1. """
2. Toward higher-order functions
3. 
4. Dan Armendariz
5. danallan@cs.harvard.edu
6. """
7. 
8. \# create an arbitrary list of numbers
nums $=[1,5,10,8,-4,432]$
9. 

print "Initial:", nums
\# now let's
result = []
for num in nums:
result.append(num + 1)
print "Add 1: ", result
19.
20. \# perhaps we now want to multiply each by some value
result = []
for num in nums:
result.append(num * 2)
24.
25. print "Double: ", result

1. """

Toward higher-order functions
3.
. Dan Armendariz
. danallan@cs.harvard.edu
6. " " "
7.
. \# create an arbitrary list of numbers
nums $=[1,5,10,8,-4,432]$
10.
print "Initial:", nums
def incrementor(num) :
""" Increment a number by 1 """
return num + 1
\# iterate over every item in the list, applying our function to it
result = []
for num in nums:
result.append(incrementor(num))
print "Add 1: ", result
def doubler(num):
""" Double a number """
return num * 2
\# iterate over every item in the list, applying our function to it
result = []
for num in nums:
result.append(doubler(num)
33.
print "Double: ", result

```
. """
Toward higher-order functions
3.
. Dan Armendariz
5. danallan@cs.harvard.edu
6. """
7.
. # create an arbitrary list of numbers
nums = [1, 5, 10, 8, -4, 432]
10.
print "Initial:", nums
def apply_to_all(f, items):
    """ Apply a function f to every item in a list """
    result = []
        for item in items
            new_item = f(item)
            result.append(new_item)
        return result
    def incrementor(num)
    """ Increment a number by 1 """
    return num + 1
def doubler(num):
    """ Double a number ""
    return num * 2
29. print "Add 1: ", apply_to_all(incrementor, nums)
30. print "Double: ", apply_to_all(doubler, nums)
```

1. """
2. Using Python's built-in Higher-Order Functions (HOFs) : map
3. 
4. Dan Armendariz
5. danallan@cs.harvard.edu
6. """
7. 

\# create an arbitrary list of numbers
nums $=[1,5,10,8,-4,432]$
10.
print "Initial:", nums
12.
def incrementor(num)
""" Increment a number by 1 """
return num + 1
def doubler(num)
""" Double a number """
return num * 2
21. \# the 'map` function is a HOF that applies a function to an entire list
print "Add 1: ", map(incrementor, nums)
23. print "Double: ", map(doubler, nums)

1. """
. Using Python's built-in Higher-Order Functions (HOFs) : reduce
2. 
3. Dan Armendariz
. danallan@cs.harvard.edu
4. """
5. 

. \# create an arbitrary list of numbers
nums $=[1,5,10,8,-4,432]$
10.
print "Initial:", nums
def add (x, y) :
""" Sum two numbers """
return $\mathrm{x}+\mathrm{y}$
def multiply(x, y)
return $x$ * $y$
\# the `reduce` function is a HOF that collapses a list to a single value
print "Sum of all values: ", reduce(add, nums, 0)
print "Multiply all values: ", reduce(multiply, nums, 1)
" " "
Reduce operates in this way:
We want to sum all of the objects, so we provide our function `sum` to reduce()
It then applies sum on the first element and the initial value, and the
output of that function is then used as input to the next iteration.
29.
30. Let's say we have a list $[1,5,8]$
. The first operation would be:
sum(0, 1) -- "0" is our initial value. The output is 1
1 is used as one of the next inputs; the other is the next item in the list
sum(1, 5) -- outputs "6"
sum(6, 8) -- outputs our final value, 14!
38.
39. Note: the order is not guaranteed!
40.

