

Data Analysis – Analyzing and Visualizing

15-110 – Friday 04/15

Announcements

- Check 6-1 was due today
 - How did it go?
- Quiz 5
 - How did it go?
 - Grades today/this weekend

Last Time

Last week we discussed the **data analysis process** and went over several methods for **representing** and **organizing** data.

This time, we'll talk more about what we can **do** with that data once we've processed it.

Learning Goals

- Perform **basic analyses** on data, including calculating **statistics** and **probabilities**, to answer simple questions
- Choose an appropriate **visualization** to create based on the number of **dimensions** and **data types**
- Create simple **matplotlib visualizations** that show the state of a dataset using **APIs** and **examples**

Analysis

Basic Data Analyses – Statistics Library

There are many basic analyses we can run on features in data to get a sense of what the data means. You've learned about some of them already in math or statistics classes, such as **mean**, **median**, and **mode**.

You can implement these in Python yourself, but you don't have to! There's already a **statistics** library that does this for you.

```
import statistics
```

```
data = [41, 65, 64, 50, 45, 13, 29, 14, 7, 14]
```

```
statistics.mean(data) # 34.2
```

```
statistics.median(data) # 35.0
```

```
statistics.mode(data) # 14
```

Example: Analyzing Ice Cream Data

We've now cleaned the ice cream dataset from last week and created categories for the ice-cream. Let's use this as a running example of how to perform analyses.

Here's a bit of code from last time to load and represent the dataset:

```
import csv
def readData(filename):
    f = open(filename, "r")
    # Semester, 3 orig, 3 cleaned, 3 categories
    data = list(csv.reader(f))
    return data
```

Example: Statistics of Ice Cream

We can start by measuring the **statistics** of the ice cream dataset.

The data is text, so we must turn it into numbers before performing analyses. Try counting the number of times a person lists a specific flavor as their favorite and putting those counts into a list to analyze.

The `count` method is handy here if we narrow down the data being counted first!

```
def getFlavorCounts(data, flavor):
    counts = []
    firstCol = data[0].index("#1 category")
    for i in range(1, len(data)): # skip header
        line = data[i]
        # only include categorized flavors
        flavorCategory = line[firstCol:firstCol+3]
        count = flavorCategory.count(flavor)
        counts.append(count)
    return counts

import statistics
d = readData("all-icecream.csv")
print(statistics.mean(getFlavorCounts(d, "chocolate")))
```

Calculating Probabilities

You'll also often want to calculate probabilities based on your data.

In general, the probability that a certain data type occurs in a dataset is the count of how often it occurred, divided by the total number of data points.

Probability:

```
lst.count(item) / len(lst)
```

Conditional probability (the probability of something occurring given another factor) is slightly more complicated. Create a modified version of the list that contains only those elements with that factor; then you can use the same equation.

```
newLst = []  
for x in lst:  
    if meetsProperty(x):  
        newLst.append(x)  
newLst.count(item) / len(newLst)
```

Example: Probabilities of Ice Cream

To calculate the probability that an ice cream flavor is someone's #1 favorite, just count the number of times that flavor shows up in the appropriate column and divide by the total number of data points.

We can easily modify this to calculate the probability that you like Flavor Y if you like Flavor X too!

```
# Probability that a flavor is chosen
def getClassProb(data, flavor):
    count = 0
    firstCol = data[0].index("#1 category ")
    for i in range(1, len(data)):
        if data[i][firstCol] == flavor:
            count += 1
    return count / (len(data) - 1)

d = readData("all-icecream.csv")
print(getClassProb(d, "chocolate"))
```

More Analysis Methods

There's plenty of other data analysis methods we could cover – bucketing, detecting outliers, dealing with missing data – but what kind of method you need will depend entirely on the context of the problem you're solving.

You should generally be able to derive an algorithm that matches the analysis you want to perform.

Visualization

Exploration vs. Presentation

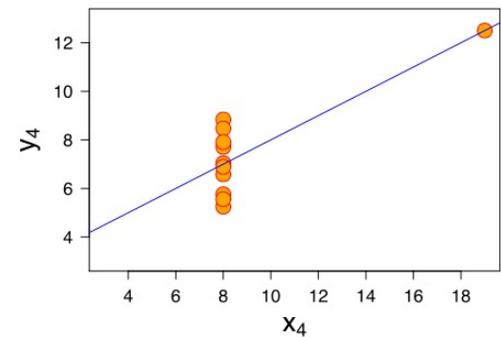
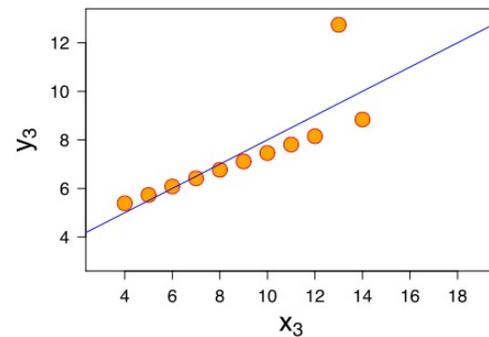
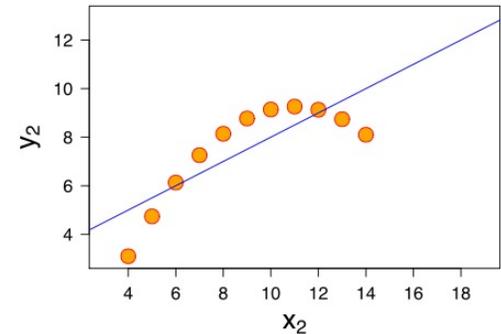
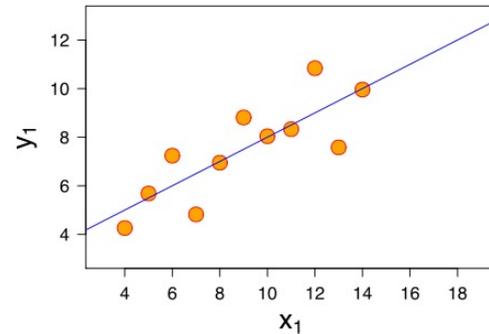
Data Visualization is the process of taking a set of data and representing it in a visual format. Whenever you've made charts or graphs in past math or science classes, you've visualized data!

