Programming Level-up

An Introduction to Matplotlib

Jay Morgan

November 5, 2022

Outline

What is Matplotlib?

In summary:

- Matplotlib is one of the defacto plotting libraries for Python. While there are many others and certainly some that are built for specific plot types, Matplotlib continues to pervade scientific plotting.
- You can create basic plots (such as line or scatter plots) to more complicated plots that include interactivity.

Installing and importing Matplotlib

Matplotlib can be installed via conda:

```
conda install matplotlib
```

or with pip:

```
pip install matplotlib
```

Remember! You can install packages in ipython REPL/juypter notebook by inserting a '!' to the beginning of a shell command.

Basic plotting

First, we will import the matplotlib module. The plotting function is located within the `pyplot` package within matplotlib. The use of this package is so common that 99% of Python users will alias this import as `plt`:

```
3 import matplotlib.pyplot as plt
```

With this package now imported, we can now use the `plot` function. To begin with, let's just plot a simple line chart. In this case, the `plot` function takes an `x` and `y` argument, where `x` denotes the values along the x-axis and `y` are the values along the y-axis.

```
4 x = np.linspace(-10, 10, 100)
5 y = np.sin(x)
6 plt.plot(x, y)
```

In this example, we have created two vectors. The first `x`, creates a

Basic plotting

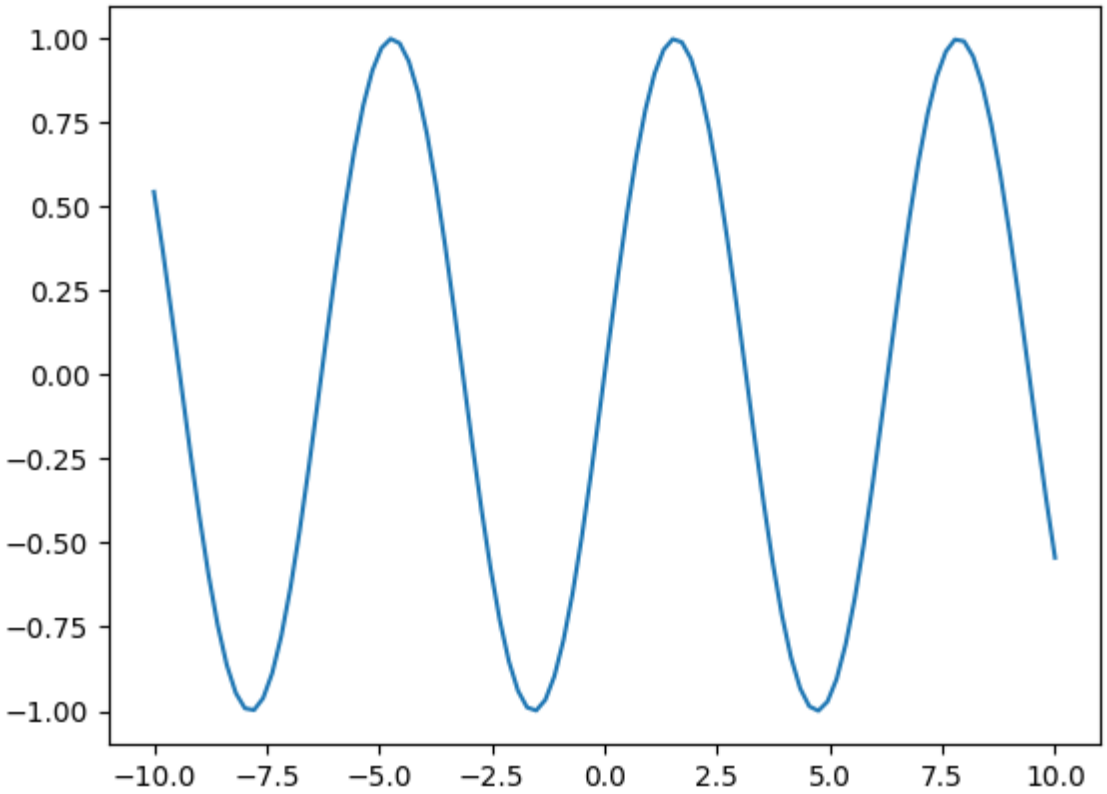

Different types of Plots

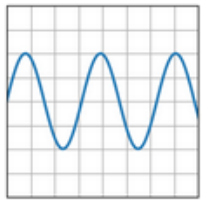

plot(x, y)

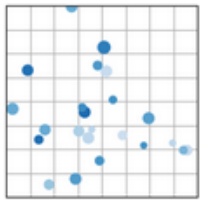

scatter(x, y)

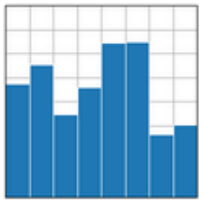

bar(x, height) / barh(y, width)

There are many different types of plots that one can make using matplotlib. These include the most popular:

- Line plots
- Scatter plots
- Bar plots
- Histograms
- Box plots
- Image plots

We're going to take a look at how we create each type of plot, examining what type of inputs they require.

Line plots

We've already seen one example of a line plot. This plot draws a line between each x,y point. For instance in the previous example, we created a sin wave by 'sampling' such wave using 100 samples from -10 to 10. Let's see what happens when we sample only 10 points:

```
7  x = np.linspace(-10, 10, 10)
8  y = np.sin(x)
9  plt.plot(x, y)
```

Line plots

We see the results are a less than ideal representation of a sin wave as `plot` will simply draw a straight line from each point.

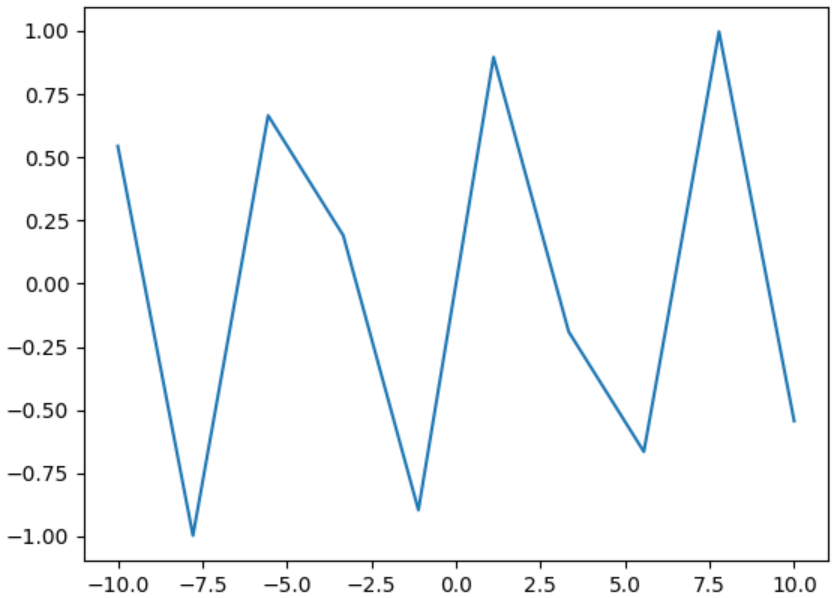

Scatter plots

If we want to see where each sample of the sin wave is, we could use instead the scatter plot, which will (by default) place a small circle at every x,y value. To create a scatter plot, we use `scatter` instead of the `plot` function. The arguments to this function are the same, however.

```
10  x = np.linspace(-10, 10, 10)
11  y = np.sin(x)
12  plt.scatter(x, y)
```

Scatter plots

Now we can see the position of each individual sample from the sin wave. If we, once again, sample 100 points from this curve, we will see better results.

