Higher Order Functions

You're used to seeing functions defined in Python:

In [2]:
```python
def meaningOfLife(x) :
    return 42 * x

print(meaningOfLife(7))
```

294

But what you may not be used to is that functions in Python are just like any other piece of data. That means that you can assign them to variables. And when you do, you can treat those variables as though they are just accessing the function:

In [3]:
```python
f = meaningOfLife
print(f(7))
```

294

Because functions act like any other data in Python, we call them *first class*. This means that we can, for example, pass them in as arguments to another function.

In [4]:
```python
def foo(fun, x) :
    return 2 * fun(x)
```

Note that `foo` here has behavior that changes based on what `fun` is. If we pass different functions in to `foo`, it will do different things. We call `foo` a *higher order function.*

In [6]:
```python
foo(f, 7)
```

Out[6]: 588

In [7]:
```python
def other(x) :
    return 39 * x

foo(other, 7)
```

Out[7]: 546

Filter

One of the best uses of higher order functions is to build generic helper functions that do different things based on the function you pass in to it.

Suppose we want to write a *filter* function, that only keeps data within a certain range:

```
In [8]:  import numpy as np
         data = np.loadtxt('inp.txt')
         print(len(data))
```

```
1000
```

```
In [9]:  print(data[:10])
```

```
[23.09471005 22.25991168 35.93632832 44.61182845 32.15680942 34.76223471
 46.67180901 31.88987773 40.08096148 37.23492821]
```

```
In [10]: def simpleFilter(data) :
             res = []
             for d in data :
                 if d >= 40 and d <= 60 :
                     res.append(d)
             return res
```

```
In [11]: filtered = simpleFilter(data)
         print(len(filtered))
         print(filtered[:10])
```

```
469
[44.611828449897416, 46.67180901057549, 40.080961476474215, 44.8976241688
8816, 46.52490829991481, 42.629775565805254, 57.721400483572964, 51.92827
054817613, 47.94710023174243, 43.82490586318856]
```

But now we want to change the filter to keep data in a different range. It looks like we have to rewrite the function:

```
In [13]: def simpleFilter(data) :
             res = []
             for d in data :
                 if d >= 60 and d <= 80 :
                     res.append(d)
             return res

         filtered = simpleFilter(data)
         print(len(filtered))
         print(filtered[:10])
```

```
22
[61.30216262756785, 70.18292815593877, 61.25057319688144, 60.252100564844
14, 61.41859000380041, 76.78503985833532, 60.008518770302615, 68.14344504
205286, 68.33551199471, 64.8951874891759]
```

One option is to add some additional parameters to our filter function. For example, we could add parameters to define the lower and upper bounds of the range we want to filter:

```
In [14]:  def simpleFilter(data, lo, hi) :
              res = []
              for d in data :
                  if d >= lo and d <= hi :
                      res.append(d)
              return res

          filtered = simpleFilter(data, 40, 60)
          print(len(filtered))
          print(filtered[:10])
```

```
469
[44.611828449897416, 46.67180901057549, 40.080961476474215, 44.8976241688
8816, 46.52490829991481, 42.629775565805254, 57.721400483572964, 51.92827
054817613, 47.94710023174243, 43.82490586318856]
```

```
In [15]:  filtered = simpleFilter(data, 60, 80)
          print(len(filtered))
          print(filtered[:10])
```

```
22
[61.30216262756785, 70.18292815593877, 61.25057319688144, 60.252100564844
14, 61.41859000380041, 76.78503985833532, 60.008518770302615, 68.14344504
205286, 68.33551199471, 64.8951874891759]
```

But that is not satisfying. What if we want to create a filter that does something entirely different, like keep numbers that are *outside* a range, or keep numbers that are *even*? We cannot use `simpleFilter` anymore. We would need to write something different each time we wanted to do a different *kind* of filtering.

So can we do better? What if we take advantage of higher order functions? Suppose we write a filter that *keeps data that pass a test* and then *pass the test to the filter*? Let's write some simple tests:

```
In [16]:  def inRange40_60(d) :
              return True if d >= 40 and d <= 60 else False
```

```
In [17]:  inRange40_60(45)
```

```
Out[17]:  True
```

```
In [18]:  inRange40_60(85)
```

```
Out[18]:  False
```

```
In [19]:  def inRange60_80(d) :
              return True if d >= 60 and d <= 80 else False
```

```
In [20]:  inRange60_80(65)
```

```
Out[20]:  True
```

In [21]:
```python
inRange60_80(85)
```

Out[21]:
```
False
```

Now we can write a filter that accepts a test `p` (`p` here stands for `predicate`):

In [22]:
```python
def higherOrderFilter(data, p) :
    res = []
    for d in data :
        if p(d) :
            res.append(d)
    return res
```

In [23]:
```python
filtered1 = higherOrderFilter(data, inRange40_60)
print(len(filtered1))
```

```
469
```

In [24]:
```python
filtered2 = higherOrderFilter(data, inRange60_80)
print(len(filtered2))
```

```
22
```

In [26]:
```python
def outOfRange(d) :
    return True if d < 40 or d > 60 else False

filtered3 = higherOrderFilter(data, outOfRange)
print(len(filtered3))
```

```
531
```

Returning functions from functions

Now we have a completely generic function. But suppose we want to simplify the process of creating tests? Instead of defining a new function from each test, what if we can write a function that *defines new functions* for us? To do this, we will take advantage of returning functions from functions:

In [27]:
```python
def createRangeP(lo, hi) :
    def p(d) :
        return True if d >= lo and d <= hi else False
    return p
```

It can be a little hard to understand what `createRangeP` is doing, so let's look at a couple of examples:

In [28]:
```python
p1 = createRangeP(40, 60)
p2 = createRangeP(60, 80)
```

```
In [29]: p1(45)
```

Out[29]: True

```
In [30]: p1(65)
```

Out[30]: False

```
In [31]: p2(45)
```

Out[31]: False

```
In [32]: p2(65)
```

Out[32]: True

Think about what happens when `createRangeP` runs. When it does, it *defines a new function called `p`*. That function has specific values for `lo` and `hi` (because we passed them in to `createRangeP`, so `p` is specialized for that particular range. We then *return* this newly created function. Note that we *have not actually run `p` yet*. Instead, `p` is now a function that runs a test on its input argument, `x`. We then run it later, as we did above.