Visualization is used for two primary purposes: **exploration** and **presentation**.

Data Exploration

In data exploration, charts created from data can provide information about that data beyond what is found in simple analyses alone.

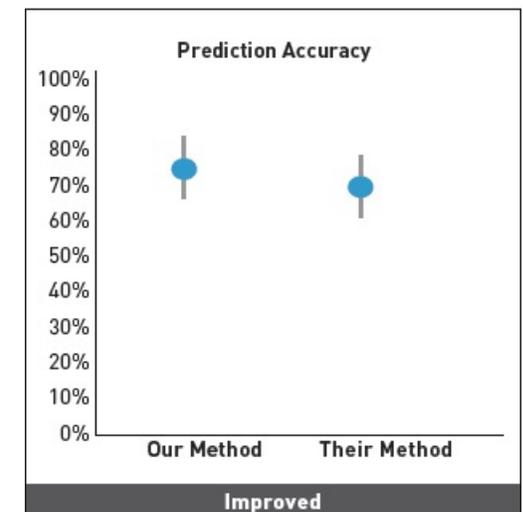
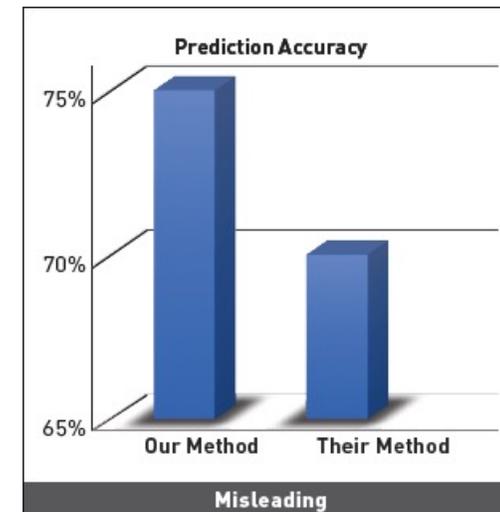
For example, the four graphs to the right all have the same mean and the same best-fit linear regression. But they tell very different stories.



Data Presentation

In data presentation, you've already found an interesting pattern in the data and you need to make that pattern easily visible to other people.

In order to choose the best visualization for the job, consider the **type** of the data you're presenting (categorical, ordinal, or numerical), and how many **dimensions** of data you need to visualize.



<https://interactions.acm.org/archive/view/july-august-2018/the-good-the-bad-and-the-biased>

One-Dimensional Data

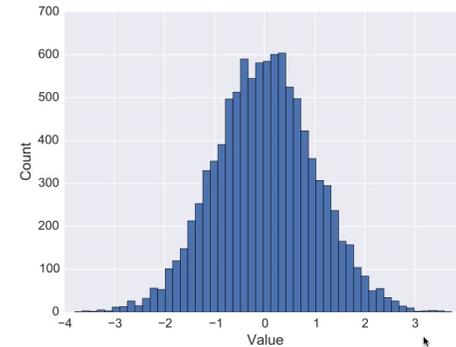
A **one-dimensional visualization** only visualizes a single feature of the dataset. For example:

"I want to know how many of each **product type** are in my data"

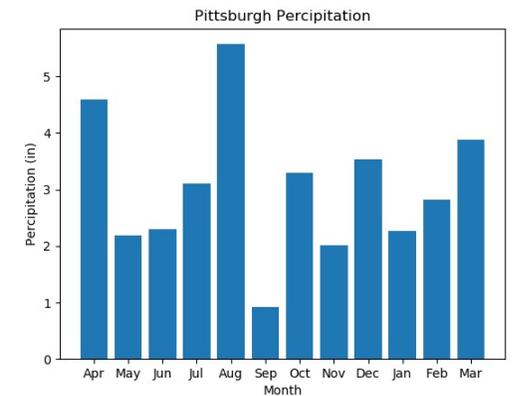
"I want to know the proportion of **people who have cats** in my data"

Charts for One-Dimensional Data

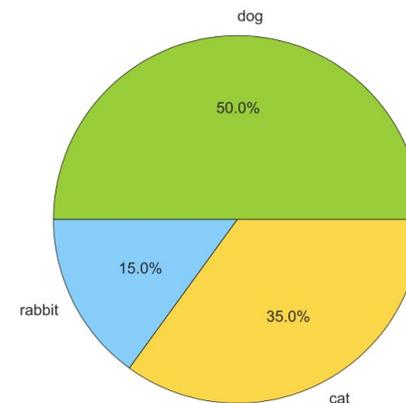
To visualize **numerical** data, use a **histogram**.



