



Part 3:

**Data Science and Machine Learning
With Amin.**

**Foundations of Numerical Analysis –
Numpy.**
















Amin Hydar Ali



Table of Contents

What NumPy is and Why It's Used in AI/ML.....	6
Installing & Importing NumPy.....	7
NumPy Arrays vs Python Lists.....	7
✓ Python List.....	8
✓ NumPy Array.....	8
⚡ Speed Comparison.....	8
Step 1: The Task 📝.....	8
Step 2: Doing it with a Python List 🐢.....	9
(Python List in Action).....	9
Step 3: Doing it with NumPy ⚡.....	9
(NumPy in Action).....	10
Step 4: Comparing the Speeds 🕒.....	10
Creating Arrays.....	10
a) From a Python list.....	10
b) Zeros and Ones.....	11
c) arange (like Python's range).....	11
d) linspace.....	11
e) Random Arrays.....	11
1. np.random.rand(3, 3).....	11
2. np.random.randn(3, 3).....	11
3. np.random.randint(0, 10, (3, 3)).....	12
f) Random Seed (Reproducibility).....	12
Array Attributes.....	12
a) .shape.....	13
b) .dtype.....	13
c) .ndim.....	13
d) .size.....	13
e) .itemsize (extra but useful).....	13
The Big Picture (Why This Matters for AI/ML).....	13
NumPy Indexing, Slicing, and Broadcasting : The Magic of Arrays.....	14
Indexing Arrays.....	14
a) 1D Arrays.....	14
b) 2D Arrays.....	15
c) 3D Arrays.....	15
Boolean Indexing & Masking.....	15
Fancy Indexing.....	16
Slicing Arrays.....	16
Breakdown of what is happening:.....	16
Step 1: Concept.....	16
Step 2: Concept for 2D slicing.....	17
Broadcasting : The Heart of NumPy.....	18
a) Adding a scalar.....	18
b) Adding arrays with compatible shapes.....	18
Breakdown of why it is compatible:.....	18

Step 1: Understand the shapes.....	18
Step 2: How NumPy broadcasts.....	18
reshape, transpose, newaxis.....	19
Adding np.newaxis.....	19
Shape After Transformation.....	20
Visualization.....	20
Mini-Project: Grayscale Image Inversion 🤖.....	20
Image:.....	21
Explanation:.....	21
1. The “image” is just numbers (matrix form).....	21
2. Displaying it with matplotlib.....	22
3. Inverting the image using broadcasting.....	22
4. Visualizing inverted.....	22
Numerical Operations in NumPy.....	23
1. Element-wise Operations.....	23
2. Matrix Multiplication.....	23
First: Why do we even need "matrix multiplication"?.....	23
1. The Three Faces of Matrix Multiplication in NumPy.....	24
1. @ (modern operator).....	24
2. np.dot().....	24
3. np.matmul().....	24
Difference.....	25
Example 1: Weighted Grades.....	25
Why Three Options?.....	25
3. 📊 Aggregations (Summaries).....	26
What Are Aggregations?.....	26
The Functions.....	26
Example in Code.....	27
Why Does This Matter in AI / Data Science?.....	27
Relatable Analogy (Friends’ Heights).....	27
4. Cumulative Operations.....	28
5. Axis Operations (Row vs Column).....	29
6. Clipping Values.....	29
Mini Project: Manual Linear Regression Step.....	29
Step by Step Break Down:.....	30
Arrays x and y.....	30
Initial weights w and b.....	30
Prediction: $y_{pred} = w * x + b$	30
Compute error: $error = ((y - y_{pred})^{**2}).mean()$	31
Gradients.....	31
Update weights.....	31
Stacking Arrays.....	32
Example:.....	32
Splitting Arrays.....	32
Example:.....	33
Repeating & Tiling.....	33
Example:.....	33
Flattening Arrays.....	33
Example:.....	33