We can now use the newly created functions in our filter:

```
In [33]: len(higherOrderFilter(data, p1))
```

Out[33]: 469

```
In [34]: len(higherOrderFilter(data, p2))
```

Out[34]: 22

We can also skip the step of assigning the result of `createRangeP` to a variable:

```
In [35]: len(higherOrderFilter(data, createRangeP(45, 75)))
```

Out[35]: 298

In class, we looked at a couple of other uses of returning functions from a function. For example, here is a function that takes in two tests (`fun1` and `fun2`) and returns a *new* test that returns true if both `fun1` and `fun2` pass:

```
In [37]: def createAnd(fun1, fun2) :
             def p(x) :
                 return fun1(x) and fun2(x)
             return p

         def keepEven(x) :
             return (int(x) % 2 == 0)

         andP = createAnd(keepEven, createRangeP(45, 75))
         len(higherOrderFilter(data, andP))

Out[37]: 150
```

As promised, here's a version of `createAnd` that takes in a whole list of functions and creates a new test that returns true if all of the functions are true. And as a bonus, a `createOr` :

```
In [39]: def createAndL(funcList) :
             def p(x) :
                 res = True
                 for f in funcList :
                     res = res and f(x)
                 return res
             return p

         def createOrL(funcList) :
             def p(x) :
                 res = False
                 for f in funcList :
                     res = res or f(x)
                 return res
             return p
```

Map and Reduce

Map and reduce are two of the most common higher-order functions. Map takes a list and a function and returns a new list where each element of the new list is an element from the first list with the function applied to it:

```
In [40]: def myMap(inp, f) :
             res = []
             for i in inp :
                 res.append(f(i))
             return res
```

```
In [41]: def sq(x) : return x * x
```

```
In [42]: small = [5, 1, 3, 7, 4, 8, 9]
```

In [43]:
```python
myMap(small, sq)
```

Out[43]: `[25, 1, 9, 49, 16, 64, 81]`

Instead of defining a new function every time we want to use it in a higher order function, we can use a *lambda* to define a function at the same time we need it:

In [44]:
```python
squared = myMap(small, lambda x : x * x)
print(squared)
```

`[25, 1, 9, 49, 16, 64, 81]`

Reduce takes a list and combines together all the elements by calling a function `f` over and over that combines the numbers (e.g., adds them together):

In [45]:
```python
def myReduce(inp, f, start) :# f(curr, i) -> curr'
    curr = start
    for i in inp :
        curr = f(curr, i)
    return curr
```

In [46]:
```python
sums = myReduce(small, lambda curr, i : curr + i, 0)
print (sums)
```

`37`

In [49]:
```python
def average(inp) :
    return (myReduce(inp, lambda curr, i : curr + i, 0) / len(inp))
```

In [50]:
```python
average(small)
```

Out[50]: `5.285714285714286`

We can combine map and reduce to compute more complicated things:

In [51]:
```python
def variance(inp) :
    avg = average(inp)
    diffs = myMap(inp, lambda x : x - avg)
    sq_diffs = myMap(diffs, lambda x : x * x)
    return average(sq_diffs)
```

In [52]:
```python
variance(small)
```

Out[52]: `7.061224489795919`

In [53]:
```python
variance(data)
```

Out[53]: `101.95142988919245`

List comprehensions

If you read a lot of Python code, you won't often see people using `map` and `filter`, because the same thing can be done more concisely using *list comprehensions*:

```
In [54]: [sq(d) for d in data]
```

```
Out[54]: [533.3656322013799,
          495.50366794291506,
          1291.4196930422397,
          1990.2152376430765,
          1034.0603921777931,
          1208.412961772856,
          2178.257756319635,
          1016.9643014619114,
          1606.4834728786102,
          1386.4398786509048,
          499.9852990840247,
          2015.79665601073,
          2164.567092315482,
          1817.2977647909267,
          978.8274925683775,
          672.2174336543578,
          1319.7606066038156,
          3331.760073785017,
          2696.5452821245763,
```

Read this "inside out": for each d in data, apply the function `sq(d)`, and put the results into an output list (note that `data` itself does not change).

We can also combine this with filter:

```
In [56]: [sq(d) for d in data if keepEven(d)]
```

```
Out[56]: [495.50366794291506,
          1990.2152376430765,
          1034.0603921777931,
          1208.412961772856,
          2178.257756319635,
          1606.4834728786102,
          499.9852990840247,
          2015.79665601073,
          2164.567092315482,
          1817.2977647909267,
          1319.7606066038156,
          603.4313675120274,
          1049.0582561292456,
          1493.1914713711317,
          788.9221757691865,
          944.2025860311232,
          1631.8857770260788,
          2534.094848066772,
          839.3604062703818,
```

Which now says: for each d in data, if `keepEven(d)` is true, compute `sq(d)` and put the result in an output list.

In [ ]: