

Lecture 12 Handout

NumPy Fundamentals and Array Operations

The Foundation of Scientific Computing in Python

INF 605 - Introduction to Programming - Python

Prof. Rongyu Lin

Quinnipiac University

School of Computing and Engineering

Fall 2024

Required Reading

Textbook: Chapter 4, Sections 4.1-4.3 (NumPy Basics)

Reference Notebooks: `ch04/04_Numpy_Basics.ipynb` for array fundamentals

Learning Objectives

By the end of this lecture, you will be able to:

1. **Understand NumPy's role** in Python's scientific computing ecosystem and why it's essential for data analysis
2. **Create NumPy arrays** from lists, tuples, and built-in functions with appropriate data types
3. **Master array indexing and slicing** including multi-dimensional arrays and boolean indexing
4. **Perform element-wise operations** and understand the difference between array and list operations
5. **Manipulate array shapes** using reshape, flatten, and transpose operations
6. **Apply mathematical operations** including statistical functions and aggregations
7. **Work with array broadcasting** to perform operations on arrays of different shapes
8. **Create professional data processing pipelines** using NumPy's efficient array operations

Prerequisites Review

Building on Your Complete Programming Foundation:

From your comprehensive foundation in Python programming (Lectures 1-11), you've mastered fundamental data types, control structures, functions, file handling, object-oriented programming with inheritance and polymorphism, and advanced string operations. You've worked

with lists, tuples, and dictionaries, understanding how Python stores and manipulates data in memory.

This lecture introduces NumPy, Python's fundamental package for scientific computing. While Python's built-in lists are flexible and easy to use, they're not optimized for numerical computations. NumPy provides a powerful N-dimensional array object that forms the foundation for nearly all scientific and data analysis libraries in Python.

Transformation Goal: Evolve from using Python lists for basic data storage to leveraging NumPy arrays for efficient numerical computing and data analysis.

1 Part 1: Introduction to NumPy

1.1 Why NumPy?

NumPy (Numerical Python) is the fundamental package for scientific computing in Python. It provides a high-performance multidimensional array object and tools for working with these arrays. Think of NumPy arrays as specialized containers optimized for numerical data - like having a Formula 1 race car instead of a family sedan when you need speed and precision.

The key advantages of NumPy over Python lists include:

- **Performance:** NumPy arrays are stored in contiguous memory and operations are implemented in C, making them 10-100x faster than Python lists for numerical operations
- **Vectorization:** Operations can be applied to entire arrays without writing loops, leading to cleaner and faster code
- **Memory Efficiency:** NumPy arrays use less memory than Python lists for numerical data
- **Broadcasting:** Sophisticated rules for performing operations on arrays of different shapes
- **Ecosystem:** NumPy is the foundation for pandas, scikit-learn, matplotlib, and most scientific Python libraries

```
1  # Comparing Python lists vs NumPy arrays
2  import numpy as np
3  import time
4
5  # Creating large datasets
6  size = 1000000
7  python_list = list(range(size))
8  numpy_array = np.arange(size)
9
10 # Timing element-wise multiplication
11 # Python list approach
12 start_time = time.time()
13 python_result = [x * 2 for x in python_list]
14 python_time = time.time() - start_time
15
16 # NumPy array approach
17 start_time = time.time()
18 numpy_result = numpy_array * 2
19 numpy_time = time.time() - start_time
20
21 print(f"Python list time: {python_time:.4f} seconds")
22 print(f"NumPy array time: {numpy_time:.4f} seconds")
23 print(f"NumPy is {python_time/numpy_time:.1f}x faster!")
```

1.2 Installing and Importing NumPy

NumPy is not part of Python's standard library and needs to be installed separately. It's typically installed using pip:

```
1 # Installation (run in terminal/command prompt)
2 pip install numpy
3
4 # Standard import convention
5 import numpy as np # 'np' is the universally accepted alias
6
7 # Verify installation
8 print(f"NumPy version: {np.__version__}")
```

