

STAT679

Computing for Data Science and Statistics

Lecture 4: Dictionaries and Tuples

Two more fundamental built-in data structures

Dictionaries

Python dictionaries generalize lists

Allow indexing by arbitrary immutable objects rather than integers

Fast lookup and retrieval

<https://docs.python.org/3/tutorial/datastructures.html#dictionaries>

Tuples

Similar to a list, in that it is a sequence of values

But unlike lists, tuples are immutable

<https://docs.python.org/3/tutorial/datastructures.html#tuples-and-sequences>

Generalized lists: Python `dict()`

Python dictionary generalizes lists

`list()`: indexed by integers

`dict()`: indexed by (almost) any data type

Dictionary contains:


a set of indices, called **keys**

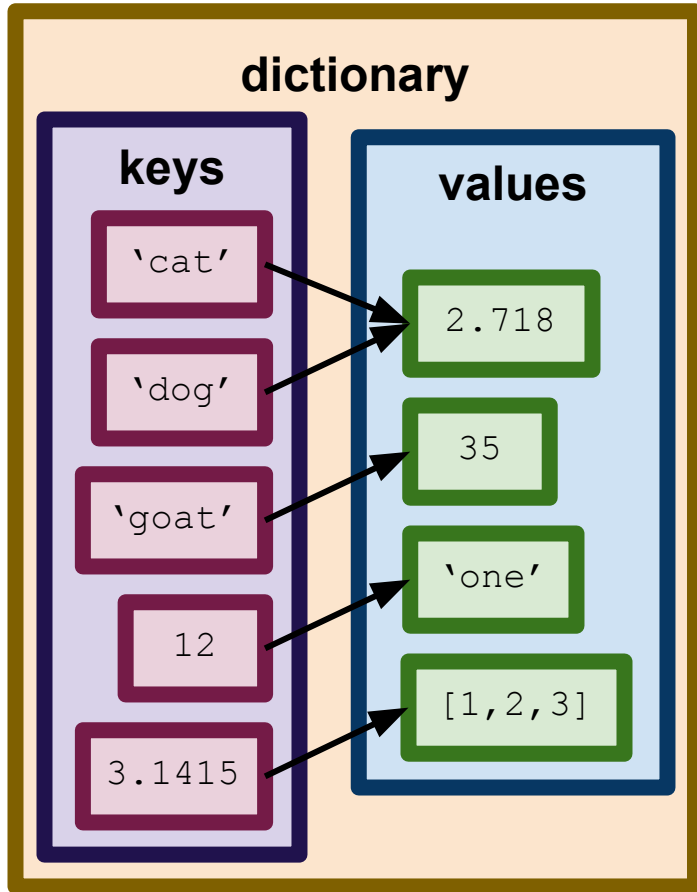
A set of values (called **values**, shockingly)

Each key associated with one (and only one) value

key-value pairs, sometimes called **items**

Like a function f : keys \rightarrow values






Dictionary **maps** keys to values.

E.g., '\cat' mapped to the float 2.718

Of course, the dictionary at the left is kind of silly. In practice, keys are often all of the same type, because they all represent a similar kind of object

Example: might use a dictionary to map UMich unique names to people



```
1 example_dict['cat']
```

```
2.718
```

```
1 example_dict['dog']
```

```
2.718
```

```
1 example_dict['goat']
```

```
35
```

```
1 example_dict[12]
```

```
'one'
```

```
1 example_dict[3.1415]
```


```
[1, 2, 3]
```

Access the value associated to key `x` by `dictionary[x]`.

Attempting to access the value associated to a non-existent key results in a `KeyError`, an error that Python supplies specifically for this situation.

Observe that `bird` is not a key in this dictionary, so when we try to index with it, we get an error.

```
1 example_dict['bird']
```



```
KeyError                                Traceback (most recent call last)
<ipython-input-12-057007a28b5b> in <module>()
----> 1 example_dict['bird']
```

```
KeyError: 'bird'
```

Creating and populating a dictionary

Example: University of Mishuges IT wants to store the correspondence between the usernames (UM IDs) of students to their actual names. A dictionary is a very natural data structure for this.

```
1 umid2name = dict()
2 umid2name['aeinstein'] = 'Albert Einstein'
3 umid2name['kyfan'] = 'Ky Fan'
4 umid2name['enoether'] = 'Emmy Noether'
5 umid2name['cshannon'] = 'Claude Shannon'
```

```
1 umid2name['cshannon']
```

```
'Claude Shannon'
```

```
1 umid2name['enoether']
```

```
'Emmy Noether'
```

```
1 umid2name['enoether'] = 'Amalie Emmy Noether'
2 umid2name['enoether']
```

```
'Amalie Emmy Noether'
```

Creating and populating a dictionary

Create an empty dictionary (i.e., a dictionary with no key-value pairs stored in it. This should look familiar, since it is very similar to list creation.

```
umid2name = dict()
```

```
2 umid2name['einstein'] = 'Albert Einstein'  
3 umid2name['kyfan'] = 'Ky Fan'  
4 umid2name['enoether'] = 'Emmy Noether'  
5 umid2name['cshannon'] = 'Claude Shannon'
```

```
1 umid2name['cshannon']
```

```
'Claude Shannon'
```

```
1 umid2name['enoether']
```

```
'Emmy Noether'
```

```
1 umid2name['enoether'] = 'Amalie Emmy Noether'  
2 umid2name['enoether']
```

```
'Amalie Emmy Noether'
```


Creating and populating a dictionary

Populate the dictionary. We are adding four key-value pairs, corresponding to four users in the system.

```
1 umid2name = dict()
2 umid2name['aeinstein'] = 'Albert Einstein'
3 umid2name['kyfan'] = 'Ky Fan'
4 umid2name['enoether'] = 'Emmy Noether'
5 umid2name['cshannon'] = 'Claude Shannon'
```

```
1 umid2name['cshannon']
```

```
'Claude Shannon'
```

```
1 umid2name['enoether']
```

```
'Emmy Noether'
```

```
1 umid2name['enoether'] = 'Amalie Emmy Noether'
2 umid2name['enoether']
```

```
'Amalie Emmy Noether'
```

Creating and populating a dictionary

```
1 umid2name = dict()
2 umid2name['aeinstein'] = 'Albert Einstein'
3 umid2name['kyfan'] = 'Ky Fan'
4 umid2name['enoether'] = 'Emmy Noether'
5 umid2name['cshannon'] = 'Claude Shannon'
```

```
1 umid2name['cshannon']
```

```
'Claude Shannon'
```

```
1 umid2name['enoether']
```

```
'Emmy Noether'
```

```
1 umid2name['enoether'] = 'Amalie Emmy Noether'
```

```
2 umid2name['enoether']
```

```
'Amalie Emmy Noether'
```

Retrieve the value associated with a key. This is called **lookup**.

Creating and populating a dictionary

```
1 umid2name = dict()
2 umid2name['aeinstein'] = 'Albert Einstein'
3 umid2name['kyfan'] = 'Ky Fan'
4 umid2name['enoether'] = 'Emmy Noether'
5 umid2name['cshannon'] = 'Claude Shannon'
```

```
1 umid2name['cshannon']
```

```
'Claude Shannon'
```

```
1 umid2name['enoether']
```

```
'Emmy Noether'
```

```
1 umid2name['enoether'] = 'Amalie Emmy Noether'
2 umid2name['enoether']
```

```
Amalie Emmy Noether'
```

Emmy Noether's actual legal name was Amalie Emmy Noether, so we have to update her record. Note that updating is syntactically the same as initial population of the dictionary.

