

# Python Basics

## Lecture 04

Practical Cloud Computing

Fall 2023

Thompson Rivers University

Python is a high-level, dynamically typed multiparadigm programming language.

Python code is often said to be almost like pseudocode, since it allows you to express very powerful ideas in very few lines of code while being very readable.

# Python

- Python is suitable for providing a solid foundation to the reader in the area of cloud computing.
- The main characteristics of Python are:
  - Multi-paradigm programming language
    - Python supports more than one programming paradigms including object-oriented programming and structured programming
  - Interpreted Language
    - Python is an interpreted language and does not require an explicit compilation step. The Python interpreter executes the program source code directly, statement by statement, as a processor or scripting engine does.
- Interactive Language
  - Python provides an interactive mode in which the user can submit commands at the Python prompt and interact with the interpreter directly.

# Python - Benefits

- Easy-to-learn, read and maintain
  - Python is a minimalistic language with relatively few keywords, uses English keywords and has fewer syntactical constructions as compared to other languages. Reading Python programs feels like English with pseudo-code like constructs. Python is easy to learn yet an extremely powerful language for a wide range of applications.
- Object and Procedure Oriented
  - Python supports both procedure-oriented programming and object-oriented programming. Procedure oriented paradigm allows programs to be written around procedures or functions that allow reuse of code. Procedure oriented paradigm allows programs to be written around objects that include both data and functionality.
- Extendable
  - Python is an extendable language and allows integration of low-level modules written in languages such as C/C++. This is useful when you want to speed up a critical portion of a program.

## Python - Benefits (cont.)

- Scalable
  - Due to the minimalistic nature of Python, it provides a manageable structure for large programs.
- Portable
  - Since Python is an interpreted language, programmers do not have to worry about compilation, linking and loading of programs. Python programs can be directly executed from source
- Broad Library Support
  - Python has a broad library support and works on various platforms such as Windows, Linux, Mac, etc.

As of January 1, 2020, Python has officially dropped support for python2.

For this class all code will use Python 3.7.

Ensure you have gone through the setup instructions and correctly installed a python3 virtual environment before proceeding with this labs.



You can double-check your Python version at the command line after activating your environment by running `python --version`.

# Basic data types

## **Numbers**

Integers and floats work  
as you would expect from other languages

```
x = 3
print(type(x)) # Prints "<class 'int'>"
print(x)      # Prints "3"
print(x + 1)  # Addition; prints "4"
print(x - 1)  # Subtraction; prints "2"
print(x * 2)  # Multiplication; prints "6"
print(x ** 2) # Exponentiation; prints "9"
x += 1
print(x)     # Prints "4"
x *= 2
print(x)     # Prints "8"
y = 2.5
print(type(y)) # Prints "<class 'float'>"
print(y, y + 1, y * 2, y ** 2) # Prints "2.5 3.5 5.0 6.25"
```

Note that unlike many languages,  
Python does not have  
unary increment (x++) or decrement (x--) operators.

Python also has built-in types for complex numbers.

You can find all of the details in the documentation:

[Built-in Types — Python 3.11.5 documentation](#)

## Booleans

Python implements all of the usual operators for Boolean logic, but uses English words rather than symbols (&&, ||, etc.)

```
t = True
f = False
print(type(t)) # Prints "<class 'bool'>"
print(t and f) # Logical AND; prints "False"
print(t or f)  # Logical OR; prints "True"
print(not t)   # Logical NOT; prints "False"
print(t != f)  # Logical XOR; prints "True"
```



# Strings

Python has great support for strings.

```
hello = 'hello'    # String literals can use single quotes
world = "world"    # or double quotes; it does not matter.
print(hello)       # Prints "hello"
print(len(hello))  # String length; prints "5"
hw = hello + ' ' + world # String concatenation
print(hw)          # prints "hello world"
hw12 = '%s %s %d' % (hello, world, 12) # sprintf style string formatting
print(hw12)        # prints "hello world 12"
```

String objects have a bunch of useful methods.

```
s = "hello"
print(s.capitalize()) # Capitalize a string; prints "Hello"
print(s.upper())      # Convert a string to uppercase; prints "HELLO"
print(s.rjust(7))     # Right-justify a string, padding with spaces; prints "  hello"
print(s.center(7))    # Center a string, padding with spaces; prints "  hello  "
print(s.replace('l', '(ell)')) # Replace all instances of one substring with another;
                                # prints "he(ell)(ell)o"
print('  world '.strip()) # Strip leading and trailing whitespace; prints "world"
```

You can find a list of all string methods in the documentation:

[Built-in Types — Python 3.11.5 documentation](#)

## Containers

Python includes several built-in container types:  
lists, dictionaries, sets, and tuples.