2 Part 2: Creating NumPy Arrays

2.1 Arrays from Python Sequences

The most straightforward way to create NumPy arrays is from Python lists or tuples. The 'np.array()' function converts sequences into arrays, automatically inferring the appropriate data type. Think of this as pouring water (your data) from a flexible container (Python list) into a rigid, optimized container (NumPy array).

```
1 # Creating 1D arrays from lists
2 numbers_list = [1, 2, 3, 4, 5]
3 numbers_array = np.array(numbers_list)
4 print(f"Array: {numbers_array}")
5 print(f"Type: {type(numbers_array)}")
6 print(f>Data type: {numbers_array.dtype}")
7
8 # Creating arrays with specific data types
9 float_array = np.array([1, 2, 3, 4, 5], dtype=np.float64)
10 print(f"Float array: {float_array}")
11
12 # Creating 2D arrays (matrices)
13 matrix_list = [[1, 2, 3],
14                [4, 5, 6],
15                [7, 8, 9]]
16 matrix_array = np.array(matrix_list)
17 print(f"\n2D Array:\n{matrix_array}")
18 print(f"Shape: {matrix_array.shape}") # (rows, columns)
19 print(f"Dimensions: {matrix_array.ndim}")
20 print(f"Total elements: {matrix_array.size}")
```

2.2 Array Creation Functions

NumPy provides numerous functions to create arrays with specific patterns or values. These are essential for initializing arrays for computations, creating test data, or setting up mathematical operations.

```
1 # Common array creation functions
2 # Zeros - often used for initialization
3 zeros_1d = np.zeros(5)
4 zeros_2d = np.zeros((3, 4)) # Note: shape as tuple
5 print(f"1D zeros: {zeros_1d}")
6 print(f"2D zeros:\n{zeros_2d}")
7
8 # Ones - useful for multiplicative operations
9 ones_1d = np.ones(5)
10 ones_2d = np.ones((2, 3), dtype=int)
```

```

11 print(f"\n1D ones: {ones_1d}")
12 print(f"2D ones:\n{ones_2d}")
13
14 # Full - create array filled with specific value
15 full_array = np.full((3, 3), 7)
16 print(f"\nFull array:\n{full_array}")
17
18 # Identity matrix - diagonal ones, zeros elsewhere
19 identity = np.eye(4)
20 print(f"\nIdentity matrix:\n{identity}")
21
22 # Empty - uninitialized array (faster but contains garbage)
23 empty_array = np.empty((2, 2))
24 print(f"\nEmpty array (uninitialized):\n{empty_array}")

```

2.3 Sequential Arrays

Creating arrays with sequential values is common in numerical computing. NumPy provides ‘arange’ (similar to Python’s range) and ‘linspace’ for different sequential patterns.

```

1 # arange - similar to range but returns array
2 integers = np.arange(10) # 0 to 9
3 evens = np.arange(0, 20, 2) # Even numbers 0 to 18
4 floats = np.arange(0, 1, 0.1) # 0.0 to 0.9 in steps of 0.1
5
6 print(f"Integers: {integers}")
7 print(f"Evens: {evens}")
8 print(f"Floats: {floats}")
9
10 # linspace - evenly spaced values over interval
11 linear = np.linspace(0, 1, 5) # 5 points from 0 to 1
12 angles = np.linspace(0, 2*np.pi, 8) # For circular calculations
13
14 print(f"\nLinspace: {linear}")
15 print(f"Angles: {angles}")
16
17 # Random arrays - essential for simulations
18 np.random.seed(42) # For reproducibility
19 random_uniform = np.random.random((3, 3)) # 0 to 1
20 random_normal = np.random.randn(5) # Standard normal
21 random_integers = np.random.randint(1, 10, size=10)
22
23 print(f"\nRandom uniform:\n{random_uniform}")
24 print(f"Random normal: {random_normal}")
25 print(f"Random integers: {random_integers}")

