latitude ( $y$ -axis) value associated with it. It should hopefully be intuitive to you now ... that what we need is a pair of matrices, of exactly the same size as the `etopo1deg` data matrix – one holding longitude values and one latitude values. There are various ways of creating the required matrices 'by hand' (or involving writing a program including a *loop*). All of them are tedious. There is a **MATLAB** function to help. But it is not entirely intuitive<sup>29</sup> ... `meshgrid`.

Spend a few minutes reading about it in `help`. In particular, look at the examples given to help you translate the **MATLAB**-speak gobbledegook of the function **Description**. You should be able to glean from all this that this function allows us to create two  $a \times b$  arrays; one with the columns all having the same values, and one with the rows all having the same values (exactly what we need for defining the (lon,lat) of all the global data points). If not, and probably not – see Box. And then lets do a simple example (adapted from `help`):

```
» [X,Y] = meshgrid(1:3,10:14)
X =
 1 2 3
 1 2 3
 1 2 3
 1 2 3
 1 2 3
Y =
 10 10 10
 11 11 11
 12 12 12
 13 13 13
 14 14 14
```

Here, we are taking 2 vectors – `[1:3]` and `[10:14]`, and asking **MATLAB** (very nicely) to create 2 matrixes, one in which `[1:3]` is replicated down, until it has the same number of rows as the length of `[10:14]`, and one in which `[10:14]` is replicated across until it has the same number of columns as the length of `[1:3]`. (Try it.)

It'll become apparent *\*why\** bother shortly. Honest.

In our Example – start my noting that the topography data is on a regular 1 degree grid starting at 0° longitude. Latitude starts (at the bottom) at -90° and goes up to +90°). We need a matrix containing all the longitude values from 0° to 359° and latitude from -90° to 89°

#### colorbar

This almost could not be simpler:

```
» colorbar
```

plots the color scale! By default, it places it to the RH side of the plot. If you wish for it to appear anywhere else, use the modified syntax:

```
» colorbar(PLACEMENT)
```

where `PLACEMENT` is one of: 'northoutside', 'southoutside', 'eastoutside', 'westoutside'. Note that these are strings and so need to be in quotation marks. (More options are summarized in a table in `help`.)

Finally, you can also add a label to the `colorbar`, but only if you get hold of its ID ('graphics handle') when you call `colorbar`, e.g.

```
» h = colorbar
```

will save the graphics handle in variable `h`, which you can then muck about with via:

```
c.Label.String = 'The
units of my lovely
colorbar';
```

(Don't fight this – use this syntax to set a label for the `colorbar` – don't worry about what it means. **MATLAB** keeps rather annoyingly changing the way it does this anyway :())

<sup>29</sup> **DON'T PANIC!**

#### meshgrid

The unholy syntax is:

```
[X,Y] = meshgrid(xv,yv)
```

Pause, and take a deep breath. On the left – the results of `meshgrid` are being returned to 2 matrixes, `X` and `Y`. These are going to be our matrixes of the longitude and latitude values (in the particular example in the text). So far so good(?)

On the right, passed into the function `meshgrid`, are two vectors – `xv` and `yv`. Pause again.

What **MATLAB** is going to do, is to take the (row) vector `xv`, and it is going to replicate it down so that there are as many rows as in the vector `yv`. This becomes the returned output matrix `X`. **MATLAB** then takes the column vector `yv`, and replicates it across so that there are as many columns as in the vector `xv`. This becomes the returned output matrix `Y`.

<sup>30</sup> These matrices need to be the same size as the data matrix.

Maybe just 'do' it and then understand what has happened after. Create the longitude and latitude grids by:

```
» [lon lat] = meshgrid([0:359],[-90:89]);
```

View (in the **Variables window**) the lon matrix first. Scan through it. Hopefully ... you'll note that it is 360 columns across, and in each column has the same value – the longitude. The matrix is 180 rows 'high', so that there is a longitude value for each latitude. Similarly, view lat. This also should make a little sense if you pause and think about it, with the one exception that the South Pole latitude is at the 'top' of the matrix – don't worry about this for now ...