To visualize **ordinal** data, use a **bar chart**.



To visualize **categorical** data, use a **pie chart**.



Two-Dimensional Data

A **two-dimensional visualization** shows how two features in the dataset relate to each other. For example:

"I want to know the **cost** of each **product category** that we have"

"I want to know the **weight** of the animals that people own, by **pet species**"

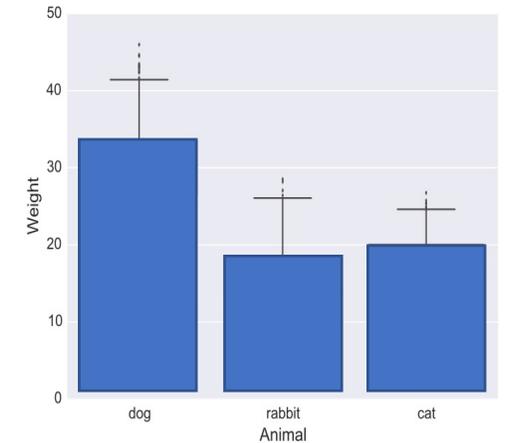
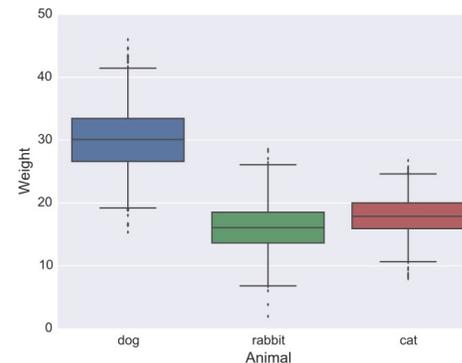
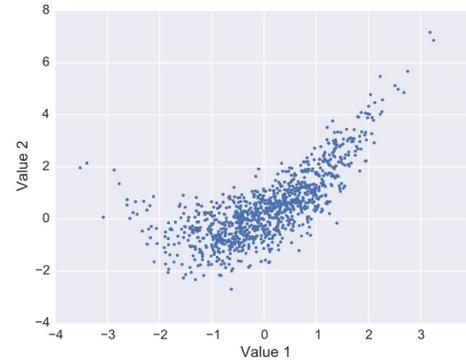
"I want to know how the **size** of the product affects **the cost of shipping**"

Charts for Two-Dimensional Data

To analyze **numerical x numerical** data, use a **scatter plot**.

To analyze **numerical x ordinal/categorical** data, use a **bar chart** for averages or a **box-and-whiskers plot** for ranges.

It is difficult to analyze **ordinal/categorical x ordinal/categorical** data visually; use a table instead.



Three-Dimensional Data

A **three-dimensional** visualization tries to show the relationship between three different features at the same time. For example:

"I want to know the **cost** and the **development time** by **product category**"

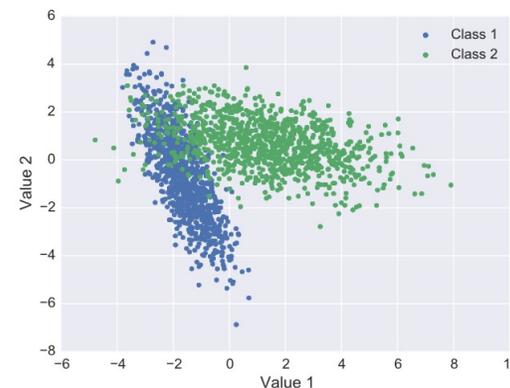
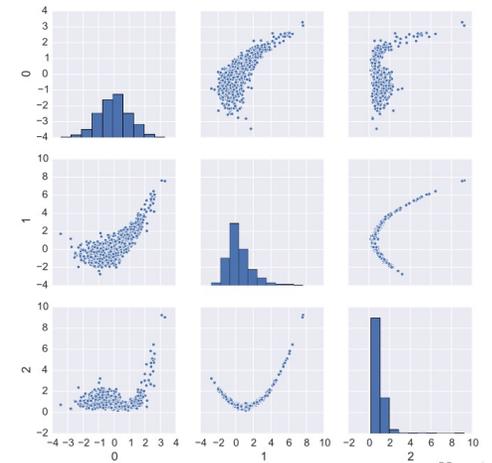
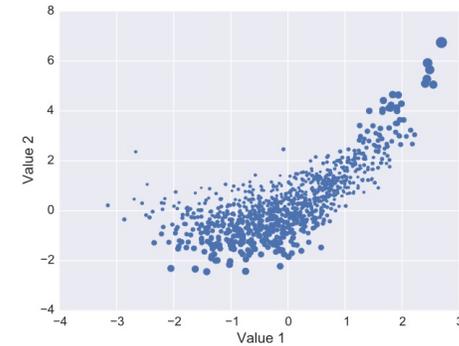
"I want to know the **weight** of the animals that people own and how much they **cost**, by **pet species**"

"I want to know how the **size** of the product and the **manufacturing location** affects **the cost of shipping**"

Charts for Three-Dimensional Data

To analyze **numerical x numerical x numerical** data, use a **bubble plot** to compare all three or a **scatter plot matrix** to compare all the pairs.