## **Lists**

A list is the Python equivalent of an array, but is resizable and can contain elements of different types.

```
xs = [3, 1, 2]      # Create a list
print(xs, xs[2])   # Prints "[3, 1, 2] 2"
print(xs[-1])     # Negative indices count from the end of the list; prints "2"
xs[2] = 'foo'     # Lists can contain elements of different types
print(xs)         # Prints "[3, 1, 'foo']"
xs.append('bar')  # Add a new element to the end of the list
print(xs)        # Prints "[3, 1, 'foo', 'bar']"
x = xs.pop()     # Remove and return the last element of the list
print(x, xs)     # Prints "bar [3, 1, 'foo']"
```



As usual, you can find all the gory details about lists in the documentation.

[Data Structures \(List\) — Python 3.11.5 documentation](#)

## Slicing

In addition to accessing list elements one at a time, Python provides concise syntax to access sublists; this is known as slicing.

```
nums = list(range(5))      # range is a built-in function that creates a list of
integers
print(nums)               # Prints "[0, 1, 2, 3, 4]"
print(nums[2:4])          # Get a slice from index 2 to 4 (exclusive); prints "[2, 3]"
print(nums[2:])           # Get a slice from index 2 to the end; prints "[2, 3, 4]"
print(nums[:2])           # Get a slice from the start to index 2 (exclusive); prints
"[0, 1]"
print(nums[:])            # Get a slice of the whole list; prints "[0, 1, 2, 3, 4]"
print(nums[:-1])          # Slice indices can be negative; prints "[0, 1, 2, 3]"
nums[2:4] = [8, 9]        # Assign a new sublist to a slice
print(nums)               # Prints "[0, 1, 8, 9, 4]"
```

We will see slicing again in the context of numpy arrays.

# Loops

```
animals = ['cat', 'dog', 'monkey']  
for animal in animals:  
    print(animal)  
# Prints "cat", "dog", "monkey", each on its own line.
```

You can loop over the elements of a list like this.

```
animals = ['cat', 'dog', 'monkey']
for idx, animal in enumerate(animals):
    print('#%d: %s' % (idx + 1, animal))
# Prints "#1: cat", "#2: dog", "#3: monkey", each on its own line
```

If you want access to the index of each element within the body of a loop, use the built-in enumerate function.

## **List comprehensions**

When programming, frequently we want to transform one type of data into another.



```
nums = [0, 1, 2, 3, 4]
squares = []
for x in nums:
    squares.append(x ** 2)
print(squares)    # Prints [0, 1, 4, 9, 16]
```

As a simple example, consider the following code that computes square numbers

```
nums = [0, 1, 2, 3, 4]
squares = [x ** 2 for x in nums]
print(squares)    # Prints [0, 1, 4, 9, 16]
```

You can make this code simpler using a list comprehension.

```
nums = [0, 1, 2, 3, 4]
even_squares = [x ** 2 for x in nums if x % 2 == 0]
print(even_squares) # Prints "[0, 4, 16]"
```

List comprehensions can also contain conditions.

## Dictionaries

A dictionary stores (key, value) pairs,  
similar to a Map in Java or an object in Javascript.

```
d = {'cat': 'cute', 'dog': 'furry'} # Create a new dictionary with some data
print(d['cat']) # Get an entry from a dictionary; prints "cute"
print('cat' in d) # Check if a dictionary has a given key; prints "True"
d['fish'] = 'wet' # Set an entry in a dictionary
print(d['fish']) # Prints "wet"
# print(d['monkey']) # KeyError: 'monkey' not a key of d
print(d.get('monkey', 'N/A')) # Get an element with a default; prints "N/A"
print(d.get('fish', 'N/A')) # Get an element with a default; prints "wet"
del d['fish'] # Remove an element from a dictionary
print(d.get('fish', 'N/A')) # "fish" is no longer a key; prints "N/A"
```

You can find all you need to know about dictionaries in the documentation.