The only way to fully make sense of things now, is to use it. Remember that use of `contour` (and `contourf`) can take matrices of  $x$  and  $y$  (here: longitude and latitude) values that correspond to the data entries in the data matrix (`etopo1deg`). Re-load the topography data in case you have flipped it about in all sorts of odd ways, and then do:

```
» [lon lat] = meshgrid([0:359],[-90:89]);
» contour(lon,lat,etopo1deg);
```

Almost! Note that the  $x$  and  $y$  axis labelling is 'correct' and particularly the  $y$ -axis, where latitude goes from -90 to 90 (although by default **MATLAB** labels in intervals of 20 starting at -80 it seems). But it also turns out that we do need to flip the data top-side-down. We can actually do this in the same line as we plot:

```
» contour(lon,lat,flipud(etopo1deg));
```

Phew! (Figure Figure 5.6.)

The final complication is that the data points in the gridded dataset (matrix `etopo1deg`), technically correspond to the mid-points of a 1 degree grid, not the corners. So if we were going to try and be formally correct<sup>31</sup>, our vectors that we'd pass into `meshgrid`, would be:

```
» xv = [0.5:359.5];
» yv = [-89.5:89.5];
```

---

OK – ANOTHER EXAMPLE on this. Previously, you downloaded and plotted monthly global distributions of surface air temperature. You plotted these simply using `pcolor` (or `image`) and the results were ... variable. Certainly not publication-quality graphics and missing appropriate longitude and latitude axes for the plots.

<sup>30</sup> There is a slight complication with this, which we'll get to shortly, but note that the data array is 360 elements ( $x$ -direction) by 180 elements ( $y$ -direction).




Figure 5.6: Usage of `contour` but with lon/lat values created by `meshgrid` function and passed in (and with the `hot colormap` (giving dark/brown colors as deep ocean, and light/white as high altitude).

<sup>31</sup> Don't worry about this for now – grids will be covered more in subsequent chapters surrounding numerical (environmental) models.

Make a copy of your original *script* (**m-file**) in which you created the animation, and give it a new name. Edit your program, and in place of `pcolor`, use `contour` or `contourf` (your choice!). Pass in just the data matrix (of monthly temperature) when calling the `contour(f)` function and don't yet worry about the lon/lat values. Get this working (i.e. debug it if not). You should end up with a contoured animation (rather than a bit-map animation).

The problem with the axis labelling should be much more apparent (than compared to the topography data, which was on a handy 1 degree grid already). So you need to make a matrix of longitude values, and one of latitude. using `meshgrid`. The grid is a little awkward:

1. The longitude grid runs from 0°E (column #1) with an increment of 1.875°; i.e., 0.000°E, 1.875°E, 3.750°E, ... up to 358.125°E (column #192).
2. Latitude runs from 88.54196°S (-88.54196°N) at row #1, to 88.54196°N (row #94) with an increment of about 1.904.

so I'll give you the answer up-front:

```
» lonv = [(1.875/2):1.875:360-(1.875/2)];
» latv = [-90+(1.904/2):1.904:90-1.904];
» [lon lat] = meshgrid(lonv,latv);
```

Now use the longitude and latitude values matrices, in conjunction with `contour(f)`, to plot the global temperature distributions 'properly'. Try plotting just one plot first, before looping through all 12 months.

At this point (before creating an animation), you might also explore some of the plotting refinements we saw earlier. For example, as per Figure 5.7. Firstly – get the units of the temperature data array into units of °C or °F rather than °K. Either: assign the `temp` array data to a new array and make the appropriate conversion from °K (all within the loop), or you can do this subtraction on the line that you actually plot the data (i.e., within the `contour/contourf` function), for example:

```
contourf(lon(:,:,month),lat(:,:,month),temp(:,:,month)-273.15);
```

would convert to °C as it plotted the data.