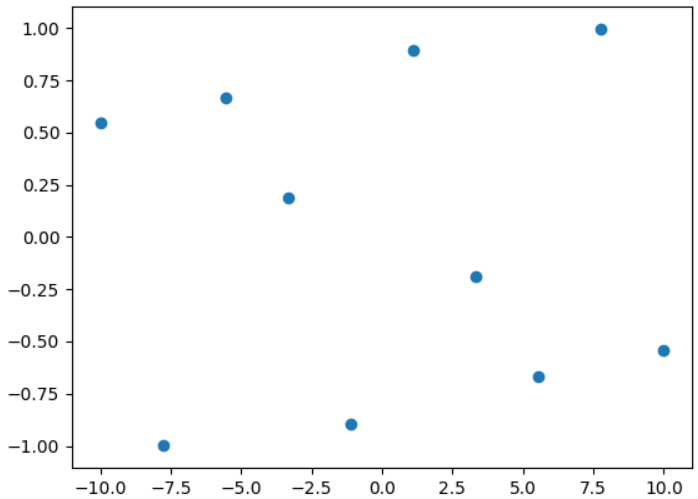

Scatter plots


```
13  x = np.linspace(-10, 10, 100)
14  y = np.sin(x)
15  plt.scatter(x, y)
```

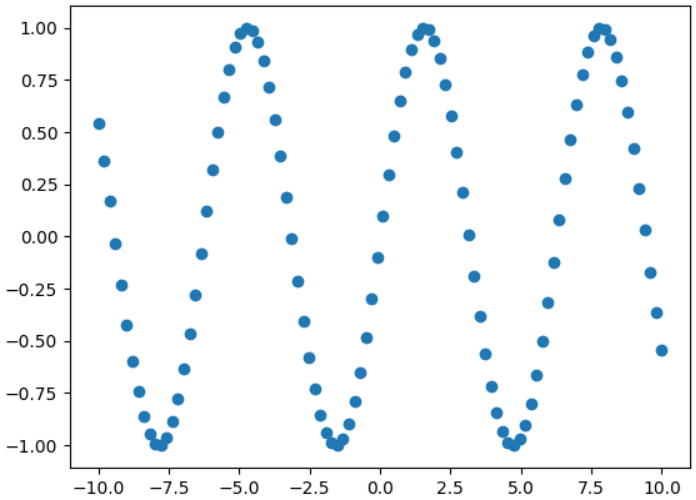

Bar plots

Bar plots are a simple plot that again takes an `x` and a `y`, where `x` is the numerical position of the bar's centre, and `y` is the height of the bar.

```
16  x = np.arange(0, 8)
17  y = np.random.uniform(2, 7, len(x))
18  plt.bar(x, y)
```

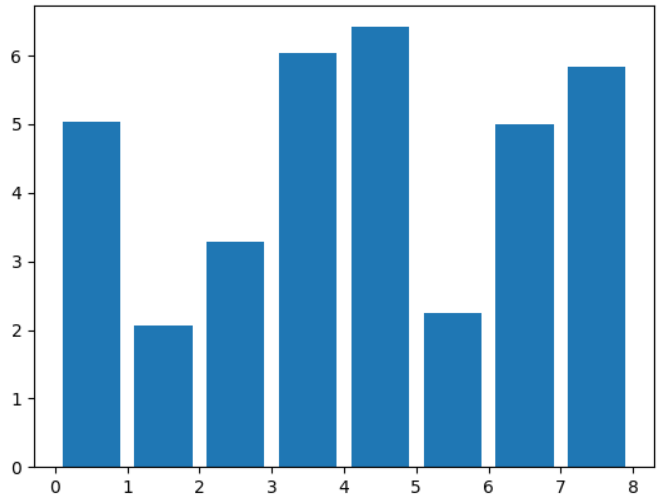

Histograms

Histograms allow us to visualise the distribution of values. In matplotlib, we can create a histogram of a vector by using the `hist` function that takes only the vector as its argument.

```
19 x = np.random.randn(1000)
20 plt.hist(x)
```

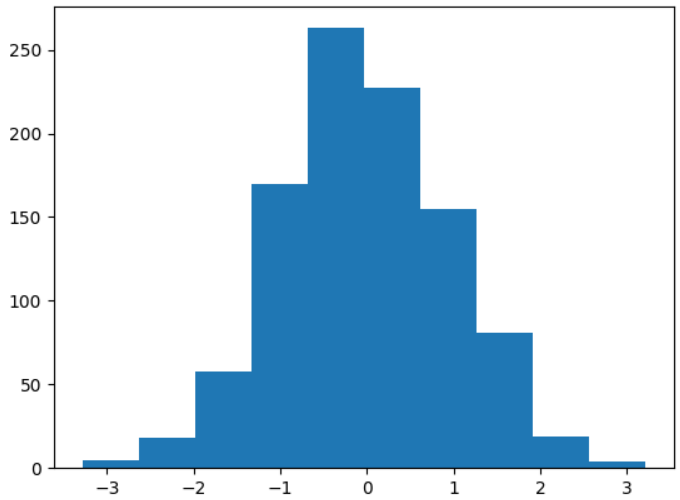

Box plots

Box plots also allow us to visualise the distribution, but the distribution of values within a group. In this example we're visualising the distribution of 3 groups. Using the `boxplot` function, we pass a matrix.

```
21 x = np.random.randn(10, 3)
22 plt.boxplot(x)
```

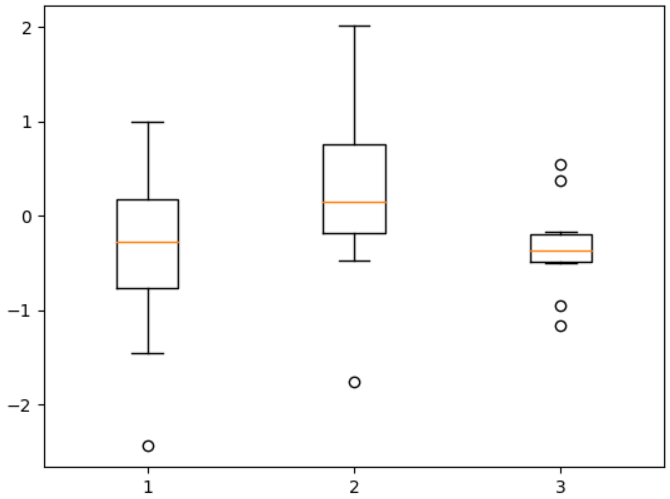

Image plots

In matplotlib, we can plot an 'image' – that is a 2D matrix – using the `imshow` function. For example:

```
23  fig = plt.figure()
24  x = np.random.randn(10, 10)
25  plt.imshow(x)
```