```

3 Part 3: Array Indexing and Slicing

3.1 Basic Indexing

NumPy arrays support all the indexing operations you learned with Python lists, plus powerful additional features. For 1D arrays, indexing works exactly like lists. For multi-dimensional arrays, you can index along each dimension separately or simultaneously.

```

1 # 1D array indexing
2 arr_1d = np.array([10, 20, 30, 40, 50])
3 print(f"Original array: {arr_1d}")
4 print(f"First element: {arr_1d[0]}")
5 print(f>Last element: {arr_1d[-1]}")
6 print(f"Third element: {arr_1d[2]}")

```

```

7
8 # 2D array indexing
9 arr_2d = np.array([[1, 2, 3],
10                   [4, 5, 6],
11                   [7, 8, 9]])
12 print(f"\n2D array:\n{arr_2d}")
13
14 # Two ways to index 2D arrays
15 print(f"Element at row 1, col 2: {arr_2d[1, 2]}") # Preferred
16 print(f"Same element: {arr_2d[1][2]}") # Works but slower
17
18 # Accessing entire rows or columns
19 print(f"Second row: {arr_2d[1]}") # Complete row
20 print(f"Second column: {arr_2d[:, 1]}") # All rows, column 1
21
22 # Modifying elements
23 arr_2d[0, 0] = 100
24 arr_2d[2] = [70, 80, 90] # Replace entire row
25 print(f"\nModified array:\n{arr_2d}")

```

3.2 Advanced Slicing

NumPy extends Python's slicing syntax to work with multiple dimensions. You can slice along each axis independently, creating powerful ways to extract subarrays. Think of this as cutting a cake - you can slice horizontally, vertically, or both to get exactly the piece you want.

```

1 # Creating a larger array for slicing demos
2 big_array = np.arange(24).reshape(4, 6)
3 print(f"Original array:\n{big_array}")
4
5 # Basic slicing
6 print(f"\nFirst two rows:\n{big_array[:2]}")
7 print(f"Last two columns:\n{big_array[:, -2:]}")
8 print(f"Middle section:\n{big_array[1:3, 2:5]}")
9
10 # Strided slicing
11 print(f"\nEvery other row:\n{big_array[::2]}")
12 print(f"Every other column:\n{big_array[:, ::2]}")
13 print(f"Reverse rows:\n{big_array[::-1]}")
14
15 # Combining slicing techniques
16 print(f"\nComplex slice - alternating rows, middle columns:\n"
17       f"{big_array[::2, 1:4]}")
18
19 # Important: Slicing creates views, not copies!
20 slice_view = big_array[1:3, 2:4]
21 print(f"\nSlice view:\n{slice_view}")
22 slice_view[0, 0] = 999
23 print(f"Original array modified:\n{big_array}")
24
25 # Creating a copy instead of view
26 slice_copy = big_array[1:3, 2:4].copy()
27 slice_copy[0, 0] = -1
28 print(f"\nCopy modified, original unchanged:\n{big_array}")

```

3.3 Boolean Indexing

Boolean indexing is one of NumPy's most powerful features, allowing you to select elements based on conditions. This creates a boolean mask array that selects only the elements where the condition is True. It's like using a filter to extract exactly the data points you need.

```

1  # Boolean indexing with 1D arrays
2  temps = np.array([72, 68, 75, 71, 69, 76, 73])
3  print(f"Temperatures: {temps}")
4
5  # Create boolean mask
6  warm_days = temps > 70
7  print(f"Warm days mask: {warm_days}")
8
9  # Use mask to select elements
10 warm_temps = temps[warm_days]
11 print(f"Warm temperatures: {warm_temps}")
12
13 # Direct boolean indexing
14 print(f"Cool temperatures: {temps[temps <= 70]}")
15
16 # Modifying elements with boolean indexing
17 temps[temps > 75] = 75 # Cap at 75
18 print(f"Capped temperatures: {temps}")
19
20 # Boolean indexing with 2D arrays
21 data = np.random.randint(0, 100, size=(5, 4))
22 print(f"\nRandom data:\n{data}")
23
24 # Multiple conditions
25 mask = (data > 30) & (data < 70) # Note: use & not 'and'
26 print(f"Values between 30 and 70: {data[mask]}")
27
28 # Set values based on condition
29 data[data < 50] = 0
30 data[data >= 50] = 1
31 print(f"Binary data:\n{data}")

