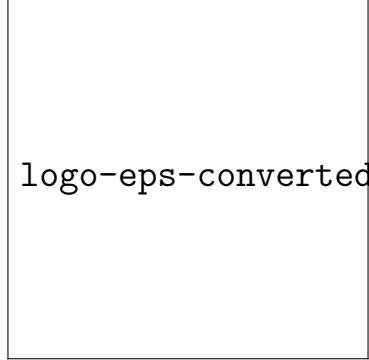


NumPy

Numerical Python for Scientific Computing



logo-eps-converted-to.pdf

August 7, 2025

Contents

1	Introduction to NumPy	4
1.1	What is NumPy?	4
1.2	Why NumPy?	4
1.3	Installation and Import	4
2	Array Creation	4
2.1	Basic Array Creation	4
2.2	Special Arrays	5
2.3	Data Types	5
3	Array Attributes and Properties	5
3.1	Essential Attributes	5
4	Array Indexing and Slicing	6
4.1	Basic Indexing	6
4.2	Advanced Indexing	6
5	Array Operations and Broadcasting	7
5.1	Element-wise Operations	7
5.2	Broadcasting	7
6	Array Manipulation	8
6.1	Reshaping and Resizing	8
6.2	Concatenation and Splitting	8
7	Views vs Copies	8
7.1	Operations that Create Views	8
7.2	Operations that Create Copies	9
7.3	Practical Example	9
8	Mathematical and Statistical Operations	9
8.1	Basic Mathematical Functions	9
8.2	Statistical Functions	10
8.3	Axis Parameter	10
9	Linear Algebra	10
9.1	Basic Operations	10
9.2	Advanced Linear Algebra	11
10	Working with Missing Data	11
10.1	NaN Handling	11
10.2	Masked Arrays	12
11	File I/O Operations	12
11.1	NumPy Native Formats	12
11.2	Text Formats	12
12	Performance Optimization	13
12.1	Vectorization vs Loops	13
12.2	Memory Optimization	13

13 Common Pitfalls and Debugging	14
13.1 Frequent Mistakes	14
13.2 Debugging Techniques	14
14 Advanced Topics	15
14.1 Custom Universal Functions (ufuncs)	15
14.2 Structured Arrays	15
15 Integration with Scientific Python Ecosystem	15
15.1 With Pandas	15
15.2 With Matplotlib	16
15.3 With SciPy	16
16 Real-World Applications	16
16.1 Image Processing	16
16.2 Data Analysis Example	17
17 Practice Exercises	18
17.1 Beginner Exercises	18
17.2 Intermediate Exercises	18
17.3 Advanced Exercises	18
18 Best Practices and Guidelines	18
19 Conclusion	19
20 Further Reading	19

1 Introduction to NumPy

1.1 What is NumPy?

NumPy (Numerical Python) is the fundamental library for numerical computing in Python. It provides:

- Powerful `ndarray` data structure for fast, vectorized computations
- Tools for linear algebra, Fourier transforms, and random number generation
- The core foundation of the scientific Python ecosystem (pandas, SciPy, scikit-learn, etc.)

1.2 Why NumPy?

Feature	Python Lists	NumPy Arrays
Memory Usage	High	Low (contiguous memory)
Speed	Slow (interpreted)	Fast (C implementation)
Vectorization	Manual loops	Built-in operations
Mathematical Operations	Limited	Extensive
Broadcasting	No	Yes

Table 1: Python Lists vs NumPy Arrays

1.3 Installation and Import

```

1 # Installation
2 pip install numpy
3
4 # Import (standard convention)
5 import numpy as np
6
7 # Check version
8 print(np.__version__)

```

2 Array Creation

2.1 Basic Array Creation

Array Creation Methods

```

1 # From Python lists
2 arr_from_list = np.array([1, 2, 3, 4, 5])
3 matrix = np.array([[1, 2, 3], [4, 5, 6]])
4
5 # Using built-in functions
6 zeros_arr = np.zeros((3, 4))           # 3x4 array of zeros
7 ones_arr = np.ones((2, 3))             # 2x3 array of ones
8 empty_arr = np.empty((2, 2))          # Uninitialized array
9 full_arr = np.full((3, 3), 7)         # Fill with specific value
10
11 # Range functions
12 range_arr = np.arange(0, 10, 2)       # [0, 2, 4, 6, 8]
13 linspace_arr = np.linspace(0, 1, 5)    # [0, 0.25, 0.5, 0.75, 1.0]