Displaying Items

Printing a dictionary lists its **items** (key-value pairs), in this rather odd format...

```
1 example_dict
{3.1415: [1, 2, 3], 12: 'one', 'cat': 2.718, 'dog': 2.718, 'goat': 35}
```

```
1 umid2name
{'aeinstein': 'Albert Einstein',
 'cshannon': 'Claude Shannon',
 'enoether': 'Amalie Emmy Noether',
 'kyfan': 'Ky Fan'}
```

...but we can use that format to create a new dictionary.

```
1 umid2name = {'aeinstein': 'Albert Einstein',
2 'cshannon': 'Claude Shannon',
3 'enoether': 'Amalie Emmy Noether',
4 'kyfan': 'Ky Fan'}
```

```
1 umid2name['kyfan']
'Ky Fan'
```

Note: the order in which items are printed isn't always the same, and (usually) isn't predictable. This is due to how dictionaries are stored in memory. More on this soon.

Dictionaries have a length

```
1 umid2name
```

```
{'aeinstein': 'Albert Einstein',  
'cshannon': 'Claude Shannon',  
'enoether': 'Amalie Emmy Noether',  
'kyfan': 'Ky Fan'}
```

Length of a dictionary is just the number of items.

```
1 len(umid2name)
```

```
4
```

Empty dictionary has length 0.

```
1 d = dict()  
2 len(d)
```

```
0
```

Note: we said earlier that all sequence objects support the length operation. But there exist objects that **aren't** sequences that also have this attribute.

Checking set membership

```
1 umid2name
```

```
{'aeinstein': 'Albert Einstein',  
'cshannon': 'Claude Shannon',  
'enoether': 'Amalie Emmy Noether',  
'kyfan': 'Ky Fan'}
```

```
1 'akolmogorov' in umid2name
```

```
False
```

```
1 'enoether' in umid2name
```

```
True
```

Suppose a new student, Andrey Kolmogorov is enrolling at UMish. We need to give him a unique name, but we want to make sure we aren't assigning a name that's already taken.

Dictionaries support checking whether or not an element is present **as a key**, similar to how lists support checking whether or not an element is present in the list.

Checking set membership: fast and slow

```
1 from random import randint
2 listlen = 1000000
3 list_of_numbers = listlen*[0]
4 dict_of_numbers = dict()
5 for i in range(listlen):
6     n = randint(1000000,9999999)
7     list_of_numbers[i] = n
8     dict_of_numbers[n] = 1
```

```
1 8675309 in list_of_numbers
```

False

```
1 1240893 in list_of_numbers
```

True

```
1 8675309 in dict_of_numbers
```

False

```
1 1240893 in dict_of_numbers
```

True

Lists and dictionaries provide our first example of how certain **data structures** are better for certain tasks than others.

Example: I have a large collection of phone numbers, and I need to check whether or not a given number appears in the collection. Both dictionaries and lists support **membership checks** of this sort, but it turns out that dictionaries are much better suited to the job.

Checking set membership: fast and slow

```
1 from random import randint
2 listlen = 1000000
3 list_of_numbers = listlen*[0]
4 dict_of_numbers = dict()
5 for i in range(listlen):
6     n = randint(1000000,9999999)
7     list_of_numbers[i] = n
8     dict_of_numbers[n] = 1
```

This block of code generates 1000000 random “phone numbers”, and creates (1) a list of all the numbers and (2) a dictionary whose keys are all the numbers.

```
1 8675309 in list_of_numbers
```

False

```
1 1240893 in list_of_numbers
```

True

```
1 8675309 in dict_of_numbers
```

False

```
1 1240893 in dict_of_numbers
```

True

Checking set membership: fast and slow

```
1 from random import randint
2 listlen = 1000000
3 list_of_numbers = listlen*[0]
4 dict_of_numbers = dict()
5 for i in range(listlen):
6     n = randint(1000000,9999999)
7     list_of_numbers[i] = n
8     dict_of_numbers[n] = 1
```

```
1 8675309 in list_of_numbers
```

False

```
1 1240893 in list_of_numbers
```

True

```
1 8675309 in dict_of_numbers
```

False

```
1 1240893 in dict_of_numbers
```

True

The `random` module supports a bunch of random number generation operations. We'll see more on this later in the course.

<https://docs.python.org/3/library/random.html>

Checking set membership: fast and slow

```
1 from random import randint
2 listlen = 1000000
3 list_of_numbers = listlen*[0]
4 dict_of_numbers = dict()
5 for i in range(listlen):
6     n = randint(1000000,9999999)
7     list_of_numbers[i] = n
8     dict_of_numbers[n] = 1
```

Initialize a list (of all zeros) and an empty dictionary.

```
1 8675309 in list_of_numbers
```

False

```
1 1240893 in list_of_numbers
```

True

```
1 8675309 in dict_of_numbers
```

False

```
1 1240893 in dict_of_numbers
```

True

Checking set membership: fast and slow

```
1 from random import randint
2 listlen = 1000000
3 list_of_numbers = listlen*[0]
4 dict_of_numbers = dict()
5 for i in range(listlen):
6     n = randint(1000000,9999999)
7     list_of_numbers[i] = n
8     dict_of_numbers[n] = 1
```

Generate `listlen` random numbers, writing them to both the list and the dictionary.

```
1 8675309 in list_of_numbers
```

False

```
1 1240893 in list_of_numbers
```

True

```
1 8675309 in dict_of_numbers
```

False

```
1 1240893 in dict_of_numbers
```

True

Checking set membership: fast and slow

```
1 from random import randint
2 listlen = 1000000
3 list_of_numbers = listlen*[0]
4 dict_of_numbers = dict()
5 for i in range(listlen):
6     n = randint(1000000,9999999)
7     list_of_numbers[i] = n
8     dict_of_numbers[n] = 1
```

```
1 8675309 in list_of_numbers
```

False

```
1 1240893 in list_of_numbers
```

True

This is slow.

```
1 8675309 in dict_of_numbers
```

False

```
1 1240893 in dict_of_numbers
```

True

This is fast.

Checking set membership: fast and slow

```
1 import time
2 start_time = time.time()
3 8675309 in list_of_numbers
4 time.time() - start_time
```

0.10922789573669434

```
1 start_time = time.time()
2 8675309 in dict_of_numbers
3 time.time() - start_time
```

0.0002219676971435547

Let's get a more quantitative look at the speed difference between lists and dictionaries.

The `time` module supports accessing the system clock, timing functions, and related operations. <https://docs.python.org/3/library/time.html>
Timing parts of your program to find where performance can be improved is called **profiling** your code. Python provides some built-in tools for more profiling, which we'll discuss later in the course, if time allows.

<https://docs.python.org/3/library/profile.html>

Checking set membership: fast and slow

```
1 import time
2 start_time = time.time()
3 8675309 in list_of_numbers
4 time.time() - start_time
0.10922789573669434
```

```
1 start_time = time.time()
2 8675309 in dict_of_numbers
3 time.time() - start_time
0.0002219676971435547
```

To see how long an operation takes, look at what time it is, perform the operation, and then look at what time it is again. The time difference is how long it took to perform the operation.