1. """
2. Use MapReduce (Hadoop) style programming to count the most common letter
unigrams (single characters) and bigrams (character pairs) in a text.
See: http://en.wikipedia.org/wiki/Bigram

Hadoop assumes that the input data is a text file. This is convenient,
because text files can be easily broken into chunks at individual lines
and distributed to many computers. It then roughly operates in this fashion:
9.
10. MAP STEP: runs a developer-provided function on one line at a time of the input 11. files. The Map step is therefore generally used to convert the data represented
12. in the line into one or more key/value pairs. Normally, we could think of
13. "key/value" pairs as a dictionary, but Hadoop lets us have duplicate keys.
14. So, in essence, if we're writing code for Hadoop (which is normally done in the
15. Java programming language, but we're using Python for familiarity), we want
16. to write a map function that will accept a line of text and output a dictionary.
17.
18. SHUFFLE STEP: As key/value pairs come in from the map step, hadoop sorts them
19. and redirects all values that have the same key to the same computer to run
20. the Reduce step. The map step must finish completely before moving on to
21. the Reduce step
22.
23. REDUCE STEP: Another developer-provided function, the reducer is given a
24. key and a list of all of the values for that key by Hadoop. It then
25. performs some computation and outputs a *single* value for that key.
26.
27. SORT STEP: Hadoop will generally sort the final output data.
28.
29. It is not unusual to have to run several Map/Reduce steps on a big data set
30. before receiving the output you want. In our example below, we'll only need a
single Map and single Reduce step to complete our goal.
32.
33. See more information on MapReduce applications, of which Hadoop is a part, here: http://en.wikipedia.org/wiki/MapReduce
35.
36. For a good exercise, consider how you could modify the code below to count
unigram WORDS and bigram WORDS. In other words, instead of counting individual characters or character pairs, counting whole words and counting word pairs.
39.
40. Dan Armendariz, Dan Bradley
41. danallan@cs.harvard.edu, dbradley@cs50.harvard.edu
42. """
43.
44. \# itemgetter will assist in sorting
45. from operator import itemgetter
46.
def unigram_mapper(line):
48. """Given a line of text, remove extraneous characters and output a key/value

```
9. pair for each character. The character will be the key and the value will
be 1
52.
53.
55.
56.
57.
58.
70.
```

73. 
74. 
75. 
76. 
```
def count_reducer(key, values):
    """Given a list of values for a key, we'll return the sum of all values."""
    total = 0
    for value in values:
        total += value
    return (key, total)
def hadoop(file, mapper, reducer):
    """A higher-order function that accepts a mapper and a reducer and
    performs Hadoop/MapReduce-style processing on a given file.
    Essentially, this code (in a very simple way) performs on a single computer
    what Hadoop and other MapReduce implementations do across a wide number
    of computers.
    """
    # a structure to store the final output
    data = []
    # MAP STEP
    pairs = []
    with open(file) as f:
        # iterate over every line in the opened file
            for line in f:
            # ask our map step to process the line
            output = mapper(line)
            # store the newly computed key/value pairs in our `pairs` structure
            pairs.extend(output)
    # SHUFFLE STEP: aggregate all values for a given key
    collapsed = {}
    for pair in pairs:
            # split the pair into its key and value
            key, value = pair
            # if the key doesn't exist yet in our dictionary, add it
            if collapsed.get(key, None) is None:
                collapsed[key] = []
            # append the value to the list of values for a given key
            collapsed[key].append(value)
    # REDUCE STEP: provide all of the values for a key to the reducer
    for key, values in collapsed.iteritems():
            # the key is different; run the reduce step on our accumulated vals
            output = reducer(key, values)
```

145. 
146. 
147. \# SORT STEP: sort the output by values in descending order
148. data.sort(key=itemgetter(1), reverse=True)
149. 
150. \# all done!
151. return data
152. 
153. \# Using "Hadoop" and our reducer and unigram mapper, count the number of
154. \# instances each character appears in Carroll's "Alice and Wonderland"
155. print "Unigram character counts:"
156. result = hadoop("aliceandwonderland.txt", unigram_mapper, count_reducer)
157. for count in result:
158. print count
159. 
160. print
161. 
162. \# On to bigrams. Our reducer is so simple we can use the same one for both!
163. print "Bigram character counts:"
164. result = hadoop("aliceandwonderland.txt", bigram_mapper, count_reducer)
165. 
166. \# there are many bigrams, so we'll only print the first 20
167. for count in result[0:10]:
168. print count