[Built-in Types \(dict\) — Python 3.11.5 documentation](#)

Loops: It is easy to iterate over the keys in a dictionary.

```
d = {'person': 2, 'cat': 4, 'spider': 8}
for animal in d:
    legs = d[animal]
    print('A %s has %d legs' % (animal, legs))
# Prints "A person has 2 legs", "A cat has 4 legs", "A spider has 8 legs"
```

Loops: It is easy to iterate over the keys in a dictionary.



```
d = {'person': 2, 'cat': 4, 'spider': 8}
for animal, legs in d.items():
    print('A %s has %d legs' % (animal, legs))
# Prints "A person has 2 legs", "A cat has 4 legs", "A spider has 8 legs"
```

If you want access to keys and their corresponding values, use the items method.

```
nums = [0, 1, 2, 3, 4]
even_num_to_square = {x: x ** 2 for x in nums if x % 2 == 0}
print(even_num_to_square) # Prints "{0: 0, 2: 4, 4: 16}"
```

Dictionary comprehensions are similar to list comprehensions, but allow you to easily construct dictionaries.

## **Sets**

A set is an unordered collection of distinct elements.

```
animals = {'cat', 'dog'}
print('cat' in animals)    # Check if an element is in a set; prints "True"
print('fish' in animals)   # prints "False"
animals.add('fish')        # Add an element to a set
print('fish' in animals)   # Prints "True"
print(len(animals))        # Number of elements in a set; prints "3"
animals.add('cat')         # Adding an element that is already in the set does nothing
print(len(animals))        # Prints "3"
animals.remove('cat')      # Remove an element from a set
print(len(animals))        # Prints "2"
```

As usual, everything you want to know about sets  
can be found in the documentation.

[Built-in Types \(set\) — Python 3.11.5 documentation](#)

```
animals = {'cat', 'dog', 'fish'}
for idx, animal in enumerate(animals):
    print('#%d: %s' % (idx + 1, animal))
# Prints "#1: fish", "#2: dog", "#3: cat"
```

Iterating over a set has the same syntax as iterating over a list.  
However since sets are unordered, you cannot make assumptions about the order in which you visit the elements of the set.

```
from math import sqrt
nums = {int(sqrt(x)) for x in range(30)}
print(nums) # Prints "{0, 1, 2, 3, 4, 5}"
```

Set comprehensions: Like lists and dictionaries, we can easily construct sets using set comprehensions.

## Tuples

A tuple is an (immutable) ordered list of values.

A tuple is in many ways similar to a list; one of the most important differences is that tuples can be used as keys in dictionaries and as elements of sets, while lists cannot.



```
d = {(x, x + 1): x for x in range(10)} # Create a dictionary with tuple keys
t = (5, 6) # Create a tuple
print(type(t)) # Prints "<class 'tuple'>"
print(d[t]) # Prints "5"
print(d[(1, 2)]) # Prints "1"
```

The documentation has more information about tuples.

[Data Structures \(tuples\) — Python 3.11.5 documentation](#)

## Functions

Python functions are defined using the **def** keyword.

```
def sign(x):  
    if x > 0:  
        return 'positive'  
    elif x < 0:  
        return 'negative'  
    else:  
        return 'zero'  
  
for x in [-1, 0, 1]:  
    print(sign(x))  
  
# Prints "negative", "zero", "positive"
```

```
def hello(name, loud=False):  
    if loud:  
        print('HELLO, %s!' % name.upper())  
    else:  
        print('Hello, %s' % name)
```

```
hello('Bob') # Prints "Hello, Bob"
```

```
hello('Fred', loud=True) # Prints "HELLO, FRED!"
```

We will often define functions to take optional keyword arguments.

There is a lot more information about Python functions in the documentation.

[Functions — Python 3.11.5 documentation](#)

Classes

```
class Greeter(object):

    # Constructor
    def __init__(self, name):
        self.name = name # Create an instance variable

    # Instance method
    def greet(self, loud=False):
        if loud:
            print('HELLO, %s!' % self.name.upper())
        else:
            print('Hello, %s' % self.name)

g = Greeter('Fred') # Construct an instance of the Greeter class
g.greet()           # Call an instance method; prints "Hello, Fred"
g.greet(loud=True) # Call an instance method; prints "HELLO, FRED!"
```

The syntax for defining classes in Python is straightforward.