```

```
14 logspace_arr = np.logspace(0, 2, 3) # [1, 10, 100]
```

2.2 Special Arrays

```
1 # Identity matrix
2 identity = np.eye(4)                      # 4x4 identity matrix
3
4 # Diagonal arrays
5 diagonal = np.diag([1, 2, 3, 4])      # Diagonal matrix
6 off_diagonal = np.diag([1, 2, 3], k=1)   # Above main diagonal
7
8 # Random arrays
9 np.random.seed(42) # For reproducibility
10 random_uniform = np.random.rand(3, 3)     # Uniform [0,1)
11 random_normal = np.random.randn(3, 3)      # Standard normal
12 random_int = np.random.randint(1, 10, (3, 3)) # Random integers
```

2.3 Data Types

Type	NumPy dtype	Description
Integer	int8, int16, int32, int64	Signed integers
Unsigned	uint8, uint16, uint32, uint64	Unsigned integers
Float	float16, float32, float64	Floating point
Complex	complex64, complex128	Complex numbers
Boolean	bool	True/False
String	U<n>	Unicode strings

Table 2: NumPy Data Types

```
1 # Specifying data types
2 int_array = np.array([1, 2, 3], dtype=np.int32)
3 float_array = np.array([1, 2, 3], dtype=np.float64)
4 bool_array = np.array([1, 0, 1], dtype=bool)
5
6 # Type conversion
7 float_arr = int_array.astype(np.float64)
```

3 Array Attributes and Properties

3.1 Essential Attributes

```
1 arr = np.random.rand(3, 4, 5)
2
3 # Shape and dimensions
4 print(f"Shape: {arr.shape}")          # (3, 4, 5)
5 print(f"Dimensions: {arr.ndim}")      # 3
6 print(f"Size: {arr.size}")           # 60 (total elements)
7
8 # Data type and memory
9 print(f"Data type: {arr.dtype}")     # float64
10 print(f"Item size: {arr.itemsize}")  # 8 bytes per element
11 print(f"Total bytes: {arr.nbytes}") # 480 bytes
```

```

12
13 # Memory layout
14 print(f"C-contiguous: {arr.flags.c_contiguous}")
15 print(f"Fortran-contiguous: {arr.flags.f_contiguous}")
16 print(f"Owns data: {arr.flags.owndata}")

```

4 Array Indexing and Slicing

4.1 Basic Indexing

1D Array Indexing

```

1 arr = np.arange(10)    # [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
2
3 # Basic indexing
4 print(arr[0])          # 0 (first element)
5 print(arr[-1])         # 9 (last element)
6 print(arr[2:5])         # [2, 3, 4] (slice)
7 print(arr[::-2])        # [0, 2, 4, 6, 8] (every 2nd element)
8 print(arr[::-1])        # [9, 8, 7, 6, 5, 4, 3, 2, 1, 0] (reverse)

```

2D Array Indexing

```

1 matrix = np.arange(12).reshape(3, 4)
2 # [[ 0  1  2  3]
3 #  [ 4  5  6  7]
4 #  [ 8  9 10 11]]
5
6 # Element access
7 print(matrix[1, 2])      # 6 (row 1, column 2)
8 print(matrix[1][2])      # 6 (alternative syntax)
9
10 # Row and column access
11 print(matrix[1, :])       # [4 5 6 7] (entire row 1)
12 print(matrix[:, 2])       # [2 6 10] (entire column 2)
13
14 # Submatrix
15 print(matrix[1:3, 1:3])    # 2x2 submatrix
16 print(matrix[:, ::2])      # Every other row and column

```

4.2 Advanced Indexing

```

1 arr = np.arange(10)
2
3 # Boolean indexing
4 mask = arr > 5
5 print(arr[mask])           # [6 7 8 9]
6
7 # Fancy indexing with arrays
8 indices = [1, 3, 5, 7]
9 print(arr[indices])        # [1 3 5 7]
10
11 # Combining conditions
12 mask = (arr > 2) & (arr < 8)
13 print(arr[mask])           # [3 4 5 6 7]
14