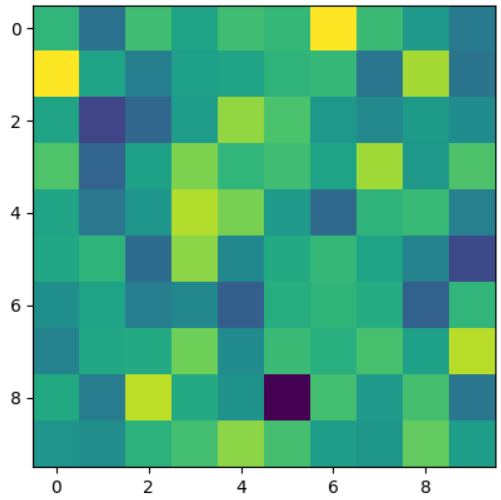

Image plots

Of course, given the name, we can then use `imshow` to plot an image as well, as long as we have the image loaded as a 2D array of values.

```
26 import PIL  # using the PIL module to read an image
27 img = np.array(PIL.Image.open("images/Lenna.png"))
28 plt.imshow(img)
```

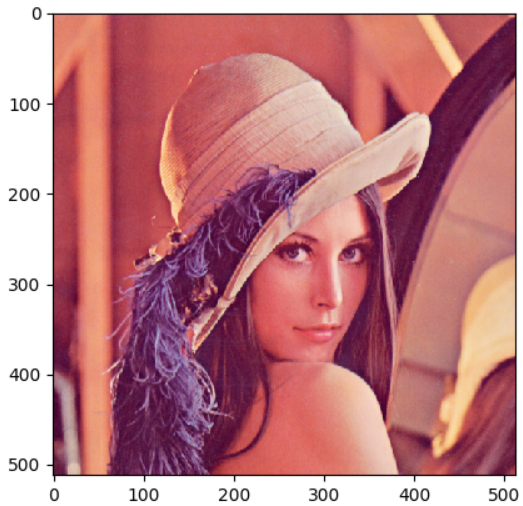

3 dimensional plots

3 dimensional plots require us to import another module from matplotlib.

```
29 from mpl_toolkits import mplot3d
```

After importing this module, we can using the *projection="3d"* and carry on plotting as normal.

```
30 fig = plt.figure()
31 ax = fig.gca(projection='3d')
32 theta = np.linspace(-4 * np.pi, 4 * np.pi, 100)
33 z = np.linspace(-2, 2, 100)
34 r = z**2 + 1
35 x = r * np.sin(theta)
36 y = r * np.cos(theta)
37 ax.plot(x, y, z, label='parametric curve')
38 ax.legend()
```

3 dimensional plots

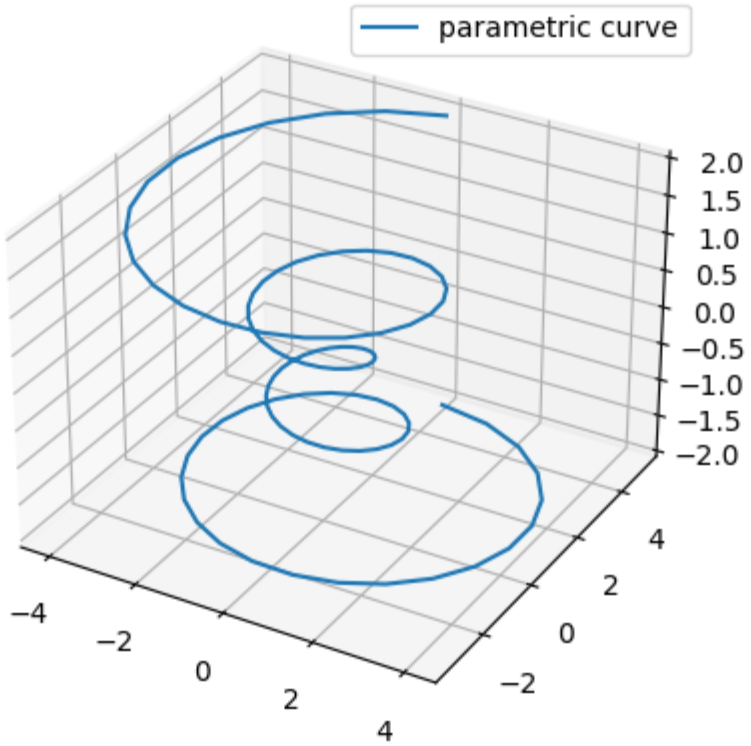

Different types of Plots

There are many more different types of plots you can make using matplotlib. You can find a comprehensive list at:

`https://matplotlib.org/stable/plot_types/index.html`

Subplots

What if we wanted to create many plots side-by-side? For this we can use the `subplots` function. This function takes the number of rows, and number of columns to create. It returns two values, the first is the figure (entire figure), and the second value is a list of sub figures. Using this list, we can place a plot of each of them.

```
39  x = np.linspace(-10, 10, 100)
40  y = np.sin(x)
41  z = np.cos(y)
42
43  fig, ax = plt.subplots(1, 2)
44  # ax is a list of sub figures
45  ax[0].plot(x, y)
46  ax[1].plot(x, z)
```

Subplots

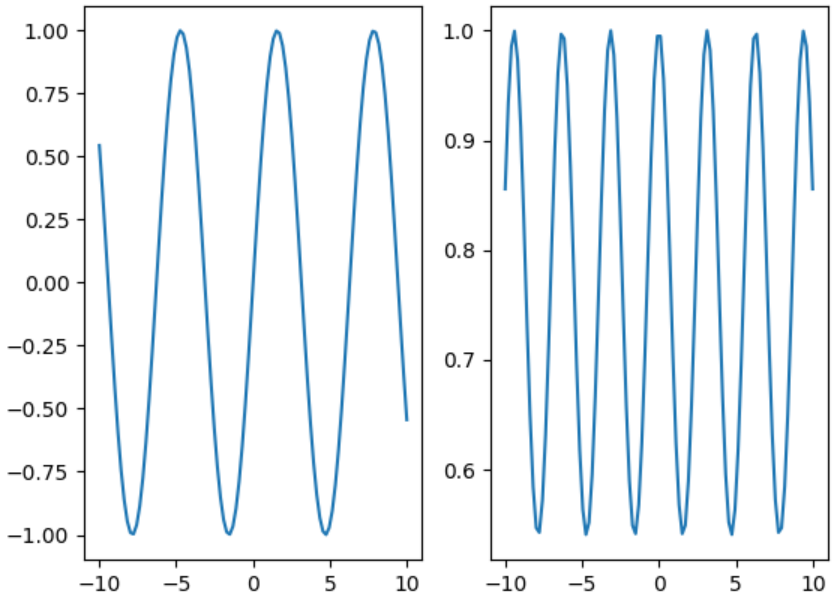
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
−0.75
−1.00
−10
−5
0
5
10
1.0
0.9
0.8
0.7
0.6
−10
−5
0
5
10