You can read a lot more about Python classes in the documentation.

[Classes — Python 3.11.5 documentation](#)

## **Numpy**

[Numpy](#) is the core library for scientific computing in Python.

It provides a high-performance multidimensional array object, and tools for working with these arrays.

## Arrays

A numpy array is a grid of values,  
all of the same type,  
and is indexed by a tuple of nonnegative integers.

The number of dimensions is the rank of the array;  
the shape of an array is a tuple of integers  
giving the size of the array along each dimension.

```
import numpy as np

a = np.array([1, 2, 3])    # Create a rank 1 array
print(type(a))           # Prints "<class 'numpy.ndarray'"
print(a.shape)           # Prints "(3,)"
print(a[0], a[1], a[2])  # Prints "1 2 3"
a[0] = 5                  # Change an element of the array
print(a)                  # Prints "[5, 2, 3]"

b = np.array([[1,2,3],[4,5,6]]) # Create a rank 2 array
print(b.shape)           # Prints "(2, 3)"
print(b[0, 0], b[0, 1], b[1, 0]) # Prints "1 2 4"
```

We can initialize numpy arrays from nested Python lists, and access elements using square brackets.

```
import numpy as np

a = np.zeros((2,2)) # Create an array of all zeros
print(a)           # Prints "[[ 0.  0.]
                  #           [ 0.  0.]]"

b = np.ones((1,2)) # Create an array of all ones
print(b)           # Prints "[[ 1.  1.]]"

c = np.full((2,2), 7) # Create a constant array
print(c)           # Prints "[[ 7.  7.]
                  #           [ 7.  7.]]"

d = np.eye(2)      # Create a 2x2 identity matrix
print(d)           # Prints "[[ 1.  0.]
                  #           [ 0.  1.]]"

e = np.random.random((2,2)) # Create an array filled with random values
print(e)           # Might print "[[ 0.91940167  0.08143941]
                  #           [ 0.68744134  0.87236687]]"
```

Numpy also provides many functions to create arrays.

You can read about other methods of array creation in the documentation.

[Array creation — NumPy v1.25 Manual](#)

## **Array indexing**

Numpy offers several ways to index into arrays.

## Slicing

Similar to Python lists, numpy arrays can be sliced.

Since arrays may be multidimensional,  
you must specify a slice for each dimension of the array.



```
import numpy as np

# Create the following rank 2 array with shape (3, 4)
# [[ 1  2  3  4]
#  [ 5  6  7  8]
#  [ 9 10 11 12]]
a = np.array([[1,2,3,4], [5,6,7,8], [9,10,11,12]])

# Use slicing to pull out the subarray consisting of the first 2 rows
# and columns 1 and 2; b is the following array of shape (2, 2):
# [[2 3]
#  [6 7]]
b = a[:2, 1:3]

# A slice of an array is a view into the same data, so modifying it
# will modify the original array.
print(a[0, 1])    # Prints "2"
b[0, 0] = 77     # b[0, 0] is the same piece of data as a[0, 1]
print(a[0, 1])    # Prints "77"
```

**Note:** A slice of an array is a view into the same data, so modifying it will modify the original array.

```
import numpy as np

# Create the following rank 2 array with shape (3, 4)
# [[ 1  2  3  4]
#  [ 5  6  7  8]
#  [ 9 10 11 12]]
a = np.array([[1,2,3,4], [5,6,7,8], [9,10,11,12]])

# Two ways of accessing the data in the middle row of the array.
# Mixing integer indexing with slices yields an array of lower rank,
# while using only slices yields an array of the same rank as the
# original array:
row_r1 = a[1, :]    # Rank 1 view of the second row of a
row_r2 = a[1:2, :]  # Rank 2 view of the second row of a
print(row_r1, row_r1.shape) # Prints "[5 6 7 8] (4,)"
print(row_r2, row_r2.shape) # Prints "[[5 6 7 8]] (1, 4)"

# We can make the same distinction when accessing columns of an array:
col_r1 = a[:, 1]
col_r2 = a[:, 1:2]
print(col_r1, col_r1.shape) # Prints "[ 2  6 10] (3,)"
print(col_r2, col_r2.shape) # Prints "[[ 2]
#                               [ 6]
#                               [10]] (3, 1)"
```

You can also mix integer indexing with slice indexing. However, doing so will yield an array of lower rank than the original array.