```

```

15 # 2D boolean indexing
16 matrix = np.random.randint(0, 10, (3, 4))
17 print(matrix[matrix > 5]) # All elements > 5 (flattened)

```

5 Array Operations and Broadcasting

5.1 Element-wise Operations

```

1 a = np.array([1, 2, 3, 4])
2 b = np.array([5, 6, 7, 8])
3
4 # Arithmetic operations
5 print(a + b)           # [6 8 10 12]
6 print(a * b)           # [5 12 21 32]
7 print(a ** 2)          # [1 4 9 16]
8 print(np.sqrt(a))      # [1.        1.414 1.732 2.      ]
9
10 # Comparison operations
11 print(a > 2)          # [False False True True]
12 print(a == b)          # [False False False False]
13
14 # Logical operations
15 print(np.logical_and(a > 1, a < 4)) # [False True True False]

```

5.2 Broadcasting

Broadcasting allows operations between arrays of different shapes following specific rules:

Broadcasting Rules

1. Arrays are aligned from the trailing dimension
2. Dimensions of size 1 can be "stretched" to match
3. Missing dimensions are assumed to be size 1

Broadcasting Examples

```

1 # Scalar and array
2 arr = np.array([[1, 2, 3], [4, 5, 6]])
3 result = arr + 10          # Adds 10 to every element
4
5 # 1D array with 2D array
6 row_vector = np.array([10, 20, 30])      # Shape: (3,)
7 result = arr + row_vector                 # Shape: (2, 3)
8
9 # Column vector with row vector
10 col_vector = np.array([[1], [2]])         # Shape: (2, 1)
11 result = col_vector + row_vector          # Shape: (2, 3)
12
13 # Mathematical example
14 x = np.linspace(0, 5, 6).reshape(6, 1)    # Column vector
15 y = np.linspace(0, 4, 5).reshape(1, 5)      # Row vector
16 distance = np.sqrt(x**2 + y**2)            # 6x5 distance matrix

```

6 Array Manipulation

6.1 Reshaping and Resizing

```

1 arr = np.arange(12)

2
3 # Reshaping (must preserve total size)
4 reshaped = arr.reshape(3, 4)          # 1D to 2D
5 reshaped = arr.reshape(2, 2, 3)        # 1D to 3D
6 reshaped = arr.reshape(-1, 4)         # Auto-calculate rows
7
8 # Flattening
9 flattened = reshaped.flatten()       # Always returns copy
10 raveled = reshaped.ravel()          # Returns view if possible
11
12 # Transpose
13 transposed = reshaped.T             # Transpose
14 transposed = np.transpose(reshaped)
15
16 # Adding/removing dimensions
17 expanded = np.expand_dims(arr, axis=1) # Add dimension
18 squeezed = np.squeeze(expanded)      # Remove size-1 dimensions

```

6.2 Concatenation and Splitting

```

1 a = np.array([[1, 2], [3, 4]])
2 b = np.array([[5, 6], [7, 8]])
3
4 # Concatenation
5 vertical = np.concatenate([a, b], axis=0)    # Stack vertically
6 horizontal = np.concatenate([a, b], axis=1)   # Stack horizontally
7
8 # Convenient functions
9 vstacked = np.vstack([a, b])                  # Vertical stack
10 hstacked = np.hstack([a, b])                  # Horizontal stack
11 dstacked = np.dstack([a, b])                  # Depth stack (3D)
12
13 # Splitting
14 large_arr = np.arange(12).reshape(4, 3)
15 split_arrays = np.split(large_arr, 2, axis=0)  # Split into 2 parts
16 hsplit_arrays = np.hsplit(large_arr, 3)        # Split horizontally
17 vsplit_arrays = np.vsplit(large_arr, 2)         # Split vertically

```

7 Views vs Copies

Understanding when NumPy creates views (shared memory) vs copies (independent memory) is crucial for performance and correctness.

Important: Views vs Copies

Views share memory with the original array, while copies create independent arrays.
Modifying a view affects the original array!