```

4 Part 4: Array Operations and Mathematics

4.1 Element-wise Operations

NumPy arrays support vectorized operations, meaning you can perform operations on entire arrays without writing loops. These operations are element-wise by default, applying the operation to each corresponding element. This is like having a team of workers each handling one element simultaneously, rather than one worker processing elements sequentially.

```

1  # Basic arithmetic operations
2  a = np.array([1, 2, 3, 4, 5])
3  b = np.array([10, 20, 30, 40, 50])
4
5  print(f"Array a: {a}")
6  print(f"Array b: {b}")
7
8  # Element-wise operations
9  print(f"\nAddition (a + b): {a + b}")
10 print(f"Subtraction (b - a): {b - a}")
11 print(f"Multiplication (a * b): {a * b}")
12 print(f"Division (b / a): {b / a}")
13 print(f"Power (a ** 2): {a ** 2}")
14
15 # Operations with scalars (broadcasting)
16 print(f"\nScalar operations:")
17 print(f"a * 10: {a * 10}")
18 print(f"b + 5: {b + 5}")
19 print(f"100 / a: {100 / a}")

```

```

20
21 # Mathematical functions
22 angles = np.array([0, np.pi/4, np.pi/2, np.pi])
23 print(f"\nAngles: {angles}")
24 print(f"Sine: {np.sin(angles)}")
25 print(f"Cosine: {np.cos(angles)}")
26
27 # More math functions
28 numbers = np.array([1, 4, 9, 16, 25])
29 print(f"\nSquare root: {np.sqrt(numbers)}")
30 print(f"Exponential: {np.exp(a[:3])}") # e^x
31 print(f"Natural log: {np.log(numbers)}")

```

4.2 Array Aggregations

Aggregation functions compute summary statistics across arrays or along specific axes. These are essential for data analysis, allowing you to quickly understand the characteristics of your data. Think of these as taking a bird's-eye view of your data landscape.

```

1 # Creating sample data
2 data = np.random.randint(1, 100, size=(4, 5))
3 print(f"Sample data:\n{data}")
4
5 # Basic aggregations on entire array
6 print(f"\nArray-wide statistics:")
7 print(f"Sum: {np.sum(data)}")
8 print(f"Mean: {np.mean(data)}")
9 print(f"Median: {np.median(data)}")
10 print(f"Standard deviation: {np.std(data):.2f}")
11 print(f"Min: {np.min(data)}, Max: {np.max(data)}")
12
13 # Aggregations along axes
14 print(f"\nAxis-specific aggregations:")
15 print(f"Sum along rows (axis=1): {np.sum(data, axis=1)}")
16 print(f"Sum along columns (axis=0): {np.sum(data, axis=0)}")
17 print(f"Mean of each row: {np.mean(data, axis=1)}")
18 print(f"Max of each column: {np.max(data, axis=0)}")
19
20 # Finding positions of min/max
21 print(f"\nPosition information:")
22 print(f"Position of minimum: {np.argmin(data)}")
23 print(f"Position of maximum: {np.argmax(data)}")
24 print(f"Position as (row, col): {np.unravel_index(np.argmax(data), data.shape)}")
25
26 # Cumulative operations
27 arr = np.array([1, 2, 3, 4, 5])
28 print(f"\nCumulative operations on {arr}:")
29 print(f"Cumulative sum: {np.cumsum(arr)}")
30 print(f"Cumulative product: {np.cumprod(arr)}")