## Integer array indexing

When you index into numpy arrays using slicing, the resulting array view will always be a subarray of the original array.

In contrast, integer array indexing allows you to construct arbitrary arrays using the data from another array.

```
import numpy as np

a = np.array([[1,2], [3, 4], [5, 6]])

# An example of integer array indexing.
# The returned array will have shape (3,) and
print(a[[0, 1, 2], [0, 1, 0]]) # Prints "[1 4 5]"

# The above example of integer array indexing is equivalent to this:
print(np.array([a[0, 0], a[1, 1], a[2, 0]])) # Prints "[1 4 5]"

# When using integer array indexing, you can reuse the same
# element from the source array:
print(a[[0, 0], [1, 1]]) # Prints "[2 2]"

# Equivalent to the previous integer array indexing example
print(np.array([a[0, 1], a[0, 1]])) # Prints "[2 2]"
```

```
import numpy as np

# Create a new array from which we will select elements
a = np.array([[1,2,3], [4,5,6], [7,8,9], [10, 11, 12]])

print(a) # prints "array([[ 1,  2,  3],
#           [ 4,  5,  6],
#           [ 7,  8,  9],
#           [10, 11, 12]])"

# Create an array of indices
b = np.array([0, 2, 0, 1])

# Select one element from each row of a using the indices in b
print(a[np.arange(4), b]) # Prints "[ 1  6  7 11]"

# Mutate one element from each row of a using the indices in b
a[np.arange(4), b] += 10

print(a) # prints "array([[11,  2,  3],
#           [ 4,  5, 16],
#           [17,  8,  9],
#           [10, 21, 12]])"
```

One useful trick with integer array indexing is selecting or mutating one element from each row of a matrix.

## **Boolean array indexing**

Boolean array indexing lets you pick out arbitrary elements of an array.

Frequently this type of indexing is used to select the elements of an array that satisfy some condition.

```
import numpy as np

a = np.array([[1,2], [3, 4], [5, 6]])

bool_idx = (a > 2)    # Find the elements of a that are bigger than 2;
                    # this returns a numpy array of Booleans of the same
                    # shape as a, where each slot of bool_idx tells
                    # whether that element of a is > 2.

print(bool_idx)      # Prints "[[False False]
                    #           [ True  True]
                    #           [ True  True]]"

# We use boolean array indexing to construct a rank 1 array
# consisting of the elements of a corresponding to the True values
# of bool_idx
print(a[bool_idx])  # Prints "[3 4 5 6]"

# We can do all of the above in a single concise statement:
print(a[a > 2])     # Prints "[3 4 5 6]"
```

For brevity we have left out a lot of details about numpy array indexing.

If you want to know more you should read the documentation.

[Indexing routines — NumPy v1.25 Manual](#)



## Datatypes

Every numpy array is a grid of elements of the same type.

Numpy provides a large set of numeric datatypes that you can use to construct arrays.

Numpy tries to guess a datatype when you create an array, but functions that construct arrays usually also include an optional argument to explicitly specify the datatype.

```
import numpy as np

x = np.array([1, 2])    # Let numpy choose the datatype
print(x.dtype)        # Prints "int64"

x = np.array([1.0, 2.0]) # Let numpy choose the datatype
print(x.dtype)        # Prints "float64"

x = np.array([1, 2], dtype=np.int64) # Force a particular datatype
print(x.dtype)        # Prints "int64"
```

You can read all about numpy datatypes in the documentation.