7.1 Operations that Create Views

```

1 original = np.arange(12).reshape(3, 4)
2
3 # These create VIEWS (share memory)
4 slice_view = original[1:3]           # Slicing
5 transpose_view = original.T          # Transpose
6 reshape_view = original.reshape(4, 3) # Reshape (when possible)
7 ravel_view = original.ravel()        # Ravel (when possible)
8
9 # Test if it's a view
10 print(slice_view.base is original)   # True for views
11 print(original.flags.owndata)         # False for views

```

7.2 Operations that Create Copies

```

1 # These create COPIES (independent memory)
2 explicit_copy = original.copy()       # Explicit copy
3 flatten_copy = original.flatten()     # Flatten always copies
4 fancy_copy = original[[0, 2]]         # Fancy indexing
5 boolean_copy = original[original > 5] # Boolean indexing
6
7 # Test memory sharing
8 print(np.shares_memory(original, explicit_copy)) # False for copies

```

7.3 Practical Example

```

1 # Demonstrate view behavior
2 matrix = np.arange(20).reshape(4, 5)
3
4 # Working with views - efficient but affects original
5 first_row = matrix[0]                 # This is a view!
6 first_row *= 10                      # Modifies original matrix!
7
8 # Working with copies - safe but uses more memory
9 safe_first_row = matrix[0].copy()
10 safe_first_row *= 10                 # Original matrix unchanged

```

8 Mathematical and Statistical Operations

8.1 Basic Mathematical Functions

```

1 arr = np.array([1, 4, 9, 16, 25])
2
3 # Trigonometric functions
4 angles = np.array([0, np.pi/4, np.pi/2, np.pi])
5 print(np.sin(angles))
6 print(np.cos(angles))
7 print(np.tan(angles))
8
9 # Exponential and logarithmic
10 print(np.exp(arr))                # e^x
11 print(np.log(arr))                # Natural log
12 print(np.log10(arr))              # Base-10 log
13 print(np.sqrt(arr))               # Square root
14
15 # Rounding
16 float_arr = np.array([1.2, 2.7, 3.1, 4.9])
17 print(np.round(float_arr))        # [1. 3. 3. 5.]

```

```

18 print(np.floor(float_arr))    # [1. 2. 3. 4.]
19 print(np.ceil(float_arr))     # [2. 3. 4. 5.]
```

8.2 Statistical Functions

```

1 data = np.random.randn(1000)    # Normal distribution
2
3 # Central tendency
4 print(f"Mean: {np.mean(data):.3f}")
5 print(f"Median: {np.median(data):.3f}")
6 print(f"Mode: {stats.mode(data)[0][0]:.3f}")  # Requires scipy
7
8 # Spread
9 print(f"Standard deviation: {np.std(data):.3f}")
10 print(f"Variance: {np.var(data):.3f}")
11 print(f"Range: {np.ptp(data):.3f}")  # Peak-to-peak
12
13 # Quantiles
14 print(f"25th percentile: {np.percentile(data, 25):.3f}")
15 print(f"75th percentile: {np.percentile(data, 75):.3f}")
16 print(f"IQR: {np.percentile(data, 75) - np.percentile(data, 25):.3f}")
17
18 # Extremes
19 print(f"Min: {np.min(data):.3f}")
20 print(f"Max: {np.max(data):.3f}")
21 print(f"Argmin: {np.argmin(data)})")  # Index of minimum
22 print(f"Argmax: {np.argmax(data)})")  # Index of maximum
```

8.3 Axis Parameter

The `axis` parameter specifies along which dimension to perform operations:

Understanding Axis Parameter

```

1 matrix = np.array([[1, 2, 3, 4],
2                  [5, 6, 7, 8],
3                  [9, 10, 11, 12]])
4
5 # axis=0: operate down rows (collapse rows)
6 col_sums = np.sum(matrix, axis=0)      # [15, 18, 21, 24]
7 col_means = np.mean(matrix, axis=0)    # [5., 6., 7., 8.]
8
9 # axis=1: operate across columns (collapse columns)
10 row_sums = np.sum(matrix, axis=1)     # [10, 26, 42]
11 row_means = np.mean(matrix, axis=1)   # [2.5, 6.5, 10.5]
12
13 # No axis: operate on entire array
14 total_sum = np.sum(matrix)           # 78
15 overall_mean = np.mean(matrix)      # 6.5
```