You can also get the plotting temperature limits and contouring consistent between months and with greater resolution by adding the following line (before the loop starts):

```
v=[-40:2:40];
```

and then to the `contour(...)` (or `contourf(...)`) function, add `,v` to the end of the list of passed parameters. This particular choice for the vector `v` tells MATLAB to do the contouring from -40 to 40 (°C), and



at a contour interval of 2 (°C).. Play around with the min and max limits of the range, and also with the contour interval to see what gives the clearest and least cluttered plot. For instance, maybe you don't want the low temperatures to go 'off' the scale (the white color in the filled contour plot).

### 5.3.1 Plotting maps

You can do some nice spatial plotting with this data using the **MATLAB Mapping Toolbox**. This should be available as part of the **MATLAB** installation in the Lab (and also if you have downloaded and installed an academic version on a personal laptop). Refer to the on-line documentation for the **Mapping Toolbox** to get you started. The key function appears to be **geoshow**. Try plotting the region encompassing the 'quake data, with a coastal outline (of land masses), and the 'quake data overlain. Explore different map projections. Remember to always ensure appropriate labelling of plots.




Figure 5.7: Example contour plot including meshgrid-generated lon/lat values. Result of `contourf(lon,lat,temp7,30)`, where the data file was `temp7.tsv`, with some embellishments.





# 6

## *Further ... Programming*

In this chapter we'll get some (more) practice building programs and crafting (often) bite-sized chunks of code that solve a specific, normally computational or numerical (rather than scientific) problem (*algorithms*) <sup>1</sup>.

<sup>1</sup> According to the all-mighty Wikipdeia (and who am I to argue?) – an "algorithm ... is a self-contained step-by-step set of operations to be performed. Algorithms perform calculation, data processing, and/or automated reasoning tasks."

## 6.1 *find!*

So – a single **MATLAB** function gets a high-level section, all to itself. Either it's really powerful and useful, or I am running out of ideas for the text<sup>2</sup>.

`find ...` finds where-ever in an array, a specific condition is met. If the specific condition occurs once, a single array location is returned. The specific condition could occur multiple times, in which case `find` will report back multiple positions in the array.

What do I mean by a 'specific condition'? Basically – exactly as per in the `if ...` construction – a conditional statement being evaluated to true.

OK – some initial Examples.

Lets say that you have a vector of numbers, e.g.:

```
A = [3 7 5 1 9 7 4 2];
```

and you want to find the maximum value in the vector – easy<sup>3</sup>

But ... you want to find \*where\* in the vector the maximum value occurs. Why might you want to do this? Rarely do you have a single vector of data on its own – generally it is always linked to at least one other vector (often time or length in scientific examples). Trivially, our second vector might be:

```
B = [0:7];
```

and is time in ms. The question then becomes: at what time did the maximum value occur? Obviously, this is easy by eye with just 8 numbers, but if you had 1000s ...

We can start by determining the maximum value.

```
c = max(A);
```

Now, we use `find` to evaluate where in the array `A` (here: a vector) the element with a value of `max(A)` occurs, or where the condition `d == c` is true, where `d` is the element in question (the maximum value). So:

```
find(A(:) == c);
```

should do it. Here, what we are saying is: take all of the elements in `A` and find where an element occurs that is equal to `c` (the maximum value which we already determined). Try it, and MATLAB should return 5 – the 5th element in the vector.

Finally, if we assign the result of `find` to `d`, we can then use `d` to determine the time at which the value of 9 occurred, i.e. `B(d)` which evaluates to 4 (ms):

In this example, `find` returned just a single element, but if we instead had:

<sup>2</sup> It is really powerful and useful.

### **find**

MATLAB defines `find`, with a basic syntax of:

```
k = find(X)
```

as 'return[ing] a vector containing the linear indices of each nonzero element in array `X`'. That means ... nothing to me. This is going to have to be a job for some Examples ... (in order to see what `find` is all about).