[Data type objects \(dtype\) — NumPy v1.25 Manual](#)

## Array math

Basic mathematical functions operate elementwise on arrays,  
and are available both  
as operator overloads and as functions in the numpy module.

```
import numpy as np

x = np.array([[1,2],[3,4]], dtype=np.float64)
y = np.array([[5,6],[7,8]], dtype=np.float64)

# Elementwise sum; both produce the array
# [[ 6.0  8.0]
#  [10.0 12.0]]
print(x + y)
print(np.add(x, y))

# Elementwise difference; both produce the array
# [[-4.0 -4.0]
#  [-4.0 -4.0]]
print(x - y)
print(np.subtract(x, y))

# Elementwise product; both produce the array
# [[ 5.0 12.0]
#  [21.0 32.0]]
print(x * y)
print(np.multiply(x, y))

# Elementwise division; both produce the array
# [[ 0.2      0.33333333]
#  [ 0.42857143  0.5      ]]
print(x / y)
print(np.divide(x, y))

# Elementwise square root; produces the array
# [[ 1.      1.41421356]
#  [ 1.73205081  2.      ]]
print(np.sqrt(x))
```

```
import numpy as np

x = np.array([[1,2],[3,4]])
y = np.array([[5,6],[7,8]])

v = np.array([9,10])
w = np.array([11, 12])

# Inner product of vectors; both produce 219
print(v.dot(w))
print(np.dot(v, w))

# Matrix / vector product; both produce the rank 1 array [29 67]
print(x.dot(v))
print(np.dot(x, v))

# Matrix / matrix product; both produce the rank 2 array
# [[19 22]
#  [43 50]]
print(x.dot(y))
print(np.dot(x, y))
```

Note that `*` is elementwise multiplication, not matrix multiplication. We instead use the `dot` function to compute inner products of vectors, to multiply a vector by a matrix, and to multiply matrices. `dot` is available both as a function in the `numpy` module and as an instance method of array objects.

```
import numpy as np
```

```
x = np.array([[1,2],[3,4]])
```

```
print(np.sum(x)) # Compute sum of all elements; prints "10"
```

```
print(np.sum(x, axis=0)) # Compute sum of each column; prints "[4 6]"
```

```
print(np.sum(x, axis=1)) # Compute sum of each row; prints "[3 7]"
```

Numpy provides many useful functions for performing computations on arrays; one of the most useful is **sum**.

You can find the full list of mathematical functions provided by numpy in the documentation.

[Mathematical functions — NumPy v1.25 Manual](#)



Apart from computing mathematical functions using arrays, we frequently need to reshape or otherwise manipulate data in arrays. The simplest example of this type of operation is transposing a matrix; to transpose a matrix, simply use the T attribute of an array object.

```
import numpy as np
```

```
x = np.array([[1,2], [3,4]])
```

```
print(x)      # Prints "[[1 2]  
              #           [3 4]]"
```

```
print(x.T)    # Prints "[[1 3]  
              #           [2 4]]"
```

*# Note that taking the transpose of a rank 1 array does nothing:*

```
v = np.array([1,2,3])
```

```
print(v)      # Prints "[1 2 3]"
```

```
print(v.T)    # Prints "[1 2 3]"
```

Numpy provides many more functions for manipulating arrays;  
you can see the full list in the documentation.

[Array manipulation routines — NumPy v1.25 Manual](#)

## Broadcasting

Broadcasting is a powerful mechanism that allows numpy to work with arrays of different shapes when performing arithmetic operations.

Frequently we have a smaller array and a larger array, and we want to use the smaller array multiple times to perform some operation on the larger array.

```
import numpy as np

# We will add the vector v to each row of the matrix x,
# storing the result in the matrix y
x = np.array([[1,2,3], [4,5,6], [7,8,9], [10, 11, 12]])
v = np.array([1, 0, 1])
y = np.empty_like(x)    # Create an empty matrix with the same shape as x

# Add the vector v to each row of the matrix x with an explicit loop
for i in range(4):
    y[i, :] = x[i, :] + v

# Now y is the following
# [[ 2  2  4]
#  [ 5  5  7]
#  [ 8  8 10]
#  [11 11 13]]
print(y)
```

Suppose that we want to add a constant vector to each row of a matrix. We could do it like this.

This works; however when the matrix  $x$  is very large, computing an explicit loop in Python could be slow.