```

5 Part 5: Array Shape Manipulation

5.1 Reshaping Arrays

Reshaping is fundamental to NumPy operations, allowing you to change the dimensions of an array without changing its data. Think of it like reorganizing a deck of cards - you can arrange them in different patterns (4 rows of 13, 13 rows of 4, etc.) but you still have the same 52 cards.

```

1  # Creating a 1D array
2  original = np.arange(12)
3  print(f"Original 1D array: {original}")
4  print(f"Shape: {original.shape}")
5
6  # Reshape to 2D
7  reshaped_2d = original.reshape(3, 4)
8  print(f"\nReshaped to 3x4:\n{reshaped_2d}")
9
10 # Reshape to different 2D
11 reshaped_alt = original.reshape(2, 6)
12 print(f"\nReshaped to 2x6:\n{reshaped_alt}")
13
14 # Reshape to 3D
15 reshaped_3d = original.reshape(2, 2, 3)
16 print(f"\nReshaped to 2x2x3:\n{reshaped_3d}")
17
18 # Using -1 for automatic dimension calculation
19 auto_reshape = original.reshape(3, -1) # NumPy calculates columns
20 print(f"\nAuto reshape (3, -1):\n{auto_reshape}")
21
22 # Flattening arrays
23 print(f"\nFlattening methods:")
24 print(f"Flatten (copy): {reshaped_2d.flatten()}")
25 print(f"Ravel (view when possible): {reshaped_2d.ravel()}")
26
27 # Important: reshape returns a view when possible
28 reshaped_view = original.reshape(3, 4)
29 reshaped_view[0, 0] = 999
30 print(f"\nOriginal modified through view: {original}")

```

5.2 Transposing and Axis Manipulation

Transposing swaps the axes of an array, which is essential for matrix operations and data alignment. For 2D arrays, this flips rows and columns. For higher dimensions, you can specify exactly how to rearrange the axes.

```

1  # Transposing 2D arrays
2  matrix = np.array([[1, 2, 3],
3                     [4, 5, 6]])
4  print(f"Original matrix:\n{matrix}")
5  print(f"Shape: {matrix.shape}")
6
7  # Simple transpose
8  transposed = matrix.T
9  print(f"\nTransposed:\n{transposed}")
10 print(f"Shape: {transposed.shape}")
11
12 # Alternative transpose methods
13 transposed_alt = np.transpose(matrix)
14 print(f"\nUsing np.transpose:\n{transposed_alt}")
15
16 # Transposing 1D arrays (no effect)
17 arr_1d = np.array([1, 2, 3, 4])
18 print(f"\n1D array: {arr_1d}")
19 print(f"1D transposed: {arr_1d.T}") # Still 1D!
20
21 # For higher dimensions
22 arr_3d = np.arange(24).reshape(2, 3, 4)
23 print(f"\n3D array shape: {arr_3d.shape}")
24 transposed_3d = np.transpose(arr_3d, axes=(1, 0, 2))

```



```

25 print(f"Transposed shape: {transposed_3d.shape}")
26
27 # Swapping specific axes
28 swapped = np.swapaxes(arr_3d, 0, 1)
29 print(f"Swapped axes shape: {swapped.shape}")

```

5.3 Stacking and Splitting

Combining and dividing arrays is essential for data preprocessing and manipulation. NumPy provides various functions to stack arrays together or split them apart, like assembling or disassembling building blocks.

```

1  # Creating sample arrays
2  a = np.array([1, 2, 3])
3  b = np.array([4, 5, 6])
4  c = np.array([7, 8, 9])
5
6  print(f"Array a: {a}")
7  print(f"Array b: {b}")
8  print(f"Array c: {c}")
9
10 # Vertical stacking (row-wise)
11 vstacked = np.vstack([a, b, c])
12 print(f"\nVertical stack:\n{vstacked}")
13
14 # Horizontal stacking (column-wise)
15 hstacked = np.hstack([a, b, c])
16 print(f"\nHorizontal stack: {hstacked}")
17
18 # Column stacking (1D to 2D columns)
19 column_stacked = np.column_stack([a, b, c])
20 print(f"\nColumn stack:\n{column_stacked}")
21
22 # Concatenate (general purpose)
23 concat_axis0 = np.concatenate([a.reshape(1, -1),
24                                b.reshape(1, -1),
25                                c.reshape(1, -1)], axis=0)
26 print(f"\nConcatenate along axis 0:\n{concat_axis0}")
27
28 # Splitting arrays
29 big_array = np.arange(12).reshape(3, 4)
30 print(f"\nArray to split:\n{big_array}")
31
32 # Split into equal parts
33 vsplit_arrays = np.vsplit(big_array, 3) # 3 equal row groups
34 print(f"\nVertical split into 3:")
35 for i, arr in enumerate(vsplit_arrays):
36     print(f"Part {i+1}: {arr}")
37
38 # Split at specific indices
39 hsplit_arrays = np.hsplit(big_array, [1, 3]) # Split at columns 1 and 3
40 print(f"\nHorizontal split at indices [1, 3]:")
41 for i, arr in enumerate(hsplit_arrays):
42     print(f"Part {i+1}:\n{arr}")