<sup>3</sup> I hope so ... check back earlier in the course on `max`.

```
A = [3 9 5 1 9 7 4 2];
```

The maximum value is still the same (9) but now ...

```
» find(A(:) == c)
ans =
 2
 5
```

What has happened is that `find` has determined that there are 2 elements in vector `A` that satisfy the condition of being equal to `c` (9) and these lie at positions (index) 2 and 5. The result vector, if you assigned it to the variable `d` again, can be used just as before to access the corresponding times in vector `B`;

```
» d = find(A(:) == c); » B(d)
ans =
 1 4
```

i.e. that the times at which the values of 9 occur are 1 and 4 (ms).

Any of the relational operators (that evaluate to *true* or *false*) can be used. In fact – looking at it this way leads us to maybe understand the **MATLAB help** text, because *true* and *false* are equivalent to 1 and 0, and `find` is defined as a function that returns the indices of the non-zero elements in a vector. By writing `A(:) == c` we are in effect creating a vector of 1s and 0s depending on whether the equality is *true* or not for each element. You can pick apart what is going on and see that this is the case, by typing:

```
» A(:) == c
ans =
 0
 1
 0
 0
 1
 0
 0
 0
```

(the statement being *true* at positions (index) 2 and 5, which is exactly what `find` told you).

For instance, we could ask `find` to tell us which elements of `A` have a value greater than 5:

```
» find(A(:) > 5)
ans =
 2
 5
 6
```

(Inspect the contents of vector A and satisfy yourself that this is the case.)

We can also use `find` to filter data. Perhaps you do not want values over 5 in the dataset. Perhaps this is above the maximum reliable range of the instrument that generated them. Having obtained a vector of locations of these values, e.g.

```
d = find(A(:) > 5);
```

we can plug this vector back into A and assign arrays of zero size to these locations – effectively, deleting the locations in the array, i.e.

```
A(d) = [];
```

They it, and note that the size<sup>4</sup> of A has shrunk to 5 – all the other elements remain, and in order, but the elements with a value greater than 5 have gone. You could apply an identical deletion (filtering) to the time array (`B(d) = []`).

Play about with some other relational operators and criteria, and make up some vectors of your own until you are comfortable with using `find`.

<sup>4</sup> Use the command `length` or view in the **Workspace** Window.

---

BACK TO THE 'QUAKE EXAMPLE: Find<sup>5</sup> how many earth quakes there were bigger than M = 8? Also determine how many quakes occurred bigger than M = 7, 6, 5, 4, and 3. Determine the day on which the magnitude 8.7 shock occurred.

<sup>5</sup> Intentional joke \*and\* clue.

In the first problem (number of quakes greater than a specified limit) – you need ask `find` to return the row numbers for all quakes satisfying the condition: `magnitude > 8.0`. `find` will return you a column vector. You don't actually need to worry about or access the contents of the vector, you just need to know how many elements there are in the vector (because there will be one element for each occurrence of `magnitude > 8.0`). This is the same as its `length` (see earlier and/or **help**).

In the second problem – you need to find the row number of the quake magnitude data which satisfies the condition: `magnitude > 8.7`. Knowing the row number, you can then access the data column containing the sate information, and hence extract the day and solve the problem.

All these problems can actually be solved in a single line of **MATLAB**, but feel free to break it down into multiple steps.

---

IN THE SEALEVEL (OXYGEN ISOTOPE) EXAMPLE, you could start by determining the maximum and minimum sea-levels that have occurred

over the last 782,000 years. Then ... because it would be helpful to know *when* the minimum and maximum sea-level heights occurred, use the `find` function to find the data row in which the minimum and maximum values occurred. Once you know the respective data rows, you can then easily pull out the ages.<sup>6</sup> Find the ages of both minimum and maximum values.

Also find all the occasions (times) on which sealevel was higher than today (modern). (Or equivalently, when the oxygen isotope value, that we are assuming directly reflects changes in level, was lower than modern<sup>7</sup>.)

You can also ask questions based in time, such as what was the sealevel (or oxygen isotope value) at 21 ka (i.e. without having to look through the data manually and determine on which row 21 ka occurs, because this is exactly what `find` can do this for you)? This can be particularly useful if the value of time is calculated or passed in from elsewhere, rather than specified as e.g. 21 ka, because you may not *a priori* know what the value will be, hence automating the script with `find` is super useful. Effectively then you are creating an *algorithm* for taking a time input and determining sealevel.