Copy vs View.....	33
Example:.....	33
Sorting.....	34
Example:.....	34
Unique Values.....	34
Example:.....	34
Index Tricks.....	34
np.ix_ → cross product of indices.....	35
That prints:.....	35
Explanation.....	35
The concrete array we'll use.....	35
Step 1 - What np.ix_([0,2], [1,2]) actually produces.....	36
Step 2 - Why those shapes? (so they can combine).....	36
Step 3 - Using it to select the submatrix.....	37
Step 4 - Contrast with arr[[0,2], [1,2]] (important!).....	37
Short visual summary.....	37
Final checklist (so you remember which to use).....	38
Related helpers (brief).....	38
np.meshgrid → make coordinate grids.....	38
Randomness & Probability.....	38
1. Random Sampling.....	39
 Uniform Distribution.....	39
 Normal Distribution (a.k.a. Bell Curve).....	39
 Binomial Distribution.....	39
 2. np.random.choice - Picking Random Items.....	40
 3. Seeds & Reproducibility.....	40
 4. Shuffling & Permutations.....	40
 5. Probability Distributions.....	41
 Why This Matters (Even for Beginners).....	41
Practical Applications for AI/ML.....	41
 1. Image Representation as NumPy Arrays.....	41
 2. Normalization & Standardization.....	42
3. One-Hot Encoding.....	42
 4. Dataset Shuffling & Batching.....	43
 5. Convolution Basics (Before Deep Learning Frameworks).....	43
6. Implementing Loss Functions (MSE, Cross-Entropy).....	44
 Summary.....	44
Connecting to AI Frameworks (Beginner Breakdown).....	45
 1. How NumPy Underlies TensorFlow & PyTorch.....	45
2. Converting Between NumPy Arrays and Tensors.....	46
Using PyTorch.....	46
Using TensorFlow.....	46
3. When to Use NumPy vs Pandas vs SciPy.....	46
Example:.....	46
4. Why NumPy Operations Look Like Tensor Operations.....	47
Example:.....	47
5. Bonus (Optional but Cool): What Are Tensors Really?.....	47
 Summary Table.....	48

The Source Code for this is available in my Github repository:

<https://github.com/alaminydar/Data-Science-and-Machine-Learning-With-Amin>

What NumPy is and Why It's Used in AI/ML

Imagine you want to build an AI system. At its heart, AI is just **math applied to data**:

- Adding, multiplying, dividing, taking averages.
- Working with large datasets (millions of numbers).
- Representing vectors (like word embeddings, pixel values in an image).
- Working with matrices (like weights in a neural network).

If you tried to do this with **pure Python lists**, you'd run into problems:

- **Slow performance** (Python loops run one step at a time).
- **High memory usage** (Python objects carry overhead).
- **No vectorization** (you can't just multiply a list by another list directly in a mathematical sense).

👉 That's why **NumPy exists**: it gives you a way to store and manipulate **numerical data efficiently**, almost like C or Fortran speed, but with Python's ease of use.

In AI/ML, every neural network weight, every image, every dataset feature is basically stored in a NumPy array (or something built on top of it, like PyTorch tensors).

So, if you understand NumPy deeply, you understand the **language that AI frameworks speak**.

Installing & Importing NumPy

To use NumPy, you first need to install it:

```
1. pip install numpy
```

Then in your Python code:

```
1. import numpy as np
```

Why as **np**?

Because you'll be calling NumPy functions a lot, and typing `numpy` everywhere would be too long. **np** has become the *standard alias* in the Python community.

So now whenever you see `np.array()` or `np.random.rand()`, just remember it's NumPy.

NumPy Arrays vs Python Lists

This is one of the most important shifts. Let's compare:

✓ Python List

```
1. py_list = [1, 2, 3, 4, 5]
```

- **Flexible:** can hold numbers, strings, objects, even mixed types.
- But each element is a full Python object → extra memory overhead.
- Operations are slow, because Python handles one element at a time.

✓ NumPy Array

```
1. np_array = np.array([1, 2, 3, 4, 5])
```

- **Homogeneous:** all elements must be the same type (e.g., all floats, all ints).
- Stored in **contiguous memory blocks**, like in C, so it's extremely efficient.
- Supports **vectorization:** you can apply operations to the whole array at once without explicit loops.

