

Plotting

Python version

Machine Learning Computer Lab
ML2022, University of Pittsburgh, Spring 2022

How can we illustrate our data and mathematical results?

Plotting allows us to visualize machine learning data. Before we do any analysis, this allows us to explore the data; after an analysis, we use various kinds of plots to analyze and publish our results. MATLAB has a rich set of plotting commands that we will explore, including

- **line plot**, a sequence of points (x_i, y_i) connected by lines;
- **scatter plot**, a set of points (x_i, y_i) indicated by markers;
- **bar plot**, a set of data y_i displayed as bar heights;
- **histogram plot**, data to be grouped into bins before bar plotting;
- **contour plot**, contour lines of a function $f(x, y)$;
- **color contour plot**, contour colors of a function $f(x, y)$;
- **surface plot**, a 3D plot of the surface $z = f(x, y)$;
- **tree plot**, representing a branching process;
- **graph plot**, a set of labeled points p_i , and a list of two-way connections $(p_i \leftrightarrow p_j)$.
- **digraph plot**, a set of labeled points p_i , and a list of one-way connections $(p_i \rightarrow p_j)$.

1 Line Plot: the BULGARIA Data

Load the text data file *bulgaria_data.txt*, which records a sequence (y_i, p_i) , of years and population:

We think of the data as samples of a smoothly varying population function $p(y)$. A line plot will suggest this behavior by connecting the data points with line segments.

Exercise: Create the Bulgaria line plot:

1. Load the text data file

```
1 > data = load ( 'bulgaria_data.txt' );
```

2. Plot column 1 (year) versus column 2 (population):

```
1 > plot ( data(:,1), data(:,2) );
```

3. Customize the plot:

```
1 > plot ( data(:,1), data(:,2), 'linewidth', 3 );
2 > grid ( 'on' )
3 > xlabel ( '<— Year —>' )
4 > ylabel ( '<— Population —>' )
5 > title ( 'Changes in Bulgarian Population' )
```

4. Save a PNG version:

```
1 > print ( '-dpng', 'bulgaria.png' );
```

2 Line Plots: the PRICE Data

The file *price_data.txt* is a table of average monthly prices for 11 consumer products, between February 2008 and February 2018. There are 241 records, and each record contains 13 items: the month, the year, and the

prices of 11 goods. Columns 3, 10 and 13 contain the prices of bananas, gas, and milk, respectively. We want a single plot that shows line plots for all three items together.

Exercise:

1. load the text data file *price_data.txt*;
2. Create a month counter:

```
1 > month = 1 : 241;
```

3. Tell MATLAB to wait for multiple plots:

```
1 > hold on;
```

4. Plot months versus banana prices (column 3):

```
1 > plot ( month, data(:,3) );
```

5. Add similar plot commands for gas and milk;
6. Tell MATLAB the multiple plots are done:

```
1 > hold off;
```

7. Save a PNG copy as *price.png*;

3 Scatter Plots: the FAITHFUL Data

The file *faithful_data.csv* contains 272 observations of the Old Faithful geyser. Each record has an index, the eruption length in minutes, and the wait in minutes until the next eruption. A line plot of index versus the second or third variable doesn't show much pattern. However, it might seem reasonable that the second and third variables are related. A long eruption might mean a long subsequent wait, for instance. Such patterns can be investigated using a scatter plot, in which we simply put a mark at each point (x, y) for which we have data.

Because the data is stored as a CSV file (comma separated values), we have the MATLAB function *csvread()* to extract the data.

Exercise:

1. Load the CSV data file *faithful_data.csv*:

```
1 > data = csvread ( 'faithful_data.csv', 1, 0 );
```

2. Plot column 2 versus column 3; the extra argument 'ro' marks data with red points, no connection:

```
1 > plot ( data(:,2), data(:,3), 'ro' );
```

3. Save a PNG copy as *faithful.png*;

What relationship between eruption time and wait time is suggested by the scatter plot of the data?