To analyze **numerical x numerical x ordinal/categorical** data, use a **colored scatter plot**.



Activity: Pick a Visualization

You do: for each of the problem prompts, determine the number of dimensions, then pick the **best** visualization to use based on the data types.

- graph the % of people who have gotten COVID vs. the % of people who have been vaccinated, separated by state
- graph the distribution of grades (72, 94, etc) in a class
- graph the ages of pets at a shelter compared to the species of pets

Coding Visualizations with Matplotlib

Matplotlib Makes Visualizations

The **matplotlib** library can be used to generate interesting visualizations in Python.

Matplotlib is **external** – you need to install it on your machine to run it. Use the **pip** command to do this.

```
pip install matplotlib
```

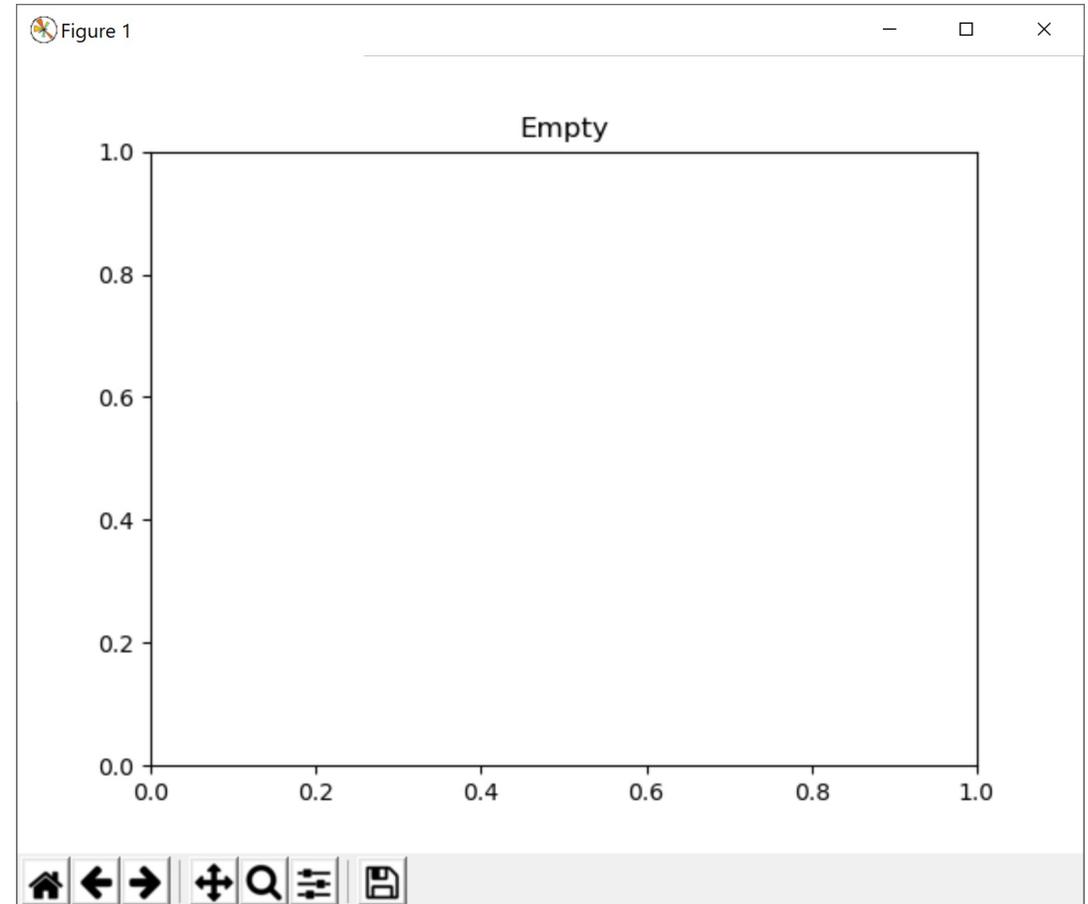
Draw Visualizations on the Plot

Matplotlib visualizations can be broken down into several components. We'll mainly care about one: the **plot** (called `plt`). This is like Tkinter's canvas, except that we'll draw visualizations on it instead of shapes.

