

Semester 2, 2018: Lab 5

S2 2018

Lab 5

Note: If you do not have time to finish all exercises (in particular, the programming problems) during the lab time, you should continue working on them later. You can do this at home (if you have a computer with python set up), in the CSIT lab rooms outside teaching hours, or on one of the computers available in the university libraries or other teaching spaces.

If you have any questions about or difficulties with any of the material covered in the course so far, ask your tutor for help during the lab.

Objectives

The purpose of this week's lab is to:

- review two of python's built-in sequence types - strings and lists;
- iterate over sequences using the `for` loop; and
- investigate some of the built-in methods for strings.

The string type

Strings (type `str`) is python's type for storing text. A string is a sequence, but unlike other sequence types (`list`, `tuple`, and `numpy.ndarray`, for example) a string can only contain characters.

The type of a value ("object") determines what we can do with it: For example, we can do arithmetic on numbers, and we can ask for the length of a sequence (using the `len` function), but we cannot ask for the length of a number, or do much arithmetic with strings. Some operations can be done on several value types, but have different effects for different types. For example, using the `+` operator on two strings concatenates them: thus, `"1.234" + "1.234"` evaluates to `"1.2341.234"`, whereas adding two floating point numbers `1.234 + 1.234` evaluates to `2.468`.

Writing string literals

Recall that a literal is an expression that represents a constant value - like a `1` represents the integer one, or `"abc"` represents the three-letter string `abc`. String literals in python are written by enclosing them in either single, double or triple quotation marks:

```
In [1]: str1 = 'a few words'
```

```
In [2]: str2 = "a few words"
```

```
In [3]: str3 = '''a few words'''
```

```
In [4]: str1 == str2
```

```
Out [4]: ...
```

```
In [5]: str2 == str3
```

```
Out [5]: ...
```

(Note that the triple quote is written by typing three single quotes.) There is no difference between single- and double-quoted strings. The triple quote has the special ability that it can stretch over several lines:

```
In [1]: many_line_string = '''This is line 1.
```

```
This is the second line.
```

```
And this is line three.'''
```

```
In [2]: many_line_string
```

```
Out [2]: ...
```

```
In [3]: print(many_line_string)
```

```
...
```

Note the difference between how `many_line_string` is displayed when you evaluate the expression and when you print it. What is the `'\n'` character?

Remember that a string enclosed in one type of quotation marks can not contain the same type of quotation marks (unless they are escaped - see the lecture slides on strings), but can contain the other types.

Exercise 0

Write string literals for each of the following sentences:

- The possessive form of 'it' is 'its' - "it's" is an abbreviation of "it is".
- A "" is three 'but two' do not make a ".

Note that the first sentence should be on two lines.

Use the `print` function to verify that your strings are displayed correctly. Try writing them both with and without using triple quotes.

Character encoding

A string is a sequence (ordered collection) of characters.

Characters (like every other type of information stored in a computer) are represented by numbers. Interpreting a number as a character requires an *encoding*; python 3 uses [the unicode standard](#) for encoding characters in strings.

Python provides the `ord` and `chr` functions for translating between numbers and the characters they represent: `ord(a_character)` returns the corresponding character code (as an `int`) while `chr(an_integer)` returns the character that the integer represents.

Exercise 1

Try the following:

```
In [1]: ord('a')
Out [1]: ...
```

```
In [2]: ord('A')
Out [2]: ...
```

```
In [3]: chr(91)
Out [3]: ...
```

```
In [4]: chr(92)
Out [4]: ...
```

```
In [5]: chr(93)
Out [5]: ...
```

```
In [6]: chr(20986)+chr(21475)
Out [6]: ...
```

```
In [7]: chr(5798) + chr(5794) + chr(5809) + chr(5835) + chr(5840) + chr(5830) + chr(5825) + chr(5823)
Out [7]: ...
```

Remember that characters outside the ASCII range (unicode numbers above 255) may not display properly, if the computer you are using does not have a font for showing those characters. Also, many of the characters below number 32 are so called “non-printable” control characters, which may not be displayed.

Sequences

Remember that:

- `str` and `list` are sequence types.
- A string (value of type `str`) can only contain characters, while a list can contain elements of any type - including a mix of elements of different types in the same list.

Exercise 2

Operations on python's built-in sequence types are summarised in [this section of the python library reference](#).

