

Lab 1

The Zen of Python

In this lab you will learn how to execute Python code from the command line instead of using Colab. In the ACME Python Essentials document

(see foundations-of-applied-mathematics.github.io/)

follow the instructions in Section 1, Getting Started, to install Python 3.7 via Anaconda on your machine. Read Running Python and do [Problem 1](#). (If your laptop is running Windows, it may be easier to do this entire lab on one of the department Linux machines.) Then read about IPython and write/execute all of the yellow code boxes in your **Sandbox** file. Now do the following problems.

Using your code from the Fall Semester Lab “Cracking the Vigenère Cipher”, create a file `vigenere.py` that decodes an encoded message in a `.txt` file (see below for instructions on passing command line arguments to a Python program.) Get the text file `ciphertext.txt` from Learning Suite. Your program should print (not return) both the key and the message using two separate print statements.

```
>>> python vigenere.py ciphertext.txt
<key>
<familiar plaintext>
```

To pass command line arguments to a Python program, use the `argparse` module. A minimal working example (saved in a python file `argparse-ex.py`):

```
import argparse
parser = argparse.ArgumentParser()
parser.add_argument("n")
args = parser.parse_args()
print int(args.n)**2
```

When we run the command `python argparse-ex.py 17` in the command line it outputs 289 since $289 = 17^2$. If you are interested in passing more arguments into your Python

program, or other options, see the tutorial

docs.python.org/2/howto/argparse.html

To read the contents of a `.txt` file in Python, use the command

```
ciphertext = open('ciphertext.txt').read()
```

Read The Zen of Python: legacy.python.org/dev/peps/pep-0020/. Now familiarize yourself with the PEP 8 Style Guide for Python Code www.python.org/dev/peps/pep-0008/ (don't read the whole thing, it's very long). I suggest you read the first 5 sections (through Whitespace in Expressions and Statements) and also Naming Conventions: Function and Variable Names. Now go back to your first few Python Labs from Fall Semester and reflect on how they might be changed to conform to this style guide. Write a few sentences listing some practices you'll try to implement this semester to make your code more readable.

Lab 2

Introduction to Plotting

This material is adapted from the ACME lab on Matplotlib*

Raw numerical data is rarely helpful unless it can be visualized. Fortunately there is a helpful Python package `matplotlib` that can help you create nice visualizations of your data. You have already seen this package in your SVD lab, but in this lab you will learn more about its capabilities.

We will focus on creating *line plots*. The following code creates an array of outputs of the function $f(x) = x^2$, then visualizes the array using `matplotlib`.

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

y = np.arange(-5,6)**2


# Visualize the plot.
plt.plot(y)                # Draw the line plot.
plt.show()                 # Reveal the resulting plot.
```

The result is shown in Figure 2.1a. Just as `np` is a standard alias for NumPy, `plt` is a standard alias for `matplotlib.pyplot` in the Python community.

The call `plt.plot(y)` creates a figure and draws straight lines connecting the entries of `y` relative to the y -axis. The x -axis is (by default) the index of the array, which in this case is the integers from 0 to 10. Calling `plt.show()` then displays the figure.

An obvious problem with Figure 2.1a is that the x -axis does not correspond correctly to the y -axis for the function $f(x) = x^2$ that is being drawn. To correct this, define an array `x` for the domain, then use it to calculate the image `y = f(x)`. The command `plt.plot(x,y)` plots `x` against `y` by drawing a line between the consecutive points `(x[i], y[i])`. Note that the arrays must have the same number of elements to be compatible.

Another problem with Figure 2.1a is its poor resolution: the curve is visibly bumpy, especially near the bottom of the curve. NumPy's `linspace()` function makes it easy to get a higher-resolution domain by creating an array of evenly-spaced values in a given interval



(a) `plt.plot(y)` uses the indices of the array for the x -axis. (b) `plt.plot(x,y)` specifies both the domain and the range.

Figure 2.1: Plots of $f(x) = x^2$ over the interval $[-5, 5]$.

where the **number of elements** is specified.

```
# Get 4 evenly-spaced values between 0 and 32 (including endpoints).
np.linspace(0, 32, 4)

# Get 50 evenly-spaced values from -5 to 5 (including endpoints).
x = np.linspace(-5, 5, 50)
y = x**2                                # Calculate the range of f(x) = x**2.
plt.plot(x, y)
plt.show()
```

The resulting plot is shown in Figure 2.1b. This time, the x -axis correctly matches up with the y -axis. The resolution is also much better because x and y have 50 entries each instead of only 10.

1. Write a function that plots the functions $\sin(x)$, $\cos(x)$, and $\arctan(x)$ on the domain $[-2\pi, 2\pi]$ (use `np.pi` for π). Call `plt.xlim(-2*np.pi, 2*np.pi)` before `plt.show()` to stretch the x -axis appropriately. Make sure the domain is refined enough to produce a figure with good resolution.

Plot Customization

The plots you created in exercise 1, are extremely basic. Most plots are greatly improved by the addition of color, legends, axis labels and titles.

`plt.plot()` receives several keyword arguments for customizing the drawing. For example, the color and style of the line are specified by the following string arguments.


| Key | Color | Key | Style |
|-----|-------|------|---------------|
| 'b' | blue | '-' | solid line |
| 'g' | green | '--' | dashed line |
| 'r' | red | '-.' | dash-dot line |
| 'c' | cyan | ':' | dotted line |
| 'k' | black | 'o' | circle marker |

Specify one or both of these string codes as the third argument to `plt.plot()` to change from the default color and style. Other `plt` functions further customize a figure.

| Function | Description |
|---|--|
| <code>legend()</code> | Place a legend in the plot |
| <code>title()</code> | Add a title to the plot |
| <code>xlim()</code> / <code>ylim()</code> | Set the limits of the x - or y -axis |
| <code>xlabel()</code> / <code>ylabel()</code> | Add a label to the x - or y -axis |

```
x1 = np.linspace(-2, 4, 100)
plt.plot(x1, np.exp(x1), 'g:', linewidth=6, label="Exponential")
plt.title("This is the title.", fontsize=18)
plt.legend(loc="upper left")      # plt.legend() uses the 'label' argument of
plt.show()                       # plt.plot() to create a legend.

x2 = np.linspace(1, 4, 100)
plt.plot(x2, np.log(x2), 'r*', markersize=4)
plt.xlim(0, 5)                   # Set the visible limits of the x axis.
plt.xlabel("The x axis")         # Give the x axis a label.
plt.show()
```




See <http://acme.byu.edu/wp-content/uploads/2020/08/Matplotlib2020.pdf> for more comprehensive lists of colors, line styles, and figure customization routines.

- Write a function to plot the curve $f(x) = \frac{1}{x-1}$ on the domain $[-2, 6]$.

- a) Although $f(x)$ has a discontinuity at $x = 1$, a single call to `plt.plot()` in the usual way will make the curve look continuous. Split up the domain into $[-2, 1)$ and $(1, 6]$. Plot the two sides of the curve separately so that the graph looks discontinuous at $x = 1$.
- b) Plot both curves with a dashed magenta line. Set the keyword argument `linewidth` (or `lw`) of `plt.plot()` to 4 to make the line a little thicker than the default setting.
- c) Use `plt.xlim()` and `plt.ylim()` to change the range of the x -axis to $[-2, 6]$ and the range of the y -axis to $[-6, 6]$.

The plot should resemble the figure below.



3. Create a plot that demonstrates your newfound knowledge of plotting in Python. It should involve at least three different functions, the axis should be labeled, the plot should have a legend, and should be labeled.

Lab 3

Solving Sudoku, Part I

In this lab we will begin to write a program to solve sudoku puzzles. This week we will write a program that will solve particularly easy sudoku puzzles; next week we will expand it to solve any sudoku.

A sudoku is a 9×9 grid containing digits in some cells, as below.

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 4 | | | | 8 | 7 | | | 2 | |
| | 8 | | | | | | 4 | | |
| | | 6 | 3 | | | | 8 | | 1 |
| 7 | | | | 1 | | | | 8 | |
| 6 | 1 | 2 | | | 9 | 8 | 7 | 3 | 4 |
| | | | | | 6 | | | 1 | 9 |
| 1 | 9 | 3 | | 4 | 2 | 7 | | 5 | |
| 8 | | 7 | | | 1 | | | 3 | |
| | 2 | | | | | 3 | | | |

The goal of the puzzle is to insert numbers into the empty squares in such a way that each digit from 1 to 9 occurs exactly once in:

- each row,
- each column, and
- each of the nine 3×3 -subsquares of the puzzle.

In this lab, we will create some basic functions to solve simple sudoku puzzles. The puzzles that I supply you will consist of a string of 81 digits from 0 to 9, with 0 indicating an empty cell, and each group of nine digits representing a row in the sudoku puzzle. For instance, the sudoku puzzle above would be entered as

400870020080000400006300801700100080612098734000060019193427500807010302020003000

1. To begin, write a function `convert(string)` that converts a string of 81 digits into a one dimensional array of 81 numbers, with the first nine numbers representing the first row of the sudoku, the second nine representing the second row, etc.
2. Create a function `entry(a,b,grid)` that takes an array `grid` of the type you returned in problem 1, and returns the (a,b) entry of the Sudoku grid (where the rows and columns of the sudoku grid are labeled with integers from 1 to 9). For example, if `S = convert(string)`, where `string` is the string of length 81 listed above, then

```
>>> grid(5, 3, S)
2
```

3. Create a function `ispossible(a,b,n,grid)` that returns `True` if n is a possible value for the (a,b) entry of `grid`, and `False` otherwise. To determine if n is a possible value, it should examine all of the nonzero digits in the same column, row, and 3×3 subsquare as (a,b) ; n is a possible value if and only if it matches none of these values. If the (a,b) entry already has a nonzero value, the function `ispossible` should return an answer as if the (a,b) -entry were zero; in other words it should tell you whether n could be placed in the (a,b) -entry if that entry were empty.

```
>>> ispossible(5, 4, 5, S)
True
>>> ispossible(7, 2, 4, S)
False
```

4. Create a function `computepossibilities(a,b,grid)` that returns a list of all possible n for which `ispossible(a,b,n,grid)` is `True`.

```
>>> computepossibilities(5, 4, S)
[5]
>>> computepossibilities(9, 3, S)
[4, 5]
```

5. Create a function `insert(a,b,n,grid)` that sets the (a,b) entry of the grid to the number n (where n is a digit from 1 to 9). If the (a,b) entry of the grid is already nonzero, or if n is not a possible value that can fit into the (a,b) entry of the grid, the function should do nothing and return a value `FALSE`, otherwise it should set the (a,b) entry of the grid to n and return a value of `TRUE`.
6. Create a function `singletons(grid)` that checks each of the entries of the grid, and if the entry is empty and has only one possible value, fills it in with the only possible value. If there are no such entries, the function should return `FALSE`, otherwise it should return `TRUE`.

7. Create a procedure `simplesolve(grid)` that runs `singletons(grid)` over and over until it returns a value of `FALSE`.

At this point, the function `simplesolve(grid)` should be able to solve very simple sudoku puzzles. In addition to the puzzle above, here is a list of ten sudokus that it should be able to solve.

```
004300209005009001070060043006002087190007400050083000600000105003508690042910300
040100050107003960520008000000000017000906800803050620090060543600080700250097100
600120384008459072000006005000264030070080006940003000310000050089700000502000190
497200000100400005000016098620300040300900000001072600002005870000600004530097061
005910308009403060027500100030000201000820007006007004000080000640150700890000420
100005007380900000600000480820001075040760020069002001005039004000020100000046352
009065430007000800600108020003090002501403960804000100030509007056080000070240090
000000657702400100350006000500020009210300500047109008008760090900502030030018206
503070190000006750047190600400038000950200300000010072000804001300001860086720005
060720908084003001700100065900008000071060000002010034000200706030049800215000090
```

- It may make your debugging easier if you write a function `printgrid(grid)` which prints the sudoku grid in a human readable form (i.e. as a 3×3 grid of 3×3 grids, with appropriate spacing). You do not need to write this function if you don't want to, and you will not be graded on it.
- All access to entries of the grid (either reading them or writing them) should be done via the function `entry(a,b,grid)` and `insert(a,b,n,grid)`. This will make the programs look more uniform, and allow for changing the data structure at a later time, without having to rewrite the program.
- Of the functions described above, the most crucial, and probably the longest, will be `impossible(a,b,n,grid)`. In particular, determining the entries of the 3×3 subsquare containing the (a,b) entry can be tricky. One possibility is to set $r=(a-1)//3$, $s=(b-1)//3$, and then the 3×3 subsquare will consist of the $(3r+i, 3s+j)$ entries with $1 \leq i, j < 4$.
- If you have extra time, write a function that looks at each column to see if there is a number that can only be placed in a single spot in that column, and if so, places it there. You can do the same for rows and 3×3 subsquares as well. Add these functions into your `simplesolve(grid)` function; they will speed up the solution to more difficult sudokus that we will do next week.

Lab 4

Visualizing Vector Valued Functions

This lab is an interactive Colab notebook.

Lab 5

Solving Sudoku, Part II

In this lab we will continue to modify our sudoku solving program from last week to solve *any* sudoku puzzles, not just easy ones. We will need some preliminary functions before writing the main routine.

1. Write a function `issolved(grid)` that checks to see if a sudoku grid is completely filled out (in other words, are all of the entries in the array nonzero). It should return `True` if all entries have been filled in, and `False` otherwise.
2. Write a function `issolvable(grid)` that checks to see if every entry that has not been filled out still has values that could possibly be put into it. In other words, check that for each entry in the grid that is still 0, `computepossibilities(a,b,grid)` returns a nonempty list.

We are now prepared to write a recursive function that will solve *any* solvable sudoku puzzle. Write a function `solve(grid)`, which takes as input a sudoku grid, and returns either a completely solved sudoku grid, or the value `False` if it fails to solve the grid. Here is some pseudocode, with a lengthier explanation below it (read all of this lab before starting to write your program).