We can construct an (almost) empty plot with the following code. Note that matplotlib comes with built-in buttons that let you zoom, move data around, and save images.

```
import matplotlib.pyplot as plt

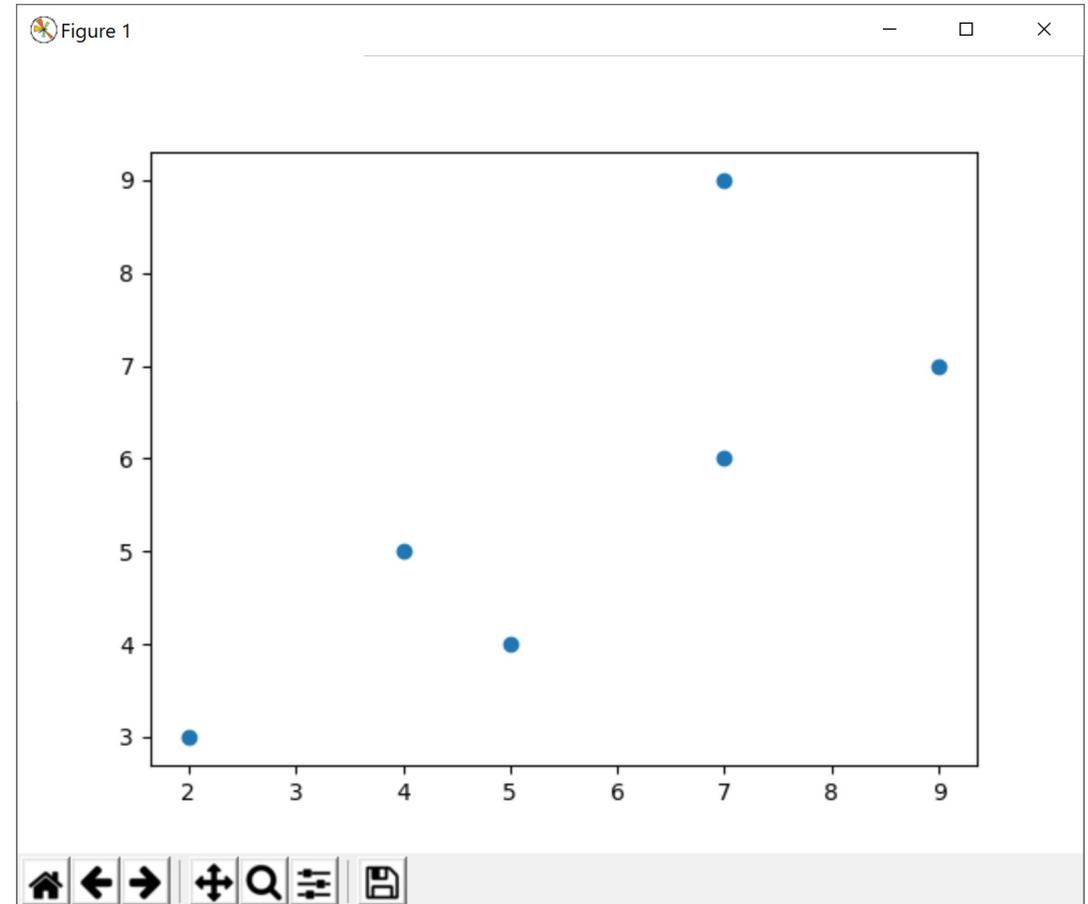
plt.title("Empty")
plt.show()
```



Add Visualizations with Methods

There are lots of built-in methods that let you construct different types of visualizations. For example, to make a scatterplot use `plt.scatter(xValues, yValues)`.

```
x = [2, 4, 5, 7, 7, 9]
y = [3, 5, 4, 6, 9, 7]
plt.scatter(x, y)
plt.show()
```