Adding a legend

Or we could put them onto the same plot. Matplotlib will automatically give them a different colour. If we use the `label` argument to `plot`, we can also give them a name that will appear when we call `legend()`.

```
47  x = np.linspace(-10, 10, 100)
48  y = np.sin(x)
49  z = np.tan(y)
50  fig, ax = plt.subplots()
51  ax.plot(x, y, label="sin(x)")
52  ax.plot(x, z, label="tan(x)")
53  ax.legend()
```

Adding a legend

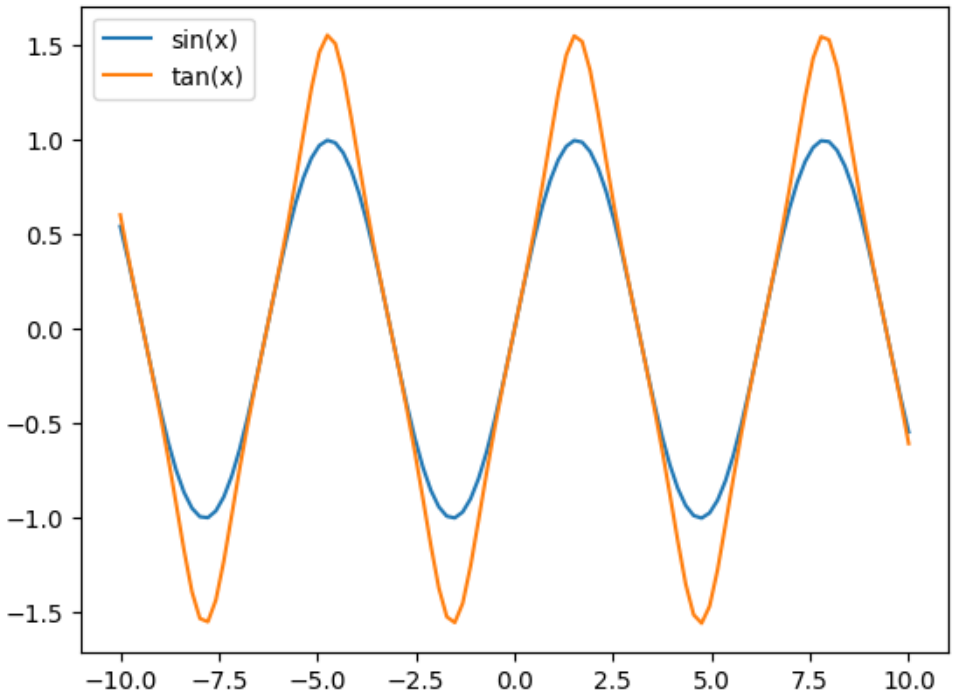

Position the legend in different places

We can change the position of the legend by specifying a different integer value for the `loc` argument (or string values such as 'upper left', 'upper right', ...). Additionally, we can change the number of columns the legend has with the `ncol` argument.

```
54  x = np.linspace(-10, 10, 100)
55  y = np.sin(x)
56  z = np.tan(y)
57
58  fig, ax = plt.subplots()
59  ax.plot(x, y, label="sin(x)")
60  ax.plot(x, z, label="tan(x)")
61  ax.legend(loc=1, ncol=2)
```

You can find the API reference for the different arguments to legend at: `https://matplotlib.org/stable/api/legend_api.html?highlight=legend#module-matplotlib.legend`

Position the legend in different places

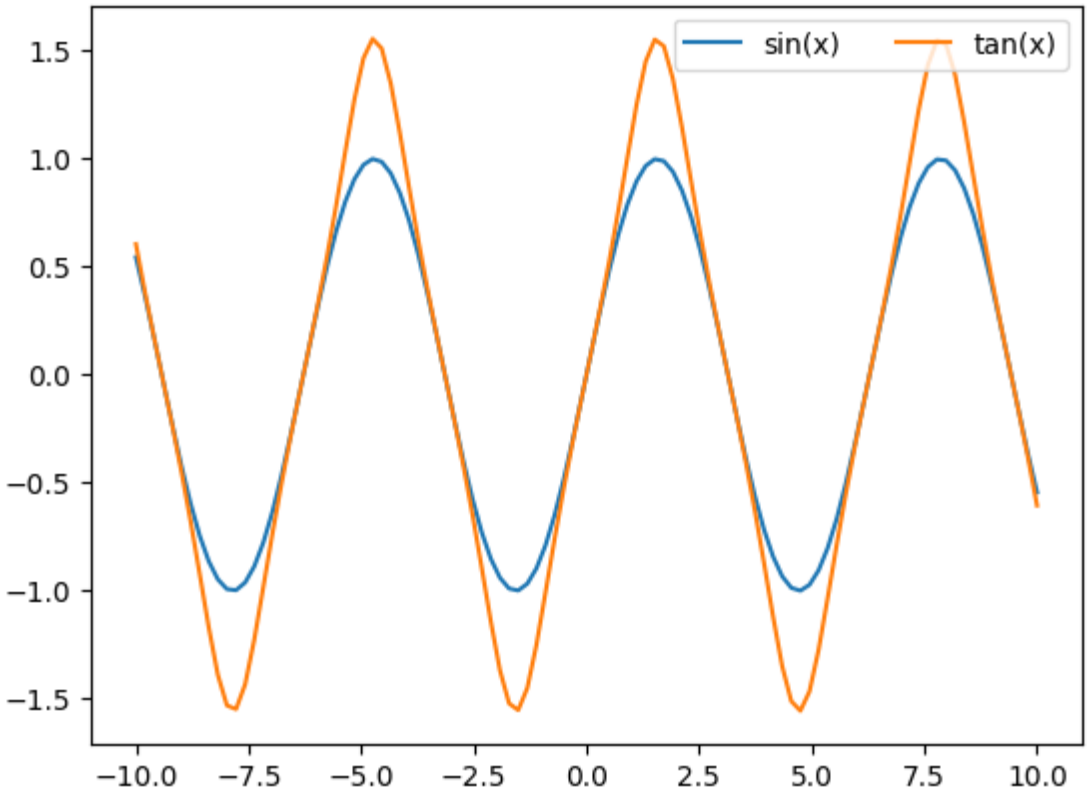

Modifying the x/y axis

Good graphs always have their axis's labelled. To do this in matplotlib, if we have a subplot object, we use `set_xlabel`, or we can use `plt.xlabel(...)`. Here is an example with an subplot object:

```
62  x = np.linspace(-10, 10, 100)
63  y = np.sin(x)
64  z = np.tan(y)
65
66  fig, ax = plt.subplots()
67  ax.plot(x, y, label="sin(x)")
68  ax.plot(x, z, label="tan(x)")
69  ax.legend(loc=1, ncol=2)
70  ax.set_xlabel("x")
71  ax.set_ylabel("y")
```