---

FOR AN EXAMPLE OF DATA-FILTERING – dig out the paleo-proxy (not ice-core) atmospheric CO<sub>2</sub> data you downloaded. One further way of plotting with `scatter` is to scale the point size by a data value. We could do with by:

```
SCATTER(data(:,1),data(:,2),data(:,2))
```

... except ... it turns out that there are atmospheric CO<sub>2</sub> values of zero or less and you cannot have an area (size) value of zero or less ...

This leads us to a new use for `find` and some basic data filtering. The simplest thing you could do to ensure no zero values, would be to add a very small number to all the values. This would defeat the 'no zero' parameter restriction, but would not help if there were negative values and you have now slightly modified and distorted the data which is not very scientific. Substituting a NaN for problem values is a useful trick, as MATLAB will simply ignore and not attempt to plot such values.

So first, lets replace any zero in the CO<sub>2</sub> column of the data with a NaN. The compact version of the command you need is:

```
data(find(data(:,2)==0),2)=NaN;
```

But as ever – perhaps break this down into separate steps and use additional arrays to store the results of intermediate steps, if it makes it easier to understand, e.g.

<sup>6</sup> HINT – if your maximum value was stored in the variable `max_value`, you found find the corresponding row by:  
`find(data(:,2) == max_value)`

What this is saying, is search the 2nd column (the sea-level values) of the array `data`, and look for a match to the value of `max_value`. The equality operator (`==`) is used in this context.

<sup>7</sup> Lower  $\delta^{18}\text{O} \Rightarrow$  less ice volume  $\Rightarrow$  higher sealevel.

#### NaN

... is **Not-a-Number** and is a representation for something that cannot be represented as a number, although if you try and divide something by zero **MATLAB** reports Inf rather than a NaN.

NaN can also be used as a function to generate arrays of NaNs. The most common/usage in this context is:

```
N = NaN(sz1,...,szN)
```

which will (according to **help**) "generate a `sz1`-by-...-by-`szN` array of NaN values where `sz1`,...,`szN` indicates the size of each dimension. For example, `NaN(3,4)` returns a 3-by-4 array of NaN values."

```
list_of_zero_locations = find(data(:,2)==0);
data(list_of_zero_locations,2) = NaN;
```

What this is saying is: first find all the locations (rows) in the 2nd column of data which are equivalent (==) to zero. Set the CO<sub>2</sub> value in all these rows, to a NaN (technically speaking: assign a value of NaN to these locations). You have now filtered out zeros, and replaced the offending values with a NaN and when **MATLAB** encounters NaNs in plotting – it ignores them and omits that row of data from the plot.

Alternatively, we could have simply deleted the entire row containing each offending zero. Breaking it down, this is similar to before in that you start by identifying the row numbers of where zeros appear in the 2nd column, but now we set the entire row to be 'empty', represented by []:

```
list_of_zero_locations = find(data(:,2)==0);
data(list_of_zero_locations,:) = [];
```

If you check the **Workspace window**<sup>8</sup>, you should notice that the size of the array data has been reduced (by 4 rows, which was the number of times a zero appeared in the 2nd column).

We are almost there with this example except it turns out that there is a CO<sub>2</sub> proxy data value less than zero(!!!) We can filter this out, just as for zeros. I'll leave this as an exercise for you<sup>9</sup> ... The plot should end up looking like Figure 6.1. As another lesson-ette, given that the circles are insanely large ... try plotting this with proportionally smaller circles<sup>10</sup>.

As a last (optional) exercise on this ... In the CO<sub>2</sub> data, there are min and max uncertainty limit values. One could color-code the points in a scatter-plot to represent either the min or the max (perhaps try this first), but one on its own is not necessarily much use. One could color-code by the difference, but this is a function of the absolute value and one would expect large uncertainty bars if the mean (central) estimate was high, and lower if it were low. Perhaps we need the *relative* range in uncertainty? Can you do this? i.e., scatter-plot the mean CO<sub>2</sub> estimate (as a function of time), but color-coding for the range in uncertainty as a proportion of the value?