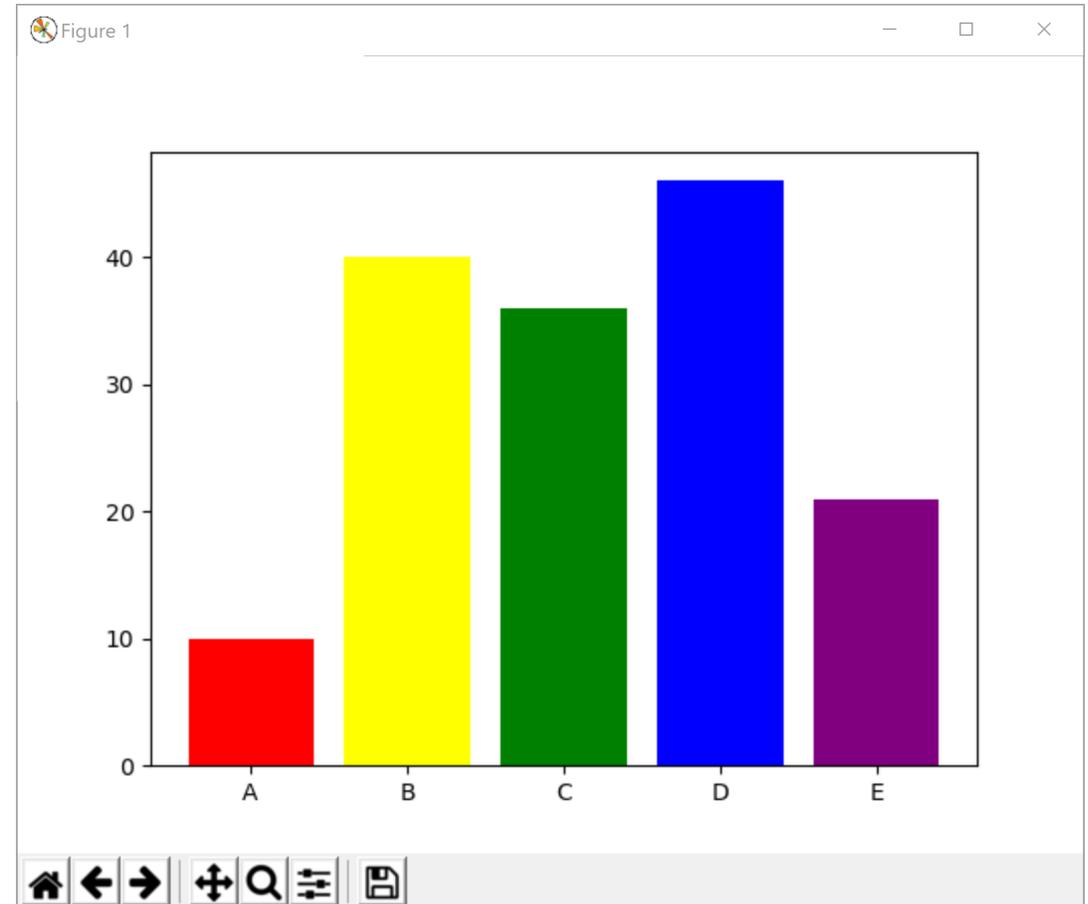


Visualization Methods have Keyword Args

You can **customize** how a visualization looks by adding **keyword arguments**. We used these in Tkinter to optionally change a shape's color or outline; in Matplotlib we can use them to add labels, error bars, and more.

For example, we might want to create a bar chart (with `plt.bar`) with a unique color for each bar. Use the keyword argument `color` to set the colors.

```
labels = [ "A", "B", "C", "D", "E" ]  
yValues = [ 10, 40, 36, 46, 21 ]  
colors = [ "red", "yellow", "green",  
          "blue", "purple" ]  
plt.bar(labels, yValues, color=colors)  
plt.show()
```



Don't Memorize – Use the Website!

There are a *ton* of visualizations you can draw in Matplotlib, and hundreds of ways to customize them. It isn't productive to try to memorize all of them.

Instead, use the documentation! Matplotlib's website is very well organized and has tons of great examples: <https://matplotlib.org/>

When you want to create a visualization, start by searching the **API** and the **pre-built examples** to find which methods might do what you need.

API Example

For example – how can we add x-axis and y-axis labels to the bar chart?

Go to the plot API:

https://matplotlib.org/stable/api/_as_gen/matplotlib.pyplot.html

Search 'label' and you'll soon find the functions `xlabel` and `ylabel`. You can [click on the function](#) to find more information. The page describes what the function does, what the required arguments are, and what it returns.

Note that the **keyword arguments** will be listed with default values. That's how we know they're optional.

```
matplotlib.pyplot.xlabel(xLabel, fontdict=None, labelpad=None, *, Loc=None, **kwargs) [source]
```

Set the label for the x-axis.

Parameters:

`xlabel`: str

The label text.

`labelpad`: float, default: `rcParams["axes.labelpad"]` (default: 4.0)

Spacing in points from the axes bounding box including ticks and tick labels. If None, the previous value is left as is.

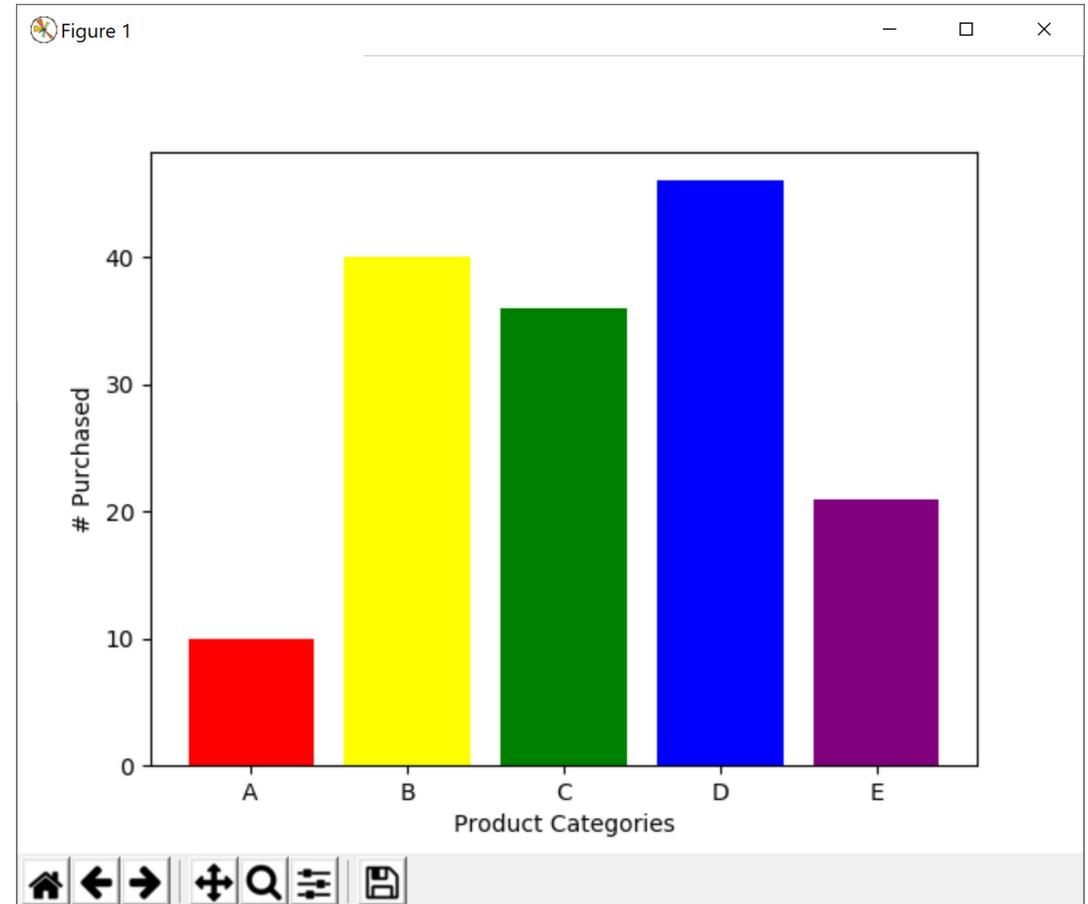
`loc`: {'left', 'center', 'right'}, default: `rcParams["xaxis.labellocation"]` (default: 'center')

The label position. This is a high-level alternative for passing parameters `x` and `horizontalalignment`.

