

Econ 520: An Introduction to Matlab

Keisuke Hirano

October 10, 2005

In this course you will be using Matlab to perform numerical computations for the problem sets. Matlab is a matrix programming language: a variable is considered a matrix by default, and matrix operations (such as matrix multiplication, transpose, inverse, etc.) are built in. Since it is a programming language, it gives you a lot of control over its calculations, but you must be rather precise in telling it what to do. This short guide is intended to help you get started using Matlab, but it only scratches the surface of what you can do with it. Much more information can be found in the Matlab manual, which also contains a very useful tutorial section.

1 Getting Started

If you are using the Windows version of Matlab, many operations can be done using the pull-down menus. However, for all platforms, you can type commands directly to the Matlab interpreter at the `>>` prompt.

A very useful command is `help`. If you just type `help` and hit return, you will be presented with a menu of help options. If you type `help command`, Matlab will provide help for that particular command. For example, `help hist` will call up information on the `hist` command, which produces histograms. In addition, you can use the `Help` menu at the top of the Matlab window to get the `Table of Contents` and `Index` of commands.

You might have to change the working directory (using the command `cd`) to tell Matlab where your files are (or will be) located. The commands `pwd` and `dir` may also be useful for file management.

To end a Matlab session, type `exit` or `quit`.

2 Working with Variables

There are two main things to remember:

First, Matlab is case sensitive. So the variable `X` is not the same as `x`.

Second, variables in Matlab are treated as matrices. For example, the command $X*Y$ is a matrix multiplication, and works only if the two matrices are conformable. To perform *array* multiplication, type $X.*Y$. This multiplies the (a,b) element of X with the (a,b) element of Y for all (a,b), and so requires the two matrices to be of the same size. X' means the transpose of X .

Suppose a and b are scalars (1×1 matrices). You can refer to the (a,b) element of the matrix X by $X(a,b)$. In a similar vein, $X(:,b)$ refers to the b'th column of X , and $X(a,:)$ refers to the a'th row of X . You can get even fancier: $X(1:10,b)$ refers to the 10 by 1 vector comprised of $X(1,b)$ through $X(10,b)$. Suppose that i is a $m \times 1$ vector of integers between 1 and n (repeats are ok), and X is an $n \times k$ matrix. Then $X(i,:)$ is the $m \times k$ matrix formed by selecting rows of X according to i .

You can also perform logical indexing of a matrix. If i is an $n \times 1$ vector of logical (true-false) values, X is an $n \times k$ matrix, $X(i,:)$ returns the matrix formed by only including the rows of X that correspond to a value of 1 in the vector i . As an example of a logical vector, suppose that a is an $n \times 1$ vector of numbers. The expression $(a > 0)$ returns an $n \times 1$ logical vector; the i 'th element of $(a > 0)$ is equal to `true` if $a(i) > 0$ and is equal to `false` otherwise. A vector of ones and zeros can be turned into a logical vector by using the `logical` command.

For small matrices, we can define them by typing something like:

```
a=[1 2 3; 4 5 6; 7 8 9]
```

which will put the following matrix in Matlab's memory:

$$a = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}.$$

To define a new matrix it will sometimes be useful to use a command along the lines of

```
X = [ones(n,1) z];
```

This will create a matrix X by putting an $n \times 1$ vector of ones in front of the matrix z . The variable z must have exactly n rows.

Note that some commands, including `sqrt`, `log`, and `exp`, perform calculations element-by-element. (There are different commands for the matrix counterparts, but you probably won't be needing those.)

3 Getting Stuff Done

There are two ways to get things done in Matlab.

1. Interactive Mode
2. Batch Mode

3.1 Interactive Mode

In this approach, you just sit at the computer, type a command at the `>>` prompt, then, when it is done, type another command, and so on, until you have the answer you want. There are a couple of things to keep in mind when doing this.

First, if you add a semicolon (`;`) at the end of a command before hitting the return key, Matlab will perform the calculation but not return the result to the screen. This can be very useful. For example, suppose `x` is a 1000 by 1 vector of numbers, and you want to create a new vector `y` as follows: `y = 3 + 2*x`. This will take each element of `x`, double it, and add 3. If you don't include the semicolon at the end, Matlab will actually list the entire `y` vector it calculates on the screen.