It turns out this is not entirely trivial because as you have seen, the data is not as well behaved as you might have hoped. In fact, it is just like real data you might encounter all the time! Before you do anything – break down into small steps what you need to do with the data, as this will inform what (if any) additional processing you might have to carry out on the data. It should be obvious, that to create a CO<sub>2</sub> difference, *relative* to the mean, you are going to have to divide by the mean value (column #2 in the array). So first off –

<sup>8</sup> Or:

```
» size(data)
```

<sup>9</sup> But you might e.g. use <=.

<sup>10</sup> HINT: you are going to want to apply a scaling factor to the vector you passed as the point size data.




Figure 6.1: Proxy reconstructed past variability in atmospheric CO<sub>2</sub> (scatter plot).

if any of the mean values are zero, it is all going to go pear-shaped. Actually, equally unhelpful, or at least, lacking in any meaning, may be negative values. If you inspect the data (in the **Variable window**), there are both zeros and negative values for mean CO<sub>2</sub> proxy estimates. We need to get rid of these. Follow the steps as before. You may also have to process the min and max values should they turn out to be the same. Likely you are going to have to delete all the rows in which (1) column #2 values are zero or below, and (2) column #3 and #4 values are equal (you could also try the NaN substitution and see if it works out). (If you need a slight hint ... one possible answer is here<sup>11</sup>, but try and work it out for yourself.)

All that is missing now, is any indication of what the color scale actually means in terms of values (and of what). MATLAB will add a colorbar to a plot with the command ... `colorbar`. Although the color scale gets automatically plotted with labels for the values, looking at the plot, we still don't know what the values are of (e.g. units). We can label the colorbar, but MATLAB needs to know what we are labelling. Each graphic object is assigned a unique ID when you create them and which normally you know nothing about. We can create a variable to store the ID, and then pass this ID to MATLAB to tell it to create a title for the colorbar. To cut a long story short:

```
colorbar_id=colorbar;
title(colorbar_id,'Relative error (%)');
```

It should end up looking something like Figure 6.2 in which you can see the high relative uncertainty (bight colors) prevail at low CO<sub>2</sub> values and 'deeper time' (ca. 200-300 Ma). The colorbar title (label) is maybe not ideal, nicer would be one aligned vertically rather than horizontally. We'll worry about that sort of refinement another time.

<sup>11</sup> In this possible solution – all rows in the array data, with mean CO<sub>2</sub> values less than or equal to zero, are deleted. Also, all rows for which the max and min values are the same, are also deleted.

```
» data=load('paleo_CO2_data.txt',
...'-ascii');
» data(find(data(:,2)<=0),:)=[];
» data(find(data(:,3)==data(:,4)),:)=[];
» data=find(data(:,3)<data(:,4));
» scatter(data(:,1),data(:,2),40,
...100*(data(:,4)-data(:,3))./data(:,2),
... 'filled');
» xlabel('Time (Ma)');
» ylabel('Atmospheric CO2 (ppm)');
» title('Proxy atmospheric CO2')
```




Figure 6.2: Proxy reconstructed past variability in atmospheric CO<sub>2</sub> (scatter plot).





## *Bibliography*

Stormy Attaway. *Matlab (Third Edition): A Practical Introduction to Programming and Problem Solving*. Butterworth-Heinemann, 2013.