Warning: this can be influenced by other processes running on your computer. See documentation for ways to mitigate that inaccuracy.

Checking set membership: fast and slow

```
1 import time
2 start_time = time.time()
3 8675309 in list_of_numbers
4 time.time() - start_time
```

0.10922789573669434

```
1 start_time = time.time()
2 8675309 in dict_of_numbers
3 time.time() - start_time
```

0.0002219676971435547

Checking membership in the dictionary is orders of magnitude faster! Why should that be?

Checking set membership: fast and slow

```
1 import time
2 start_time = time.time()
3 8675309 in list_of_numbers
4 time.time() - start_time
```

0.10922789573669434

```
1 start_time = time.time()
2 8675309 in dict_of_numbers
3 time.time() - start_time
```

0.0002219676971435547

The time difference is due to how the `in` operation is implemented for lists and dictionaries.

Python compares `x` against each element in the list until it finds a match or hits the end of the list. So this takes time **linear** in the length of the list.


Python uses a **hash table**. This lets us check if `x` is in the dictionary in (almost) the same amount of time, regardless of how many items are in the dictionary. How is that possible?

Crash course: hash tables

Universe of objects



Let's say I have a set of 4 items:




I want to find a way to know **quickly** whether or not an item is in this set.

Crash course: hash tables

Universe of objects



Let's say I have a set of 4 items:



I want to find a way to know **quickly** whether or not an item is in this set.

Hash function f maps objects to “buckets”
(in this example, let's say there are four buckets)

$$f : \{ \text{red}, \text{orange}, \text{yellow}, \text{green}, \text{blue}, \text{cyan}, \text{purple} \} \rightarrow \{1, 2, 3, 4\}$$

Crash course: hash tables

Hash function f maps objects to “buckets”



Let's say I have a set of 4 items:





Assign objects to buckets based on the outputs of the hash function.

$$f(\text{red circle}) = 1$$

$$f(\text{orange circle}) = 3$$






$$f(\text{yellow circle}) = 2$$

$$f(\text{green circle}) = 1$$



Crash course: hash tables

Hash function maps objects to “buckets”



Crash course: hash tables

Hash function maps objects to “buckets”



Let's say I have a set of 4 items:







Q: is this item in the set?








$$f(\text{cyan circle}) = 4$$

Look in bucket 4. Nothing's there, so the item wasn't in the set.



Crash course: hash tables

Hash function maps objects to “buckets”



Crash course: hash tables

Hash function maps objects to “buckets”



Let's say I have a set of 4 items:







Q: is this item in the set?



$$f(\text{yellow circle}) = 2$$

Look in bucket 2, and we find the object, so it's in the set.




Crash course: hash tables

Hash function maps objects to “buckets”



Let's say I have a set of 4 items:



Q: is this item in the set?



Bucket 1



Bucket 2



Bucket 3



Bucket 4

When more than one object falls in the same bucket, we call it a **hash collision**.

Crash course: hash tables

Hash function maps objects to “buckets”







Let's say I have a set of 4 items:

Q: is this item in the set?






$$f(\text{green circle}) = 1$$

Look in bucket 1, and there's more than one thing. Compare against each of them, eventually find a match.



Crash course: hash tables

Hash function maps objects to “buckets”



Worst possible case: have to check everything in the bucket only to conclude there's no match.

Crash course: hash tables

Hash function maps objects to "buckets"







Let's say I have a set of 4 items:

Q: is this item in the set?

$$f(\text{purple circle}) = 1$$

Look in bucket 1, and there's more than one thing. Compare against each of them, no match, so it's not in the set.



Crash course: hash tables

Hash function maps objects to “buckets”

Key point: hash table lets us avoid comparing against every object in the set (provided we pick a good hash function that has few collisions)

More information:

Downey Chapter B.4

https://en.wikipedia.org/wiki/Hash_table

https://en.wikipedia.org/wiki/Hash_function

For the purposes of this course, it suffices to know that dictionaries (and the related **set** object, which we'll see soon), have faster membership checking than lists because they use hash tables.

Common pattern: dictionary as counter

Example: counting word frequencies

Naïve idea: keep one variable to keep track of each word

We're gonna need a lot of variables!

Better idea: use a dictionary, keep track of only the words we see

```
1 wdcounts = dict()
2 for w in list_of_words:
3     wdcounts[w] += 1
```

This code as written won't work! It's your job in one of your homework problems to flesh this out. You may find it useful to read about the `dict.get()` method: <https://docs.python.org/3/library/stdtypes.html#dict.get>

Traversing a dictionary

Suppose I have a dictionary representing word counts...

...and now I want to display the counts for each word.

```
1 for w in wdcnt:  
2     print(w, wdcnt[w])
```

```
half 3  
a 3  
league 3  
onward 1  
all 1  
in 1  
the 2  
valley 1  
of 1  
death 1  
rode 1  
six 1  
hundred 1
```

Traversing a dictionary yields the keys, in no particular order. Typically, you'll get them in the order they were added, but this is not guaranteed, so don't rely on it.

This kind of traversal is, once again, a very common pattern when dealing with dictionaries. Dictionaries support iteration over their keys. They, like sequences, are **iterators**. We'll see more of this as the course continues.

<https://docs.python.org/dev/library/stdtypes.html#iterator-types>

Common Pattern: Reverse Lookup and Inversion

Returning to our example, what if I want to map a (real) name to a username?

E.g., I want to look up Emmy Noether's username from her real name

```
1 umid2name
```

```
{'aeinstein': 'Albert Einstein',  
'cshannon': 'Claude Shannon',  
'enoether': 'Amalie Emmy Noether',  
'kyfan': 'Ky Fan'}
```

```
1 name2umid = dict()  
2 for uname in umid2name:  
3     truename = umid2name[uname]  
4     name2umid[truename] = uname  
5 name2umid
```

```
{'Albert Einstein': 'aeinstein',  
'Amalie Emmy Noether': 'enoether',  
'Claude Shannon': 'cshannon',  
'Ky Fan': 'kyfan'}
```

The keys of `umid2name` are the values of `name2umid` and vice versa. We say that `name2umid` is the **reverse lookup table** (or the **inverse**) for `umid2name`.

Common Pattern: Reverse Lookup and Inversion

Returning to our example, what if I want to map a (real) name to a username?

E.g., I want to look up Emmy Noether's username from her real name

```
1 umid2name
```

```
{'aeinstein': 'Albert Einstein',  
'cshannon': 'Claude Shannon',  
'enoether': 'Amalie Emmy Noether',  
'kyfan': 'Ky Fan'}
```

```
1 name2umid = dict()  
2 for uname in umid2name:  
3     truename = umid2name[uname]  
4     name2umid[truename] = uname  
5 name2umid
```

```
{'Albert Einstein': 'aeinstein',  
'Amalie Emmy Noether': 'enoether',  
'Claude Shannon': 'cshannon',  
'Ky Fan': 'kyfan'}
```

The keys of `umid2name` are the values of `name2umid` and vice versa. We say that `name2umid` is the **reverse lookup table** (or the **inverse**) for `umid2name`.

What if there are duplicate values? In the word count example, more than one word appears 2 times in the text... How do we deal with that?

Common Pattern: Reverse Lookup and Inversion

```
1 print(wdcnt)
```

```
{'half': 3, 'a': 3, 'league': 3, 'onward': 1, 'all': 1, 'in': 1, 'the': 2, 'vall  
1, 'six': 1, 'hundred': 1}
```

```
1 wdcnt_reverse = dict()  
2 for w in wdcnt:  
3     c = wdcnt[w]  
4     if c in wdcnt_reverse:  
5         wdcnt_reverse[c].append(w)  
6     else:  
7         wdcnt_reverse[c] = [w]  
8 wdcnt_reverse
```

Here's our original word count dictionary (cropped for readability). Some values (e.g., 1 and 3) appear more than once.

Solution: map values with multiple keys to a list of all keys that had that value.

```
{1: ['onward', 'all', 'in', 'valley', 'of', 'death', 'rode', 'six', 'hundred'],  
2: ['the'],  
3: ['half', 'a', 'league']}
```

What if there are duplicate values? In the word count example, more than one word appears 2 times in the text... How do we deal with that?

Common Pattern: Reverse Lookup and Inversion

```
1 print(wdcnt)
```

```
{'half': 3, 'a': 3, 'league': 3, 'onward': 1, 'all': 1, 'in': 1, 'the': 2, 'vall  
1, 'six': 1, 'hundred': 1}
```

```
1 wdcnt_reverse = dict()
```

```
2 for w in wdcnt:
```

```
3     c = wdcnt[w]
```

```
4     if c in wdcnt_reverse:
```

```
5         wdcnt_reverse[c].append(w)
```

```
6     else:
```

```
7         wdcnt_reverse[c] = [w]
```

```
8 wdcnt_reverse
```

```
{1: ['onward', 'all', 'in', 'valley', 'of', 'death', 'rode', 'six', 'hundred'],
```

```
2: ['the'],
```

```
3: ['half', 'a', 'league']}
```

Here's our original word count dictionary (cropped for readability). Some values (e.g., 1 and 3) appear more than once.

Note: there is a more graceful way to do this part of the operation, mentioned in homework 2.

...es with multiple keys that had that value.

What if there are duplicate values? In the word count example, more than one word appears 2 times in the text... How do we deal with that?

Keys Must be Hashable

```
1 d = dict()
2 animals = ['cat', 'dog', 'bird', 'goat']
3 d[animals] = 1.61803
```

```
-----
TypeError                                 Traceback (most recent call last)
<ipython-input-77-9fa9089d27b7> in <module>()
      1 d = dict()
      2 animals = ['cat', 'dog', 'bird', 'goat']
----> 3 d[animals] = 1.61803
```

```
TypeError: unhashable type: 'list'
```

From the documentation: “All of Python’s immutable built-in objects are hashable; mutable containers (such as lists or dictionaries) are not.”

<https://docs.python.org/3/glossary.html#term-hashable>

Dictionaries can have dictionaries as values!

Suppose I want to map pairs (x,y) to numbers.

```
1 times_table = dict()
2 for x in range(1,13):
3     if x not in times_table:
4         times_table[x] = dict()
5     for y in range(1,13):
6         times_table[x][y] = x*y
7 times_table[7][9]
```

Each value of x maps to another dictionary.

Note: We're putting this if-statement here to illustrate that in practice, we often don't know the order in which we're going to observe the objects we want to add to the dictionary.

Dictionaries can have dictionaries as values!

Suppose I want to map pairs (x,y) to numbers.

```
1 times_table = dict()
2 for x in range(1,13):
3     if x not in times_table:
4         times_table[x] = dict()
5     for y in range(1,13):
6         times_table[x][y] = x*y
7 times_table[7][9]
```

In a few slides we'll see a more natural way to perform this mapping in particular, but this "dictionary of dictionaries" pattern is common enough that it's worth seeing.

Common pattern: memoization

```
1 def naive_fibo(n):
2     if n < 0:
3         raise ValueError('Negative Fibonacci number?')
4     if n==0:
5         return 0
6     elif n==1:
7         return 1
8     else:
9         return naive_fibo(n-1) + naive_fibo(n-2)
10 for i in range(8,13):
11     print(naive_fibo(i))
```

Raise an error. You'll need this in many of your future homeworks.
<https://docs.python.org/3/tutorial/errors.html#raising-exceptions>

21
34
55
89
144

Common pattern: memoization

```
1 def naive_fibo(n):
2     if n < 0:
3         raise ValueError('Negative Fibonacci number?')
4     if n==0:
5         return 0
6     elif n==1:
7         return 1
8     else:
9         return naive_fibo(n-1) + naive_fibo(n-2)
10 for i in range(8,13):
11     print(naive_fibo(i))
```

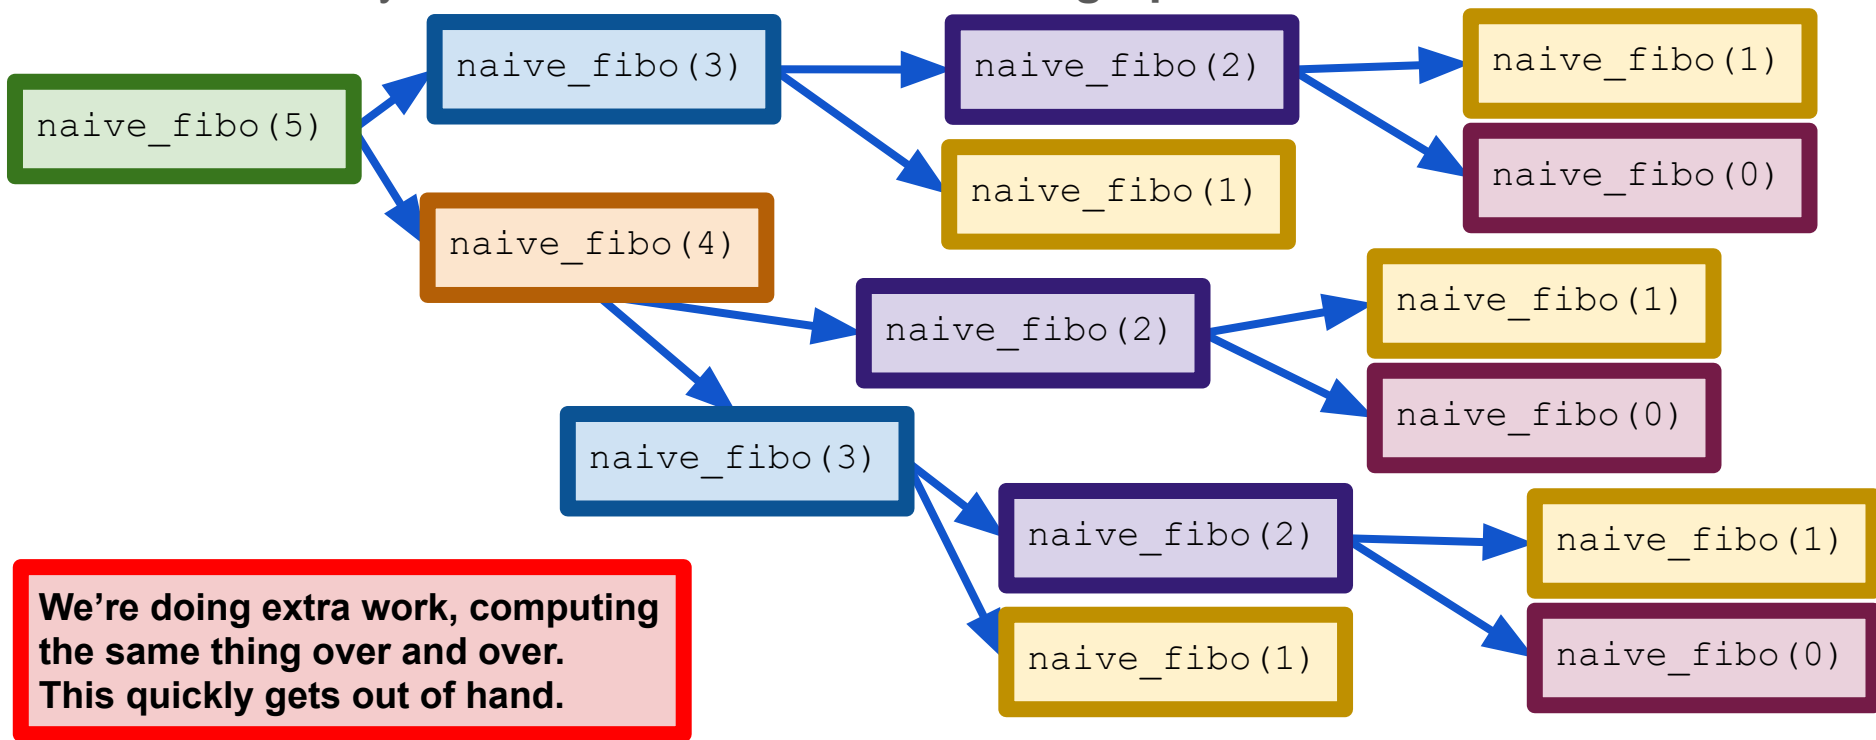
Raise an error. You'll need this in many of your future homeworks.
<https://docs.python.org/3/tutorial/errors.html#raising-exceptions>

21
34
55
89
144

This gets slow as soon as the argument gets even moderately big. **Why?**

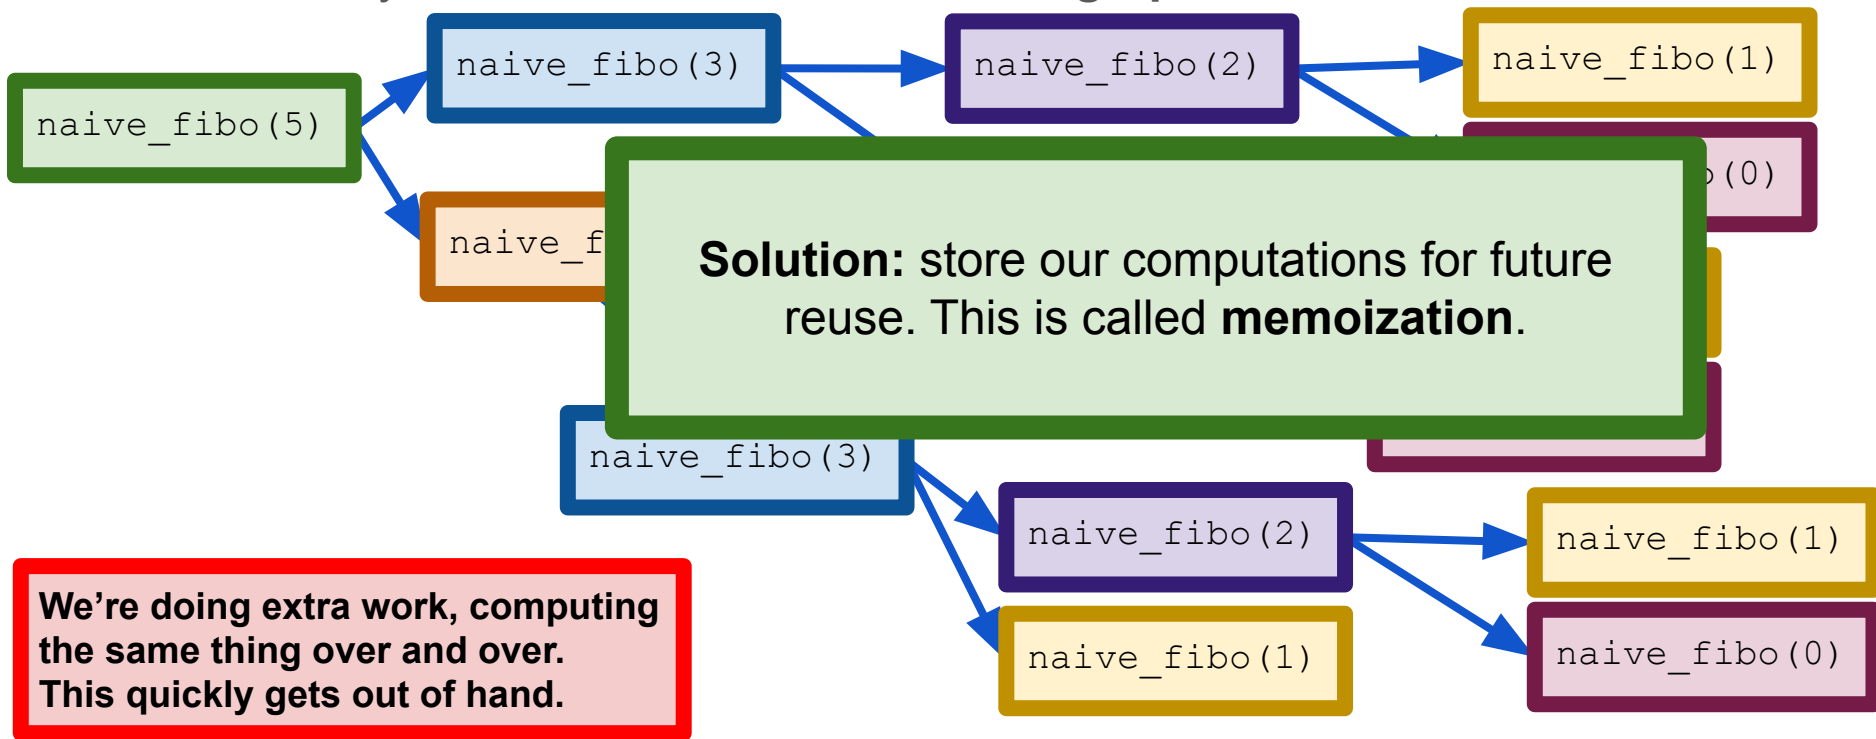
Common pattern: memoization

The inefficiency is clear when we draw the **call graph** of the function



Common pattern: memoization

The inefficiency is clear when we draw the **call graph** of the function



Common pattern: memoization

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return f
9 fibo(30)
```

This is the dictionary that we'll use for memoization. We'll store `known[n] = fibo(n)` the first time we compute `fibo(n)`, and every time we need it again, we just look it up!

832040

Common pattern: memoization

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return(f)
9 fibo(30)
```

If we already know the n-th Fibonacci number, there's no need to compute it again. Just look it up!

832040

Common pattern: memoization

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return f
9 fibo(30)
```

If we don't already know it, we have to compute it, but before we return the result, we memoize it in `known` for future reuse.

832040

Common pattern: memoization

```
1 import time
2 start_time = time.time()
3 naive_fibo(30)
4 time.time() - start_time
```

0.8452379703521729

```
1 start_time = time.time()
2 fibo(30)
3 time.time() - start_time
```

0.00015687942504882812

```
1 fibo(100)
```

354224848179261915075

The time difference is enormous!

Note: this was done with `known` set to its initial state, so this is a fair comparison.

If you try to do this with `naive_fibo`, you'll be waiting for quite a bit!

```
1 fibo(100)
```

```
354224848179261915075
```

Our memoized Fibonacci function can compute some truly huge numbers!

```
1 fibo(1000)
```

```
434665576869374564356885276750406258025646605173717804024817290895  
347752096896232398733224711616429964409065331879382989696499285160
```

I cropped this huge number for readability.

```
1 fibo(10000)
```

```
-----  
RecursionError                                Traceback (most recent call last)  
<ipython-input-116-94cf333eb78c> in <module>()  
----> 1 fibo(10000)
```

I cropped some of the error message for readability.

```
----> 6         f = fibo(n-1) + fibo(n-2)  
      7         known[n] = f  
      8         return(f)
```

```
RecursionError: maximum recursion depth exceeded
```

Python runs out of levels of recursion. You can change this maximum recursion depth, but it can introduce instability:
<https://docs.python.org/3.5/library/sys.html#sys.setrecursionlimit>

1 f

35422

1 f

43466

34775

1 f

Recur

<ipy

----->

I cropp

----->

Recur

LEONARDO DICAPRIO

WATANABE GORDON-LEVITT COTILLARD PAGE HARDY MURPHY BERENGER CCAINE

YOUR MIND IS THE SCENE OF THE CRIME.

INCEPTION

FROM THE DIRECTOR OF THE DARK KNIGHT

Warner Bros. Pictures presents... IN ASSOCIATION WITH LEGENDARY PICTURES... A SYNOPSIS PRODUCTION... A FILM BY CHRISTOPHER NOLAN... LEONARDO DICAPRIO "INCEPTION" KEN WATANABE JOSEPH GORDON-LEVITT MARION COTILLARD ELLEN PAGE TOM HARDY CILLIAN MURPHY TOM BERENGER AND MICHAEL CAINE... PRODUCED BY CLOY HENDRIX LOVAS... WRITTEN BY WALLY PISTER... DIRECTED BY CHRISTOPHER NOLAN... inceptiomovie.com

EXPERIENCE IT JULY 16 IN THEATERS AND IMAX

Fibonacci function can only handle huge numbers!

605173717804024817290895
331879382989696499285160

I cropped this huge number for readability.

Traceback (most recent call last)
File "inception.py", line 1, in
e>()

File "inception.py", line 1, in
e>()

Python runs out of levels of recursion. You can change this maximum recursion depth, but it can introduce instability: <https://docs.python.org/3.5/library/sys.html#sys.setrecursionlimit>

RecursionExceeded

Common pattern: memoization

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return(f)
9 fibo(30)
```

832040



Congratulations! You've seen your first example of **dynamic programming!** Lots of popular interview questions fall under this purview.
E.g., https://en.wikipedia.org/wiki/Tower_of_Hanoi

Common pattern: memoization

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return(f)
9 fibo(30)
```

832040

Note: the dictionary `known` is declared **outside** the function `fibo`. There is a good reason for this: we don't want `known` to disappear when we finish running `fibo`! We say that `known` is a **global variable**, because it is defined in the "main" program.

Name Spaces

A **name space** (or **namespace**) is a context in which code is executed

The “outermost” namespace (also called a **frame**) is called `__main__`

Running from the command line or in Jupyter? You’re in `__main__`

Often shows up in error messages, something like,

```
“Error ... in __main__: blah blah blah”
```

Variables defined in `__main__` are said to be **global**

Function definitions create their own **local** namespaces

Variables defined in such a context are called **local**

Local variables cannot be accessed from outside their frame/namespace

Similar behavior inside for-loops, while-loops, etc

Name Spaces

Example: we have a program simulating a light bulb

Bulb state is represented by a global Boolean variable, `lightbulb_on`

```
1 lightbulb_on = False
2 def lights_on():
3     lightbulb_on = True
4 lights_on()
5 lightbulb_on
```

False

Bulb is initially off.

Calling this function sets the bulb to the "on" state.

But after calling `lights_on`, the state variable is still `False`. What's going on?

Name Spaces

```
1 lightbulb_on = False
2 def flip_switch():
3     lightbulb_on = not lightbulb_on
4 flip_switch()
```

The fact that this code causes an error shows what is really at issue. By default, Python treats the variable `lightbulb_on` inside the function definition as being a **different** variable from the `lightbulb_on` defined in the main namespace. This is, generally, a good design. It prevents accidentally changing global state information.

UnboundLocalError

Traceback (most recent call last)

<ipython-input-125-b39d1f83dc2a> in <module>()

2 def flip_switch():

3 lightbulb_on = not lightbulb_on

----> 4 flip_switch()

<ipython-input-125-b39d1f83dc2a> in flip_switch()

1 lightbulb_on = False

2 def flip_switch():

----> 3 lightbulb_on = not lightbulb_on

4 flip_switch()

UnboundLocalError: local variable 'lightbulb_on' referenced before assignment

Name Spaces

We have to tell Python that we want `lightbulb_on` to mean the *global* variable

```
1 lightbulb_on = False
2 def flip_switch():
3     global lightbulb_on
4     lightbulb_on = not lightbulb_on
5 flip_switch()
6 lightbulb_on
```

True

```
1 flip_switch()
2 lightbulb_on
```

False

Tell Python that we want `lightbulb_on` to refer to the global variable of the same name.

Now, when we call `flip_switch`, the value of `lightbulb_on` is changed successfully.

Warning: this is all well and good, but it is considered best practice to avoid global variables in large programs, as they can make debugging hard. This isn't so crucial for our course, since we won't be building anything especially large, but you should be aware of it.

Important note

Why is this okay, if `known` isn't declared `global`?

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return(f)
9 fibo(30)
```

832040

Important note

Why is this okay, if `known` isn't declared `global`?

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return(f)
9 fibo(30)
```

`known` is a dictionary, and thus mutable. Maybe mutable variables have special powers and don't have to be declared as `global`?

832040

Important note

Why is this okay, if `known` isn't declared `global`?

```
1 known = {0:0, 1:1}
2 def fibo(n):
3     if n in known:
4         return known[n]
5     else:
6         f = fibo(n-1) + fibo(n-2)
7         known[n] = f
8         return(f)
9 fibo(30)
```

832040

`known` is a dictionary, and thus mutable. Maybe mutable variables have special powers and don't have to be declared as `global`?

Correct answer: global vs local distinction is only important for **variable assignment**. We aren't performing any variable assignment in `fibo`, so no need for the global declaration. Contrast with `lights_on`, where we were reassigning `lightbulb_on`. Variable assignment is **local** by default.

Tuples

Similar to a list, in that it is a sequence of values

But unlike lists, tuples are immutable

Because they are immutable, they are hashable

So we can use tuples where we wanted to key on a list

Documentation:

<https://docs.python.org/3/tutorial/datastructures.html#tuples-and-sequences>

<https://docs.python.org/3/library/stdtypes.html#tuples>

Creating Tuples

```
1 t = 1,2,3,4,5
2 t
(1, 2, 3, 4, 5)
```

Tuples created with “comma notation”.
Parentheses are optional, but I recommend them for ease of reading.

```
1 t = (1,2,3,4,5)
2 t
(1, 2, 3, 4, 5)
```

Python always displays tuples with parentheses.

```
1 t = 'cat',
2 t
('cat',)
```

Creating a tuple of one element requires a trailing comma. Failure to include this comma, even with parentheses, yields... not a tuple.

```
1 t = ('cat')
2 t
'cat'
```

Creating Tuples

```
1 t1 = tuple()  
2 t1
```

```
()
```

```
1 t2 = tuple(range(5))  
2 t2
```

```
(0, 1, 2, 3, 4)
```

```
1 t3 = tuple('goat')  
2 t3
```

```
('g', 'o', 'a', 't')
```

```
1 tuple([[1,2,3],[4,5,6]])
```

```
([1, 2, 3], [4, 5, 6])
```

```
1 print(type(t2))
```

```
<class 'tuple'>
```

Can also create a tuple using the `tuple()` function, which will cast any sequence to a tuple whose elements are those of the sequence.

Tuples are Sequences

```
1 t = ('a', 'b', 'c', 'd', 'e')
2 t[0]
```

'a'

As sequences, tuples support indexing, slices, etc.

```
1 t[1:4]
```

('b', 'c', 'd')

```
1 t[-1]
```

'e'

And of course, sequences have a length.

```
1 len(t)
```

5

Reminder: sequences support all the operations listed here:
<https://docs.python.org/3.3/library/stdtypes.html#typeseq>

Tuple Comparison

Tuples support comparison, which works analogously to string ordering.

```
1 (1,2,3) < (2,2,3)
```

True

```
1 (2,2,20) <= (2,2,2)
```

False

```
1 (1,2,3) < (1,2,3,4)
```

True

```
1 ('cat', 'dog', 'goat') > ('dog', 'cat', 'goat')
```

False

```
1 (1, 'cat', (1,2,3)) > (0, 'bird', (1,2,0))
```

True

0-th elements are compared. If they are equal, go to the 1-th element, etc.

Just like strings, the “prefix” tuple is ordered first.

Tuple comparison is element-wise, so we only need that each element-wise comparison is allowed by Python.

Tuples are Immutable

```
1 fruits = ('apple', 'banana', 'orange', 'kiwi')
2 fruits[2] = 'grapefruit'
```

Tuples are immutable, so changing an entry is not permitted.

```
-----
TypeError                                 Traceback (most recent call last)
<ipython-input-48-c40a1905a6e9> in <module>()
      1 fruits = ('apple', 'banana', 'orange', 'kiwi')
----> 2 fruits[2] = 'grapefruit'
```

TypeError: 'tuple' object does not support item assignment

As with strings, have to make a new assignment to the variable.

```
1 fruits = fruits[0:2] + ('grapefruit',) + fruits[3:]
2 fruits
```

```
('apple', 'banana', 'grapefruit', 'kiwi')
```

```
1 fruits = fruits[0:2] + 'grapefruit', + fruits[3:]
```

Note: even though 'grapefruit', is a tuple, Python doesn't know how to parse this line. Use parentheses!

```
-----
TypeError                                 Traceback (most recent call last)
<ipython-input-50-f62749483e65> in <module>()
----> 1 fruits = fruits[0:2] + 'grapefruit', + fruits[3:]
```

TypeError: can only concatenate tuple (not "str") to tuple

Useful trick: tuple assignment

```
1 a = 10
2 b = 5
3 print(a, b)
```

10 5

Tuples in Python allow us to make many variable assignments at once. Useful tricks like this are sometimes called **syntactic sugar**.

https://en.wikipedia.org/wiki/Syntactic_sugar

```
tmp = a
a = b
b = tmp
print(a, b)
```

5 10

Common pattern: swap the values of two variables.

```
1 a = 10
2 b = 5
3 (a,b) = (b,a)
4 print(a, b)
```

5 10

This line achieves the same end, but in a single assignment statement instead of three, and without the extra variable `tmp`.

Useful trick: tuple assignment

```
1 (x,y,z) = (2*'cat', 0.57721, [1,2,3])
```

```
2 (x,y,z)
```

```
('catcat', 0.57721, [1, 2, 3])
```

← Tuple assignment requires one variable on the left for each expression on the right.

```
1 (x,y,z) = ('a','b','c','d')
```

ValueError

Traceback (most recent call last)

<ipython-input-68-e118c50f83dd> in <module>()
----> 1 (x,y,z) = ('a','b','c','d')

ValueError: too many values to unpack (expected 3)

If the number of variables doesn't match the number of expressions, that's an error.

```
1 (x,y,z) = ('a','b')
```

ValueError

Traceback (most recent call last)

<ipython-input-69-875f95cea434> in <module>()
----> 1 (x,y,z) = ('a','b')

ValueError: not enough values to unpack (expected 3, got 2)

Useful trick: tuple assignment

```
1 email = 'klevin@umich.edu'  
2 email.split('@')
```

```
['klevin', 'umich.edu']
```

The `string.split()` method returns a list of strings, obtained by splitting the calling string on the characters in its argument.

```
1 (user, domain) = email.split('@')  
2 user
```

```
'klevin'
```

Tuple assignment works so long as the right-hand side is **any** sequence, provided the number of variables matches the number of elements on the right. Here, the right-hand side is a list, `['klevin', 'umich.edu']`.

```
1 domain
```

```
'umich.edu'
```

```
1 (x,y,z) = 'cat'  
2 print(x, y, z)
```

A string is a sequence, so tuple assignment is allowed. Sequence elements are characters, and indeed, `x`, `y` and `z` are assigned to the three characters in the string.

```
c a t
```

Tuples as Return Values

```
1 import random
2 def five_numbers(t):
3     t.sort()
4     n = len(t)
5     return (t[0], t[n//4], t[n//2], t[(3*n)//4], t[-1])
6 five_numbers([1,2,3,4,5,6,7])
```

This function takes a list of numbers and returns a tuple summarizing the list.
https://en.wikipedia.org/wiki/Five-number_summary

(1, 2, 4, 6, 7)

```
1 randnumlist = [random.randint(1,100) for x in range(60)]
2 (mini,lowq,med,upq,maxi) = five_numbers(randnumlist)
3 (mini,lowq,med,upq,maxi)
```

(3, 27, 54, 73, 98)

Test your understanding: what does this list comprehension do?

Tuples as Return Values

More generally, sometimes you want more than one return value

```
1 t = divmod(13,4)
2 t
```

(3, 1)

```
1 (quotient,remainder) = divmod(13,4)
2 quotient
```

3

```
1 remainder
```

1

`divmod` is a Python built-in function that takes a pair of numbers and outputs the quotient and remainder, as a tuple. Additional examples can be found here: <https://docs.python.org/3/library/functions.html>

Useful trick: variable-length arguments

```
1 def my_min( *args ):  
2     return min(args)  
3 my_min(1,2,3)
```

1

```
1 my_min(4,5,6,10)
```

4

```
1 def print_all( *args ):  
2     print(args)  
3 print_all('cat', 'dog', 'bird')
```

```
('cat', 'dog', 'bird')
```

```
1 print_all()
```

()

A parameter name prefaced with * **gathers** all arguments supplied to the function into a tuple. This is also called **argument packing**.

Note: this is also one of several ways that one can implement optional arguments, though we'll see better ways later in the course.

Gather and Scatter

The opposite of the gather operation is **scatter**

```
1 t = (13, 4)
2 divmod(t)
```

divmod takes two arguments, so this is an error.