If nothing obvious shows up, you can do a broader internet search of 'matplotlib x-axis label', which will often point you to the right place.

Adding xlabel and ylabel

```
labels = [ "A", "B", "C", "D", "E" ]  
yValues = [ 10, 40, 36, 46, 21 ]  
colors = [ "red", "yellow", "green",  
          "blue", "purple" ]  
plt.bar(labels, yValues, color=colors)  
  
plt.xlabel("Product Categories")  
plt.ylabel("# Purchased")  
  
plt.show()
```



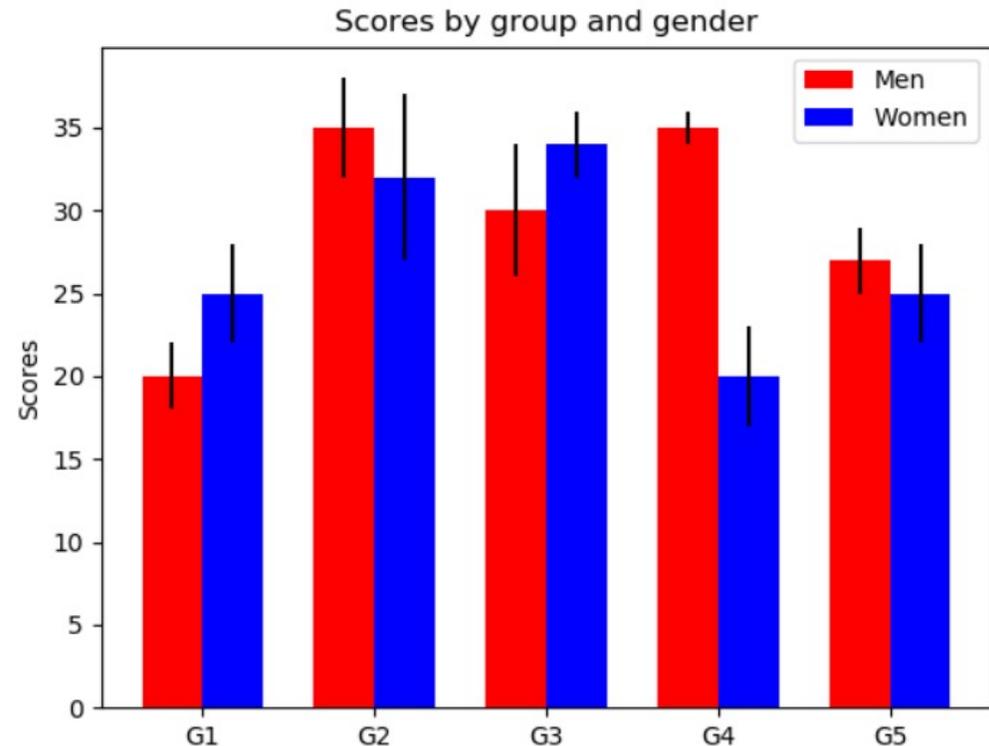
'Example' Example

Alternatively, you can browse the Matplotlib examples page to find visualizations and features that might prove useful:

<https://matplotlib.org/stable/gallery/index.html>

Perhaps we're interested in using [grouped bar charts](#) to show the breakdown between products purchased in different countries. The example code provides a starting place for which functions to use.

Try copying the example code into your editor and running it. Then try changing some things to see how the results are affected.



Sidebar: Plot vs Axis

You might have noticed that the grouped bar chart example looks slightly different than the code we've written so far. It sets

```
fig, ax = plt.subplots()
```

and calls methods on `ax` for the rest of the code.

This is an alternate way to write code in Matplotlib. Instead of drawing on the plot, break the plot into one of more axes with `plt.subplots`, then draw directly on the axis.

This is mainly useful if you want to draw more than one visualization in a single window. For the visualizations we'll do in this class, `plt` will work fine.

Going from Example to Our Own Code

```
import matplotlib
import matplotlib.pyplot as plt
import numpy as np

labels = ['G1', 'G2', 'G3', 'G4', 'G5']
men_means = [20, 34, 30, 35, 27]
women_means = [25, 32, 34, 20, 25]

x = np.arange(len(labels)) # the label locations
width = 0.35 # the width of the bars

fig, ax = plt.subplots()
rects1 = ax.bar(x - width/2, men_means,
                width, label='Men')
rects2 = ax.bar(x + width/2, women_means,
                width, label='Women')

...
```

```
labels = [ "A", "B", "C", "D", "E" ]
yValuesA = [ 10, 40, 36, 46, 21 ]
yValuesB = [ 20, 45, 35, 62, 32 ]

xValuesA = []
xValuesB = []
w = 0.35
for i in range(len(labels)):
    xValuesA.append(i - w / 2)
    xValuesB.append(i + w / 2)

plt.bar(xValuesA, yValuesA, width=w)
plt.bar(xValuesB, yValuesB, width=w)

plt.xlabel("Product Categories")
plt.ylabel("# Purchased")

plt.show()
```

Example: Visualizing Ice Cream

Let's use Matplotlib to visualize how popular the ice cream flavors were. We're visualizing counts of categorical data, so we can use a **pie chart**.