Note that adding the vector  $v$  to each row of the matrix  $x$  is equivalent to forming a matrix  $vv$  by stacking multiple copies of  $v$  vertically, then performing elementwise summation of  $x$  and  $vv$ .

```
import numpy as np

# We will add the vector v to each row of the matrix x,
# storing the result in the matrix y
x = np.array([[1,2,3], [4,5,6], [7,8,9], [10, 11, 12]])
v = np.array([1, 0, 1])
vv = np.tile(v, (4, 1))    # Stack 4 copies of v on top of each other
print(vv)                  # Prints "[[1 0 1]
                           #           [1 0 1]
                           #           [1 0 1]
                           #           [1 0 1]]"

y = x + vv # Add x and vv elementwise
print(y)   # Prints "[[ 2  2  4
           #           [ 5  5  7]
           #           [ 8  8 10]
           #           [11 11 13]]"
```

We could implement this approach like this.

Numpy broadcasting allows us to perform this computation without actually creating multiple copies of  $v$ .



```
import numpy as np

# We will add the vector v to each row of the matrix x,
# storing the result in the matrix y
x = np.array([[1,2,3], [4,5,6], [7,8,9], [10, 11, 12]])
v = np.array([1, 0, 1])
y = x + v # Add v to each row of x using broadcasting
print(y) # Prints "[[ 2  2  4]
          #          [ 5  5  7]
          #          [ 8  8 10]
          #          [11 11 13]]"
```

Consider this version, using broadcasting.

## Broadcasting two arrays together follows these rules:

1. If the arrays do not have the same rank, prepend the shape of the lower rank array with 1s until both shapes have the same length.
2. The two arrays are said to be compatible in a dimension if they have the same size in the dimension, or if one of the arrays has size 1 in that dimension.
3. The arrays can be broadcast together if they are compatible in all dimensions.
4. After broadcasting, each array behaves as if it had shape equal to the elementwise maximum of shapes of the two input arrays.
5. In any dimension where one array had size 1 and the other array had size greater than 1, the first array behaves as if it were copied along that dimension.

If this explanation does not make sense,  
try reading the explanation from the documentation

[Broadcasting — NumPy v1.25 Manual](#)

or this explanation

[Eric's Broadcasting Doc - SciPy wiki dump](#)

Functions that support broadcasting are known as universal functions.

You can find the list of all universal functions in the documentation.

[Universal functions \(ufunc\) — NumPy v1.25 Manual](#)

```
import numpy as np

# Compute outer product of vectors
v = np.array([1,2,3]) # v has shape (3,)
w = np.array([4,5])   # w has shape (2,)
# To compute an outer product, we first reshape v to be a column
# vector of shape (3, 1); we can then broadcast it against w to yield
# an output of shape (3, 2), which is the outer product of v and w:
# [[ 4  5]
#  [ 8 10]
#  [12 15]]
print(np.reshape(v, (3, 1)) * w)

# Add a vector to each row of a matrix
x = np.array([[1,2,3], [4,5,6]])
# x has shape (2, 3) and v has shape (3,) so they broadcast to (2, 3),
# giving the following matrix:
# [[2 4 6]
#  [5 7 9]]
print(x + v)
```

Here are some applications of broadcasting.

```
# Add a vector to each column of a matrix
# x has shape (2, 3) and w has shape (2,).
# If we transpose x then it has shape (3, 2) and can be broadcast
# against w to yield a result of shape (3, 2); transposing this result
# yields the final result of shape (2, 3) which is the matrix x with
# the vector w added to each column. Gives the following matrix:
# [[ 5  6  7]
#  [ 9 10 11]]
print((x.T + w).T)
# Another solution is to reshape w to be a column vector of shape (2, 1);
# we can then broadcast it directly against x to produce the same
# output.
print(x + np.reshape(w, (2, 1)))

# Multiply a matrix by a constant:
# x has shape (2, 3). Numpy treats scalars as arrays of shape ();
# these can be broadcast together to shape (2, 3), producing the
# following array:
# [[ 2  4  6]
#  [ 8 10 12]]
print(x * 2)
```

Here are some applications of broadcasting (continue).

Broadcasting typically makes your code more concise and faster,  
so you should strive to use it where possible.

## Numpy Documentation

This brief overview has touched on many of the important things that you need to know about numpy, but is far from complete.

Check out the numpy reference to find out much more about numpy.

[NumPy reference — NumPy v1.25 Manual](#)



## SciPy

Numpy provides a high-performance multidimensional array and basic tools to compute with and manipulate these arrays.