4 Bar Plot: the ALBUM Data

Line plots treat the data as representing a continuously varying numeric quantity. But sometimes, it is more satisfying to think of our data as grouped into bins or buckets or lumps. The job of the plot is then to exhibit the number of items in each bin by drawing a bar of the corresponding height.

As an example, we consider the sales of music albums over a period of years. It is natural to want to see this data reported by separate years.

Exercise:

1. Load the text data file *album_data.txt*:

```
1 > data = load ( 'album_data.txt' );
```

2. Extract the "year" and "sales" variables as columns 1 and 2:

```
1 > year = data(:,1);  
2 > sales = data(:,2);
```

3. Make a bar plot:

```
1 > bar ( year , sales );
```

4. Save a PNG copy as *album.png*;

5 Horizontal Bar Plot: the SCHOOLYEAR Data

Usually, a bar plot exhibits the bars vertically. The length of the bar represents a count, but what does the position mean? Frequently, the bars represent counts of data over bins of equal size, and evenly spaced. But other times, the data is associated with different people, or different countries, or other items, so that each bar needs a readable label.

For such data, a horizontal bar plot might be best. The labels might appear on the left, with a horizontal bar to the right, as suggested by this simple typewritten example:

```
Votes in favor:  xxxxxxxxxx  
Votes against:  xxxxxxxxxxxxxxxx  
Votes withheld:  xxxx
```

For this exercise, our data will involve the length of the school year in days, for various countries. MATLAB has some trouble reading data files that contain a mixture of text and numeric data, so instead, we will get our data from a MATLAB function.

Exercise:

1. Access the MATLAB M file *schoolyear_data.m*:

```
1 > [ country , days ] = schoolyear_data ( );
```

2. Make a horizontal bar plot of the school days:

```
1 > barh ( days );
```

3. Label each bar with this horrible command:

```
1 > set ( gca , 'yticklabel' , country );
```

4. Save a PNG copy as *schoolyear.png*;

6 Histogram: the SNOWFALL Data

Suppose we want to display a large set of observations of some numerical quantity, $(x_1, x_2, \dots, x_n)$. If the data has some slowly varying trend, a line plot might be suitable. If the data is not too numerous, a bar plot might work. But if the data is numerous and somewhat chaotic, we are better off with a *histogram*, which simplifies the data by grouping it into a small number of bins. A histogram then shows how many data items fall into each range.

Exercise:

1. Access the data file *snowfall_data.txt*:

```
1 > data = load ( 'snowfall_data.txt' );
```

2. Copy out column 10 as “inches”:

```
1 > inches = data(:,10);
```

3. Create a bar plot of the inches, and convince yourself it is not a good presentation of the data:

```
1 > bar ( inches ) ;
```

4. Instead, create a histogram, allowing MATLAB to choose the number of bins:

```
1 > histogram ( inches ) ;
```

5. Save a PNG copy as *snowfall.png*;

What features of the data does the histogram allow you to see right away?

7 Contours: the VOLCANO Data

We may have a collection of measurements of temperature, or elevation, or pollution levels, sampled at regular points on a grid. One way to visualize such data is to use a contour plot. Contours can be indicated by a sequence of curves, each of which connects points with the same value of the measurement; a more vivid image can be done using colors.

Exercise:

1. Access the CSV data file *volcano_data.csv*:

```
1 > z = csvread ( 'volcano_data.csv' );
```

2. Make a line contour plot:

```
1 > contour ( z ) ;
```

3. Make a color contour plot:

```
1 > contourf ( z ) ;
```

4. Save a PNG copy as *volcano.png*;

8 Surfaces: the MEXICAN_HAT Function

A contour plot can also be useful to examine the behavior of a function $z = f(x, y)$; even better is a 3D surface plot. In this case, we have to generate the (x, y) grid values, and set up a way to evaluate the function.