Modifying the x/y axis

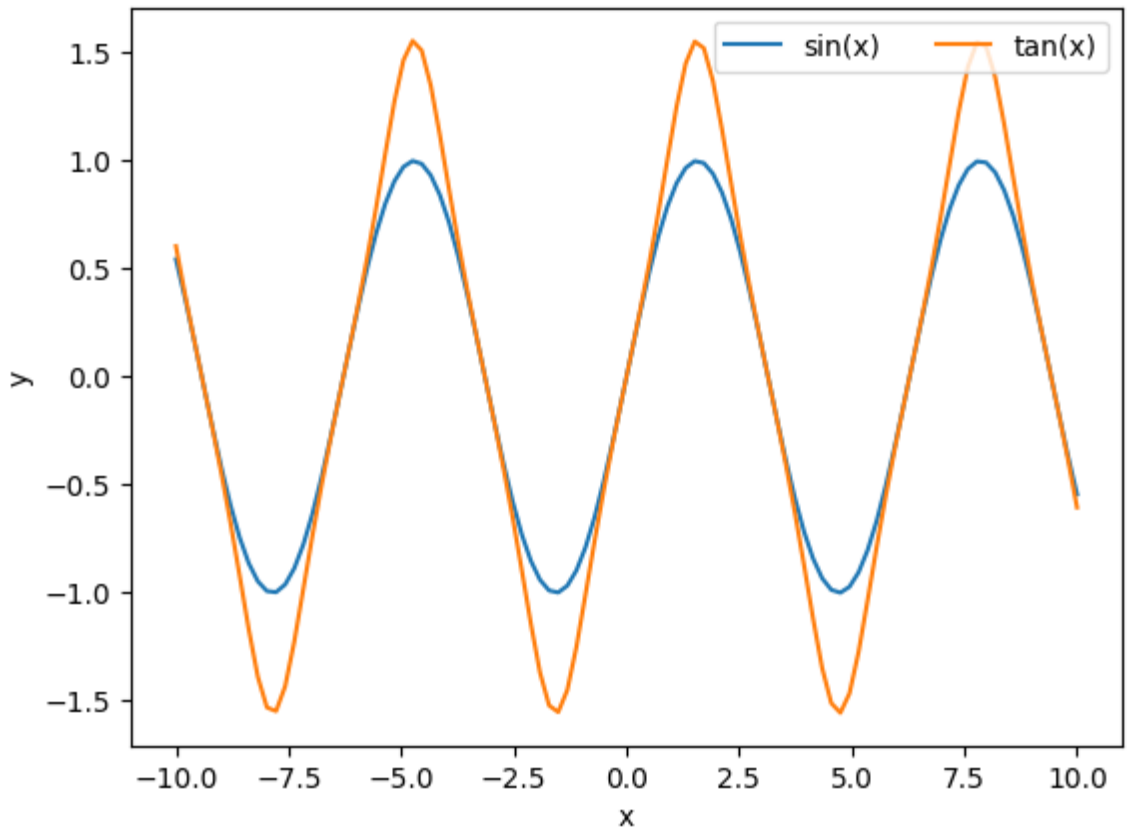

Changing figure size

A common change you may want to make to your figure is to change its size or aspect ratio. `figure()` or `subplots()` take an optional argument called `figsize`. This argument expects a tuple representing the width and height of the figure in inches.

```
72  fig = plt.figure(figsize=(8, 2.5))
73
74  # or most likely
75  fig, ax = plt.subplots(figsize=(8, 2.5))
76  x = np.linspace(-10, 10, 100)
77  y = np.sin(x)
78  z = np.tan(y)
79  ax.plot(x, y, label="sin(x)")
80  ax.plot(x, z, label="tan(x)")
81  ax.legend(loc=1, ncol=2)
82  ax.set_xlabel("x")
83  ax.set_ylabel("y")
```

Here we are creating a figure with 8 inches of width, and 2.5 inches of height.

Changing figure size

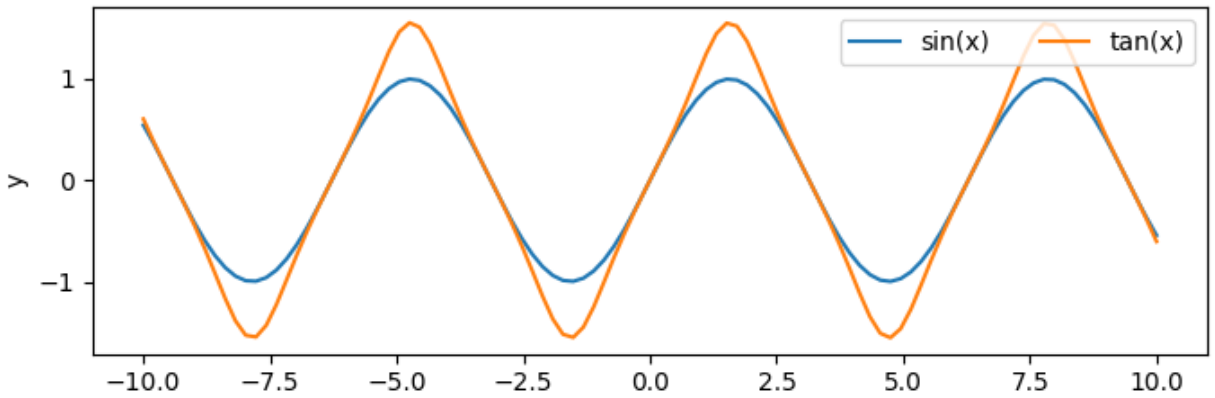

Changing figure size

This is especially useful when you have many sub-figures, as by default, they will be 'squashed' into the default aspect ratio. We can 'give them more space' by modifying this `figsize` argument when creating the many sub-figures.

```
84  fig, ax = plt.subplots(1, 2, figsize=(8, 2.5))
85  x = np.linspace(-10, 10, 100)
86  y = np.sin(x)
87  z = np.tan(y)
88  ax[0].plot(x, y, label="sin(x)")
89  ax[1].plot(x, z, label="tan(x)")
```