9 Linear Algebra

9.1 Basic Operations

```

1 A = np.array([[1, 2], [3, 4]])
2 B = np.array([[5, 6], [7, 8]])
```

```

3 v = np.array([1, 2])
4
5 # Matrix multiplication
6 C = np.dot(A, B)          # Matrix multiplication
7 C = A @ B                 # Alternative syntax (Python 3.5+)
8
9 # Vector operations
10 dot_product = np.dot(v, v)    # Scalar result
11 outer_product = np.outer(v, v) # Matrix result
12
13 # Matrix properties
14 det_A = np.linalg.det(A)      # Determinant
15 trace_A = np.trace(A)        # Trace (sum of diagonal)

```

9.2 Advanced Linear Algebra

```

1 # Create a symmetric positive definite matrix
2 A = np.array([[4, 2, 1],
3               [2, 5, 3],
4               [1, 3, 6]])
5
6 # Eigenvalues and eigenvectors
7 eigenvals, eigenvecs = np.linalg.eig(A)
8 print(f"Eigenvalues: {eigenvals}")
9
10 # Matrix decompositions
11 U, s, Vt = np.linalg.svd(A)       # Singular Value Decomposition
12 L = np.linalg.cholesky(A)         # Cholesky decomposition
13 Q, R = np.linalg.qr(A)           # QR decomposition
14
15 # Matrix inverse and pseudo-inverse
16 A_inv = np.linalg.inv(A)          # Inverse (for square matrices)
17 A_pinv = np.linalg.pinv(A)        # Pseudo-inverse (Moore-Penrose)
18
19 # Solving linear systems: Ax = b
20 b = np.array([1, 2, 3])
21 x = np.linalg.solve(A, b)          # Solve Ax = b
22 print(f"Solution: {x}")
23 print(f"Verification: {np.allclose(A @ x, b)}")

```

10 Working with Missing Data

10.1 NaN Handling

```

1 # Creating arrays with NaN
2 data = np.array([1.0, 2.0, np.nan, 4.0, 5.0])
3
4 # Detecting NaN
5 nan_mask = np.isnan(data)          # Boolean array
6 has_nan = np.any(np.isnan(data))   # True if any NaN
7
8 # NaN-aware functions
9 print(f"Mean (ignoring NaN): {np.nanmean(data)}")
10 print(f"Sum (ignoring NaN): {np.nansum(data)}")
11 print(f"Std (ignoring NaN): {np.nanstd(data)}")
12
13 # Removing NaN values
14 clean_data = data[~np.isnan(data)]
15

```

```
16 # Replacing NaN values
17 filled_data = np.where(np.isnan(data), 0, data)    # Replace with 0
18 filled_data = np.nan_to_num(data, nan=0.0)        # Replace with 0
```

10.2 Masked Arrays

```
1 # Masked arrays for handling missing/invalid data
2 temperatures = np.array([20.5, -999, 22.1, -999, 19.8])
3
4 # Create masked array (mask invalid values)
5 masked_temp = np.ma.masked_where(temperatures == -999, temperatures)
6
7 # Operations automatically ignore masked values
8 print(f"Mean temperature: {masked_temp.mean():.1f}")
9 print(f"Valid measurements: {masked_temp.count()}")
10
11 # Fill masked values
12 filled = masked_temp.filled(fill_value=20.0)
```

11 File I/O Operations

11.1 NumPy Native Formats

```
1 # Save and load single arrays
2 data = np.random.randn(100, 50)
3
4 # Binary format (fast, preserves exact data)
5 np.save('data.npy', data)
6 loaded_data = np.load('data.npy')
7
8 # Compressed format
9 np.savez_compressed('data.npz', data=data)
10
11 # Multiple arrays
12 sales = np.random.randint(100, 1000, (12, 4))
13 costs = sales * 0.6
14 np.savez('company_data.npz', sales=sales, costs=costs)
15
16 # Load multiple arrays
17 loaded = np.load('company_data.npz')
18 sales_loaded = loaded['sales']
19 costs_loaded = loaded['costs']
```