```

6 Part 6: Broadcasting

6.1 Understanding Broadcasting Rules

Broadcasting is NumPy's powerful mechanism for performing operations on arrays of different shapes. It follows specific rules to "stretch" smaller arrays across larger ones, enabling element-wise operations without explicitly creating copies. Think of broadcasting like using a stamp - you can apply the same pattern across a larger surface without manually copying it.

Broadcasting rules:

1. Arrays are compatible for broadcasting if their dimensions match or if one dimension is 1
2. Arrays are broadcast together by adding dimensions of size 1 to the left
3. After broadcasting, all dimensions must match

```
1  # Broadcasting scalars
2  array = np.array([[1, 2, 3],
3                   [4, 5, 6]])
4  scalar = 10
5
6  print(f"Array:\n{array}")
7  print(f"Array + scalar: \n{array + scalar}")  # Scalar broadcasts to all
8  elements
9
10 # Broadcasting 1D array to 2D
11 row_array = np.array([10, 20, 30])
12 print(f"\nRow array: {row_array}")
13 print(f"Array + row array:\n{array + row_array}")  # Broadcasts across rows
14
15 # Broadcasting with reshape for column operations
16 col_array = np.array([[100], [200]])  # Shape (2, 1)
17 print(f"\nColumn array:\n{col_array}")
18 print(f"Array + column array:\n{array + col_array}")  # Broadcasts across
19 columns
20
21 # More complex broadcasting
22 a = np.ones((3, 4))
23 b = np.arange(4)
24 c = np.arange(3).reshape(3, 1)
25
26 print(f"\nShapes: a={a.shape}, b={b.shape}, c={c.shape}")
27 print(f"a + b:\n{a + b}")
28 print(f"a + c:\n{a + c}")
29
30 # Broadcasting in practice - normalization
31 data = np.random.randint(0, 100, size=(5, 3))
32 print(f"\nOriginal data:\n{data}")
33
34 # Normalize each column (subtract mean, divide by std)
35 col_means = data.mean(axis=0)  # Shape (3,)
36 col_stds = data.std(axis=0)    # Shape (3,)
37 normalized = (data - col_means) / col_stds
38 print(f"\nNormalized data:\n{normalized}")
39 print(f"Normalized means: {normalized.mean(axis=0)}")  # Should be ~0
40 print(f"Normalized stds: {normalized.std(axis=0)}")    # Should be ~1
```

6.2 Practical Broadcasting Applications

Broadcasting enables elegant solutions to common data manipulation tasks. Here we explore practical applications that demonstrate the power and efficiency of broadcasting in real-world