Vectorization is the magic: instead of writing a loop to add 2 to every number, you just write:

```
1. np_array + 2
```

and NumPy applies it to the whole array internally, in compiled C speed.

⚡ Speed Comparison

```
1. import time
2. import numpy as np
3.
4. # Python list
5. lst = list(range(1000000))
6. start = time.time()
7. lst = [x + 2 for x in lst]
8. print("Python list:", time.time()-start)
9.
10. # NumPy array
11. arr = np.arange(1000000)
12. start = time.time()
13. arr = arr + 2
14. print("NumPy array:", time.time()-start)
```

You'll see NumPy is many times faster. Let go in depth on the code analysis

When people first hear about NumPy, they often wonder:

“Why can't I just use normal Python lists?”

Let's break it down slowly using the above code as example.

Step 1: The Task

Suppose we have **1 million numbers** and we want to **add 2** to each number.

This is a very common operation in AI/ML we often need to adjust or transform entire datasets.

Step 2: Doing it with a Python List 🐢

```
1. import time
2.
3. # Create a Python list with 1 million numbers
4. lst = list(range(1000000))
5.
6. # Start the timer
7. start = time.time()
8.
9. # Add 2 to each element (using a loop under the hood)
10. lst = [x + 2 for x in lst]
11.
12. # Stop the timer
13. print("Python list:", time.time() - start)
```

Breakdown of what's happening here:

- Python has to **go through each element one by one**.
- For every number, it takes it out, adds 2, then puts it back in the new list.
- That's like having a **million boxes in a row** and opening each box, adding 2 to the number inside, then closing the box again.

It works... but it's slow.

(Python List in Action)

```
1. Before: [0, 1, 2, 3, 4, 5, ...]
2.         ↑ ↑ ↑ ↑ ↑
3.         | | | | |
4. loop through each, add 2
5.
6. After: [2, 3, 4, 5, 6, 7, ...]
```

Step 3: Doing it with NumPy ⚡

```
1. import numpy as np
2.
3. # Create a NumPy array with 1 million numbers
4. arr = np.arange(1000000)
5.
6. # Start the timer
7. start = time.time()
8.
9. # Add 2 to the entire array (vectorized operation)
10. arr = arr + 2
11.
12. # Stop the timer
13. print("NumPy array:", time.time() - start)
```

Breakdown of what's happening here:

- NumPy doesn't loop over each element like Python.

- Instead, it uses **vectorization**: it sends the whole operation (+2) directly to optimized C code under the hood.
- Think of it like telling a factory machine:
“Add 2 to every box in this row, all at once.”
- No more opening and closing boxes one by one.

This makes it **blazingly fast**.

(NumPy in Action)

```
1. Before: [0, 1, 2, 3, 4, 5, ...]
2. Apply "+2" to all at once (machine-like efficiency ⚡)
3.
4. After: [2, 3, 4, 5, 6, 7, ...]
```

Step 4: Comparing the Speeds

Typical output:

```
1. Python list: 0.12 seconds
2. NumPy array: 0.004 seconds
```

Even though both do the **same job**, NumPy is usually **30x+ faster** (sometimes even 100x depending on the task).

Creating Arrays

Creating arrays is the first step to fluency. NumPy gives you many tools:

a) From a Python list

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

Converts a Python list into a NumPy array.

b) Zeros and Ones

```
1. np.zeros((3, 3)) # 3x3 matrix of zeros
2. np.ones((2, 4)) # 2x4 matrix of ones
```

Useful for initializing weights or empty data structures in ML.

c) arange (like Python's range)

```
1. np.arange(0, 10, 2) # [0, 2, 4, 6, 8]
```

Creates evenly spaced values with a step size.

d) linspace

```
1. np.linspace(0, 1, 5) # [0. , 0.25, 0.5 , 0.75, 1. ]
```

Creates evenly spaced numbers between start and stop. Very useful in plotting or when you need normalized values.