To remind yourself what you can do with a sequence, run the following in the python shell:

```
In [1]: aseq = "abcd"

In [2]: type(aseq)          # 1. what type of sequence is this?
Out [1]: ...

In [3]: aseq + aseq        # 2. concatenation
Out [1]: ...

In [4]: aseq * 4           # 3. repetition
Out [1]: ...

In [5]: aseq[0]           # 4. Indexing
Out [1]: ...

In [6]: type(aseq[0])
Out [1]: ...

In [7]: aseq[-2]
Out [1]: ...

In [8]: aseq[1:-2]        # 5. Slicing
Out [1]: ...

In [9]: aseq[1:2]         # 5b.
Out [1]: ...

In [10]: bseq = "abdc"

In [11]: aseq < bseq      # 6. Comparison
Out [1]: ...
```

```
In [12]: for elem in aseq:      # 7. Iteration, using a for loop.
        print(elem)
```

```
In [13]: min(aseq)           # 8. built-in functions: min, max, sorted
Out [1]: ...
```

```
In [14]: max(aseq)
Out [1]: ...
```

```
In [15]: sorted(aseq)
Out [1]: ...
```

Next, try the operations above with a different sequence type, this time a list:

```
In [1]: aseq = [1,2,3,4]
```

```
In [2]: bseq = [1,3,2,4]
```

```
In [3]: type(aseq)          # 1. what type of sequence is this?
Out [1]: ...
        # 2 - 8 as above
```

This experiment will show you some important differences, but also some similarities, between strings and lists:

- Slicing a list returns a list (test 5b).
- Ordering comparisons are defined the same way for lists as for strings.

The built-in functions `min` and `max` and `sorted` are applicable to any sequence type (in fact, to any iterable type). Look them up using python's `help` function. What happens if you apply them to an empty sequence?

Iteration over sequences

Python has two kinds of loop statements: the `while` loop, which repeatedly executes a suite as long as a condition is true, and the `for` loop, which executes a suite once for every element of a sequence. (To be precise, the `for` loop works not only on sequences but on any type that is *iterable*. All sequences are iterable, but later in the course we will see examples of types that are iterable but not sequences.)

Both kinds of loop can be used to iterate over a sequence. Which one is most appropriate to implement some function depends on what the function needs to do with the sequence. The `for` loop is simpler to use, but only allows you to look at one element at a time. The `while` loop is more complex to use (you must initialise and update an index variable, and specify the loop condition correctly) but allows you greater flexibility; for example, you can skip elements in the sequence (increment the index by more than one) or look at elements in more than one position in each iteration.

Exercise 3

(a) In [lab 4](#), you wrote a function to count the number of negative numbers in a list of numbers. Here is an example of a solution to this problem, implemented using a `while` loop:

```
def count_negative(array):
    count = 0
    index = 0
    while index < len(array):
        if array[index] < 0:
            count = count + 1
        index = index + 1
    return count
```

Rewrite this function so that it uses a `for` loop instead.

(b) Write a function called `count_capitals` that takes a string argument and returns the number of capital (upper case) letters in the string. The function should look very similar to the one you wrote for the previous exercise.

To do this, you will need to determine if a letter is a capital. It will be helpful to know that in the unicode character encoding, the capital letters (of the English alphabet) are ordered sequentially; that is `ord('A') + 1 == ord('B')`, `ord('A') + 2 == ord('C')`, etc, up to `ord('A') + 25 == ord('Z')`. (Alternatively, have a look at the documentation of python's string methods; there are several that help you do things with letter case.)

(c) (*advanced*) Is it possible to write a general function `count` that takes a sequence and some property X and counts how many elements in it have that property?

Hint: Functions are values in python. You can pass a function as an argument to another function.

(d) Two other functions that you were asked to write in [lab 4](#) were `most_average` and `is_increasing`. Can both of these be written using a `for` loop? Is the `while` loop more suitable for either one of the two problems?

Slicing

Python's built-in sequence types (which include type `str`) provide a mechanism, called slicing, to select parts of a sequence (that is, substrings if the sequence happens to be a string). It is done using the notation `sequence[start:end]`. There is also an extended form of slicing, which takes three arguments, written `sequence[start:end:step]`.

(Punch & Enbody's book has a detailed description of slicing, including its extended form, in Section 4.1.5 (page 183). Downey's book discusses slicing in Section "String Slices" in Chapter 8; the extended form of slicing is only briefly mentioned in Exercise 8-3.)