```
def solve(grid):

    run simplesolve(grid)

    if grid is solved:
        return grid

    if grid is not solvable:
        return False

    find an entry (a,b) of the grid which is still 0

    possibilities = list of all possible entries for the (a,b) slot
```

```
for each possibility n:

    make a copy of grid, called newgrid
    insert n into the (a,b) entry of newgrid

    if newgrid is solvable:
        newgrid = solve(newgrid)
        if not(newgrid is False):
            return newgrid

return False
```

The following notes clarify some of the steps above.

- (a) After running `simplesolve` and checking if the grid is solvable, we have a grid for which each unfilled entry has multiple possibilities (since if any entry had only one possibility, `simplesolve(grid)` would have filled it in). The function should find some entry of the grid which has not been filled out, and store the coordinates of the entry in the variables `a` and `b`. It should also compute the possible numbers that could fit into the (a, b) entry, and store them in the variable `possibilities`.

Note: If you want your program to run faster, don't just find an arbitrary entry that hasn't been filled out; instead find the one with the smallest number of possibilities. This will reduce the size of the next loop.

- (b) We will now loop through the entries of `possibilities`, one at a time. For each entry:
- (i) Make a copy of the sudoku grid, which we will call `newgrid` (be sure not to just use the command `newgrid=grid`, since this will only make a new reference to `grid`, not actually a new copy).
 - (ii) Use the `insert` command (from last week) to set the (a, b) entry of `newgrid` to equal the k th entry of `possibilities`.
 - (iii) Check if `newgrid` is solvable. If so, run the command `newgrid=solve(newgrid)`. Otherwise do nothing.
 - (iv) After running `newgrid=solve(newgrid)`, either `newgrid` will be a solved sudoku grid, or the value `False`. If `newgrid` is not `False`, then `return newgrid`. If `newgrid` is `False`, there is no need to do anything; this means that this possibility was not the correct entry; it does not lead to a solution.

There are two ways that we can exit this loop; we may have returned a solved sudoku (in which case, any code after the loop will never be reached), or we may have tried each possible value in the (a, b) entry, and found that none of them lead to a solution, in which case the code after the loop will be run.

At this point, the function `solve(grid)` should return a solved sudoku grid for any sudoku puzzle that you put in that has a valid solution. Here are some puzzles for you to try it on. Some of them may take your program a while to solve.

```

003000900020904060700050003010305080006080300050209070500010008080703010009000400
980100000501000000067089000300206400004000800005708006000350260000000301000002045
010702080300050001006000700600070004090423060400080009003000900800040002050607040
903050400060008010100000007000070090300809005050060000200000006080400030006030801
4000300000006008000000000001000050090080000600070200000000102700503000040900000000,
708000300000201000500000000040000026300080000000100090090600004000070500000000000
708000300000601000500000000040000026300080000000100090090200004000070500000000000
3070400000000000091800000000400000700000160000000250000000000380090000500020600000
50070060000380000000000020062040000000000091700000000000035080400000100000090000
4007006000038000000000002006205000000000009170000000000043080500000100000090000
040010200000009070010000000000430600800000050000200000705008000000600300900000000
705000002000401000300000000010600400200050000000000090000370000080000600090000080
000000410900300000300050000048007000000000062010000000600200005070000800000090000
705000002000401000300000000010600400200050000000000090000370000090000800080000060
080010000005000030000000400000605070890000200000300000200000109006700000000400000
809000300000701000500000000070000026300090000000100040060200004000080500000000000
708000300000601000400000000060000025300080000000100090090500002000070400000000000

```

Remarks:

- Unlike `simplesolve(grid)`, the function `solve(grid)` must return a value, not just change the entries of `grid`. This is because many (perhaps most) of the values that it inserts into `grid` will be guesses that are incorrect. If we actually changed the original grid with these guesses, we would have to keep track of them and be able to erase them.
- The logic of the recursive `solve(grid)` command does not actually require the function `simplesolve(grid)` at all (however, it will run faster with `simplesolve`). Try commenting out the `simplesolve(grid)` command at the beginning of the function, and see what happens. (This will work better if you programmed `solve(grid)` to look for the unfilled entry with the fewest possibilities).
- Be sure that you understand the logic of the program. Bonus: can your program be modified so that it could find multiple solutions to a badly formed sudoku puzzle? Try it on the following puzzle, which has two solutions.

```

407200000100400005000016098620300040300900000001072600002005870000600004530097061

```

Lab 6

Planes, Trains, and Automobiles

(Or how what was once “useless math” makes beautifully curved objects). Over a century ago Bernstein introduced his now famous set of polynomials defined as

Definition 1. Given $n \in \mathbb{N}$ the Bernstein polynomials $\{B_j^n(x)\}_{j=0}^n$ of degree n are defined as

$$B_j^n(x) = \binom{n}{j} x^j (1-x)^{n-j}, \quad \text{where} \quad \binom{n}{j} = \frac{n!}{j!(n-j)!}. \quad (6.1)$$

These were originally designed to give a constructive proof of the Stone–Weierstrass approximation theorem, which guarantees that any continuous function on a closed interval can be approximated as closely as desired by polynomials – essentially deep theoretical math with limited application. However, in the early 1960s two French engineers, de Casteljau and Bézier, independently used Bernstein polynomials to address a perplexing problem in the automotive design industry, namely how to specify “free-form” shapes consistently. Their solution was to use Bernstein polynomials to describe parametric curves and surfaces with a finite number of points that control the shape of the curve or surface. This enabled the use of computers to sketch the curves and surfaces. The resulting curves are called *Bézier curves* and are of the form

$$r(t) = \sum_{j=0}^n \mathbf{p}_j B_j^n(t), \quad t \in [0, 1],$$

where $\mathbf{p}_0, \dots, \mathbf{p}_n$ are the control points and $B_j^n(t)$ are the Bernstein polynomials. In this lab we will play with these Bézier curves, and a slick algorithm to create them called the de Casteljau algorithm.

At first glance, the formula for $r(t)$ looks complicated to calculate. De Casteljau gave a geometric algorithm to trace out a Bézier curve for an ordered set of n control points. We will parameterize this curve between the first and last point with respect to t , letting t go from 0 to 1. To evaluate the curve at a given time T_0 do the following:

- Parameterize the lines between each of the points letting t go from 0 to 1. For the points P_n and P_{n+1} this parameterization is $P_n(1-t) + P_{n+1}t$.
- Evaluate each of these parameterizations at t_0 . Note that this gives $n-1$ new points.




Figure 6.1: An illustration of the de Casteljau algorithm for $t_0 = 0.25$. The blue lines are what is called the *control polygon*. The green lines are the first parametrization, the red line is the second, and the red point is the point on the curve when $t_0 = 0.25$






Figure 6.2: $t_0 = .5$

- Repeat this process on the list of points until only one point is left. This is the value of the Bézier curve for these control points at the parameter t_0 .

This is what is known as De Casteljau's Algorithm. It can be illustrated as follows:

Figures 6.1 through 6.4 are illustrations of De Casteljau's Algorithm for 4 points at $t_0 = .25, .5, .75$, and 1 respectively.

1. Implement De Casteljau's algorithm by writing a function that returns the location of a point on the Bézier curve. Your function should accept a $m \times n$ array representing m control points, and a time t_0 at which to evaluate the curve. This looks much more scary than it actually is!
2. Create a function to plot a Bézier curve in 2-D. Your function should accept a $m \times 2$ array representing m control points. Using what you learned the first week in lab (i.e., `linspace()`, plotting tools) to create a curve that looks smooth.

Figure 6.3: $t_0 = .75$ Figure 6.4: $t_0 = 1$.

Lab 7

Modular Exponentiation

In today's lab we will write a program that computes a huge power of an integer modulo n for some n . The big idea is to use the binary representation of the exponent to do this very quickly. For example,

$$73 = 1 \cdot 2^6 + 0 \cdot 2^5 + 0 \cdot 2^4 + 1 \cdot 2^3 + 0 \cdot 2^2 + 0 \cdot 2^1 + 1 \cdot 2^0 = 2^6 + 2^3 + 2^0 = 64 + 8 + 1,$$

so to compute, say, 6^{73} modulo 11, we can do the following procedure:

$$\begin{aligned} 6^1 &\equiv 6 \pmod{11} * \\ 6^2 &\equiv 36 \equiv 3 \pmod{11} \\ 6^4 &\equiv (6^2)^2 \equiv 3^2 \equiv 9 \pmod{11} \\ 6^8 &\equiv (6^4)^2 \equiv 9^2 \equiv 81 \equiv 4 \pmod{11} * \\ 6^{16} &\equiv (6^8)^2 \equiv 4^2 \equiv 16 \equiv 5 \pmod{11} \\ 6^{32} &\equiv (6^{16})^2 \equiv 5^2 \equiv 25 \equiv 3 \pmod{11} \\ 6^{64} &\equiv (6^{32})^2 \equiv 3^2 \equiv 9 \pmod{11} * \end{aligned}$$

In each step, we square the result of the previous step and reduce modulo 11 so that we are never working with numbers that are too large. Then using the results from the $*$ equations above we find that

$$6^{73} \equiv 6^{64} \cdot 6^8 \cdot 6^1 \equiv 9 \cdot 4 \cdot 6 \equiv 9 \cdot 24 \equiv 9 \cdot 2 \equiv 18 \equiv 7 \pmod{11}.$$

Note that this method uses only 8 multiplications mod 11 to compute 6^{73} mod 11.

1. Write a function `power(a,b,n)` that will compute a^b in $\mathbb{Z}/n\mathbb{Z}$. Have it return the answer as an integer k with $0 \leq k < n$. It should work for very large values of a , b and n . It should **not** use the built in python function `pow(a,b,n)`.

Use it to compute $2^{1234567890}$ in $\mathbb{Z}/n\mathbb{Z}$ for $n = 1234567891$

Note: To compute the binary representation of a number, you can either use the function you wrote last semester or use the Python built-in `bin(n).replace("0b", "")`

2. Recall that the set

$$(\mathbb{Z}/n\mathbb{Z})^\times = \{\bar{a} : 1 \leq a < n : \gcd(a, n) = 1\}$$

is a group. (If you don't know what a group is, that's fine, you can still do this problem.) Define a function $\varphi : \mathbb{N} \rightarrow \mathbb{N}$ by the rule $\varphi(n) = |(\mathbb{Z}/n\mathbb{Z})^\times|$. This function is called the Euler phi function.

Write a function `eulerphi(n)` that determines the size of \mathbb{Z}_n^\times . Test your code by finding $|\mathbb{Z}_{11}^\times| = 10$, $|\mathbb{Z}_{56}^\times| = 24$, and $|\mathbb{Z}_{120}^\times| = 32$. What is $\varphi(p)$ when p is a prime?

3. Write a function `order(a,n)` that returns the order of a in the multiplicative group $(\mathbb{Z}/n\mathbb{Z})^\times$. (This is the smallest positive integer k such that $\bar{a}^k = \bar{1}$ in $\mathbb{Z}/n\mathbb{Z}$.) If a is not relatively prime to n , your function should raise the error "Number not relatively prime to modulus".

```
>>> order(4,11)
5
>>> order(33,56)
6
>>> order(7,120)
4
```

4. Write a function `gen(n)` that determines if $(\mathbb{Z}/n\mathbb{Z})^\times$ is cyclic and returns a generator g if it is, or `False` otherwise. (It is cyclic if there is an integer g , called a generator, such that the order of \bar{g} in $(\mathbb{Z}/n\mathbb{Z})^\times$ is the size of $(\mathbb{Z}/n\mathbb{Z})^\times$.) For each number n between 2 and 100, use your `gen` function to determine whether $(\mathbb{Z}/n\mathbb{Z})^\times$ is cyclic, and print your results in the following format:

```
2 1
3 2
4 3
5 2
6 5
7 3
8 False
:
```

where each line is of the form `n g`, where g is a generator of $(\mathbb{Z}/n\mathbb{Z})^\times$ if $(\mathbb{Z}/n\mathbb{Z})^\times$ is cyclic, and just `n False` if $(\mathbb{Z}/n\mathbb{Z})^\times$ is not cyclic.

Lab 8

Plotting surfaces in 3-D

It has often been said that a picture is worth 1,000 words. This is still true in Math 314, however it is often no longer quite as easy to draw the pictures we need since our data is often in 3 dimensions. Visualizing a 3-dimensional surface or a function $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ (two inputs, one output) requires a different kind of plot than the line plots we have seen thus far. The good news, however, is that the process to create such a plot is similar to creating a line plot but requires slightly more setup: first construct an appropriate domain, then calculate the image of the function on that domain.

NumPy's `np.meshgrid()` function is the standard tool for creating a 2-dimensional domain in the Cartesian plane. Given two 1-dimensional coordinate arrays, `np.meshgrid()` creates two corresponding coordinate matrices. See Figure 8.1.




Figure 8.1: `np.meshgrid(x, y)`, returns the arrays `X` and `Y`. The returned arrays give the x - and y -coordinates of the points in the grid formed by `x` and `y`. Specifically, the arrays `X` and `Y` satisfy $(X[i, j], Y[i, j]) = (x[i], y[j])$.

```

x, y = [0, 1, 2], [3, 4, 5]      # A rough domain over [0,2]x[3,5].
X, Y = np.meshgrid(x, y)        # Combine the 1-D data into 2-D data.
for xrow, yrow in zip(X,Y):
    ...      print(xrow, yrow, sep='\t')
    ...
[0 1 2]      [3 3 3]
[0 1 2]      [4 4 4]
[0 1 2]      [5 5 5]

```

With a 2-dimensional domain, $g(x, y)$ is usually visualized with one of three kinds of plots.

- A *heat map* assigns a color to each point in the domain, producing a 2-dimensional colored picture describing a 3-dimensional shape. Darker colors typically correspond to lower values while lighter colors typically correspond to higher values.

Use `plt.pcolormesh()` to create a heat map.

- A *contour map* draws several *level curves* of g on the 2-dimensional domain. A level curve corresponding to the constant c is the collection of points $\{(x, y) \mid c = g(x, y)\}$. Coloring the space between the level curves produces a discretized version of a heat map. Including more and more level curves makes a filled contour plot look more and more like the complete, blended heat map.

Use `plt.contour()` to create a contour plot and `plt.contourf()` to create a filled contour plot. Specify either the number of level curves to draw, or a list of constants corresponding to specific level curves.

- A *wireframe map* creates the surface in 3-dimensions using a wire grid. The wireframe allows one to "look through" the plot to see different characteristics of the surface

Use `ax.plot_wireframe()` to create a wireframe plot.

These functions each receive the keyword argument `cmap` to specify a color scheme (some of the better schemes are "viridis", "magma", and "coolwarm"). For the list of all Matplotlib color schemes, see http://matplotlib.org/examples/color/colormaps_reference.html.

Finally, `plt.colorbar()` draws the color scale beside the plot to indicate how the colors relate to the values of the function.

```

# Create a 2-D domain with np.meshgrid().
x = np.linspace(-np.pi, np.pi, 100)
y = x.copy()
X, Y = np.meshgrid(x, y)
Z = np.sin(X) * np.sin(Y)      # Calculate g(x,y) = sin(x)sin(y).