Exercise:

1. Create X and Y grid arrays:

```
1 > [X,Y] = meshgrid ( -8 : 0.5 : 8 ) ;
```

2. Create the Z grid array:

```
1 > R = sqrt ( X.^2 + Y.^2 + eps ) ;  
2 > Z = sin ( R ) ./ R ;
```

3. Start with a line contour plot:

```
1 > contour ( X, Y, Z );
```

4. Compare a color contour plot:

```
1 > contourf ( X, Y, Z );
```

5. Finish with a 3D surface plot:

```
1 > surf ( X, Y, Z );
```

6. Save a PNG copy as *mexican_hat.png*;

9 Trees: the GENEALOGY data

A tree plot is way to illustrate a process which involves branching or dependence. A tree plot can illustrate all the paths that start at the initial point and branch one way or another at decision points.

The locations on a tree are called *nodes*, and the node representing the initial state is called the *root node*. If we assign each node an identifying index, then the structure of the tree can be described if each node identifies its “parent node”, which connects it back to the root node.

Exercise:

1. Create a tree description:

```
1 > nodes = [ 0, 1, 2, 2, 4, 4, 4, 1, 8, 8, 10, 10 ];
```

2. Display the tree:

```
1 > treeplot ( nodes );
```

3. Create labels for the nodes:

```
1 > labels = { ...  
2   'Adam',  
3   'Bert',  
4   'Carl',  
5   'Dale',  
6   'Eddy',  
7   'Fred',  
8   'Gina',  
9   'Hank',  
10  'Inez',  
11  'Jane',  
12  'Kurt',  
13  'Leah' };
```

4. Find out where the nodes are:

```
1 > [ x, y ] = treelayout ( nodes );
```

5. Attach labels to the nodes:

```
1   for i = 1 : length ( x )  
2       text ( x(i) + 0.1, y(i), labels{i} );  
3   end
```


6. Save a PNG copy as *genealogy.png*;

10 Graphs: the NETWORK data

Mathematically, a graph is a set of *nodes*, some of which are connected by *edges*. The edges usually represent two-way connections; one-way connections arise in “directed graphs”.

The MATLAB function `graph()` can define a graph G using two lists of equal length, containing the starting node s and terminal node t of each edge.

For your exercise, consider the following network:



Exercise:

1. Use the picture above to describe the graph of the network by filling in the s and t vectors:

```
1 > s = [ 1, 1, 2, ..., 10 ];  
2 > t = [ 2, 3, 4, ..., 11 ];
```

2. Create a graph called G :

```
1 > G = graph ( s, t );
```

3. Plot the graph


```
1 > plot ( G );
```

4. save a PNG copy as `network.png`;

11 Directed Graphs: the WEB data

A directed graph displays nodes connected by one-way edges. There may be a value associated with each edge, representing a weight, length or probability of access. A good display should include that information.

For your exercise, consider the following network, and assume that nodes 1 through 6 are labeled Alpha, Beta, Gamma, Delta, Epsilon and Zeta, respectively.



Exercise:

1. Use the picture above to describe the graph of the network by filling in the s and t vectors:

```

1 > s = [ 1, 1, 2, ..., 6 ];
2 > t = [ 5, 6, 1, ..., 4 ];

```

2. Define the digraph

```

1 > G = digraph ( s, t, nodenames );

```

3. Set a label for each node:

```

1 > node_labels = { 'Alpha', 'Beta', 'Gamma', 'Delta', 'Epsilon', 'Zeta' };

```

4. Set a label for each edge:

```

1 > edge_labels = { '0.1', '0.9', '0.3', ..., '1.0' };

```

5. Plot the graph:

```

1 > plot ( G, ...
2   'Layout', 'circle', ...
3   'EdgeLabel', edge_labels, ...
4   'NodeLabel', node_labels, ...
5   'ArrowSize', 15 );

```

6. Save a PNG copy as *web.png*;