e) Random Arrays

ML thrives on randomness (weight initialization, shuffling data, sampling).

```
1. np.random.rand(3, 3) # Uniform distribution between 0 and 1
2. np.random.randn(3, 3) # Normal distribution (mean=0, std=1)
3. np.random.randint(0, 10, (3, 3)) # Integers between 0 and 10
```

When we use **np.random**, we're basically asking the computer to “roll dice” for us. But the type of dice we roll depends on the function we choose. Let's break it down:

1. np.random.rand(3, 3)

Think of it as rolling a **special die** that always gives you a number between **0 and 1**.

- Each roll is equally likely just like spinning a fair wheel from 0.0 → 1.0.
- You ask for a 3x3 grid, so you get 9 rolls arranged in a little table.

 Example Output:

```
1. [[0.12, 0.75, 0.55],
2. [0.89, 0.34, 0.01],
3. [0.67, 0.43, 0.92]]
```

All numbers are between **0 and 1**.

2. np.random.randn(3, 3)

Now imagine rolling a **different die**: this one produces numbers centered around **0**.

- Most results will be close to 0, but sometimes you'll get **negative** or **positive** numbers farther away.
- This follows the **bell curve** (normal distribution, mean = 0, std = 1).

 Example Output:

```
1. [[-0.45, 0.11, 1.23],  
2. [ 0.88, -1.02, -0.33],  
3. [ 0.09, -0.72, 0.55]]
```

Notice values spread around **0**, not limited to 0–1.

3. `np.random.randint(0, 10, (3, 3))`

Finally, this is like rolling a **regular dice**, but you can **choose the range**.

- Here: numbers between **0 and 9** (10 is excluded).
- Again, you get a 3x3 table filled with integers.

 Example Output:

```
1. [[2, 7, 1],  
2. [0, 9, 5],  
3. [3, 6, 4]]
```

f) Random Seed (Reproducibility)

If you don't set a seed, NumPy will generate different random numbers each time you run your code. But in ML experiments, we need **reproducibility** (to compare results).

```
1. np.random.seed(42)  
2. print(np.random.rand(3))
```

Run this multiple times → you'll always get the same numbers.

This is **critical in AI research**. Imagine you train two models and they give different results, just because random numbers changed. Setting the seed makes experiments consistent.

Array Attributes

Once you have arrays, you need to understand their properties. Think of it like metadata about your data.

Let's create an example:

```
1. arr = np.random.rand(3, 4)
```

This creates a 3x4 array. Now:

a) .shape

```
1. arr.shape
```

- Returns the dimensions of the array.
- (3, 4) → 3 rows, 4 columns.
- In AI, this is how we know the **input size, output size, batch size, etc.**

b) .dtype

```
1. arr.dtype
```

- Data type of the elements: int32, float64, etc.
- Why important? Neural networks usually require **floats** for computation.
- Using wrong dtype (like int) could cause errors in training.

c) .ndim

```
1. arr.ndim
```

- Number of dimensions.
- 1D = vector, 2D = matrix, 3D = tensor (e.g., images with color channels).
- Deep learning often uses **high-dimensional tensors**.

d) .size

```
1. arr.size
```

- Total number of elements in the array.
- Example: (3,4) shape → size = 12.

e) .itemsize (extra but useful)

```
1. arr.itemsize
```

- Size in bytes of each element.
- Helps you estimate memory usage (important when dealing with huge datasets).

The Big Picture (Why This Matters for AI/ML)

- **Arrays are the language of AI.**
Every dataset, every image, every word embedding, every layer weight is stored as an array.

- **Speed matters.**
If you're training a model with millions of parameters, you cannot afford slow Python loops. NumPy gives you speed close to C.
- **Shapes and dimensions are critical.**
A single mistake in `.shape` can break your whole neural network. Understanding **shape**, **ndim**, and **broadcasting** will make you a much stronger ML engineer.
- **Randomness is part of intelligence.**
Without random initialization and random sampling, most modern AI algorithms wouldn't work.