# Plot the heat map of f over the 2-D domain.
plt.subplot(131)
plt.pcolormesh(X, Y, Z, cmap="viridis")
plt.colorbar()
plt.xlim(-np.pi, np.pi)
plt.ylim(-np.pi, np.pi)

# Plot a contour map of f with 10 level curves.
plt.subplot(132)
plt.contour(X, Y, Z, 10, cmap="coolwarm")
plt.colorbar()

#plot a wireframe map, specifying the strides
fig = plt.figure()
ax = fig.add_subplot(133, projection='3d')
# Plot a basic wireframe.
ax.plot_wireframe(X, Y, Z, rstride=5, cstride=5)
plt.show()

```




Figure 8.2: heatmap, contour, and wireframe plots for the function $\sin(x)\sin(y)$.

1. Make a heatmap plot to help you visualize the following functions:

a) $f(x, y) = \frac{x^2 - y^2}{x^2 + y^2}$

b) $f(x, y) = xy^2 - x^3$ (monkey saddle)

c) $f(x, y) = \frac{2x^2 + 3xy + 4y^2}{3x^2 + 5y^2}$

2. Repeat exercise 1, but this time make contour plots of the functions.
3. Repeat exercise 1, but this time make wireframe plots of the functions.
4. Select a function of two variables that you find interesting and create a heatmap, contour, and wireframe plot for the function.
5. The following function represents the continuous topography of a mountain range, giving the elevation $h(x, y)$ at each point (x, y) :

$$h(x, y) = \left(5000 - \frac{x^2 + y^2 + xy}{200} + \frac{25(x + y)}{2} \right) e^{-\left| \frac{x^2 + y^2}{1000000} - \frac{3(x + y)}{2000} + \frac{7}{10} \right|}.$$

Here is some Python-friendly code you can copy and paste (you'll need to modify `exp` and `abs` depending on how you imported numpy):

```
( 5000-0.005*(x**2+y**2+x*y)+12.5*(x+y) )
* exp( -abs(0.000001*(x**2+y**2)-0.0015*(x+y)+0.7) )
```

Make a wireframe plot to help you visualize this function. For this whole problem, keep the x and y values in the box

$$0 \leq x, y \leq 1600.$$

6. (bonus) A mosquito wants to fly from the starting point $A = (200, 200, h(200, 200))$ to the ending point $B = (1400, 1400, h(1400, 1400))$. It will do this by first flying straight up to the point $(200, 200, H)$, then, maintaining a constant elevation of H , it will fly around any obstacles in its way until it reaches $(1400, 1400, H)$, at which point it will fly straight down to B . What is the minimum value of H so that this is possible? Keep in mind that the x, y values of the mosquito's position must always satisfy the condition $0 \leq x, y \leq 1600$. (You may assume that the mosquito takes up no physical space, so can fit through gaps of arbitrarily small size.)

Hint: what is the maximum value of $h(0, y)$ (or $h(x, 0)$)?

Lab 9

Maximal Paths

In this lab, we will write programs to find paths through arrays of numbers that maximize certain sums. As an example, we present the triangle below:

| | | | | | | | | | | | | |
|----|----|----|----|----|-----------|-----------|-----------|-----------|-----------|-----------|-----------|----|
| | | | | | | 19 | | | | | | |
| | | | | | 41 | | 30 | | | | | |
| | | | 28 | | 39 | | 62 | | | | | |
| | | 96 | | 42 | | 43 | | 19 | | | | |
| | | 56 | 31 | | 09 | | 14 | | 22 | | | |
| | 81 | | 87 | | 72 | | 25 | 44 | | 98 | | |
| | 08 | 40 | 15 | | 08 | | 49 | | 98 | 27 | | |
| 69 | | 02 | 23 | | 79 | | 26 | | 34 | 67 | 41 | |
| 72 | 57 | 06 | 29 | | 63 | | 39 | | 57 | | 58 | 76 |
| 64 | 11 | 39 | 82 | 97 | | 61 | 83 | 82 | | 53 | | 97 |

Our goal is to move downward from the top entry of the triangle to the bottom, choosing to go left or right as we go from each row to the next, and maximize the sum of the numbers on our path. For instance, the path going 19-41-39-43-14-44-98-67-58-82 (indicated in boldface above) has a sum of 505. It is not, however the path having the highest sum.

Triangles like the one above with 10, 15, 20, 25, 30, 100, and 200 rows are available online, at

math.byu.edu/~nick/triangles.txt

This file consists of n rows of positive numbers, each two digits, with a space between the numbers. The maximum path sums for the first four triangles are 581, 1060, 1437, 1918, respectively.

1. Write a function `all_sums(string)` that will take as input a string representing a triangle of at most 10 rows and return a list of all the sums obtained by valid paths (sorted in increasing order, without repetition).

The first few lines of your function should be:

```
lines = string.split('\n')
str_grid = [line.split(' ') for line in lines]
grid = [[int(entry) for entry in line] for line in str_grid]
```

2. Write a function `max_path(string)` that will take as input a string representing any of the triangles given, and return the largest sum of a valid path. For small triangles, you can do this using the brute-force method of problem 1. But in order for this function to run efficiently for larger triangles, you will need to find the maximum path size more cleverly. *Hint:* consider the second row up from the bottom in the triangle above. For each entry in that row the maximum path size from that entry to the bottom row can be determined easily; you either move left or right. Now replace each entry in that row with the maximum path size from that entry to the bottom row:

72 + 64 57 + 39 06 + 82 29 + 97 63 + 97 39 + 83 57 + 83 58 + 82 76 + 97

You can now ignore the bottom row and think of this as a triangle with one fewer row. Now iterate.

3. Write a function `max_path_rect(string)` to find the path with largest sum for a rectangular array of numbers, where you wish to go from the upper left to the lower right corner, and you can only move right or down. Test your program on the rectangles of various sizes found at

math.byu.edu/~nick/rectangles.txt

The maximum path sums for the first two rectangles are 1401 and 2904, respectively.

| | | | | | | | | | | | | | | | |
|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| 47 | 54 | 58 | 48 | 66 | 30 | 06 | 06 | 33 | 76 | 22 | 35 | 77 | 22 | 01 | 26 |
| 56 | 31 | 22 | 31 | 74 | 98 | 39 | 67 | 21 | 58 | 71 | 90 | 74 | 80 | 63 | 85 |
| 36 | 27 | 70 | 06 | 50 | 58 | 62 | 88 | 58 | 13 | 25 | 80 | 13 | 36 | 14 | 51 |
| 40 | 89 | 57 | 75 | 76 | 55 | 78 | 71 | 35 | 20 | 64 | 23 | 35 | 60 | 82 | 10 |
| 62 | 21 | 95 | 58 | 69 | 53 | 81 | 63 | 52 | 75 | 82 | 65 | 58 | 66 | 49 | 47 |
| 79 | 82 | 29 | 03 | 37 | 87 | 15 | 93 | 48 | 50 | 49 | 91 | 92 | 18 | 39 | 02 |
| 35 | 35 | 77 | 10 | 30 | 43 | 45 | 46 | 85 | 42 | 25 | 80 | 10 | 55 | 71 | 01 |
| 75 | 19 | 63 | 56 | 10 | 82 | 56 | 19 | 48 | 58 | 48 | 20 | 83 | 93 | 89 | 50 |
| 18 | 01 | 95 | 10 | 14 | 69 | 53 | 40 | 28 | 34 | 41 | 24 | 26 | 60 | 23 | 59 |
| 13 | 73 | 56 | 07 | 49 | 05 | 99 | 27 | 16 | 11 | 79 | 53 | 48 | 47 | 44 | 45 |
| 61 | 17 | 40 | 63 | 15 | 49 | 92 | 06 | 68 | 11 | 21 | 68 | 55 | 65 | 97 | 08 |
| 56 | 54 | 41 | 75 | 02 | 74 | 32 | 70 | 85 | 39 | 83 | 51 | 21 | 29 | 89 | 17 |
| 78 | 05 | 89 | 45 | 66 | 19 | 75 | 53 | 55 | 54 | 84 | 57 | 18 | 19 | 80 | 85 |
| 68 | 74 | 54 | 41 | 95 | 79 | 74 | 76 | 30 | 83 | 88 | 16 | 52 | 54 | 90 | 75 |
| 71 | 34 | 83 | 38 | 91 | 81 | 34 | 66 | 39 | 68 | 38 | 74 | 46 | 47 | 42 | 94 |
| 63 | 62 | 33 | 37 | 77 | 52 | 08 | 42 | 22 | 65 | 36 | 80 | 17 | 54 | 25 | 98 |

The path above:

47-56-36-40-89-57-75-76-69-53-87-43-82-69-53-99-
-92-32-75-74-76-66-42-22-65-36-80-17-54-25-98

has sum 1885.

Lab 10

Symbolic Python and Partial Differentiation

Chances are you have used a tool similar to Wolfram Alpha to simplify, integrate, or differentiate a complicated algebraic expression. Although Python is primarily used for crunching, relating, or visualizing numerical data, using the SymPy module one can also do symbolic mathematics in Python, including algebra, differentiation, integration, and more. In this lab we introduce SymPy syntax and emphasize how to use symbolic algebra for numerical computing.

Symbolic Variables and Expressions

Most variables in Python refer to a number, string, or data structure. Doing computations on such variables results in more numbers, strings, or data structures. A *symbolic variable* is a variable that represents a mathematical symbol, such as x or θ , not a number or another kind of data. Operating on symbolic variables results in an *expression*, representative of an actual mathematical expression. For example, if a symbolic variable Y refers to a mathematical variable y , the multiplication $3*Y$ refers to the expression $3y$. This is all done without assigning an actual numerical value to Y .

SymPy is Python's library for doing symbolic algebra and calculus. It is typically imported with `import sympy as sy`, and symbolic variables are usually defined using `sy.symbols()`.

```

import sympy as sy
x0 = sy.symbols('x0')                # Define a single variable.

# Define multiple symbolic variables simultaneously.
x2, x3 = sy.symbols('x2, x3')        # Separate symbols by commas,
m, a = sy.symbols('mass acceleration') # by spaces,
x, y, z = sy.symbols('x:z')          # or by colons.
x4, x5, x6 = sy.symbols('x4:7')

# Combine symbolic variables to form expressions.
expr = x**2 + x*y + 3*x*y + 4*y**3
force = m * a
print(expr, force, sep='\n')

```

SymPy has its own version for each of the standard mathematical functions like $\sin(x)$, $\log(x)$, and \sqrt{x} , and includes predefined variables for special numbers such as π . The naming conventions for most functions match NumPy, but some of the built-in constants are named slightly differently.

| | | | | | | |
|-----------|-----------------------|---------------------------|-------------------------|-------------------|-----------------------|-------------------------|
| Functions | $\sin(x)$ sy.sin() | $\arcsin(x)$ sy.asin() | $\sinh(x)$ sy.sinh() | e^x sy.exp() | $\log(x)$ sy.log() | \sqrt{x} sy.sqrt() |
| Constants | π sy.pi | e sy.E | $i = \sqrt{-1}$ sy.I | ∞ sy.oo | | |

Other trigonometric functions like $\cos(x)$ follow the same naming conventions. For a complete list of SymPy functions, see <http://docs.sympy.org/latest/modules/functions/index.html>.

1. Write a function that returns the expression $\frac{2}{5}e^{x^2-y}\cosh(x+y) + \frac{3}{7}\log(xy+1)$ symbolically. Make sure that the fractions remain symbolic.

Evaluating Expressions

Every SymPy expression has a `subs()` method that substitutes one variable for another. The result is usually still a symbolic expression, even if a numerical value is used in the substitution. The `evalf()` method actually evaluates the expression numerically after all symbolic variables have been assigned a value. Both of these methods can accept a dictionary to reassign multiple symbols simultaneously.

```

>>> x,y = sy.symbols('x y')
>>> expr = sy.expand((x + y)**3)
>>> print(expr)
x**3 + 3*x**2*y + 3*x*y**2 + y**3

# Replace the symbolic variable y with the expression 2x.
>>> expr.subs(y, 2*x)
27*x**3

# Replace x with pi and y with 1.
>>> new_expr = expr.subs({x:sy.pi, y:1})
>>> print(new_expr)
1 + 3*pi + 3*pi**2 + pi**3
>>> new_expr.evalf() # Numerically evaluate the expression
71.0398678443373

# Evaluate the expression by providing values for each variable.
>>> expr.evalf(subs={x:1, y:2})
27.000000000000000

```

These operations are good for evaluating an expression at a single point, but it is typically more useful to turn the expression into a reusable numerical function. To this end, `sy.lambdify()` takes in a symbolic variable (or list of variables) and an expression, then returns a callable function that corresponds to the expression.

```

# Turn the expression sin(x)^2 into a function with x as the variable.
>>> f = sy.lambdify(x, sy.sin(x)**2)
>>> print(f(0), f(np.pi/2), f(np.pi), sep=' ')
0.0 1.0 1.4997597826618576e-32

# Lambdify a function of several variables.
>>> f = sy.lambdify((x,y), sy.sin(x)**2 + sy.cos(y)**2)
>>> print(f(0,1), f(1,0), f(np.pi, np.pi), sep=' ')
0.2919265817264289 1.708073418273571 1.0

```

It is almost always computationally cheaper to lambdify a function than to use substitutions.

1. The Maclaurin series up to order N for e^x is defined as

$$e^x \approx \sum_{n=0}^N \frac{x^n}{n!}. \quad (10.1)$$

Write a function that accepts an integer N . Define an expression for (10.1), then substitute in $-y^2$ for x to get a truncated Maclaurin series of e^{-y^2} . Lambdify the

resulting expression and plot the series on the domain $y \in [-2, 2]$. Plot e^{-y^2} over the same domain for comparison.

(Hint: use `sy.factorial()` to compute the factorial.)

Call your function with increasing values of N to check that the series converges correctly.

Calculus

SymPy is also equipped to perform standard calculus operations, including derivatives, integrals, and taking limits. Like other elements of SymPy, calculus operations can be temporally expensive, but they give exact solutions whenever solutions exist.

Differentiation

The command `sy.Derivative()` creates a closed form, unevaluated derivative of an expression. This is like putting $\frac{d}{dx}$ in front of an expression without actually calculating the derivative symbolically. The resulting expression has a `doit()` method that can be used to evaluate the actual derivative. Equivalently, `sy.diff()` immediately takes the derivative of an expression.

Both `sy.Derivative()` and `sy.diff()` accept a single expression, then the variable or variables that the derivative is being taken with respect to.