Changing figure size

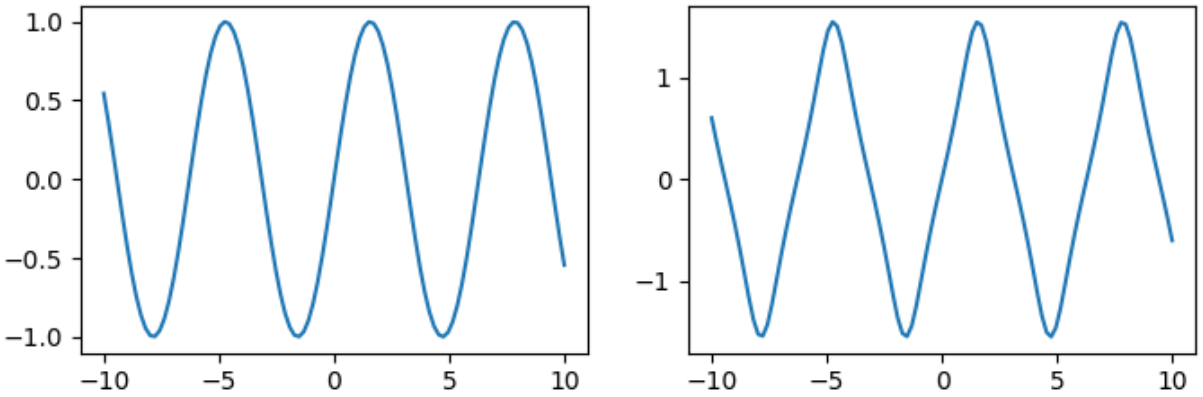

Line properties

When creating a plot, there are many different properties you can change. Some of these include:

- color – the colour of the line
- alpha – the amount of transparency (1.0 is opaque, 0.0 is transparent)
- linewidth, lw – the width of the stroke width
- linestyle, ls – the style of the line (i.e. a dotted line)

There are also some properties for the markers, i.e. the circles in the scatter plot. These properties are:

- marker – the type of marker (you can use different shapes instead of a circle
- markersize – the size of the mark
- markerfacecolor – colour of the marker
- markeredgewidth – outline width of the marker.

Line properties

If this example we are modifying some of the line properties that include the color (c), setting it to a string value of "green". The linewidth (lw) to be thicker, and making the line to be a dotted line by specifying the linestyle (ls) to "=–={".

```
90  fig = plt.figure()
91  x = np.linspace(-5, 5, 100)
92  y = np.sin(x)
93  plt.plot(x, y,
94           c="green", # or color
95           lw=3, # or linewidth
96           ls="--")
```

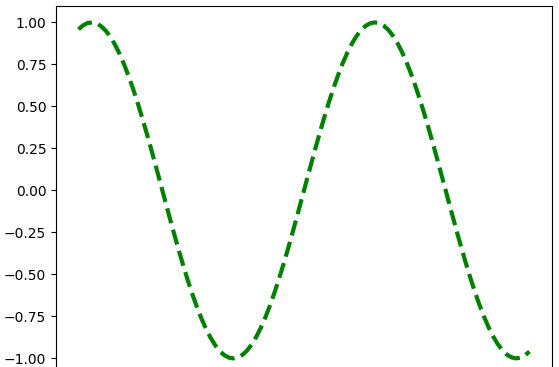

Colormaps

When we create a heatmap using `imshow`, the gradients of colour are automatically set. Yet, we can control the colour gradient using a colour map. First we must import `cm` from matplotlib:

```
97 from matplotlib import cm
```

Then we can get a colour map with 10 levels using `get_cmap`:

```
98 blues = cm.get_cmap("Blues", 10) # 10 levels
99 reds = cm.get_cmap("Reds", 2) # 2 levels
```

You can find a full list of different colour maps at: `https://matplotlib.org/stable/tutorials/colors/colormaps.html`

Colourmaps

Now that we have our new colour maps, we can pass it as an `cmap` argument when we create a plot.

```
100 x = np.random.randn(10, 10)
101 y = np.random.randn(10, 10)
102 fig, ax = plt.subplots(1, 2, figsize=(8, 3))
103 p1 = ax[0].imshow(x, cmap=blues)
104 p2 = ax[1].imshow(y, cmap=reds)
105 fig.colorbar(p1, ax=ax[0])
106 fig.colorbar(p2, ax=ax[1])
```

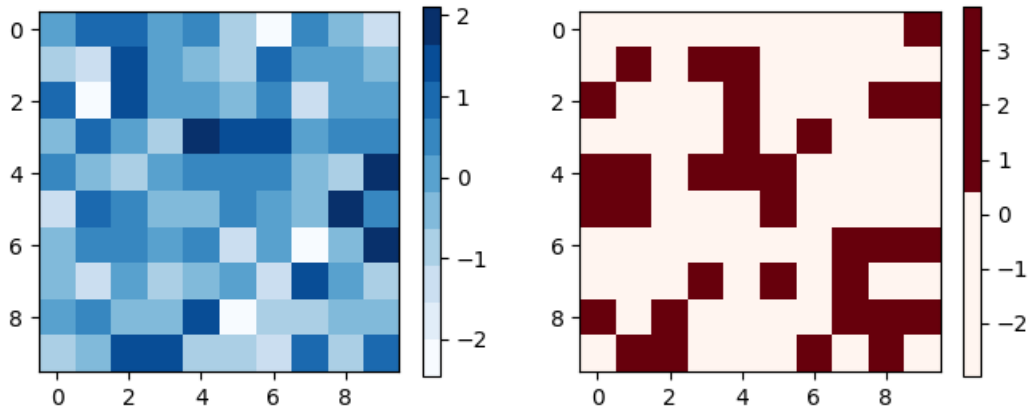

Ticks

If we want to customise the numbers along each axis, we use the `set_xticks` for the x-axis and `set_yticks` for the y-axis. These functions take the list of locations for each 'tick', and optionally a list of labels to use instead of the numbers.

```
107  x = np.linspace(-2, 2, 100)
108  y = np.sin(x)
109
110  bx = np.arange(2, 7)
111  by = np.random.uniform(2, 7, len(bx))
112
113  fig, ax = plt.subplots(1, 2, figsize=(8, 3))
114  ax[0].plot(x, y)
115  ax[0].set_xticks([-2, 0, 2])
116  ax[1].bar(bx, by)
117  ax[1].set_xticks(bx, ["a", "b", "c", "d", "e"])
```

Ticks

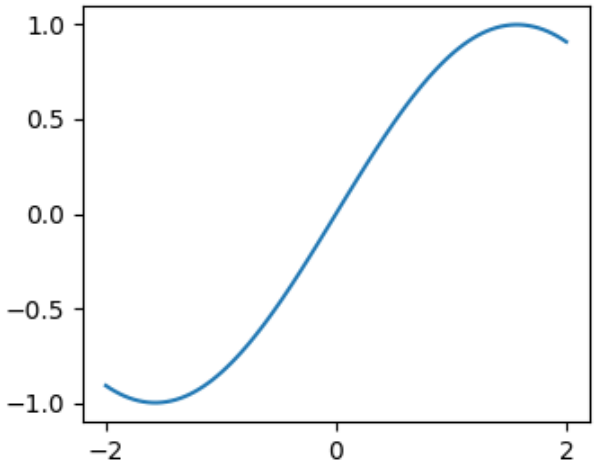

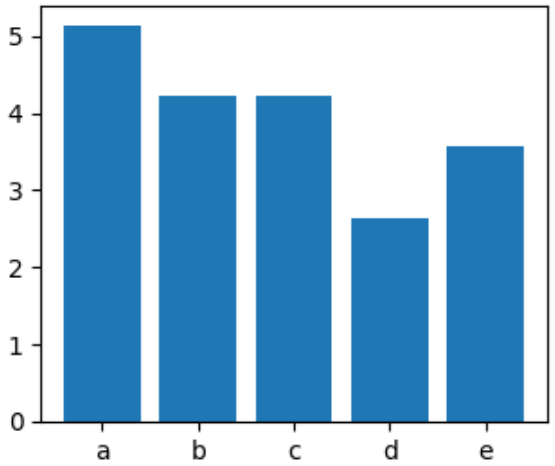

Grids

In all of the previous plots, the background has no grids, they are simply white. If we wanted to add grid lines to the plot we use the `.grid()` method. This function, by default, adds the major grid lines.

```
118  x = np.linspace(-2, 2, 100)
119  y = np.sin(x)
120  z = np.tan(x)
121  fig, ax = plt.subplots(1, 2, figsize=(8, 3))
122  ax[0].plot(x, y)
123  ax[0].grid()
124  ax[1].plot(x, z)
125  ax[1].grid(which="both", color="r")
```