NumPy Indexing, Slicing, and Broadcasting : The Magic of Arrays

NumPy isn't just about storing numbers it's about **accessing, transforming, and combining them efficiently.**

Indexing, slicing, and broadcasting are the tools that make NumPy **so powerful**, especially for AI and data science.

Indexing Arrays

Indexing is how we pick out elements from an array. Think of it as pointing at a specific box in a huge warehouse.

a) 1D Arrays

```
1. import numpy as np
2.
3. arr = np.array([10, 20, 30, 40, 50])
4. print(arr[0]) # 10 → first element
5. print(arr[-1]) # 50 → last element
```

- `arr[0]` → the first "box"
- `arr[-1]` → the last "box" (Python lets you count backwards)

Visualization:

```
1. Index: 0  1  2  3  4
2. Value:10 20 30 40 50
```

Think of arrows pointing to the number you want.

b) 2D Arrays

```
1. arr2d = np.array([[1, 2, 3],
2.                [4, 5, 6],
3.                [7, 8, 9]])
4. print(arr2d[0, 1]) # 2 → row 0, col 1
5. print(arr2d[2, -1]) # 9 → row 2, last column
```

- `arr2d[row, col]` → pick a cell in a grid
- Negative indices → count from the bottom/right

Visualization:

```
1. Col Row 0 1 2
2.
3. 0    1 2 3
4. 1    4 5 6
5. 2    7 8 9
```

Arrows point to the selected element.

c) 3D Arrays

3D arrays = “cubes” of numbers (think video frames or RGB images).

```
1. arr3d = np.arange(27).reshape(3,3,3)
2. print(arr3d[0, 1, 2]) # pick element at depth 0, row 1, col 2
```

- First index → which “slice” (depth)
- Second → row
- Third → column

Visualization:

- Imagine layers of 3x3 grids stacked like floors in a building.

Boolean Indexing & Masking

Sometimes, you don't want a single element, but all **elements satisfying a condition**.

```
1. arr = np.array([10, 15, 20, 25, 30])
2. mask = arr > 20
3. print(mask) # [False False False True True]
4. print(arr[mask]) # [25 30]
```

- Create a **mask** (True/False) for every element
- Use it to **filter elements quickly**

Visualization:

```
1. arr: 10 15 20 25 30
2. mask: F F F T T
3. result:      25 30
```

This is **super useful in AI**: e.g., selecting all pixels brighter than a threshold.

Fancy Indexing

Fancy indexing allows you to pick **specific arbitrary elements** using a list of indices:

```
1. arr = np.array([10, 20, 30, 40, 50])
2. indices = [0, 2, 4]
3. print(arr[indices]) # [10 30 50]
```

Instead of consecutive slices, you can **jump around** to specific elements.

Slicing Arrays

Slicing lets you select **ranges of elements** using [start:stop:step].

```
1. arr = np.array([10, 20, 30, 40, 50, 60])
2. print(arr[::2]) # every other element → [10 30 50]
```

2D Slicing Example:

```
1. arr2d = np.arange(1,10).reshape(3,3)
2. print(arr2d[:, 1]) # all rows, column 1 → [2 5 8]
3. print(arr2d[1, :]) # row 1, all columns → [4 5 6]
```

Breakdown of what is happening:

Step 1: Concept

- `arr[start:stop:step]` → pick elements from start to stop-1, stepping by step
- `::2` → start to end, step 2 → every other element

Visualization:

```
1. Index: 0 1 2 3 4 5
2. Value: 10 20 30 40 50 60
3. Pick: * * *
4. Result: 10 30 50
5.
6.
7. Index: 0 1 2 3 4 5 6 7 8
8. Value: 10 20 30 40 50 60 70 80 90
9. Pick: * * * * *
10. Result: 10 30 50 70 90
```

- Start at index 0 → pick 10
- Jump 2 → index 2 → pick 30
- Jump 2 → index 4 → pick 50

```
1. arr2d = np.arange(1,10).reshape(3,3)
2. print(arr2d[:, 1]) # all rows, column 1 → [2 5 8]
3. print(arr2d[1, :]) # row 1, all columns → [4 5 6]
```