```
>>> x, y = sy.symbols('x y')
>>> f = sy.sin(y)*sy.cos(x)**2

# Make an expression for the derivative of f with respect to x.
>>> df = sy.Derivative(f, x)
>>> print(df)
Derivative(sin(y)*cos(x)**2, x)

>>> df.doit()                                # Perform the actual differentiation.
-2*sin(x)*sin(y)*cos(x)

# Alternatively, calculate the derivative of f in a single step.
>>> sy.diff(f, x)
-2*sin(x)*sin(y)*cos(x)

# Calculate the derivative with respect to x, then y, then x again.
>>> sy.diff(f, x, y, x)
2*(sin(x)**2 - cos(x)**2)*cos(y)    # Note this expression could be simplified.
```

1. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a smooth function. A *critical point* of f is a number $x_0 \in \mathbb{R}$ satisfying $f'(x_0) = 0$. The second derivative test states that a critical point x_0 is a

local minimum of f if $f''(x_0) > 0$, or a local maximum of f if $f''(x_0) < 0$ (if $f''(x_0) = 0$, the test is inconclusive).

Now consider the polynomial

$$p(x) = 2x^6 - 51x^4 + 48x^3 + 312x^2 - 576x - 100.$$

Use SymPy to find all critical points of p and classify each as a local minimum or a local maximum. Plot $p(x)$ over $x \in [-5, 5]$ and mark each of the minima in one color and the maxima in another color. Return the collections of local minima and local maxima as sets.

2. Calculate the partial derivatives with respect to x and y of the following functions

a) $f(x, y) = \frac{x^2 - y^2}{x^2 + y^2}$

b) $f(x, y) = xy^2 - x^3$ (monkey saddle)

c) $f(x, y) = \frac{2x^2 + 3xy + 4y^2}{3x^2 + 5y^2}$

A more thorough treatment of symbolic python can be found at <http://acme.byu.edu/wp-content/uploads/2020/08/Sympy2020.pdf>.

Bonus material (it will mean a lot more to you later in the semester)

The *Jacobian matrix* of a multivariable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ at a point $\mathbf{x}_0 \in \mathbb{R}^n$ is the $m \times n$ matrix J whose entries are given by

$$J_{ij} = \frac{\partial f_i}{\partial x_j}(\mathbf{x}_0).$$

For example, the Jacobian for a function $f : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is defined by

$$J = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \frac{\partial f_1}{\partial x_3} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \frac{\partial f_2}{\partial x_3} \end{bmatrix}, \quad \text{where} \quad f(\mathbf{x}) = \begin{bmatrix} f_1(\mathbf{x}) \\ f_2(\mathbf{x}) \end{bmatrix}, \quad \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}.$$

To calculate the Jacobian matrix of a multivariate function with SymPy, define that function as a symbolic matrix (`sy.Matrix()`) and use its `jacobian()` method. The method requires a list of variables that prescribes the ordering of the differentiation.

```
# Create a matrix of symbolic variables.
>>> r, t = sy.symbols('r theta')
>>> f = sy.Matrix([r*sy.cos(t), r*sy.sin(t)])

# Find the Jacobian matrix of f with respect to r and theta.
>>> J = f.jacobian([r,t])
>>> J
Matrix([
  [cos(theta), -r*sin(theta)],
  [sin(theta),  r*cos(theta)]]

# Evaluate the Jacobian matrix at the point (1, pi/2).
>>> J.subs({r:1, t:sy.pi/2})
Matrix([
  [0, -1],
  [1,  0]])
```

Lab 11

One-dimensional Optimization and Gradient Descent

Most mathematical optimization problems involve estimating the minimizer(s) of a scalar-valued function. Many algorithms for optimizing functions with a high-dimensional domain depend on routines for optimizing functions of a single variable. There are many techniques for optimization in one dimension, each with varying degrees of precision and speed. In this lab we first implement Newton's method and then we attack the problem of gradient descent

Newton's Method

Newton's method is an important root-finding algorithm that can also be used for optimization. Given $f : \mathbb{R} \rightarrow \mathbb{R}$ and a good initial guess x_0 , the sequence $(x_k)_{k=1}^{\infty}$ generated by the recursive rule

$$x_{k+1} = x_k - \frac{f(x_k)}{f'(x_k)}$$

converges to a point \bar{x} satisfying $f(\bar{x}) = 0$. The first-order necessary conditions from elementary calculus state that if f is differentiable, then its derivative evaluates to zero at each of its local minima and maxima. Therefore using Newton's method to find the zeros of f' is a way to identify potential minima or maxima of f . Specifically, starting with an initial guess x_0 , set

$$x_{k+1} = x_k - \frac{f'(x_k)}{f''(x_k)} \tag{11.1}$$

and iterate until $|x_k - x_{k-1}|$ is satisfactorily small. Note that this procedure does not use the actual function f at all, but it requires many evaluations of its first and second derivatives. As a result, Newton's method converges in few iterations, but it can be computationally expensive.

Newton's method for optimization works well to locate minima when $f''(x) > 0$ on the entire domain. However, it may fail to converge to a minimizer if $f''(x) \leq 0$ for some portion of the domain. If f is not unimodal, the initial guess x_0 must be sufficiently close to a local minimizer x^* in order to converge.

1. Let $f : \mathbb{R} \rightarrow \mathbb{R}$. Write a function that accepts f' , f'' , a starting point x_0 , a stopping tolerance `tol`, and a maximum number of iterations `maxiter`. Implement Newton's method using (11.1) to locate a local optimizer. Return the approximate optimizer, whether or not the algorithm converged, and the number of iterations computed.
2. Test your function from problem 1 by minimizing $f(x) = x^2 + \sin(5x)$ with an initial guess of $x_0 = 0$. Compare your results to `scipy.optimize.newton()`, which implements the root-finding version of Newton's method.

```
>>> df = lambda x : 2*x + 5*np.cos(5*x)
>>> d2f = lambda x : 2 - 25*np.sin(5*x)
>>> opt.newton(df, x0=0, fprime=d2f, tol=1e-10, maxiter=500)
-1.4473142236328096
```

Note that other initial guesses can yield different minima for this function.

Descent Methods

Consider now a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. *Descent methods*, also called *line search methods*, are optimization algorithms that create a convergent sequence $(x_k)_{k=1}^\infty$ by the following rule.

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k \quad (11.2)$$

Here $\alpha_k \in \mathbb{R}$ is called the *step size* and $\mathbf{p}_k \in \mathbb{R}^n$ is called the *search direction*. The choice of \mathbf{p}_k is usually what distinguishes an algorithm; in the one-dimensional case ($n = 1$), $p_k = f'(x_k)/f''(x_k)$ results in Newton's method.

To be effective, a descent method must also use a good step size α_k . If α_k is too large, the method may repeatedly overstep the minimum; if α_k is too small, the method may converge extremely slowly. See Figure 11.1.

Given a search direction \mathbf{p}_k , the best step size α_k minimizes the function $\phi_k(\alpha) = f(\mathbf{x}_k + \alpha \mathbf{p}_k)$. Since f is scalar-valued, $\phi_k : \mathbb{R} \rightarrow \mathbb{R}$, so Newton's method (or any other 1-D optimization method) can be used to minimize ϕ_k .

The Method of Steepest Descent

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with first derivative $Df : \mathbb{R}^n \rightarrow \mathbb{R}^n$. The following iterative technique is a common template for methods that aim to compute a local minimizer \mathbf{x}^* of f .

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k \quad (11.3)$$

Here \mathbf{x}_k is the k th approximation to \mathbf{x}^* , α_k is the *step size*, and \mathbf{p}_k is the *search direction*. Newton's method and its relatives follow this pattern, but they require the calculation (or approximation) of the inverse Hessian matrix $Df^2(\mathbf{x}_k)^{-1}$ at each step. The following idea is a simpler and less computationally intensive approach than Newton and quasi-Newton methods.




Figure 11.1: If the step size α_k is too large, a descent method may repeatedly overstep the minimizer.

The derivative $Df(\mathbf{x})^T$ (often called the *gradient* of f at \mathbf{x} , sometimes notated $\nabla f(\mathbf{x})$) is a vector that points in the direction of greatest **increase** of f at \mathbf{x} . It follows that the negative derivative $-Df(\mathbf{x})^T$ points in the direction of steepest **decrease** at \mathbf{x} . The *method of steepest descent* chooses the search direction $\mathbf{p}_k = -Df(\mathbf{x}_k)^T$ at each step of (11.3), resulting in the following algorithm.

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k Df(\mathbf{x}_k)^T \quad (11.4)$$

Setting $\alpha_k = 1$ for each k is often sufficient for Newton and quasi-Newton methods. However, a constant choice for the step size in (11.4) can result in oscillating approximations or even cause the sequence $(\mathbf{x}_k)_{k=1}^\infty$ to travel away from the minimizer \mathbf{x}^* . To avoid this problem, the step size α_k can be chosen in a few ways.

- Start with $\alpha_k = 1$, then set $\alpha_k = \frac{\alpha_k}{2}$ until $f(\mathbf{x}_k - \alpha_k Df(\mathbf{x}_k)^T) < f(\mathbf{x}_k)$, terminating the iteration if α_k gets too small. This guarantees that the method actually descends at each step and that α_k satisfies the Armijo rule, without endangering convergence.
- At each step, solve the following one-dimensional optimization problem.

$$\alpha_k = \underset{\alpha}{\operatorname{argmin}} f(\mathbf{x}_k - \alpha Df(\mathbf{x}_k)^T)$$

Using this choice is called *exact steepest descent*. This option is more expensive per iteration than the above strategy, but it results in fewer iterations before convergence.

3. Write a function that accepts an objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, its derivative $Df : \mathbb{R}^n \rightarrow \mathbb{R}^n$, an initial guess $\mathbf{x}_0 \in \mathbb{R}^n$, a convergence tolerance `tol` defaulting to $1e^{-5}$, and a maximum number of iterations `maxiter` defaulting to 100. Implement the exact

method of steepest descent, using a one-dimensional optimization method to choose the step size (use `opt.minimize_scalar()` or your own 1-D minimizer). Iterate until $\|Df(\mathbf{x}_k)\|_\infty < \text{tol}$ or $k > \text{maxiter}$. Return the approximate minimizer \mathbf{x}^* , whether or not the algorithm converged (`True` or `False`), and the number of iterations computed.

4. Test your function from Problem 3 on $f(x, y, z) = x^4 + y^4 + z^4$ (easy) and the Rosenbrock function (hard). It should take many iterations to minimize the Rosenbrock function, but it should converge eventually with a large enough choice of `maxiter`.

Lab 12

The RSA Cryptosystem

In this lab we will implement the RSA public key cryptosystem. So that we are all using the same alphabet, and so that our messages can include both uppercase and lowercase letters and some punctuation, run the following code.

```
>>> def str_to_int(s):
>>>     return int(''.join(str(ord(c)-32).rjust(2,'0') for c in s))
>>> def int_to_str(n):
>>>     return '' if n==0 else int_to_str(n//100)+chr((n%100)+32)
```

To check that everything is working:

```
>>> str_to_int("Hello, world?")
24069767679120087798276683102
>>> int_to_str(24069767679120087798276683102)
"Hello, world?"
```

We will need the following fact from group theory: for any positive integer N , let $\varphi(N)$ denote the size of the group $(\mathbb{Z}/N\mathbb{Z})^\times$, as in Lab 7. If a is relatively prime to N then

$$a^{\varphi(N)} \equiv 1 \pmod{N}.$$

(If G is a finite group of order m then $a^m = e$ for all $a \in G$.) Hence, if $k \equiv 1 \pmod{\varphi(N)}$, we can write $k = 1 + \ell\varphi(N)$, and we have

$$a^k = a^1(a^{\varphi(N)})^\ell \equiv a \pmod{N}.$$

In order to create an RSA key pair, we will find an integer N that is a product of two large primes p and q . For this value of $N = pq$, we know that $\varphi(N) = (p-1)(q-1)$. We will choose an integer e that is relatively prime to $\varphi(N)$ (usually the prime 1234577 will work). Then we compute d such that $de \equiv 1 \pmod{\varphi(N)}$. The public key is the pair

$$[e, N]$$

and the private key (which must be kept secret) is the pair

$$[d, N].$$

To encrypt a message to a person, you first look up their public key $[e, N]$. Convert your message to an integer M using `str_to_int`. Then compute

$$C \equiv M^e \pmod{N} \quad \text{with } 0 \leq C < N.$$

This number is the encrypted message that is sent.

To decrypt a message C sent to you that is encrypted as above with your public key $[e, N]$, you would compute

$$P \equiv C^d \pmod{N} \quad \text{with } 0 \leq P < N.$$

Since $C \equiv M^e \pmod{N}$, we see that $P \equiv (M^e)^d \equiv M^{de} \equiv M \pmod{N}$, since $de \equiv 1 \pmod{\phi(N)}$. Hence, you would retrieve the original message. Note that you should be the only one who knows your private key, so you should be the only one able to read the message.

If two people, Alice and Bob, each have a public/private key pair (say Alice's is $[e_A, N_A]$ and $[d_A, N_A]$ and Bob's is $[e_B, N_B]$ and $[d_B, N_B]$), then Alice can send Bob an encrypted and signed message as follows.

Let M be a number smaller than N representing the message. Then Alice encrypts it using Bob's public key $[e_B, N_B]$ to get

$$C \equiv M^{e_B} \pmod{N_B} \quad \text{with } 0 \leq C < N_B.$$

Alice will then take the encrypted text C and run it through her decryption key (which is something that only she can do) to get a signature:

$$S \equiv C^{d_A} \pmod{N_A} \quad \text{with } 0 \leq S < N_A.$$

She then sends Bob the pair of numbers (C, S) . Using his decryption key (which only he can do), Bob can decrypt C to obtain and read M . Using Alice's encryption key (which is publicly known) on S will result in a number $V \equiv S^{e_A} \pmod{N_A}$. If $V \equiv C \pmod{N_A}$, this verifies that the message comes from Alice. If $V \not\equiv C \pmod{N_A}$, then the message is not from Alice.

Your assignment:

1. Create three functions:

- (a) A function `create_key(p,q)` that takes two large primes `p` and `q`, and creates a public key $[e, N]$ and a private key $[d, N]$ where $N = pq$ and e is the smallest number larger than 1234576 that is relatively prime to $\phi(N) = (p-1)(q-1)$. The decryption exponent is unique modulo $\phi(N)$, so d should be between 0 and $\phi(N)$.
- (b) A function `encrypt(message, e_key)` that takes a string and a public encryption key, and returns an integer C that is the RSA encrypted message.
- (c) A function `decrypt(C, d_key)` that takes an integer and a private decryption key, and returns the decrypted message as a string.

2. A list of large primes (over 100 digits) has been placed online at

math.byu.edu/~doud/RSA/largeprimes

Using the first two primes from this list, create a public/private key pair. Encrypt the message “I have finished part 2 of the assignment!” using the public key. You should probably check that the private key, together with your `decrypt` function, correctly decrypts the message.

3. You have been sent an encrypted message:

math.byu.edu/~doud/RSA/keypair

Note that this page contains both your public (encryption) key, EKEY, and your private (decryption) key, DKEY. It also contains a number C , which is the encrypted message that has been sent to you, encrypted with EKEY.

Decrypt the message!

4. A message has been encrypted using a key created in exactly the way specified in Problem 2; by choosing two primes from the list of large primes, and creating a key. The public key used (and the number C representing the encrypted message) are here:

math.byu.edu/~doud/RSA/publickey

Write a function `RSA_crack(primes,C,e_key)` which takes a list of primes from which `e_key` was generated (you may assume that N is a product of two distinct primes from this list), an encrypted message C , and a public encryption key `e_key`, and returns the decrypted message.

Use your function to find the private (decryption) key used above and decrypt the message.

5. Write a function `verify_signature(C,S,e_key)` which takes as input an encrypted message C , a signature S , and the encryption key of the purported sender. Both C and S will be integers. The function should return `True` if the signature is valid, and `False` otherwise.

As before, your RSA key is found at

math.byu.edu/~doud/RSA/keypair

Two people, each claiming to be Alice, have sent you a message encrypted with your public key. Both messages (and their signatures) are contained in the web page

math.byu.edu/~doud/RSA/SignedMessages

Assuming that Alice keeps her private key secret, and using the list of public keys at

math.byu.edu/~doud/RSA/PublicKeyDatabase

decrypt both messages and check the signatures to determine which (if any) of the two messages is really from Alice. Can you tell who the other one is from?

Lab 13

Solving the Rubik's cube, Part I

The set of all possible “moves” of the Rubik's cube is a good example of a group. Each possible move can be constructed as a product of moves consisting of rotating a single face of the Rubik's cube through a 90° rotation, either clockwise or counterclockwise. We will use the following notation for these moves:

One face of the Rubik's cube will be specified as the “Up” (top) face, and another as the “Front” face. These specifications will fix a “Down” (bottom) face (the face opposite the top face), a “Back” face (opposite the front face), and a “Right” and “Left” face (on the right or left as you look at the front face, with the top face on the top). We will use the following letters to denote the different faces:

| | |
|---|-------|
| U | Top |
| F | Front |
| R | Right |
| L | Left |
| B | Back |
| D | Down |

Representing a turn of a face will be indicated by giving the letter of the face that we wish to turn, in upper case for a clockwise turn (as we look directly at the face), and in lower case for a counterclockwise turn. The turns will be listed in order from left to right in the order in which we perform them. For instance, the sequence **RFr_f** would correspond to the following:

1. Turn the right face 90 degrees clockwise.
2. Turn the front face 90 degrees clockwise.
3. Turn the right face 90 degrees counterclockwise.
4. Turn the front face 90 degrees counterclockwise.

To keep the programming easier, we will not allow ourselves to turn the entire cube (in particular the squares in the center of each face will never change position). All moves that we allow will be moves of individual faces.

We will number the smaller faces of the Rubik's cube as indicated on the handout at

<https://math.byu.edu/~doud/Math495R/Rubik-handout.pdf>

With this numbering, you should check that turning the right face 90 degrees clockwise (the move denoted by **R**) would move the small face in spot 2 to spot 18, the face in spot 18 to spot 23, the face in spot 23 to spot 7, and the face in spot 7 to spot 2. The move **R** will actually involve five such 4-cycles: we have that as a permutation of the faces,

$$R = (2, 18, 23, 7)(3, 10, 22, 15)(6, 14, 19, 11)(26, 37, 46, 35)(30, 34, 42, 38)$$

We will define a variable `cube` that keeps track of the position of each sub-face as we manipulate the cube. When the cube is solved, `cube` will be a list of length 49, with the i th entry equal to i (note that there are only 48 subfaces that can move, but in order to keep the numbering simple, we will leave the zero entry always equal to 0).

1. Find the cycle structures for the moves **R**, **L**, **U**, **D**, **F**, and **B**. As an example, the cycle representing **R** would be

`[[2, 18, 23, 7], [3, 10, 22, 15], [6, 14, 19, 11], [26, 37, 46, 35], [30, 34, 42, 38]]`.

Store these cycle structures in variables named `Right`, `Left`, `Up`, `Down`, `Front`, and `Back`. I suggest you crowdsource this part of the assignment.

2. Write a function `apply(move, cube)` that takes a list representing a movement of the cube (written as a product of cycles as in problem 1) and a list representing the cube (as described above), and performs the indicated move. It should actually change the values of the list `cube`, so that it does not actually need to return a value.
3. Write a function `execute(moves, cube)` that takes a string containing moves of the cube, and a position of the cube, and executes those moves on the cube.

The function should apply all of the indicated moves to the cube, actually changing the values of the list. For instance, for a capital letter **R**, it should apply the permutation `Right` that you found above. For a lowercase **r**, it should apply the permutation `Right` three times (since a 90° counterclockwise move is the same as three 90° clockwise moves). After applying all moves specified in the string, the function should then return the sequence of moves that it made. If `moves` contains characters other than "`FBRLUDfbrlud`", these characters should be ignored, and not included in the returned variable.

After this function is written and debugged, you should not change `cube` other than through this function (except to debug your code).

In order to test the `execute` function, use it to find the orders of the following sequences of moves: '`RF`' should have order 105 (i.e., applying the move `RF` 105 times should return the cube to its initial position, but fewer than 105 times should not), `Ur` should have order 63, `rdL` should have order 180, and `RDFLBURdf1bu` should have order 360. Note that this last sequence of moves uses all twelve possible moves, so it should be a good test for your function.

4. Write functions: `align25(cube)`, `align26(cube)`, `align27(cube)`, `align28(cube)`, that put the subfaces numbered 25, 26, 27, and 28 in their correct positions. Each one should perform its task without changing the positions of the other subfaces in the set {25, 26, 27, 28}. Each function should return the sequence of moves that it makes.

Here is my suggestion for how to write these functions.

For `align25`, create a list of 24 sequences of moves. Entry 0 of this list would be the sequence of moves necessary to move subface 25 from position 25 to position 25; entry 1 of this list would be the sequence of moves necessary to move subface 25 from position 26 to position 25, and so on. Find the current position of subface 25, lookup the correct sequence of moves in the list, and execute the moves.

For `align25`, here is what my function looks like:

```
def align25(cube):
    i=cube.index(25)
    move=['25', 'u26', 'uu27', 'U28', 'BLU29', 'RB30', 'uRB31', 'lb32', ...][i-25]
    return execute(move, cube)
```

In order to help me keep track of which entry of the list is which, I have included the starting position of subface 25 (so, for instance, “Bru41” would indicate the sequence of moves needed to move the subface in position 41 to position 25). If the `execute` function is correctly written, these extra digits will be ignored when the move is executed.

The functions for `align26`, `align27`, `align28` can be constructed similarly.

5. Create four functions, `align1(cube)`, `align2(cube)`, `align3(cube)`, and `align4(cube)`, that return sequences of moves that put subfaces 1 through 4 in their correct positions without disturbing subfaces 25 to 28. These functions should work similarly to the ones written in part 4.
6. Create a function `solvetop(cube)` that has a scrambled cube as input, and returns a sequence of moves that solves the top level of the cube. The easiest way to do this would be to run each of the functions `align25(cube)`, `align26(cube)`, `align27(cube)`, `align28(cube)`, `align1(cube)`, `align2(cube)`, `align3(cube)`, and `align4(cube)` in the indicated order; put together their output into a single string, and return that string. Done efficiently, this can be a one or two line function.

You should test your `solvetop` program using some of the 100 scrambled cubes found:

math.byu.edu/~doud/Math495R/Cubes.txt

If your routines work properly, subfaces 1–12 and 24–32 should all be put into the correct positions by `solvetop(cube)`.

If you need to see a graphical simulation of a cube, you can find one online at

math.byu.edu/~doud/RubiksCube/

Enter a sequence of moves into the textbox at the top of the page, click the button labeled “Turn,” and the program will execute the specified moves.

Lab 14

Riemann Sums

In this lab we will study Riemann sums in higher dimensions.

1. Write a function `riemann_sum_2D` which takes 7 parameters `f`, `xMin`, `xMax`, `yMin`, `yMax`, `N`, and `method` and returns the Riemann sum

$$\sum_{j=1}^N \sum_{i=1}^N f(x_i^*, y_j^*) \Delta x \Delta y$$

where $\Delta x = (\text{xMax} - \text{xMin})/N$, $\Delta y = (\text{yMax} - \text{yMin})/N$ and `method` determines whether we are using the lower left, upper right, or midpoints of the partition. (The options for the method should be `left`, `right`, `mid`).

2. Using your function from problem 1, use the midpoint method to calculate the Riemann sums for $N = 10$ and $N = 20$ for the following functions and domains:
 - a) $f(x, y) = x \sin(xy)$ on the rectangle $[0, \pi] \times [0, \pi]$.
 - b) $f(x, y) = y^2 e^{-x-y}$ on the rectangle $[0, 1] \times [0, 1]$.
 - c) $f(x, y) = x^3 y^2 + xy$ on the rectangle $[0, 1] \times [1, 2]$.
3. Consider the integral of $f(x, y) = x \sin(x + y)$ on the rectangle $[0, \pi/6] \times [0, \pi/3]$. First calculate the value of this integral analytically. Then make a plot that shows the error of the midpoint Riemann integral approximation as N ranges from 1 to 100.
4. Write a function `riemann_sum_3D` which takes 9 parameters `f`, `xMin`, `xMax`, `yMin`, `yMax`, `zMin`, `zMax`, `N`, and `method` and returns the Riemann sum

$$\sum_{k=1}^N \sum_{j=1}^N \sum_{i=1}^N f(x_i^*, y_j^*) \Delta x \Delta y \Delta z,$$

where $\Delta x = (\text{xMax} - \text{xMin})/N$, $\Delta y = (\text{yMax} - \text{yMin})/N$, $\Delta z = (\text{zMax} - \text{zMin})/N$ and `method` determines whether we are using the lower left, upper right, or midpoint of the partition. (*Hint*: you can copy much of the code you wrote for Problem 1.)

5. Using your function from problem 4, use the midpoint method to calculate the Riemann sums for $N = 10$ and $N = 20$ for the following function and domain:

$$f(x, y, z) = xy + z^2 \text{ on the rectangle } [0, 2] \times [0, 1] \times [0, 3].$$

Lab 15

Solving the Rubik's cube, Part II

In this lab, we will continue programming our solution to the Rubik's cube. We will also construct several functions to help us debug our program.

Before beginning this lab, you should be sure that your functions from Lab 13 work correctly. They will be needed for this lab.

We will do the following in this lab: putting subfaces 33–40 in the correct position, and getting the edge faces 45–48 all on the bottom face. In the next (and final) Rubik's cube lab, we will complete the solution of the cube.

1. As a preliminary, we will create a function which scrambles a Rubik's cube. The function should have the following syntax: `scramble(cube,n=50)`. As input, it will take a vector representing a state of the Rubik's cube, and an (optional) integer n . It will then execute n random moves on the cube to scramble it. It should move the elements of `cube`, and return only the list of moves that it made.

Note that there are 12 possible moves: "RLUDFBr1udfb". To choose one at random use "import random" at the beginning of your program, and then `random.randint(0,11)` will choose a random integer from 0 to 11 (inclusive).

2. We now create a function that will translate sequences of moves designed to move subfaces on the front of the cube so that they move subfaces on other sides of the cube.

For instance, if we have a sequence of moves that takes the contents of the bottom front edge (location 43) and moves it to the right front edge (location 35), we want to be able to adjust it to move location 42 to 34 (on the right side), location 41 to 33 (on the back), or location 44 to 36, on the left.

To do this, hold up your cube handout, with the "Up" face on top, and the "Front" face facing you. Now rotate the cube 90° counterclockwise so that the "Left" face faces you. The face on the right hand side is now the face labeled "Front", so if you were to do an **R** move without reference to the labels on the cube, it would actually be a **F** move. Similarly, **L** would actually perform **B**, **F** would actually perform **L**, and **B** would actually perform **R**. The moves **T** and **D** would have the same effects as if the cube had not been rotated.

Note that if we do a sequence of moves that transfer the contents of location 43 to location 35, then after the rotation it would rotate the contents of location 44 to location 36.

Construct a function `toprotate(moves)` that will translate a series of moves by 90°, as above: in other words, it will make the following substitutions:

$$\begin{array}{cccc} R \rightarrow F, & F \rightarrow L, & L \rightarrow B, & B \rightarrow R \\ r \rightarrow f, & f \rightarrow l, & l \rightarrow b, & b \rightarrow r \end{array}$$

Probably the easiest way to do this is with the `translate` method on strings. To illustrate this method, run the following lines of code:

```
word="Convert abcde to numbers."
translation=str.maketrans("abcde","12345")
print(word.translate(translation))
```

Notice that each a has been turned into a 1, each b into a 2, and so on. Use this technique to make the translation function described above. Once you have `toprotate` working properly, the command `toprotate("RFLBrflb")` should return `"FLBRflbr"`.

3. Create a function `edgesfrombottom(cube)` that will (assuming that the top face of the cube is solved) look at which faces are in locations 41 to 44, and if any face numbered between 33 and 40 is in one of these four spots will place it in its proper location. It should exit when all of the subfaces in locations 41 to 44 have numbers higher than 40. It should initialize a variable `result=""`, and use this variable to record and return any moves that it executes.

The first step to placing a subface in its proper place is to rotate it to the proper side of the cube. Note that all the edges on the front of the cube have numbers congruent to 3 mod 4; all the ones on the right have numbers congruent to 2 mod 4, all the ones on the back have numbers congruent to 1 mod 4, and all the ones on the left have numbers congruent to 0 mod 4. So, if we find a face numbered 33–40 in one of the spots 41–44, it is on the correct side of the cube if its number is congruent mod 4 to the position that it is in (i.e. if $(i - \text{cube}[i]) \% 4 == 0$). If it is in the wrong position, we need to use a `d`, `dd`, or `D` move to put it on the correct side.

Once we have it on the correct side, we need to move it to the correct position in the middle layer of the cube.

To do this, note that the sequence of moves `"drDRDfdf"` will move the contents of face 43 (the bottom edge of the front) to face 35 (the right edge of the front) without disturbing the top face or any of faces 33–40 except for 35 and 38. The sequence of moves `"DLdldfDF"` will move the contents of face 43 to face 39 (the left edge of the front) under similar restrictions.

These moves adjust the “Front” face of the cube. If we apply `toprotate` to them, they will have the desired effect on the “Left” face of the cube; namely moving the contents of position 44 to either position 36 or 40. Applying `toprotate` twice will affect the “Back” face of the cube, and applying `toprotate` three times will affect the “Right” face of the cube. (Try to find a formula for how many times you need to apply `toprotate`, instead of using multiple if statements.)

After checking each of locations 41 to 44, if no moves were executed (i.e., if `result` is still empty), the function should return the empty string. If any moves were executed, then we need to run the code again (since we may have disturbed a location that we had previously checked). Do this by returning `result+edgesfrombottom(cube)`.

In writing the code for this lab, the functions on the next page that test your code may be useful. If `testall(10000)` yields a positive result, then your code has worked correctly on 10000 scrambled cubes. These functions can be downloaded at

math.byu.edu/~doud/Math495R/Lab16Debug.py

After you finish this lab, you should go to Lab 20 and start working on the next function.

To test (and debug) the function `edgesfrombottom`, you can use the following function:

```
def test(n):
    flag=True
    for i in range(n):
        cube=list(range(49))
        t=scramble(cube)
        t=solvetop(cube)
        t=edgesfrombottom(cube)
        for position in range(41,45):
            if cube[position]<41:
                flag=False
                print("Your function has a problem")
    if flag:
        print("Your function seems to be fine.")
```

To test and debug the function `edgesfromsides`, use the following function.