scenarios.

```
1 # Creating a multiplication table using broadcasting
2 rows = np.arange(1, 11).reshape(10, 1) # Column vector
3 cols = np.arange(1, 11)                 # Row vector
4
5 multiplication_table = rows * cols
6 print("Multiplication table (1-10):")
7 print(multiplication_table)
8
9 # Distance calculation between points
10 # Points in 2D space
11 points = np.array([[0, 0], [1, 0], [0, 1], [1, 1]])
12 # Calculate distances from origin [0, 0]
13 origin = np.array([0, 0])
14 distances = np.sqrt(np.sum((points - origin)**2, axis=1))
15 print(f"\nDistances from origin: {distances}")
16
17 # Image-like data manipulation
18 # Simulate RGB image data (height=3, width=4, channels=3)
19 image = np.random.randint(0, 256, size=(3, 4, 3))
20 print(f"\nOriginal 'image' shape: {image.shape}")
21
22 # Adjust brightness by scaling all pixels
23 brightness_factor = 0.5
24 dimmed = (image * brightness_factor).astype(int)
25 print(f"Dimmed image sample:\n{dimmed[0]}") # First row
26
27 # Apply different scaling to each color channel
28 channel_scales = np.array([1.2, 0.8, 0.9]) # R, G, B scales
29 adjusted = (image * channel_scales).clip(0, 255).astype(int)
30 print(f"Channel-adjusted sample:\n{adjusted[0]}")
```

7 Part 7: Practical NumPy Applications

7.1 Data Analysis Pipeline

Let's build a complete data analysis pipeline using NumPy, demonstrating how these concepts work together in practice. We'll analyze temperature data from multiple weather stations.

```
1 # Simulating temperature data from 5 weather stations over 7 days
2 np.random.seed(42)
3 stations = ['Station_A', 'Station_B', 'Station_C', 'Station_D', 'Station_E']
4 days = ['Mon', 'Tue', 'Wed', 'Thu', 'Fri', 'Sat', 'Sun']
5
6 # Generate temperature data (Celsius)
7 # Base temperatures with daily variations
8 base_temps = np.array([20, 22, 19, 21, 23]) # Per station
9 daily_variation = np.random.randn(7, 5) * 3 # Random variation
10 temperatures = base_temps + daily_variation
11
12 print("Temperature data (\textdegree{}C):")
13 print(temperatures)
14
15 # Analysis 1: Basic statistics per station
16 print("\nStation Statistics:")
17 print(f"Mean temperatures: {temperatures.mean(axis=0)}")
18 print(f"Max temperatures: {temperatures.max(axis=0)}")
19 print(f"Min temperatures: {temperatures.min(axis=0)}")
20 print(f"Temperature ranges: {temperatures.max(axis=0) - temperatures.min(axis=0)}")
21
```

```

22 # Analysis 2: Find hottest and coldest days
23 daily_means = temperatures.mean(axis=1)
24 hottest_day = days[np.argmax(daily_means)]
25 coldest_day = days[np.argmin(daily_means)]
26 print(f"\nHottest day: {hottest_day} ({daily_means.max():.1f}\textdegree{}C
    average)")
27 print(f"Coldest day: {coldest_day} ({daily_means.min():.1f}\textdegree{}C
    average)")
28
29 # Analysis 3: Identify extreme temperatures
30 threshold_high = 25
31 threshold_low = 17
32 extreme_high = temperatures > threshold_high
33 extreme_low = temperatures < threshold_low
34
35 print(f"\nDays with temperatures above {threshold_high}\textdegree{}C:")
36 high_days, high_stations = np.where(extreme_high)
37 for day, station in zip(high_days, high_stations):
38     print(f"    {days[day]} at {stations[station]}: {temperatures[day, station]
        :.1f}\textdegree{}C")
39
40 # Analysis 4: Temperature anomalies
41 station_means = temperatures.mean(axis=0, keepdims=True)
42 anomalies = temperatures - station_means
43 print(f"\nLargest positive anomaly: {anomalies.max():.1f}\textdegree{}C")
44 anomaly_pos = np.unravel_index(np.argmax(anomalies), anomalies.shape)
45 print(f"    Occurred on {days[anomaly_pos[0]]} at {stations[anomaly_pos[1]]}")