```
-----
TypeError                                 Traceback (most recent call last)
<ipython-input-106-c7c0a10eef7e> in <module>()
      1 t = (13, 4)
----> 2 divmod(t)
```

```
TypeError: divmod expected 2 arguments, got 1
```

```
1 divmod(*t)
(3, 1)
```

Instead, we have to “untuple” the tuple, using the **scatter** operation. This makes the elements of the tuple into the arguments of the function.

```
1 *t
```

```
File "<ipython-input-109-f9912a2ca07d>", line 1
```

```
*t
  ^
```

```
SyntaxError: can't use starred expression here
```

Note: gather/scatter only works in certain contexts (e.g., for function arguments).

Combining lists: `zip`

Python includes a number of useful functions for combining lists and tuples

```
1 t1 = ['apple', 'orange', 'banana', 'kiwi']
2 t2 = [1, 2, 3, 4]
3 zip(t1,t2)
```

```
<zip at 0x10c95d5c8>
```

`zip()` returns a zip object, which is an **iterator** containing as its elements tuples formed from its arguments.
<https://docs.python.org/3/library/functions.html#zip>

```
1 for tup in zip(t1,t2):
2     print(tup)
```

```
('apple', 1)
('orange', 2)
('banana', 3)
('kiwi', 4)
```

Iterators are, in essence, objects that support for-loops. All sequences are iterators. Iterators support, crucially, a method `__next__()`, which returns the “next element”. We’ll see this in more detail later in the course.
<https://docs.python.org/3/library/stdtypes.html#iterator-types>

Combining lists: `zip`

`zip()` returns a zip object, which is an **iterator** containing as its elements tuples formed from its arguments.

<https://docs.python.org/3/library/functions.html#zip>

```
1 for tup in zip(['a','b','c'],[1,2,3,4]):  
2     print(tup)
```

```
('a', 1)  
( 'b', 2)  
( 'c', 3)
```

Given arguments of different lengths, `zip` defaults to the shortest one.

```
1 for tup in zip(['a','b','c','d'],[1,2,3]):  
2     print(tup)
```

```
('a', 1)  
( 'b', 2)  
( 'c', 3)
```

`zip` takes any number of arguments, so long as they are all **iterable**. Sequences are iterable.

```
1 for tup in zip([1,2,3],['a','b','c'],'xyz'):  
2     print(tup)
```

```
(1, 'a', 'x')  
(2, 'b', 'y')  
(3, 'c', 'z')
```


Iterables are, essentially, objects that can *become* iterators. We'll see the distinction later in the course.

<https://docs.python.org/3/library/stdtypes.html#typeiter>

Combining lists: zip

```
1 def count_matches(s, t):
2     cnt = 0
3     for (a,b) in zip(s,t):
4         if a==b:
5             cnt += 1
6     return( cnt )
7 count_matches([1,1,2,3,5],[1,2,3,4,5])
```

zip is especially useful for iterating over several lists in lockstep.



2

```
1 count_matches([1,2,3,4,5],[1,2,3])
```

Test your understanding: what should this return?

Combining lists: zip

```
1 def count_matches(s, t):
2     cnt = 0
3     for (a,b) in zip(s,t):
4         if a==b:
5             cnt += 1
6     return( cnt )
7 count_matches([1,1,2,3,5],[1,2,3,4,5])
```

zip is especially useful for iterating over several lists in lockstep.

2

```
1 count_matches([1,2,3,4,5],[1,2,3])
```

3

Test your understanding: what should this return?

Related function: `enumerate()`

```
1 for t in enumerate('goat'):  
2     print(t)
```

```
(0, 'g')  
(1, 'o')  
(2, 'a')  
(3, 't')
```

```
1 s = 'goat'  
2 for i in range(len(s)):  
3     print((i,s[i]))
```

```
(0, 'g')  
(1, 'o')  
(2, 'a')  
(3, 't')
```

`enumerate` returns an **enumerate object**, which is an iterator of (index,element) pairs. It is a more graceful way of performing the pattern below, which we've seen before.
<https://docs.python.org/3/library/functions.html#enumerate>

Dictionaries revisited

```
1 hist = {'cat':3,'dog':12,'goat':18}
2 hist.items()
```

```
dict_items([('cat', 3), ('dog', 12), ('goat', 18)])
```

```
1 for (k,v) in hist.items():
2     print(k, ':', v)
```

```
cat : 3
dog : 12
goat : 18
```

`dict.items()` returns a `dict_items` object, an iterator whose elements are (key,value) tuples.

Dictionaries revisited

```
1 hist = {'cat':3,'dog':12,'goat':18}
2 hist.items()
```

```
dict_items([('cat', 3), ('dog', 12), ('goat', 18)])
```

```
1 for (k,v) in hist.items():
2     print(k, ':', v)
```

```
cat : 3
dog : 12
goat : 18
```

`dict.items()` returns a `dict_items` object, an iterator whose elements are (key,value) tuples.

```
1 d = dict([(0,'zero'),(1,'one'),(2,'two')])
2 d
```

```
{0: 'zero', 1: 'one', 2: 'two'}
```

Conversely, we can create a dictionary by supplying a list of (key,value) tuples.

```
1 dict(zip('cat','dog'))
```

```
{'a': 'o', 'c': 'd', 't': 'g'}
```

Tuples as Keys

```
1 name2umid = {('Einstein', 'Albert'): 'aeinstein',  
2   ('Noether', 'Emmy'): 'enoether',  
3   ('Shannon', 'Claude'): 'cshannon',  
4   ('Fan', 'Ky'): 'kyfan'}  
5 name2umid
```

```
{('Einstein', 'Albert'): 'aeinstein',  
 ('Fan', 'Ky'): 'kyfan',  
 ('Noether', 'Emmy'): 'enoether',  
 ('Shannon', 'Claude'): 'cshannon'}
```

In (most) Western countries, the family name is said last (hence “last name”), but it is frequently useful to key on this name before keying on a given name.

```
1 name2umid[('Einstein', 'Albert')]
```

```
'aeinstein'
```

```
1 sparsemx = dict()  
2 sparsemx[(1,4)] = 1  
3 sparsemx[(3,5)] = 1  
4 sparsemx[(12,13)] = 2  
5 sparsemx[(11,13)] = 3  
6 sparsemx[(19,13)] = 5  
7 sparsemx
```

Keying on tuples is especially useful for representing sparse structures. Consider a 20-by-20 matrix in which most entries are zeros. Storing all the entries requires 400 numbers, but if we only record the entries that are nonzero...

```
{(1, 4): 1, (3, 5): 1, (11, 13): 3, (12, 13): 2, (19, 13): 5}
```

Data Structures: Lists vs Tuples

Use a **list** when:

- Length is not known ahead of time and/or may change during execution
- Frequent updates are likely

Use a **tuple** when:

- The set is unlikely to change during execution
- Need to key on the set (i.e., require immutability)
- Want to perform multiple assignment or use in variable-length argument list

Most code you see will use lists, because mutability is quite useful