```
def test2(n):
    flag=True
    for i in range(n):
        cube=list(range(49))
        t=scramble(cube)
        t=solvetop(cube)
        t=edgesfrombottom(cube)
        t=edgesfromsides(cube)
        if cube[33:41]!=list(range(33,41)):
            flag=False
            print("Your function has a problem")
            print(cube[33:41])
    if flag:
        print("Your functions seem to be fine.")
```

The following function tests all of the functions that we have written so far.

```
def testall(n):
    flag=True
    for i in range(n):
        cube=list(range(49))
        t=scramble(cube)
        t=solvetop(cube)
        t=edgesfrombottom(cube)
        t=edgesfromsides(cube)
        t=flipbottomedges(cube)
        if cube[1:13]!=list(range(1,13)):
            flag=False
        if cube[25:41]!=list(range(25,41)):
            flag=False
        if min(cube[45:])<45:
            flag=False
    if flag:
        print("Your code seems to be working")
```

Lab 16

Monte Carlo Integration

Integration is challenging, even for a computer. Although Riemann sums can give fairly accurate answers they do not always work well in higher dimensional settings. In today's lab we will talk about another technique for numerical integration, called Monte Carlo integration.

Volume Estimation

To give you an idea of how Monte Carlo integration works, let's consider the concrete example of a circle in \mathbb{R}^2 . It is well known that the area of a circle of radius r is $A = \pi r^2$ —in fact one way to numerically estimate π is to compute the area of the unit circle. The way that Monte Carlo integration works is to randomly choose points in a domain whose area/volume is easy to calculate that encompasses the object in question. For this concrete example we will randomly choose points from the square domain $\Omega = [-1, 1] \times [-1, 1]$. We can then estimate the area of the circle by determining the percentage of points that lie in the circle and multiplying this percentage by the total area of the sample domain gives an estimate for the area of the circle.

The following code samples 2000 uniformly distributed random points in Ω , determines what percentage of those points are within the unit circle, then multiplies that percentage by 4 (the area of Ω) to get an estimate for π .

```
>>> import numpy as np
>>> from scipy import linalg as la

# Get 2000 random points in the 2-D domain [-1,1]x[-1,1].
>>> points = np.random.uniform(-1, 1, (2,2000))

# Determine how many points are within the circle.
>>> lengths = la.norm(points, axis=0)
>>> num_within = np.count_nonzero(lengths < 1)

# Estimate the circle's area.
>>> 4 * (num_within / 2000)
3.198
```


The estimate of the area of the circle 3.198 isn't perfect, but it only differs from the true area by about 0.0564. On average, increasing the number of sample points decreases the estimate error.




Figure 16.1: Estimating the area of the unit circle using random points.

1. The n -dimensional *open unit ball* is the set $U_n = \{\mathbf{x} \in \mathbb{R}^n \mid \sum_{i=1}^n x_i^2 < 1\}$. Write a function that accepts an integer n and a keyword argument N defaulting to 10^4 . Estimate the volume of U_n by drawing N points over the n -dimensional domain $[-1, 1] \times [-1, 1] \times \cdots \times [-1, 1]$. (Hint: the volume of $[-1, 1] \times [-1, 1] \times \cdots \times [-1, 1]$ is 2^n .)

When $n = 2$, this is the same experiment outlined above so your function should return an approximation of π . The volume of the U_3 is $\frac{4}{3}\pi \approx 4.18879$, and the volume of U_4 is $\frac{\pi^2}{2} \approx 4.9348$. Try increasing the number of sample points N to see if your estimates improve.

Integral Estimation

The strategy for estimating the area of a circle can be formulated as an integral problem. Define $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ by

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } \sum_i x_i^2 < 1 \text{ (}\mathbf{x} \text{ is within the unit circle)} \\ 0 & \text{otherwise,} \end{cases}$$

and let $\Omega = [-1, 1] \times [-1, 1]$ as before. Then

$$\int_{-1}^1 \int_{-1}^1 f(x, y) dx dy = \int_{\Omega} f(\mathbf{x}) dV = \pi.$$

To estimate the integral we chose N random points $\{\mathbf{x}_i\}_{i=1}^N$ in Ω . Since f indicates whether or not a point lies within the unit circle, the total number of random points that lie in the

circle is the sum of the $f(\mathbf{x}_i)$. Then the average of these values, multiplied by the volume $V(\Omega)$, is the desired estimate:

$$\int_{\Omega} f(\mathbf{x}) dV \approx V(\Omega) \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_i). \quad (16.1)$$

This remarkably simple equation can be used to estimate the integral of any integrable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ over any domain $\Omega \subset \mathbb{R}^n$ and is called the general formula for *Monte Carlo integration*.

The intuition behind (16.1) is that $\frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_i)$ approximates the average value of f on Ω , and multiplying the approximate average value by the volume of Ω yields the approximate integral of f over Ω . This is a little easier to see in one dimension: for a single-variable function $f : \mathbb{R} \rightarrow \mathbb{R}$, the Average Value Theorem states that the average value of f over an interval $[a, b]$ is given by

$$f_{avg} = \frac{1}{b-a} \int_a^b f(x) dx.$$

Then using the approximation $f_{avg} \approx \frac{1}{N} \sum_{i=1}^N f(x_i)$, the previous equation becomes

$$\int_a^b f(x) dx = (b-a) f_{avg} \approx V(\Omega) \frac{1}{N} \sum_{i=1}^N f(x_i), \quad (16.2)$$

which is (16.1) in one dimension. In this setting $\Omega = [a, b]$ and hence $V(\Omega) = b - a$.

2. Write a function that accepts a function $f : \mathbb{R} \rightarrow \mathbb{R}$, bounds of integration a and b , and an integer N defaulting to 10^4 . Use `np.random.uniform()` to sample N points over the interval $[a, b]$, then use (16.2) to estimate the integral

$$\int_a^b f(x) dx.$$

Test your function on the following integrals, or on other integrals that you can check by hand.

$$\int_{-4}^2 x^2 dx = 24 \quad \int_{-2\pi}^{2\pi} \sin(x) dx = 0 \quad \int_1^{10} \frac{1}{x} dx = \log(10) \approx 2.30259$$

Be careful not to use Monte Carlo integration to estimate integrals that do not converge. For example, since $1/x$ approaches ∞ as x approaches 0 from the right, the integral

$$\int_0^1 \frac{1}{x} dx$$

does not converge. Even so, attempts at Monte Carlo integration still return a finite value. Use various numbers of sample points to see whether or not the integral estimate is converging.

```

\begin{lstlisting}
>>> for N in [5000, 7500, 10000]:
...     print(np.mean(1. / np.random.uniform(0, 1, N)), end='\t')
...
11.8451683722    25.5814419888    7.64364735049    # No convergence.

```

Integration in Higher Dimensions

The implementation of (16.1) for a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with $n > 1$ introduces a few tricky details, but the overall procedure is the same for the case when $n = 1$. We consider only the case where $\Omega \subset \mathbb{R}^n$ is an n -dimensional box $[a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_n, b_n]$.

- (i) If $n = 1$ then Ω is a line, so $V(\Omega) = b_1 - a_1$. If $n = 2$ then Ω is a rectangle, and hence $V(\Omega) = (b_1 - a_1)(b_2 - a_2)$, the product of the side lengths. The volume of a higher-dimensional box Ω is also the product of the side lengths,

$$V(\Omega) = \prod_{i=1}^n (b_i - a_i) \quad (16.3)$$

- (ii) It is easy to sample uniformly over an interval $[a, b]$ with `np.random.uniform()`, or even over the n -dimensional cube $[a, b] \times [a, b] \times \cdots \times [a, b]$ (such as in Problem ??). However, if $a_i \neq a_j$ or $b_i \neq b_j$ for any $i \neq j$, the samples need to be constructed in a slightly different way.

The interval $[0, 1]$ can be transformed to the interval $[a, b]$ by scaling it so that it is the same length as $[a, b]$, then shifting it to the appropriate location.

$$[0, 1] \xrightarrow{\text{scale by } b-a} [0, b-a] \xrightarrow{\text{shift by } a} [a, b]$$

This suggests a strategy for sampling over $[a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_n, b_n]$: sample uniformly from the n -dimensional box $[0, 1] \times [0, 1] \times \cdots \times [0, 1]$, multiply the i th component of each sample by $b_i - a_i$, then add a_i to that component.

$$[0, 1] \times \cdots \times [0, 1] \xrightarrow{\text{scale}} [0, b_1 - a_1] \times \cdots \times [0, b_n - a_n] \xrightarrow{\text{shift}} [a_1, b_1] \times \cdots \times [a_n, b_n] \quad (16.4)$$

3. Write a function that accepts a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, a list of lower bounds $[a_1, a_2, \dots, a_n]$, a list of upper bounds $[b_1, b_2, \dots, b_n]$, and an integer N defaulting to 10^4 . Use (16.1), (16.3), and (16.4) with N sample points to estimate the integral

$$\int_{\Omega} f(\mathbf{x}) dV,$$

where $\Omega = [a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_n, b_n]$.

Test your function on the following integrals.

$$\int_0^1 \int_0^1 x^2 + y^2 \, dx \, dy = \frac{2}{3} \quad \int_{-2}^1 \int_1^3 3x - 4y + y^2 \, dx \, dy = 54$$

$$\int_{-4}^4 \int_{-3}^3 \int_{-2}^2 \int_{-1}^1 x + y - wz^2 \, dx \, dy \, dz \, dw = 0$$

Note carefully how the order of integration defines the domain; in the last example, the x - y - z - w domain is $[-1, 1] \times [-2, 2] \times [-3, 3] \times [-4, 4]$, so the lower and upper bounds passed to your function should be $[-1, -2, -3, -4]$ and $[1, 2, 3, 4]$, respectively.

Lab 17

Solving the Rubik's cube, Part III

In this lab we will solve the remaining parts of the Rubik's cube.

1. Create a function `edgesfromsides(cube)`. It should look at the subfaces in positions 33 to 36, and if one of them is incorrect, it should move this face to the bottom edge, and then run `edgesfrombottom(cube)` to place the edge in its correct position. Once all four positions from 33 to 36 have been dealt with, the second layer of the cube should be solved.

To move a face from a side edge to the bottom, use the basic move “`drDRDFdf`” (for position 35), with `toprotate` applied the appropriate number of times to make it apply to the other three positions.

At this point, you have code that solves the top two layers of the Rubik's cube.

2. We will now work on putting the subfaces numbered 45 to 48 on the bottom face of the cube. Many instruction manuals for solving the cube call this “creating a yellow cross”, since they typically have the bottom face be yellow, and when all four bottom edges are on the bottom, they form the shape of a cross.

In order to do this, we will create a function called `flipbottomedges(cube)`. It should change the entries of cube, and return the sequence of moves that it executes.

It turns out that there are only four possibilities for how the bottom faces can be positioned.

- (1) None of them are on the bottom. In this case, the sequence of moves “`FLDldfBRDrdrDRdb`” will flip them all to be on the bottom.
- (2) Two of them are on the bottom, opposite from each other. In this case, after rotating the bottom so that the two correct ones are in position 46 and 48, the sequence of moves “`FLDldf`” will leave all four bottom faces on the bottom.
- (3) Two of them are on the bottom, adjacent to each other. In this case, after rotating the bottom so that the two correct ones are in positions 45 and 48, the sequence of moves “`RDFdfr`” will leave all four bottom faces on the bottom.
- (4) All four of the bottom edges are already on the bottom; in this case, nothing needs to be done.

Count the number of subfaces in position 45–48 that are numbered above 45–48. If this number is 0, we are in case 1 above. If the number is 4, we are in case 4 above. Otherwise, we are in either case 2 or case 3, and we should perform a sequence of D moves until location 48 contains a bottom face and location 47 does not. Then, depending on which of positions 45 and 46 contain bottom faces, we perform the appropriate sequence of moves.

3. We will now write a function `positionbottomedges(cube)` that will put the bottom edges in the correct positions. It should change `cube`, and return the sequence of moves that it executes.

At this point, all four bottom edges are on the bottom of the cube, but they are probably not positioned correctly. It will always be possible to rotate the bottom face in such a way that either 2 or 4 of the bottom faces are positioned correctly. To do this, count the number of subfaces in positions 45–48 that are numbered correctly. As long as that number is less than two, execute a D move, and then repeat.

If all four faces are positioned correctly, we are finished with this step. If only two faces are positioned correctly, we have more to do.

If the two faces that are positioned correctly are opposite each other, they are either in positions 46 and 48 or positions 45 and 47. In the first case execute the sequence of moves: “FDfDFDDfRDrDRDDrD”. In the second case, `toprotate` this sequence of moves once, and execute the result. This will put all four edge pieces in the correct positions.

If the two faces that are positioned correctly are adjacent to each other, if they are in positions 45 and 48, execute the sequence of moves “LD1DLDD1D. If they are positioned differently, `toprotate` this sequence of commands the correct number of times to have the desired effect. This will put all four edge pieces in the correct positions. If the correct ones are in positions 45 and 46 a single `toprotate` should suffice; if they are in position 46 and 47, two `toprotates` will be needed, if they are in positions 47 and 48, three `toprotates` will be needed.

4. We will now write a function `positionbottomcorners(cube)` that will put the bottom corners in the correct positions. It should change `cube`, and return the sequence of moves that it executes. We will identify the bottom corner pieces by looking at positions 21, 22, 23, and 24. Note that if the face in one of these positions is congruent modulo 4 to its position, (i.e. if $(\text{cube}[i] - i) \% 4 = 0$) then it is in the correct position, but might be twisted incorrectly.

To position the bottom corners, we have a basic move, “rDLdRD1d” that fixes the corner containing position 24, and moves the other three corners in a clockwise pattern (twisting them as they move). If the cube has been solved to this point, it should be in one of the states below:

- (1) If none of the corners are in the correct position, execute the basic move. One will now be in the correct position.

- (2) If one of the corners is correct, use `toprotate` on the basic move above the correct number of times so that it leaves the correct corner fixed. Then applying the `toprotated` basic move either once or twice will put all the corners in the correct position.
 - (3) It should never happen that two or three of the corners are in the correct position.
 - (4) If all four corners are in the correct position, we are done with this step.
5. We are now ready for the final step. This step will rotate each corner cube into the correct orientation.
- Write a function called `rotatebottomcorners(cube)`. It should change the cube, and return the sequence of moves that it makes.
- The function should consist of a loop that does the following four times.
- (1) While the face in position 21 is numbered less than 21, execute the following move “`buBUbuBU`”. (You may have to do this twice.)
 - (2) Execute the move `D`.
6. Finally, we put together all of the functions that we have produced so far, to get a single function that solves the Rubik’s cube. Your function might look something like the one below:

```
def solvecube(cube):
    result=""
    result+=solvetop(cube)
    result+=edgesfrombottom(cube)
    result+=edgesfromsides(cube)
    result+=flipbottomedges(cube)
    result+=positionbottomedges(cube)
    result+=positionbottomcorners(cube)
    result+=rotatebottomcorners(cube)
    return result
```

Once you have written this function, use the functions at the website
math.byu.edu/~doud/Math495R/Lab16Debug.py
to test it on a large number of cubes.

Lab 18

Plotting Vector Valued Functions

In this lab you will create visualizations to help you understand vector fields. Before beginning the lab, please review how to create quiver plots in python.

[problemsolvingwithpython.com/06-Plotting-with-Matplotlib/
06.15-Quiver-and-Stream-Plots/](https://problemsolvingwithpython.com/06-Plotting-with-Matplotlib/06.15-Quiver-and-Stream-Plots/)

A quiver plot is a plot that shows vectors as arrows and is especially useful when talking about vector fields.

1. Plot the vector field $\mathbf{F}(x, y) = (y^2 - 2xy)\mathbf{i} + (3xy - 6x^2)\mathbf{j}$.
2. Let $\mathbf{F}(\mathbf{x}) = (r^2 - 2r)\mathbf{x}$, where $\mathbf{x} = \langle x, y \rangle$ and $r = |\mathbf{x}|$. You may have to vary your domain to see what is happening. First plot for $-1 \leq x, y \leq 1$ with a step size of 0.1. Then plot again for $-5 \leq x, y \leq 5$ with a step size of 0.5.
3. Plot the gradient vector field of f together with a contour map of f for:
 - a) $f(x, y) = \ln(1 + x^2 + 2y^2)$
 - b) $f(x, y) = \cos x - 2 \sin y$
4. Of course there are also vector fields in higher dimensions. See matplotlib.org/3.1.1/gallery/mplot3d/quiver3d.html

Plot the following vector fields:

- a) $\mathbf{F}(x, y, z) = \mathbf{i} + 2\mathbf{j} + 3\mathbf{k}$
- b) $\mathbf{F}(x, y, z) = \mathbf{i} + 2\mathbf{j} + z\mathbf{k}$
- c) $\mathbf{F}(x, y, z) = x\mathbf{i} + y\mathbf{j} + z\mathbf{k}$
- d) $\mathbf{F}(x, y, z) = \sin y\mathbf{i} + (x \cos y + \cos z)\mathbf{j} - y \sin z\mathbf{k}$

Lab 19

K-Means Clustering

One of the current areas of high interest is data science and machine learning. Machine learning can roughly be divided into three main types, supervised learning, unsupervised learning, and reinforcement learning. Unsupervised learning, or pattern recognition involves identifying useful patterns or structure from data that is unlabeled. One example of this is clustering, where the goal is to determine points that should be clustered together. The goal of this lab is to introduce the k-means algorithm, a simple and popular clustering method.

The objective of clustering is to find a partition of the data such that points in the same subset will be “close” according to some way of measuring. There are many different ways of measuring how close two points are but we will use the Euclidean distance introduced in your text.

More formally, suppose we have a collection of \mathbb{R}^K -valued observations $X = \{x_1, x_2, \dots, x_n\}$. Let $N \in \mathbb{N}$ and let \mathcal{S} be the set of all N -partitions of X , where an N -partition is a partition with exactly N nonempty elements. We can represent a typical partition in \mathcal{S} as $S = \{S_1, S_2, \dots, S_N\}$, where

$$X = \bigcup_{i=1}^N S_i$$

and

$$|S_i| > 0, \quad i = 1, 2, \dots, N.$$

We seek the N -partition S^* that minimizes the within-cluster sum of squares, i.e.

$$S^* = \arg \min_{S \in \mathcal{S}} \sum_{i=1}^N \sum_{x_j \in S_i} \|x_j - \mu_i\|_2^2,$$

where μ_i is the mean of the elements in S_i , i.e.

$$\mu_i = \frac{1}{|S_i|} \sum_{x_j \in S_i} x_j.$$

The K-Means Algorithm

Finding the global minimizing partition S^* is generally intractable since the set of partitions can be very large indeed, but the *k-means* algorithm is a heuristic approach that can often provide reasonably accurate results.

We begin by specifying an initial cluster mean $\mu_i^{(1)}$ for each $i = 1, \dots, N$ (this can be done by random initialization, or according to some heuristic). For each iteration, we adopt the following procedure. Given a current set of cluster means $\mu^{(t)}$, we find a partition $S^{(t)}$ of the observations such that

$$S_i^{(t)} = \{x_j : \|x_j - \mu_i^{(t)}\|_2^2 \leq \|x_j - \mu_l^{(t)}\|_2^2, \quad l = 1, \dots, N\}.$$

We then update our cluster means by computing for each $i = 1, \dots, N$. We continue to iterate in this manner until the partition ceases to change.

Figure 19.1 shows two different clusterings of the iris data produced by the *k-means* algorithm. Note that the quality of the clustering can depend heavily on the initial cluster means. We can use the within-cluster sum of squares as a measure of the quality of a clustering (a lower sum of squares is better). Where possible, it is advisable to run the clustering algorithm several times, each with a different initialization of the means, and keep the best clustering. Note also that it is possible to have very slow convergence. Thus, when implementing the algorithm, it is a good idea to terminate after some specified maximum number of iterations. The algorithm can be summarized as follows.




Figure 19.1: Two different K-Means clusterings for the iris dataset. Notice that the clustering on the left predicts the flower species to a high degree of accuracy, while the clustering on the right is less effective.

1. Choose k initial cluster centers.
2. For $i = 0, \dots, \text{max_iter}$,
 - a) Assign each data point to the cluster center that is closest, forming k clusters.

- b) Recompute the cluster centers as the means of the new clusters.
- c) If the old cluster centers and the new cluster centers are sufficiently close, terminate early.

Those students planning on enrolling in the ACME program or who are completing a degree in computer science will likely have the opportunity to code up the k-means algorithm as part of the program. For this lab we will use the built in version of k-means clustering that comes with the sklearn package.

Detecting Active Earthquake Regions

Suppose we are interested in learning about which regions are prone to experience frequent earthquake activity. We could make a map of all earthquakes over a given period of time and examine it ourselves, but this, as an unsupervised learning problem, can be solved using our k -means clustering tool.




Figure 19.2: Earthquake epicenters over a 6 month period.

The file `earthquake_coordinates.npy` contains earthquake data throughout the world from January 2010 through June 2010. Each row represents a different earthquake; the columns are scaled longitude and latitude measurements. We want to cluster this data into active earthquake regions. For this task, we might think that we can regard any epicenter as a point in \mathbb{R}^2 with coordinates being their latitude and longitude. This, however, would be incorrect, because the earth is not flat. Instead, latitude and longitude should be viewed in *spherical coordinates* in \mathbb{R}^3 , which could then be clustered.

A simple way to accomplish this transformation is to first transform the latitude and longitude values to spherical coordinates, and then to Euclidean coordinates. Recall that

a spherical coordinate in \mathbb{R}^3 is a triple (r, θ, φ) , where r is the distance from the origin, θ is the radial angle in the xy -plane from the x -axis, and φ is the angle from the z -axis. In our earthquake data, once the longitude is converted to radians it is an appropriate θ value; the latitude needs to be offset by 90° degrees, then converted to radians to obtain φ . For simplicity, we can take $r = 1$, since the earth is roughly a sphere. We can then transform to Euclidean coordinates using the following relationships.

$$\begin{aligned}\theta &= \frac{\pi}{180} (\text{longitude}) & \varphi &= \frac{\pi}{180} (90 - \text{latitude}) \\ r &= \sqrt{x^2 + y^2 + z^2} & x &= r \sin \varphi \cos \theta \\ \varphi &= \arccos \frac{z}{r} & y &= r \sin \varphi \sin \theta \\ \theta &= \arctan \frac{y}{x} & z &= r \cos \varphi\end{aligned}$$

There is one last issue to solve before clustering. Each earthquake data point has norm 1 in Euclidean coordinates, since it lies on the surface of a sphere of radius 1. Therefore, the cluster centers should also have norm 1. Otherwise, the means can't be interpreted as locations on the surface of the earth, and the *k-means* algorithm will struggle to find good clusters. A solution to this problem is to normalize the mean vectors at each iteration, so that they are always unit vectors.

1. Add a keyword argument `normalize=False` to your `KMeans` constructor. Modify `fit()` so that if `normalize` is `True`, the cluster centers are normalized at each iteration.

Cluster the earthquake data in three dimensions by converting the data from raw data to spherical coordinates to euclidean coordinates on the sphere.

- a) Convert longitude and latitude to radians, then to spherical coordinates. (Hint: `np.deg2rad()` may be helpful.)
- b) Convert the spherical coordinates to euclidean coordinates in \mathbb{R}^3 .
- c) Use your `KMeans` class with normalization to cluster the euclidean coordinates.
- d) Translate the cluster center coordinates back to spherical coordinates, then to degrees. Transform the cluster means back to latitude and longitude coordinates. (Hint: use `numpy.arctan2()` for arctan, so that that correct quadrant is chosen).
- e) Plot the data, coloring by cluster. Also mark the cluster centers.

With 15 clusters, your plot should resemble the Figure [19.3](#).

- (bonus) Add a keyword argument `3d=False` to your `kmeans` function, and add code to show the three-dimensional plot instead of the two-dimensional scatter plot should this argument be set to `True`. Maintain the same color-coding scheme as before. Use `mpl_toolkits.mplot3d.Axes3D` to make your plot.




Figure 19.3: Earthquake epicenter clusters with $k = 15$.

Color Quantization

The k -means algorithm uses the euclidean metric, so it is natural to cluster geographic data. However, clustering can be done in any abstract vector space. The following application is one example.

Images are usually represented on computers as 3-dimensional arrays. Each 2-dimensional layer represents the red, green, and blue color values, so each pixel on the image is really a vector in \mathbb{R}^3 . Clustering the pixels in RGB space leads a one kind of image segmentation that facilitate memory reduction. (If you would like to read more – go to https://en.wikipedia.org/wiki/Color_quantization

- 2 Write a function that accepts an image array (of shape $(m, n, 3)$), an integer number of clusters k , and an integer number of samples S . Reshape the image so that each row represents a single pixel. Choose S pixels to train a k -means model on with k clusters. Make a copy of the original picture where each pixel has the same color as its cluster center. Return the new image. You may use `sklearn.cluster.KMeans` instead of your implementation from the previous problem.

Test your code with different values of k and S on some of the images from <https://www.nasa.gov/multimedia/imagegallery/index.html> Include one of your color-quantized pictures in your submitted lab.

Lab 20

Machine Learning with Nim, Part I

In this Lab, we will program a computer to learn to play the game of Nim. Note that we will *not* program the computer to play perfectly—indeed, we will begin with a computer that plays randomly (and probably loses most of the time). We will then train the computer to recognize winning moves, until the computer plays perfectly.

Rules of Nim

A game of Nim begins with three piles of objects; one having 7 objects, the other having 5 objects, and the last having 3 objects. Two players take turns removing objects from the piles. A turn consists of removing any positive number of objects from one of the piles. The winner is the last player that is able to make a valid move (i.e. the player who removes the last objects from the last nonempty pile).

There are many variations of Nim on the internet; we could use different numbers of objects in each pile, or have a different number of piles. In some versions, the person who takes the last object loses. For our game, we will use the rules listed above.

Nim is a completely solved game: with perfect play, depending on the starting configuration, either the first player will always win, or the second player will always win. For our starting configuration, unless the first player plays poorly, the first player will always win.

A template for this lab is available at

math.byu.edu/~doud/Math495R/Lab20template.py

1. The first function that we will write is `checkmove(position,pile,remove)`. The first input is a list `position` indicating the current position of the board. The second, `pile`, is an integer 1, 2, or 3, and the third, `remove`, is a positive integer. The function should check whether taking `remove` objects from pile `pile` is a valid move, and return `True` or `False`. For example, if `position=[7,5,3]`, then `pile=3`, `remove=1` would be a valid move, but `pile=3`, `remove=5` would not (since there are not five objects on pile 3 to remove).
2. The next function is `player(n,position)`. The input to this function is an integer `n` indicating whether the human player is player one or player two, and a list `position`

indicating the position of the board prior to the player's move.

The function should ask the player "Player **n**, what move would you like to make?" (with "n" filled in with the playernumber) and accept an answer as a string. The string should be of the form **3 1** (two integers separated by a space). The first integer indicates which pile the player wishes to remove objects from, and the second indicates how many objects should be removed. The program should parse this string into two integers, and check whether it is a valid move for the indicated position.

If the move is not valid, the program should indicate this, and ask the player again for a valid move.

If the move is valid, the program should subtract the given number of objects from the correct component of **position**. Note that the program does not need to return a value; **position** is a vector that was passed to the function, so changing a component will change the original vector.

3. The second function that we need to write is **playgame(strategy)**. For the moment, we will not use the input variable **strategy**; it will be used to store the computer's strategy, once we have developed a strategy for the computer.

This function should begin by asking "Would you like to play a game? (Y/N)". It should accept input in the form of a string. If the string is either "Y" or "y", it should proceed; if the input is "N" or "n" it should terminate, and other input should lead to the question being asked again.

If the program proceeds, it should ask

"Player 1: Computer (C) or Human (H)? (C/H) "

and accept a string as an answer. It should store this answer in the variable **P1**. If the answer is neither "C" nor "H", it should ask the question again.

It should then repeat the process for player 2, storing the response in the variable **P2**.

If both players are computers, the program should ask "How many games?" and accept an integer as a response. If either player is human, this question should not be asked—a single game will be played.

If more than one game is to be played (which should only happen with two computer players), the program should print

"—Beginning Game 1—".

The program should then initialize the board, setting the position to **[7,5,3]**. Beginning with player one, it should print the board position, in the format

"The piles have 7, 5, 3 objects"

(where the numbers are taken from the vector indicating the position) and then call the function `player(n,position)` or `computer(n,position,moves,strategy)` (depending on whether player `n` is a human or computer) to allow the player to make a move. It should then change the player number, and repeat (printing “The piles have *, *, * objects” each time), until the position becomes `[0,0,0]`.

Once the game is over, if both players are computers and multiple games have been requested, it should print

“—Beginning Game 2—”

and initialize and play a new game with two computer players. It should repeat this until the desired number of games have been played. Once all of the computer vs. computer games are over, it should print

“—Completed `n` games—”

where `n` is the number of games requested.

If either player is human, then once the game is over, the program should ask “Would you like to play again? (Y/N)”, and depending on the answer, start over (asking if the players are human or computers), or terminate.

At this point, you should be able to play a two-person game (with no computer players) by running the program.

4. In order to have a computer play, we need to tell it how to play. We will create a vector called `strategy`, indexed by integers from 0 to 753. If the integer `n` is a valid board position for the game (i.e. the first digit is from 0 to 7, the second digit from 0 to 5, and the third digit from 0 to 3), then `strategy[n]` will itself be a vector containing all possible moves from the given position. Each move will be represented by a triple of numbers, representing the pile, number of objects to remove, and the quality of the move (to be described later). For instance, if the integer `n` is equal to 312, the board position will be `[3,1,2]`, and `strategy[312]` will be the vector

`[[1, 3, 50], [1, 2, 50], [1, 1, 50], [2, 1, 50], [3, 2, 50], [3, 1, 50]]`.

For now, we just set the third number in each triple to 50 (it will be used to determine which moves are best, later).

Create a function `initialize()` that will create and return the desired vector `strategy`.

Lab 21

Machine Learning with Nim, Part II

In this part of the lab, we will train the computer to learn from its mistakes and become a perfect Nim player. We can only hope that it does not learn more than this.

1. We first teach the computer the rules of the game, so that it can play very badly. The function `computer(n,position,moves,strategy)` takes as input an integer `n` indicating the player number, a vector `position` giving the board configuration, a vector `moves` which we describe below, and the vector `strategy` that tells the computer all valid moves.

The variable `moves` should be a vector of length two. The first component of this vector will consist of moves made so far by player one, and the second will consist of moves made so far by player 2. As the computer makes a move, it will add information describing that move to `moves[n-1]`.

This routine should first convert `position` into a three-digit integer `pos`. It will then look at the possible moves in `strategy[pos]`. Set `mx` equal to the largest value of the third component of the possible moves, and then create a vector `bestmoves` containing the indices in `strategy[pos]` of all the possible moves whose third component matches that maximum value. As an example, if `position=[3,1,2]` then `pos=312` and we might have