SciPy builds on this, and provides a large number of functions that operate on numpy arrays and are useful for different types of scientific and engineering applications.

The best way to get familiar with SciPy is to browse the documentation.  
We will highlight some parts of SciPy that you might find useful.

[SciPy API — SciPy v1.11.2 Manual](#)

## Image operations

SciPy provides some basic functions to work with images.

For example, it has functions to read images from disk into numpy arrays, to write numpy arrays to disk as images, and to resize images.

Here is a simple example that showcases these functions.

```
from scipy.misc import imread, imsave, imresize

# Read an JPEG image into a numpy array
img = imread('assets/cat.jpg')
print(img.dtype, img.shape) # Prints "uint8 (400, 248, 3)"

# We can tint the image by scaling each of the color channels
# by a different scalar constant. The image has shape (400, 248, 3);
# we multiply it by the array [1, 0.95, 0.9] of shape (3,);
# numpy broadcasting means that this leaves the red channel unchanged,
# and multiplies the green and blue channels by 0.95 and 0.9
# respectively.
img_tinted = img * [1, 0.95, 0.9]

# Resize the tinted image to be 300 by 300 pixels.
img_tinted = imresize(img_tinted, (300, 300))

# Write the tinted image back to disk
imsave('assets/cat_tinted.jpg', img_tinted)
```



Left: The original image. Right: The tinted and resized image.

Matplotlib

Matplotlib is a plotting library.

## Plotting


The most important function in matplotlib is `plot`, which allows you to plot 2D data.

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

# Compute the x and y coordinates for points on a sine curve
x = np.arange(0, 3 * np.pi, 0.1)
y = np.sin(x)

# Plot the points using matplotlib
plt.plot(x, y)
plt.show() # You must call plt.show() to make graphics appear.
```






Running this previous code produces this plot.

With just a little bit of extra work we can easily plot multiple lines at once, and add a title, legend, and axis labels.

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

# Compute the x and y coordinates for points on sine and cosine curves
x = np.arange(0, 3 * np.pi, 0.1)
y_sin = np.sin(x)
y_cos = np.cos(x)

# Plot the points using matplotlib
plt.plot(x, y_sin)
plt.plot(x, y_cos)
plt.xlabel('x axis label')
plt.ylabel('y axis label')
plt.title('Sine and Cosine')
plt.legend(['Sine', 'Cosine'])
plt.show()
```



Running this previous code produces this plot.

You can read much more about the `plot` function in the documentation.

[pyplot — Matplotlib 2.0.2 documentation](#)

## Subplots

You can plot different things in the same figure using the `subplot` function.

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


# Compute the x and y coordinates for points on sine and cosine curves
x = np.arange(0, 3 * np.pi, 0.1)
y_sin = np.sin(x)
y_cos = np.cos(x)

# Set up a subplot grid that has height 2 and width 1,
# and set the first such subplot as active.
plt.subplot(2, 1, 1)

# Make the first plot
plt.plot(x, y_sin)
plt.title('Sine')

# Set the second subplot as active, and make the second plot.
plt.subplot(2, 1, 2)
plt.plot(x, y_cos)
plt.title('Cosine')

# Show the figure.
plt.show()
```



Running this previous code produces this plot.



You can read much more about the `subplot` function in the documentation.

[pyplot — Matplotlib 2.0.2 documentation](#)

Check out the Matplotlib reference to find out more.

[Matplotlib](#)

## Images

You can use the **imshow** function to show images.

```
import numpy as np
from scipy.misc import imread, imresize
import matplotlib.pyplot as plt
```

```
img = imread('assets/cat.jpg')
img_tinted = img * [1, 0.95, 0.9]
```

```
# Show the original image
```

```
plt.subplot(1, 2, 1)
plt.imshow(img)
```

```
# Show the tinted image
```

```
plt.subplot(1, 2, 2)
```

```
# A slight gotcha with imshow is that it might give strange results
# if presented with data that is not uint8. To work around this, we
# explicitly cast the image to uint8 before displaying it.
```

```
plt.imshow(np.uint8(img_tinted))
plt.show()
```



Running this previous code produces this plot.

Check out the Python Docs to find out more.

[Python Docs](#)