11.2 Text Formats

```
1 # Save to text (human-readable)
2 data = np.random.rand(5, 3)
3 np.savetxt('data.csv', data, delimiter=',', fmt='%.3f')
4
5 # Load from text
6 loaded = np.loadtxt('data.csv', delimiter=',')
7
8 # Advanced loading with headers and specific columns
9 # File format: "Name,Age,Score1,Score2,Score3"
10 student_scores = np.loadtxt('students.csv',
11                             delimiter=',',
12                             skiprows=1,                      # Skip header
```

```

13         usecols=(2, 3, 4), # Load only scores
14         dtype=float)
15
16 # Handle missing data
17 data_with_missing = np.genfromtxt('messy_data.csv',
18                                 delimiter=',',
19                                 missing_values='N/A',
20                                 filling_values=0)

```

12 Performance Optimization

12.1 Vectorization vs Loops

Performance Tip

Always prefer vectorized operations over Python loops for numerical computations. NumPy operations are implemented in C and are typically 10-100x faster.

```

1 import time
2
3 # Large dataset
4 large_array = np.random.rand(1_000_000)
5
6 # Python loop (slow)
7 def python_sum(arr):
8     total = 0.0
9     for value in arr:
10         total += value
11     return total
12
13 # Time comparison
14 start = time.time()
15 py_result = python_sum(large_array)
16 py_time = time.time() - start
17
18 start = time.time()
19 np_result = np.sum(large_array)
20 np_time = time.time() - start
21
22 print(f"Python loop: {py_time:.4f} seconds")
23 print(f"NumPy vectorized: {np_time:.4f} seconds")
24 print(f"Speedup: {py_time/np_time:.1f}x")

```

12.2 Memory Optimization

```

1 # Use appropriate data types
2 small_ints = np.arange(1000, dtype=np.int8)      # 1 byte per element
3 large_ints = np.arange(1000, dtype=np.int64)      # 8 bytes per element
4
5 print(f'int8 memory: {small_ints.nbytes} bytes')
6 print(f'int64 memory: {large_ints.nbytes} bytes')
7
8 # In-place operations to save memory
9 arr = np.random.rand(1000, 1000)
10 arr += 5           # In-place addition (saves memory)
11 arr *= 2           # In-place multiplication
12
13 # Instead of:

```

```

14 # arr = arr + 5    # Creates new array
15 # arr = arr * 2    # Creates new array
16
17 # Pre-allocate arrays when possible
18 result = np.empty(1000) # Faster than appending to lists
19 for i in range(1000):
20     result[i] = expensive_computation(i)

```

13 Common Pitfalls and Debugging

13.1 Frequent Mistakes

Common Pitfall: Unexpected Broadcasting

```

1 # This can create unexpected results
2 a = np.array([[1, 2, 3]])          # Shape: (1, 3)
3 b = np.array([[1], [2], [3]])      # Shape: (3, 1)
4 result = a + b                  # Shape: (3, 3) - often unexpected!
5
6 # Be explicit about your intentions
7 a_explicit = np.array([1, 2, 3])        # Shape: (3,)
8 b_explicit = np.array([1, 2, 3])        # Shape: (3,)
9 result_explicit = a_explicit + b_explicit # Shape: (3,)

```

Common Pitfall: View Modifications

```

1 # Views can cause unexpected side effects
2 matrix = np.arange(12).reshape(3, 4)
3 row = matrix[0]      # This is a view!
4 row[0] = 999         # Modifies the original matrix!
5
6 # Solution: Use copy() when you need independence
7 safe_row = matrix[0].copy()
8 safe_row[0] = 999   # Original matrix unchanged

```

13.2 Debugging Techniques

```

1 def debug_array(arr, name="Array"):
2     """Comprehensive array debugging information"""
3     print(f"\n==== {name} Debug Info ====")
4     print(f"Shape: {arr.shape}")
5     print(f"Dtype: {arr.dtype}")
6     print(f"Size: {arr.size}")
7     print(f"Dimensions: {arr.ndim}")
8     print(f"Memory (bytes): {arr.nbytes}")
9     print(f"Strides: {arr.strides}")
10    print(f"C-contiguous: {arr.flags.c_contiguous}")
11    print(f"F-contiguous: {arr.flags.f_contiguous}")
12    print(f"Owns data: {arr.flags.owndata}")
13
14    if arr.size > 0:
15        print(f"Min: {arr.min()}")
16        print(f"Max: {arr.max()}")
17        print(f"Mean: {arr.mean():.3f}")
18
19    if arr.size <= 20:

```

```

20         print(f"Values:\n{arr}")
21     else:
22         print(f"First 5 values: {arr.flat[:5]}")
23
24 # Usage
25 problematic_array = np.random.rand(3, 4)
26 debug_array(problematic_array, "My Array")

```

14 Advanced Topics

14.1 Custom Universal Functions (ufuncs)

```

1 # Create custom vectorized functions
2 def sigmoid(x):
3     """Sigmoid activation function"""
4     return 1 / (1 + np.exp(-x))
5
6 # Vectorize for NumPy arrays
7 vectorized_sigmoid = np.vectorize(sigmoid)
8
9 # Usage
10 x = np.linspace(-10, 10, 100)
11 y = vectorized_sigmoid(x)
12
13 # Alternative: Use existing NumPy functions
14 def fast_sigmoid(x):
15     """Faster implementation using NumPy"""
16     return 1 / (1 + np.exp(-np.clip(x, -500, 500))) # Prevent overflow

```

14.2 Structured Arrays

```

1 # Define custom data types
2 dtype = np.dtype([
3     ('name', 'U20'),          # Unicode string, max 20 chars
4     ('age', 'i4'),            # 32-bit integer
5     ('salary', 'f8'),         # 64-bit float
6     ('department', 'U10')    # Unicode string, max 10 chars
7 ])
8
9 # Create structured array
10 employees = np.array([
11     ('Alice', 30, 75000.0, 'Engineering'),
12     ('Bob', 25, 50000.0, 'Marketing'),
13     ('Charlie', 35, 85000.0, 'Engineering')
14 ], dtype=dtype)
15
16 # Access fields
17 print(employees['name'])      # All names
18 print(employees['salary'])    # All salaries
19 print(employees[employees['age'] > 30]) # Employees over 30

```

15 Integration with Scientific Python Ecosystem

15.1 With Pandas

```

1 import pandas as pd
2

```

```
3 # NumPy to Pandas
4 data = np.random.randn(100, 4)
5 df = pd.DataFrame(data, columns=[‘A’, ‘B’, ‘C’, ‘D’])
6
7 # Pandas to NumPy
8 numpy_data = df.values           # All data
9 numpy_subset = df[['A', 'B']].values # Specific columns
```

15.2 With Matplotlib

```
1 import matplotlib.pyplot as plt
2
3 # Generate data
4 x = np.linspace(0, 2*np.pi, 100)
5 y1 = np.sin(x)
6 y2 = np.cos(x)
7
8 # Plot
9 plt.figure(figsize=(10, 6))
10 plt.plot(x, y1, label='sin(x)', linewidth=2)
11 plt.plot(x, y2, label='cos(x)', linewidth=2)
12 plt.xlabel('x')
13 plt.ylabel('y')
14 plt.title('Trigonometric Functions')
15 plt.legend()
16 plt.grid(True, alpha=0.3)
17 plt.show()
```

15.3 With SciPy

```
1 from scipy import stats, optimize, integrate
2
3 # Statistical analysis
4 data = np.random.normal(100, 15, 1000)
5 mean_est, std_est = stats.norm.fit(data)
6
7 # Optimization
8 def objective(x):
9     return (x - 3)**2 + 1
10
11 result = optimize.minimize(objective, x0=0)
12 print(f"Minimum at x = {result.x[0]:.3f}")
13
14 # Numerical integration
15 def integrand(x):
16     return np.exp(-x**2)
17
18 integral, error = integrate.quad(integrand, 0, np.inf)
19 print(f"Integral: {integral:.6f}    {error:.2e}")
```

16 Real-World Applications

16.1 Image Processing

```
1 # Create synthetic image
2 def create_image(size=256):
3     x = np.linspace(-2, 2, size)
```

```

4     y = np.linspace(-2, 2, size)
5     X, Y = np.meshgrid(x, y)
6
7     # Create pattern
8     image = np.sin(2*np.pi*X) * np.cos(2*np.pi*Y)
9     image = (image + 1) / 2 * 255 # Normalize to 0-255
10    return image.astype(np.uint8)
11
12 image = create_image()
13
14 # Image operations
15 def image_operations(img):
16     # Horizontal flip
17     flipped = np.fliplr(img)
18
19     # Rotation (90 degrees)
20     rotated = np.rot90(img)
21
22     # Brightness adjustment
23     brighter = np.clip(img + 50, 0, 255)
24
25     # Contrast enhancement
26     enhanced = np.clip(img * 1.5, 0, 255)
27
28     return flipped, rotated, brighter, enhanced
29
30 processed_images = image_operations(image)