```
strategy[pos]=[[1,3,80],[1,2,20],[1,1,40],[2,1,80],[3,2,60],[3,1,80]]
```

The maximum value of the third component is 80, and we should end up with the vector `bestmoves=[0,3,5]`, indicating that the 0, 3, and 5 entries of the vector achieve the maximum.

Randomly choose an element `m` of `bestmoves` (this can be done, for instance, with the command `m=bestmoves[random.randint(0,len(bestmoves)-1)]`), and execute the indicated move (i.e. the move described in `strategy[pos][m]` by subtracting the appropriate number of objects from the appropriate pile). Print out “Computer player `n` takes `a` objects from pile `b`”, and append the vector `[pos,m]` to the vector `moves[n]`.

Finally, if the current value of `position` is equal to `[0,0,0]`, print “Computer player `n` wins” and exit the function. Note that no value needs to be returned.

Your function should now allow you to play a computer. The computer should, at this point, be selecting random moves from all possible moves (because every possible move has a “quality” of 50). You should be able to defeat this random computer easily.

2. Now copy your function `computer` into the function `computertrainer`. We will modify it slightly to make it more suitable to training the computer. To do this, we will want to do two things differently.

First, eliminate the print statements. We will train the computer by having it play thousands of games against itself, and we don’t want to see the output of those games.

Second, when selecting the best moves, choose all the moves with a quality greater than `mx-90`. If we only select the best moves, then the computer will only ever try one winning move of all the possible moves; namely the first one that leads it to a victory. There could be other, even better moves available, and we want the computer to try them out. Note that the value 90 is arbitrary, after the next step you may want to experiment with it, to see how best to train the computer. Basically a value closer to 100 reserves judgement on whether a move is good or bad until later; a smaller value judges a move to be permanently bad as soon as there is one that is marginally better.

3. Our final function will be `traincomputer(strategy)`. This function will teach the computer which moves are the best, so that it can consistently win.

We want the computer to play thousands of games against itself. Each time the computer wins, it should increase the quality of the moves that it made; each time it loses, it should decrease the quality. Eventually, the best moves will become apparent. Although this is the basic idea, we will make some modifications so that the best moves become apparent faster.

The function `traincomputer(strategy)` will consist of a large loop, which will cause the computer to play against itself over and over, becoming a little bit better each time.

The first thing to do inside the loop is to set the starting position, `position=[7,5,3]`, the player number `playernumber=2`, and initialize the list of moves that have been made, `moves=[[],[]]`. Note that the moves made are divided between moves made by player one and moves made by player two.

Next, have the computer play a game against itself, using the `computertrainer` function. When the game ends, keep track of which player was the winner, and which was the loser.

Now, examine the moves made by the winner. Each of these moves will consist of a pair, `[pos,move]`, where `pos` is the position before the move, and `mov` is the number of the move that was made from that position. For each of them increase the value of `strategy[pos][mov][2]` by 29 (but only up to a maximum of 100).

Now examine the moves made by the loser. Each of these moves will consist of a pair, `[pos,move]`, where `pos` is the position before the move, and `mov` is the number of the move that was made from that position. For each of them decrease the value of `strategy[pos][mov][2]` by 11 (but only down to a minimum of 0).

As a further innovation, the last move made by the winner is a certain win; adjust the quality of that move to 1000 (we always want to make that move if we are in this position, regardless of how good the other moves are). The last move made by the loser leads immediately to a win by the other player, decrease the quality of that move to -1000 (since we would prefer to never make that move again, and don't need any more practice to know that it is a bad move).

With the parameters listed above, the function `traincomputer(strategy)` should produce a good strategy after about 15,000 games have been played. Changing the parameters (the 90 in the previous part, and the 29 and 11 in this part) will adjust the speed at which the strategy approaches a good strategy. Feel free to adjust these parameters to experiment, and see if you can get a good strategy in less than 15,000 games.

Lab 22

Finding Areas of Arbitrary Polygons

You may remember from elementary or middle school the formulas for special polygons such as triangles, trapezoids, and regular hexagons. In this lab we will use the theorems of Green and Stokes to find the area of arbitrary polygons in the plane and in 3-space. In theory we could use Monte Carlo integration to find the area, but for an arbitrary polygon it can be hard to determine whether a point is inside or outside the polygon. Recall Green's Theorem which says:

Green's Theorem: Let C be a positively oriented, piecewise-smooth, simple closed curve in the plane and let D be the region bounded by C . If P and Q have continuous partial derivatives on an open region that contains D , then

$$\int_C P \, dx + Q \, dy = \iint_D \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dA.$$

Since we know that $\iint_D 1 \, dA$ gives the area of a region D if we can find a P and Q such that $\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} = 1$ then we could use Green's theorem to find the area. There are several ways to do this but one way would be to set $P(x, y) = -\frac{1}{2}y$ and $Q(x, y) = \frac{1}{2}x$, then $A = \frac{1}{2} \oint_C x \, dy - y \, dx$. A few moments of thought (and recalling problem 25 in Section 16.4), will show that if C is the line segment connection the point (x_1, y_1) to the point (x_2, y_2) then

$$\int_C x \, dy - y \, dx = x_1 y_2 - x_2 y_1.$$

1. Write a function that takes in a list of vertices of a polygon in the plane, listed in counter clockwise order and determines the area of the polygon.
2. Try your code on the following two polygons:
 - a) $(0, 0)$, $(2, 1)$, $(1, 3)$ $(0, 2)$, and $(-1, 1)$
 - b) $(3, 0)$, $(5, 3)$, $(1, 7)$ $(-1, 4)$, $(-5, 7)$, $(-5, -2)$, $(-2, -6)$ $(5, -6)$, $(2, -3)$ and $(5, -2)$

Suppose, now you wish to find the area of the face of polyhedra in \mathbb{R}^3 . One way you could accomplish this would be to rotate and translate the polyhedra so that the face lies in the xy plane. Alternatively you could use Stokes' theorem, which we state below.

Stokes' Theorem: Let S be an oriented piecewise-smooth surface that is bounded by a simple, closed piecewise-smooth boundary curve C with positive orientation. Let \mathbf{F} be a vector field whose components have continuous partial derivatives on an open region in \mathbb{R}^3 that contains S . Then

$$\int_C \mathbf{F} \cdot d\mathbf{R} = \iint_S \text{curl } \mathbf{F} \cdot d\mathbf{S} = \iint_S \text{curl } \mathbf{F} \cdot \mathbf{n} dS.$$

Recalling that the surface area is equal to $\iint_S 1 dS$, we need only find appropriate values for \mathbf{F} and \mathbf{n} so that $\text{curl } \mathbf{F} \cdot \mathbf{n} = 1$. Let's begin with finding an appropriate value for \mathbf{n} . Since we are interested in finding the surface area of a face of a polyhedra, we know that the face is a plane. The normal to the plane (and hence the surface) can be determined by finding the equation of the plane.

3. Write a function that takes in three points in \mathbb{R}^3 and returns the coefficients a , b , c , and d such that the 3 points lie on the plane $ax + by + cz = d$. (If you are not enrolled in Math 314, or you want a review of how to do this, there is an example at the end of the lab.)

Now that we have found $\mathbf{n} = \langle a, b, c \rangle$ it remains to find an appropriate value of \mathbf{F} so that $\text{curl } \mathbf{F} \cdot \mathbf{n} = 1$. In general, you cannot always find the anti-curl. That is to say given a vector field \mathbf{G} it is not always possible to find a vector field \mathbf{F} such that $\text{curl } \mathbf{F} = \mathbf{G}$. However for this specific problem it can be done. Here's how:

Suppose that $\mathbf{G} = \left\langle \frac{1}{3a}, \frac{1}{3b}, \frac{1}{3c} \right\rangle$. A straight forward calculation shows $\mathbf{G} \cdot \mathbf{n} = 1$. Since \mathbf{G} is divergence free (i.e., the divergence of $\mathbf{G} = 0$), we can find an \mathbf{F} such that $\text{curl } \mathbf{F} = \mathbf{G}$. Unsurprisingly, this \mathbf{F} is not unique and one such example of an \mathbf{F} is:

$$\mathbf{F} = \frac{1}{6} \left\langle \frac{z}{b} - \frac{y}{c}, \frac{x}{c} - \frac{z}{a}, \frac{y}{a} - \frac{x}{b} \right\rangle.$$

(Those of you who are enrolled in Math 314 should check this, it is good practice for your final.)

If C is the line segment connection the point (x_1, y_1, z_1) to the point (x_2, y_2, z_2) then

$$\int_C \mathbf{F} \cdot d\mathbf{r} = \frac{1}{6} \left(\frac{x_1 y_2 - x_2 y_1}{c} + \frac{x_2 z_1 - x_1 z_2}{b} + \frac{y_1 z_2 - y_2 z_1}{a} \right)$$

4. Write a function that takes in a list of vertices of a face of a polyhedron in \mathbb{R}^3 , listed in counter clockwise order, and determines the surface area of that face. Your code should:
 - a) Calculate the normal vector to the plane (and the equation of the plane).

- b) Verify that all of the given points lie on the plane.
 - c) Determine the surface area of the face.
5. Check your code by calculating the surface area of the top face of the unit cube (i.e., insert as vertices $(0, 0, 1)$, $(1, 0, 1)$, $(1, 1, 1)$, and $(0, 1, 1)$).
6. Once you have verified that your code works find the surface area of the following polygons in 3-space.
- a) $(1, 1, 3)$, $(1 - 7/\sqrt{2}, 4, 3 - 7/\sqrt{2})$, $(1 - 7/\sqrt{2}, -4, 3 - 7/\sqrt{2})$, $(1 + 5/\sqrt{2}, -4, 2 + 5/\sqrt{2})$, and $(1 + 5/\sqrt{2}, 6, 2 + 5/\sqrt{2})$
 - b) $(0.2367, -0.5266, -0.8165)$, $(-6.6449, 2.7101, -1.2247)$, $(-4.7513, -1.5027, -7.7567)$, $(5.8285, -4.3431, -2.8577)$, and $(3.4614, 0.9229, 5.3072)$

Finding the equation of a plane example. Suppose you want to find the equation of the plane that passes through the points $A = (3, 0, -1)$, $B = (-2, -2, 3)$ and $C = (7, 1, -4)$. The first thing is to create two vectors from the points. We determine the vectors $\mathbf{AB} = \langle -5, -2, 4 \rangle$ and $\mathbf{AC} = \langle 4, 2, -3 \rangle$. These are two vectors that lie on the plane. The next step is to find a vector orthogonal to these two vectors, which can be done by finding the cross product. Hence $\langle a, b, c \rangle = \mathbf{n} = \mathbf{AB} \times \mathbf{AC} = \langle -2, 1, -2 \rangle$. Once we have found this normal vector we can use the equation of the plane

$$a(x - x_0) + b(y - y_0) + c(z - z_0) = 0 \quad \text{where } (x_0, y_0, z_0) \text{ is a point on the plane.}$$

Hence the equation of our plane is $-2(x - 3) + 1(y - 0) - 2(z - (-1)) = 0$ or writing it in the desired form $2x - y + 2z = 4$.