# Index

- .mat environment, 41
  - ; environment, 26
  - = environment, 25, 26
- addition environment, 26
- addpath environment, 40
- and environment, 27
- assignment operator environment, 27
- axis environment, 34
- break environment, 74
- cell array environment, 88
- cell2mat environment, 87, 88
- clabel environment, 98, 99
- clear all environment, 28
- clear environment, 28
- close environment, 28
- colon operator environment, 30–32, 37
- colorbar environment, 100, 111
- colormap environment, 97
- Command Window, 22
- comments environment, 88
- contour environment, 96, 98
- contourf environment, 96
- disp environment, 52, 53, 71
- division environment, 26
- else environment, 60
- elseif environment, 60
- end environment, 31, 32
- environments
  - .mat, 41
  - ;, 26
  - =, 25, 26
  - addition, 26
  - addpath, 40
  - and, 27
  - assignment operator, 27
  - axis, 34
  - break, 74
  - cell array, 88
  - cell2mat, 87, 88
  - clabel, 98, 99
  - clear, 28
  - clear all, 28
  - close, 28
  - colon operator, 30–32, 37
  - colorbar, 100, 111
  - colormap, 97
  - comments, 88
  - contour, 96, 98
  - contourf, 96
  - disp, 52, 53, 71
  - division, 26
  - else, 60
  - elseif, 60
  - end, 31, 32
  - equality, 27
  - exist, 75, 78
  - exit, 28
  - exponentiation, 26
  - fclose, 86
  - figure, 33
  - find, 106, 107, 109
  - fliplr, 31, 38
  - flipud, 31
  - flipup, 38
  - fopen, 86–88
  - for, 66
  - fprintf, 41
  - functions, 28
  - geoshow, 103
  - getframe, 73
  - greater than, 27
  - greater than or equal to, 27
  - hist, 49
  - hold, 46
  - if ... end, 60
  - image, 50, 96
  - imagesc, 96
  - imread, 50
  - inequality, 27
  - input, 60, 61, 78
  - legend, 46
  - length, 31, 80
  - less than, 27
  - less than or equal to, 27
  - load, 40–42
  - ls, 40
  - m-file, 52
  - m-files, 34
  - meshgrid, 100
  - movie2avi, 73
  - multiplication, 26
  - NaN, 109
  - not, 27
  - num2str, 72
  - or, 27
  - pcolor, 50, 79
  - pi, 28
  - plot, 33
  - print, 35
  - rotate, 38
  - save, 41
  - scatter, 33, 47
  - sin, 34
  - size, 31, 37
  - sort, 42
  - sortrows, 42
  - strcmp, 61
  - subplot, 35
  - subtraction, 26
  - sum, 38

- switch ... case ... end, 65
- textscan, 86–88
- title, 34
- transpose, 38
- transpose operator, 31
- while, 66
- xlabel, 34
- xlsread, 90
- ylabel, 34
- equality environment, 27
- exist environment, 75, 78
- exit environment, 28
- exponentiation environment, 26
- fclose environment, 86
- figure environment, 33
- find environment, 106, 107, 109
- fliplr environment, 31, 38
- flipud environment, 31
- flipup environment, 38
- fopen environment, 86–88
- for environment, 66
- fprintf environment, 41
- functions environment, 28
- geoshow environment, 103
- getframe environment, 73
- greater than environment, 27
- greater than or equal to environment, 27
- hist environment, 49
- hold environment, 46
- if ... end environment, 60
- image environment, 50, 96
- imagesc environment, 96
- imread environment, 50
- inequality environment, 27
- input environment, 60, 61, 78
- legend environment, 46
- length environment, 31, 80
- less than environment, 27
- less than or equal to environment, 27
- license, 2
- load environment, 40–42
- ls environment, 40
- m-file environment, 52
- m-files environment, 34
- meshgrid environment, 100
- movie2avi environment, 73
- multiplication environment, 26
- NaN environment, 109
- not environment, 27
- num2str environment, 72
- or environment, 27
- pcolor environment, 50, 79
- pi environment, 28
- plot environment, 33
- print environment, 35
- rotate environment, 38
- save environment, 41
- scatter environment, 33, 47
- sin environment, 34
- size environment, 31, 37
- sort environment, 42
- sortrows environment, 42
- strcmp environment, 61
- subplot environment, 35
- subtraction environment, 26
- sum environment, 38
- switch ... case ... end environment, 65
- textscan environment, 86–88
- The command line, 22
- title environment, 34
- transpose environment, 38
- transpose operator environment, 31
- variable, 24
- while environment, 66
- xlabel environment, 34
- xlsread environment, 90
- ylabel environment, 34