Start by looking up how to make a pie chart in the plot API:

https://matplotlib.org/stable/api/as_gen/matplotlib.pyplot.pie.html#matplotlib.pyplot.pie

Example: Reformat the Data

A pie chart requires a list `x` that holds the size of each portion. We can optionally provide a list labels with the labels.

Let's find all the flavors in the dataset and get the counts of each one using `getFlavorCounts` from before.

```
data = readData("all-icecream.csv")
firstCol = data[0].index("#1 category")
flavors = []
portions = []
for i in range(1, len(data)):
    flavor = data[i][firstCol]
    if flavor not in flavors: # haven't counted yet
        flavors.append(flavor)
        counts = getFlavorCounts(data, flavor)
        portions.append(sum(counts))
```

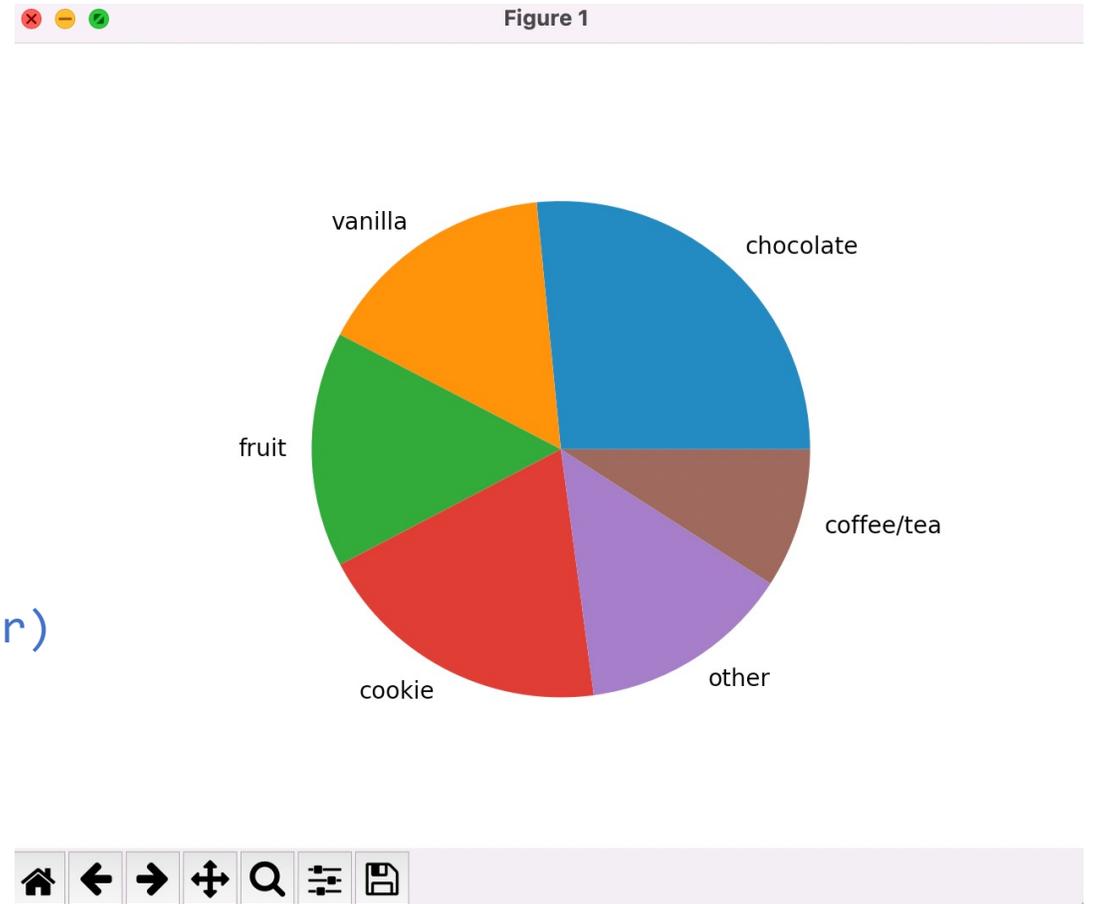
Example: Create the Pie Chart

Now we can combine it all together into one pie chart!

```
import matplotlib.pyplot as plt

data = readData("all-icecream.csv")
firstCol = data[0].index("#1 category")
flavors = []
portions = []
for i in range(1, len(data)):
    flavor = data[i][firstCol]
    if flavor not in flavors:
        flavors.append(flavor)
        counts = getFlavorCounts(data, flavor)
        portions.append(sum(counts))

plt.pie(portions, labels=flavors)
plt.show()
```



Learning Goals

- Perform **basic analyses** on data, including calculating **statistics** and **probabilities**, to answer simple questions
- Choose an appropriate **visualization** to create based on the number of **dimensions** and **data types**
- Create simple **matplotlib visualizations** that show the state of a dataset using **APIs** and **examples**

Feedback: <https://bit.ly/110-s22-feedback>