Grids

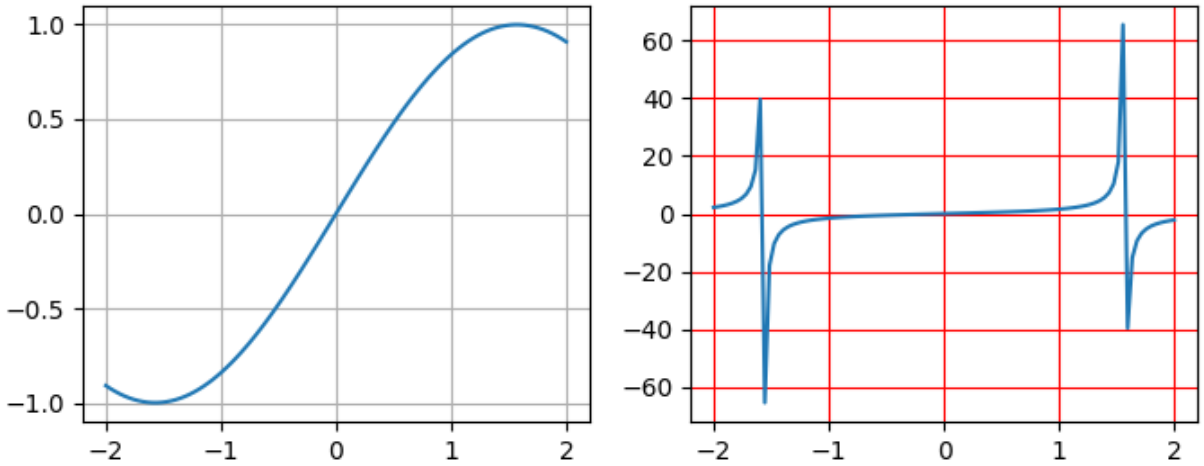

Scale

The default behaviour of matplotlib is to plot using a linear scale. In certain situations, we want view the plot using a different scale. For this we can use `set_yscale`.

```
126 x = np.linspace(-2, 10, 100)
127 y = np.exp(x)
128 fig, ax = plt.subplots(1, 2, figsize=(8, 3))
129 ax[0].plot(x, y)
130 ax[0].grid()
131 ax[1].plot(x, y)
132 ax[1].set_yscale('log')
133 ax[1].grid()
```

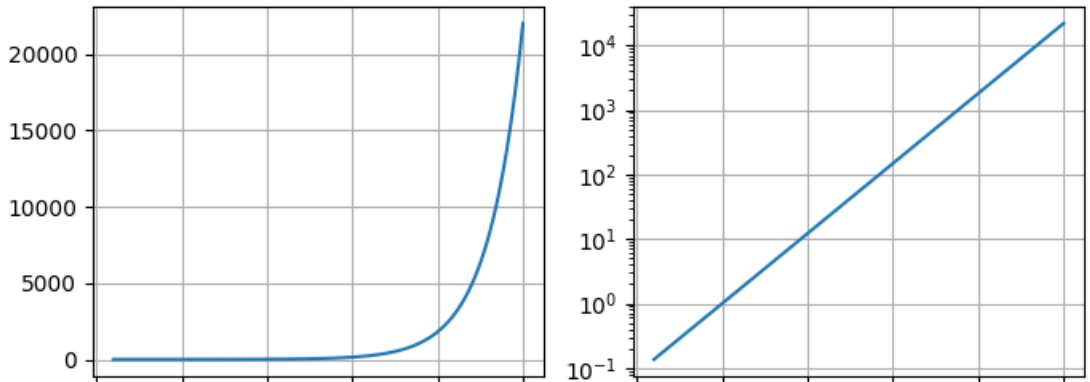

Setting the plot limits

By default, matplotlib will calculate the minimum and maximum values of the data, and use those values to set the limits of the plot. Using `set_xlim` and `set_ylim` we can change this default behaviour.

```
134 x = np.linspace(-2, 2, 100)
135 y = np.sin(x)
136 fig, ax = plt.subplots(1, 2, figsize=(8,3))
137 ax[0].plot(x, y)
138 ax[0].set_ylim(-1, 2)
139 ax[1].plot(x, y)
140 ax[1].set_xlim(-3, 3)
```

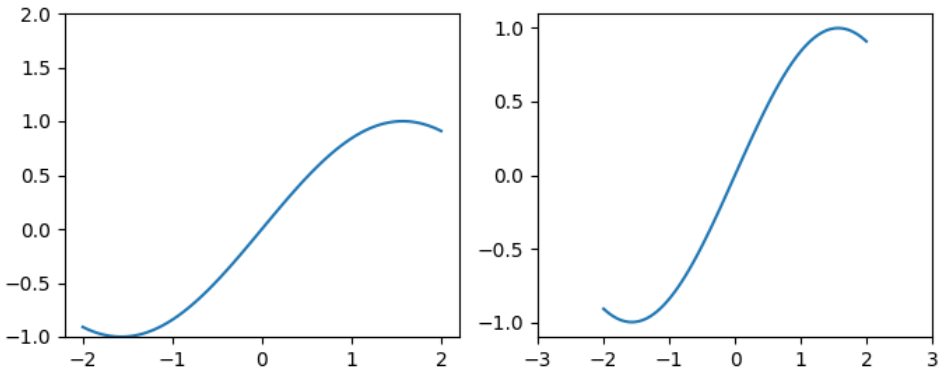

Annotations

We can annotate our plot in a number of way:

- .axhline – plot a horizontal line (axvline for vertical lines)/
- .annotate – add text to the plot at a certain position.

```
141 x = np.linspace(-2, 2, 100)
142 y = np.sin(x)
143 fig, ax = plt.subplots()
144 ax.plot(x, y)
145 ax.axhline(0, c='gray', ls='--')
146 ax.annotate("0th line", (-2, 0), xytext=(-1.5, 0.25),
147             arrowprops=dict(facecolor='black', shrink=0.05,
148                             width=0.5, headwidth=5.0))
```