Exercise 4(a)

To make sure you understand what the arguments in a slicing expression mean, run through the following examples. For each expression, try to work out what the output will be before you evaluate the expression.

```
In [1]: my_string = "Angry Public Swamp Methods"
```

```
In [2]: L = len(my_string)
```

```
In [3]: my_string[1:L]
```

```
Out [3]: ...
```

```
In [4]: my_string[0:L - 1]
```

```
Out [4]: ...
```

```
In [5]: my_string[0:L:2]
```

```
Out [5]: ...
```

```
In [6]: my_string[L:0:-1]
```

```
Out [6]: ...
```

```
In [7]: my_string[6:6+6]
```

```
Out [7]: ...
```

```
In [8]: my_string[11:11-6:-1]
```

```
Out [8]: ...
```

```
In [9]: my_string[2*L]
```

```
Out [9]: ...
```

```
In [10]: my_string[0:2*L]
```

```
Out [10]: ...
```

(“Angry Public Swamp Methods” - Why such a strange, non-sense string? It’s designed to make it easier for you to see what happens when you try different slicing expressions. Can you see what’s special about it? Hint: Remember from the earlier exercise that `ord('S')` is not equal to `ord('s')` - ‘S’ and ‘s’ are different characters.)

Exercise 4(b)

Use the statement

```
In [1]: my_list = [i + 1 for i in range(26)]
```

to assign `my_list` a list of integers of the same length as the string used in the previous exercise, and try the slicing expressions on `my_list` instead of `my_string`. Is the result of all expressions what you expect?

String methods

A “method” is the same thing as a function, but it uses a slightly different call syntax. A method is always called on a particular object (value). A method call,

```
object.method(...)
```

can be thought of as “on this object, perform that method”.

For now, we will just explore some of the rich set of methods that python provides for performing operations on built-in data types, such as strings. You can find the documentation of python's string methods at <http://docs.python.org/3/library/stdtypes.html#text-sequence-type-str>, or by using the built-in help function in the python shell.

Try the following:

```
In [1]: sentence = "the COMP1730 lectures are boring, but the labs are great"
```

```
In [2]: sentence.capitalize()
```

```
Out [2]: ...
```

```
In [3]: sentence.swapcase()
```

```
Out [3]: ...
```

```
In [4]: sentence.count("e")
```

```
Out [4]: ...
```

```
In [5]: sentence.find("lectures")
```

```
Out [5]: ...
```

```
In [6]: sentence.find("exciting")
```

```
Out [6]: ...
```

```
In [7]: sentence.split(" ")
```

```
Out [7]: ...
```

```
In [8]: sentence.upper().find("LECTURES")
```

```
Out [8]: ...
```

```
In [9]: sentence.find("LECTURES").upper()
```

```
Out [9]: ...
```

After each method call, try `print(sentence)`. Did any of the methods modify the initial string?

Exercise 5

(From Punch & Enbody, Chapter 4: Question 14, on page 221.) There are five string methods that modify case: `capitalize`, `title`, `swapcase`, `upper` and `lower`.

- Look up each of these methods in the [python on-line documentation](#) or using the built-in `help` function. (Note: To find the documentation of a string method using `help`, you must prefix the method name with `str.`; that is, use `help(str.capitalize)` instead of `help(capitalize)`.)
- Based on your understanding, can you predict what will be the effect of each of these methods on the following strings:

```
s1 = "turner"
s2 = "north lyneham"
s3 = "AINSLIE"
s4 = "NewActon"?
```

To check if your understanding is correct, write python code to perform the operation, run it and see.

Programming problems

Note: We don't expect everyone to finish all these problems during the lab time. If you do not have time to finish these programming problems in the lab, you should continue working on them later (at home, in the CSIT labs after teaching hours, or on one of the computers available in the university libraries or other teaching spaces).

Cesar cipher

A Caesar cipher is based on simply shifting the letters in the alphabet by a fixed amount. For example we might do the following:

```
Plain:    ABCDEFGHIJKLMNOPQRSTUVWXYZ
Cipher:   DEFGHIJKLMNOPQRSTUVWXYZABC
```

So each 'A' in the message is replaced by a 'D', each 'M' by a 'P', and so on. That is, there is a shift of 3 letters. Note that the alphabet wraps around at the end: An 'X' (third from the end) is replaced by an 'A', etc.