Second, as you go along making calculations and defining new variables as in the previous example, Matlab will keep these variables in its memory. To see which variables are currently in Matlab's memory, type `who`. The alternative command `whos` lists the variables in memory plus some additional information about the variables.

Third, you may want to keep track of the operations you perform and the things that get printed to the screen. To do this, type `diary filename`. Anything that appears on the Matlab screen after you issue the `diary` command will also get written to the text file `filename`, until you type the command `diary off`. If there is already a file called `filename` in the working directory, the screen output will get appended to the end of that file.

Fourth, after performing a set of calculations, you may end up with some possibly large variables that you wish to save for later use. The command `save filename y Y X` will save the variables `y Y X` to the file `filename.MAT`. The `.MAT` file is a Matlab storage file. It is a binary file, not a text file, so you cannot look at the contents with a text editor. Instead, you will have to issue the Matlab command `load filename`, to tell Matlab to read the file and put it back in its memory.

Below are some other commands you may find useful. Remember, you can always use the help facility or consult the manual to see how to use these and any other commands.

<code>randn</code>	<code>log</code>	<code>clear</code>	<code>sqrt</code>
<code>rand</code>	<code>exp</code>	<code>size</code>	<code>plot</code>
<code>if</code>	<code>ones</code>	<code>length</code>	<code>hist</code>
<code>for</code>	<code>zeros</code>	<code>mean</code>	<code>min</code>
<code>while</code>	<code>eye</code>	<code>median</code>	<code>max</code>
<code>kron</code>	<code>diag</code>	<code>cov</code>	<code>sort</code>

```
chol          cd          std          find
inv           dir         pwd
```

3.2 Batch Mode

You can automate things by typing in all the commands you want to perform, in order, in a text file that has the extension `.m`. (Such files are generally called “batch files.”) In Matlab, these are called “M-files,” and you should use them extensively when working on the problem sets. Suppose you write the following program in an M-file called `prog.m`:

```
% A short sample program
x=randn(100,1);
y=3+2*x;
hist(y,20)
```

At the Matlab prompt, you would type `prog` and hit return. Matlab would read `prog.m`, then perform the commands in order: first generate 100 independent standard normal draws, then multiply those by 2 and add 3, then graph the histogram of these draws (using 20 bins). The first line is ignored because it starts with `%`; it’s a good idea to comment your code extensively with `%`.¹

Keep in mind that the M-file must be a text file, not, for example, a Word file. (If you write the M-file in a word processor, just remember to save it as a text file.) If you are using Matlab for Windows, you can use the **New M-file** and **Open M-file** options in the **File** menu to work with M-files.

When handing in problem sets, you should include a copy of the M-files used to solve them.

4 Functions

Matlab lets you define new functions, which can then be called from within a Matlab session or from a batch file. The function needs to reside in the current working directory or along Matlab’s search path (which can be displayed and modified with the command `path`).

Here is an example of a function in Matlab. The file containing the function should be called `bernoulli.m`.

¹If you are experiencing difficulties printing an M-file that starts with `%`, try adding an extra blank line at the top of the file. Postscript printers may misinterpret the file type if the M-file starts with `%`.

```

function y=bernoulli(p,n);

% BERNOULLI.M
% This function generates n independent draws of a Bernoulli
% random variable with probability of success p.

% first, draw n uniform random variables
x = rand(n,1);

% set y=1 if x is less than p. This gives probability p of success
y = (x <= p);
% end function definition

```

This function takes two arguments, p and n , and returns a vector of n independent Bernoulli random variables with probability of success equal to p . Notice that I have commented the file extensively. The function can then be called from a Matlab interactive session, or from within another batch file or another function. For example, the following interactive session calls the `bernoulli` function:

```

>> p=.7;
>> y=bernoulli(p,1000);
>> mean(y)
ans = 0.70800

```

This calls `bernoulli.m` to generate 1000 independent Bernoulli random variables with $p = .7$. The average of these 1000 variables is .708, quite close to the expected value .7.

Note: when defining functions, you can have it return more than one variable. For example, you could define a function `func1.m` that takes as inputs x_1 and x_2 , and returns y_1 and y_2 . The first line of the function definition would be:

```

function [y1, y2] = func1(x1,x2);

y1=x1-x2;
y2=x1*(x2-x1);
% end function definition

```