```

16.2 Data Analysis Example

```

1 # Sales data analysis
2 np.random.seed(42)
3
4 # Generate synthetic sales data
5 months = 12
6 products = 5
7 sales_data = np.random.normal(1000, 200, (months, products))
8 sales_data = np.maximum(sales_data, 0) # Ensure non-negative
9
10 product_names = ['Product A', 'Product B', 'Product C', 'Product D', 'Product E',
11                   '']
11 month_names = ['Jan', 'Feb', 'Mar', 'Apr', 'May', 'Jun',
12                  'Jul', 'Aug', 'Sep', 'Oct', 'Nov', 'Dec']
13
14 # Analysis
15 total_sales = np.sum(sales_data)
16 monthly_totals = np.sum(sales_data, axis=1)
17 product_totals = np.sum(sales_data, axis=0)
18
19 # Find best and worst performers
20 best_month_idx = np.argmax(monthly_totals)
21 worst_month_idx = np.argmin(monthly_totals)
22 best_product_idx = np.argmax(product_totals)
23
24 print(f"Best month: {month_names[best_month_idx]} ({monthly_totals[
25           best_month_idx]:.2f})")
25 print(f"Worst month: {month_names[worst_month_idx]} ({monthly_totals[
26           worst_month_idx]:.2f})")
26 print(f"Best product: {product_names[best_product_idx]} ({product_totals[
27           best_product_idx]:.2f})")
27
28 # Growth analysis

```

```
29 monthly_growth = np.diff(monthly_totals) / monthly_totals[:-1] * 100
30 average_growth = np.mean(monthly_growth)
31 print(f"Average monthly growth: {average_growth:.2f}%")
```

17 Practice Exercises

17.1 Beginner Exercises

1. Create a 5×5 matrix with values ranging from 0 to 24
2. Extract the border elements of the matrix
3. Replace all odd numbers with -1
4. Calculate the sum of each row and column

17.2 Intermediate Exercises

1. Implement a function to normalize a dataset (zero mean, unit variance)
2. Create a function to find the k-nearest neighbors in a 2D dataset
3. Implement moving average smoothing for a 1D signal
4. Calculate correlation matrix for a multi-variate dataset

17.3 Advanced Exercises

1. Implement Principal Component Analysis (PCA) from scratch
2. Create a function for convolution of two signals
3. Implement k-means clustering algorithm
4. Design a neural network layer using only NumPy

18 Best Practices and Guidelines

NumPy Best Practices

1. **Always use vectorized operations** instead of Python loops
2. **Choose appropriate data types** to minimize memory usage
3. **Be aware of views vs copies** to avoid unexpected behavior
4. **Use axis parameter wisely** for multi-dimensional operations
5. **Preallocate arrays** when size is known in advance
6. **Use in-place operations** when possible to save memory
7. **Handle NaN and infinity values** appropriately
8. **Profile your code** to identify bottlenecks

19 Conclusion

NumPy is the foundation of scientific computing in Python, providing efficient array operations, mathematical functions, and seamless integration with the broader scientific Python ecosystem. Mastering NumPy concepts like broadcasting, views vs copies, and vectorization will significantly improve your data analysis and scientific computing capabilities.

Key takeaways:

- NumPy arrays are faster and more memory-efficient than Python lists
- Understanding broadcasting enables elegant solutions to complex problems
- Views and copies behavior affects both performance and correctness
- Vectorized operations should always be preferred over loops
- NumPy integrates seamlessly with pandas, matplotlib, SciPy, and scikit-learn

20 Further Reading

- Official NumPy Documentation: <https://numpy.org/doc/>
 - NumPy User Guide: <https://numpy.org/doc/stable/user/>
 - SciPy Lecture Notes: <https://scipy-lectures.org/>
 - Python Data Science Handbook: <https://jakevdp.github.io/PythonDataScienceHandbook/>
-