```

7.2 Financial Data Processing

NumPy is extensively used in financial analysis. Here's an example processing stock price data and calculating various financial metrics.

```

1 # Simulating stock price data
2 np.random.seed(42)
3 days = 252 # Trading days in a year
4 initial_price = 100
5 daily_returns = np.random.randn(days) * 0.02 # 2% daily volatility
6
7 # Calculate price series using cumulative product
8 price_multipliers = 1 + daily_returns
9 prices = initial_price * np.cumprod(price_multipliers)
10
11 print(f"Stock price statistics:")
12 print(f"Starting price: ${initial_price:.2f}")
13 print(f"Ending price: ${prices[-1]:.2f}")
14 print(f"Max price: ${prices.max():.2f}")
15 print(f"Min price: ${prices.min():.2f}")
16
17 # Calculate moving averages
18 window_short = 20
19 window_long = 50
20
21 # Simple moving averages
22 sma_short = np.convolve(prices, np.ones(window_short)/window_short, mode='valid',
    ')
23 sma_long = np.convolve(prices, np.ones(window_long)/window_long, mode='valid')
24
25 # Calculate daily returns from prices
26 price_returns = np.diff(prices) / prices[:-1]
27
28 # Risk metrics

```

```

29 volatility = np.std(price_returns) * np.sqrt(252) # Annualized
30 sharpe_ratio = np.mean(price_returns) * 252 / volatility
31 max_drawdown = np.min(prices / np.maximum.accumulate(prices) - 1)
32
33 print(f"\nRisk Metrics:")
34 print(f"Annual volatility: {volatility:.1%}")
35 print(f"Sharpe ratio: {sharpe_ratio:.2f}")
36 print(f"Maximum drawdown: {max_drawdown:.1%}")
37
38 # Find best and worst periods
39 rolling_returns = np.convolve(price_returns, np.ones(5)/5, mode='valid')
40 best_week = np.argmax(rolling_returns)
41 worst_week = np.argmin(rolling_returns)
42
43 print(f"\nBest 5-day period: Days {best_week}-{best_week+4}")
44 print(f"Worst 5-day period: Days {worst_week}-{worst_week+4}")

```

Summary and Best Practices

Key Takeaways:

1. NumPy arrays are the foundation of scientific computing in Python, offering 10-100x performance improvements over lists for numerical operations
2. Arrays can be created from sequences, built-in functions (zeros, ones, arange, linspace), or random generators
3. Indexing and slicing work similarly to lists but extend naturally to multiple dimensions
4. Vectorized operations eliminate the need for explicit loops, making code cleaner and faster
5. Broadcasting enables operations between arrays of different shapes following consistent rules
6. Shape manipulation (reshape, transpose, stack, split) is essential for data preprocessing
7. Always consider memory views vs. copies when slicing and reshaping arrays

Common Pitfalls to Avoid

- Forgetting that slicing creates views, not copies - modifications affect the original array
- Using Python's 'and', 'or', 'not' instead of NumPy's '&', '—', '~' for boolean operations
- Not understanding broadcasting rules, leading to unexpected results or errors
- Creating unnecessary copies of large arrays, causing memory issues
- Using loops instead of vectorized operations, resulting in slow code
- Mixing NumPy arrays with Python lists without understanding the performance implications

Practice Exercises

Exercise 1: Grade Analysis System

Create a program that analyzes student grades across multiple subjects:

- Generate random grades (60-100) for 30 students across 5 subjects
- Calculate average grade per student and per subject
- Identify students with any failing grades (≤ 70)
- Find the top 5 performing students
- Calculate grade distribution statistics

Exercise 2: Image Processing Basics

Simulate basic image operations using NumPy:

- Create a 100x100 pixel "image" with random grayscale values (0-255)
- Apply brightness and contrast adjustments
- Create a simple blur effect using averaging
- Detect edges by calculating pixel differences
- Generate a histogram of pixel intensities

Exercise 3: Sales Data Analysis

Build a sales analysis system:

- Create sales data for 12 months across 10 products
- Calculate month-over-month growth rates
- Identify best and worst performing products
- Find seasonal patterns using moving averages
- Generate a sales forecast based on trends

Next Week Preview

In Lecture 13, we'll explore advanced NumPy features including structured arrays, memory layout optimization, advanced indexing techniques, and integration with pandas for data analysis. We'll also cover performance optimization strategies and real-world applications in data science and machine learning.