Annotations

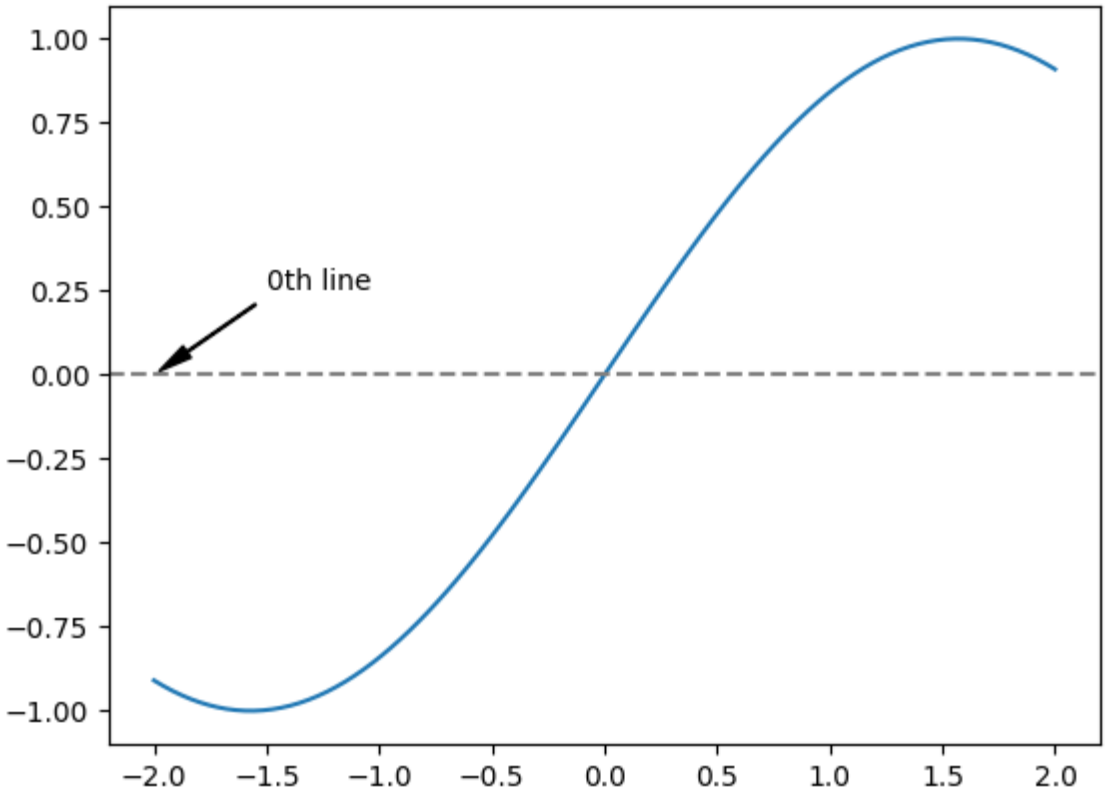

Creating a twin axes plot

Sometimes you will want to display multiple sub-plots on the same plot, but where each have a very different range in values. Instead of having a single y-axis, with `twinx()` we can create a two y-axis plot.

```
149 x = np.arange(10, 100)
150 y = np.exp(x)
151 z = np.log(x)
152
153 fig, ax = plt.subplots(1, 2)
154 ax[0].plot(x, y, label="exp(x)")
155 ax[0].plot(x, z, label="log(x)")
156 ax[0].legend()
157
158 ax2 = ax[1].twinx()
159 ax[1].plot(x, y)
160 ax2.plot(x, z, color="orange")
161 ax2.tick_params(axis="y", labelcolor="orange")
```

Creating a twin axes plot

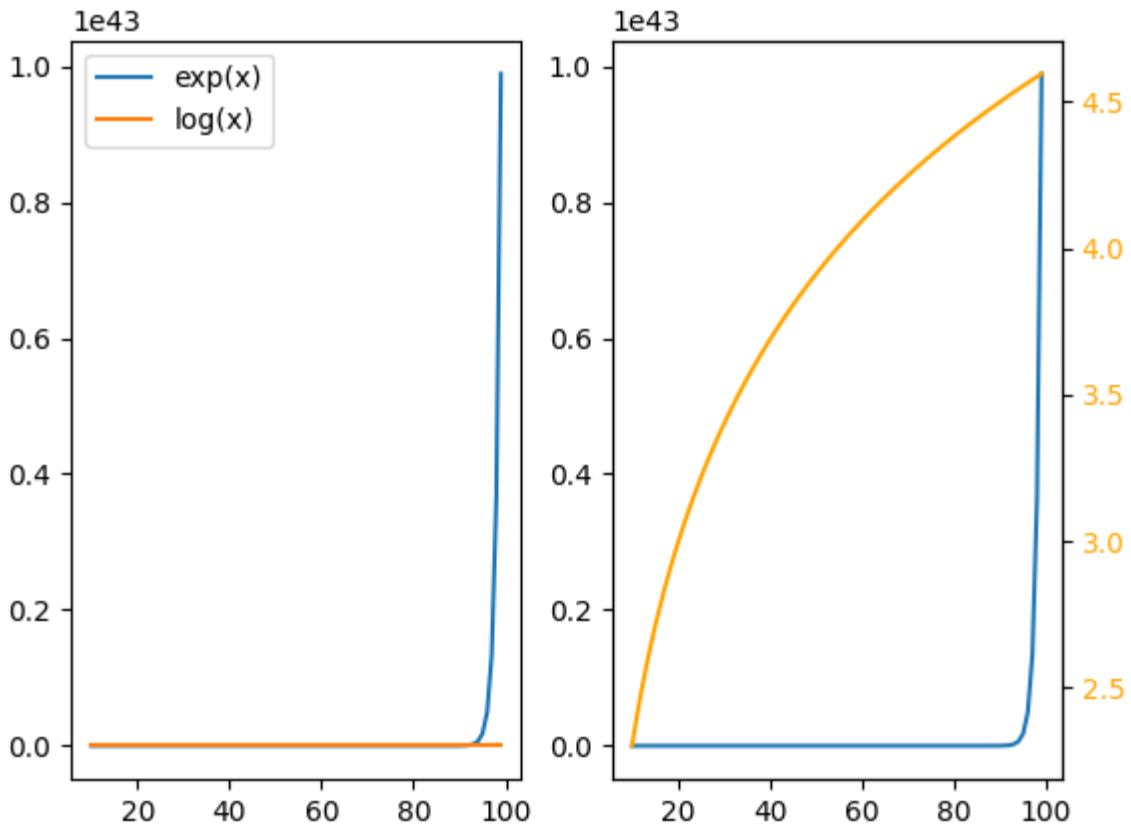

Learn more

There are many many more types of plots you can create with matplotlib. I would recommend that you read the documentation to fully appreciate everything that it can visualise:

- Gallery – `https://matplotlib.org/stable/gallery/index.html`
- Plotting tutorials – `https://matplotlib.org/stable/tutorials/index.html`
- Basic plot types – `https://matplotlib.org/stable/plot_types/index.html`