(a) Write a function that takes two arguments, a string to encrypt and a shift value (an integer) to use for encrypting it, and returns the encrypted string. Apply the same shift to both lower and upper case letters. Do not alter the non-alphabetical characters (like space, comma etc).

Examples:

- Encrypting “Et tu, Brutus!” with a shift of 3 should return “Hw wx, Euxwxv!”
(Wikipedia has a [long list of Latin phrases](#) if you want to encrypt more examples in Latin.)
- Encrypting “IBM” with a shift of -1 should return “HAL”.

To decrypt a string, you can call the same function with the negative of the shift that was used to encrypt it. Also note that the function is invariant of the shift modulo 26: that is, a shift of 29 is the same as a shift of 3 ($3 == 29 \% 26$), and a shift of -1 is the same as a shift of 25 ($25 == -1 \% 26$).

(b) To break the Caesar cipher you just need to guess the shift. Try the following three approaches:

- There are only 25 possible different shifts to decrypt the code. Write a function that takes an encrypted string and prints out the first five or so words decrypted using successively larger shifts (up to 25). See if you can guess based on this which is the right shift value. (How many words do you need to check to tell the right shift value from the others?)
- Repeat the above, but this time decrypt the entire message using increasing shift values. For each shift value search the decrypted message to find how many of the following 40 “common” three letter words exist:

```
the, and, for, are, but, not, you, all, any, can,
her, was, one, our, out, day, get, has, him, his,
how, man, new, now, old, see, two, way, who, boy,
did, its, let, put, say, she, too, use, dad, mom
```

Return the shift value that gives the highest number of different three letter words. Does this automatic process succeed in breaking the cipher? (How did you deal with upper and lower case letters?)

- Next, consider the occurrence of individual letters. The [most common letter in English is “e”, which occurs nearly 13% of the time](#). For the encrypted text, determine the most frequently occurring letter, assume it should be an “e”, from this determine the shift, and decrypt the message. Does this automatic process correctly break the cipher? (Again - how did you deal with upper and lower case?)

Tests Here are some encrypted strings that you can try your decryption methods on:

- `''Awnhu pk neoa wjz awnhu pk xaz lwgao w iwj dawhpdu, xqp okyewhhu zawz''`
(Note the use of triple quotes because the string contains a line break.)
- `''Jcstghipcsxcv xh iwpi etctigpixcv fjpaxin du zcdlatsvt iwpi vgdh ugdb iwtdgn, egpirt, rdckxrixdc, phhtgixdc, tggdg pcs wjbxaxipixdc.'' (Gjat 7: Jht p rdadc puitg pc xcstetctci rapjht id xcigdsjrt p axhi du epgixrjapgh, pc peedhxixkt, pc pbeauxrxpidxc dg pc xaajhigpixkt fjdiipixdc. Ugdb Higjcz & Lwxit, "Iwt Tatbtcih du Hinat".)''`

(Note that the string is enclosed in single quotes because it contains double quotes.)

- "Cywo cmsoxdscdc gybu cy rkbm drobo sc xy dswv vopd pyb cobsyec drsxusxq. (kddbslledon dy Pbkxmsc Mbsmu)"

You should also make up your own tests (use the encryption function you wrote to encrypt sentences, then test if your cipher-breaking functions find the correct shift!)

Pig Latin

Pig Latin is a game of alterations played on words. To translate an English word into Pig Latin, the initial consonant sound is transposed to the end of the word and an “ay” is affixed. Specifically, there are two rules:

- If a word begins with a vowel, append “yay” to the end of the word.
- If a word begins with a consonant, remove all the consonants from the beginning up to the first vowel and append them to the end of the word. Finally, append “ay” to the end of the word.

For example,

- dog => ogday
- scratch => atchscray
- is => isyay
- apple => appleyay

Write a function that takes one argument, a string, and returns a string with each word in the argument translated into Pig Latin.

Hints:

- The `split` method, when called with no additional arguments, breaks a string into words, and returns the words in a list.
- Slicing is your friend: it can pick off the first character for checking, and you can slice off pieces of a string and use string concatenation (the `+` operator) to make a new word.
- Making a string of vowels allows use of the `in` operator: `vowels="aeiou"` (how do you make this work with both upper and lower case?)
- Test your function with a diverse range of examples. Your tests should cover all cases (for example, test words beginning with a vowel and words beginning with a consonant). Pay particular attention to edge cases (for example, what happens if the word consists of just one vowel, like “a”? what happens if the string is empty?).