Step 2: Concept for 2D slicing

arr2d shape (3,3):

```
1. Col Row 0 1 2
2.
3. 0    1 2 3
4. 1    4 5 6
5. 2    7 8 9
```

arr2d[:,1]

- : → all rows
- 1 → column 1

Pick column 1 from every row: 2,5,8

arr2d[1,:] →

- 1 → row 1
- : → all columns

Pick row 1 completely: 4,5,6

Visualization:

```
1. arr2d[:,1] -> pick vertical slice (column):
2. [[1 2 3]  2
3. [4 5 6]  → 5
4. [7 8 9]  8]
5.
6. arr2d[1,:] -> pick horizontal slice (row):
7. [4 5 6]
```

- : → “pick all elements in this direction”
- Step (::2) → “jump every N elements”

Broadcasting : The Heart of NumPy

Broadcasting allows NumPy to **perform operations on arrays of different shapes** without explicit loops.

a) Adding a scalar

```
1. arr = np.array([1,2,3])
2. print(arr + 5) # [6 7 8]
```

NumPy “stretches” the scalar 5 to match the array length.

b) Adding arrays with compatible shapes

```
1. A = np.array([[1,2,3],
2.           [4,5,6]])
3. B = np.array([10,20,30])
4. print(A + B)
5. Output:
6. [[11 22 33]
7.  [14 25 36]]
```

Breakdown of why it is compatible:

Step 1: Understand the shapes

- A shape $\rightarrow (2, 3) \rightarrow$ 2 rows, 3 columns

```
1. Row 0: 1 2 3
2. Row 1: 4 5 6
```

- B shape $\rightarrow (3,) \rightarrow$ 1D array of length 3

10 20 30

Step 2: How NumPy broadcasts

Rule: NumPy compares shapes **from the right**. If dimensions match or one is 1 \rightarrow compatible.

- A.shape = (2,3)
- B.shape = (3,) \rightarrow treated as (1,3) \rightarrow can stretch along row dimension.

Visualization:

```
1. A:      B (broadcasted):
2. [[1 2 3]  [[10 20 30]
3. [4 5 6]  [10 20 30]]
```

- NumPy “stretches” B along rows \rightarrow now both arrays have shape (2,3)
- Then element-wise addition happens:

Result:

```
1. [[1+10, 2+20, 3+30] → [11, 22, 33]
2. [4+10, 5+20, 6+30] → [14, 25, 36]
```

- B is “copied” virtually along the missing dimension
- No actual memory copy happens, NumPy handles it efficiently under the hood.

Takeaway: shapes are compatible if **from the right: sizes match or 1**.

- Here, (2,3) and (1,3) → broadcastable → (2,3) result.
- B is “broadcast” over each row of A.
- No loops, memory-efficient, extremely fast.

reshape, transpose, newaxis

- **reshape:** changes the array’s shape without changing data.

```
1. arr = np.arange(6)
2. print(arr.reshape(2,3)) # 2 rows, 3 cols
```

- **transpose:** flips rows ↔ columns

```
1. arr2d = np.array([[1,2,3],[4,5,6]])
2. print(arr2d.T)
```

- **newaxis:** adds a dimension

```
1. arr = np.array([1,2,3])
2. print(arr[:, np.newaxis].shape) # (3,1)
```

Breakdown of how it works:

- This is a **1D array** with 3 elements.
- Shape (3,) means:
 - 3 elements,
 - only **one axis** (like a line).

Visualization:

```
1. Index: 0 1 2
2. Value: 1 2 3
```

Adding np.newaxis

```
1. arr[:, np.newaxis]
```

Here's what happens:

- `:` → means **take all elements** from the original array.
- `np.newaxis` → means **add a new dimension (axis)** at that position.

So we're turning a (3,) array into a **column vector** (3,1).

Shape After Transformation

```
1. print(arr[:, np.newaxis].shape)
2. # (3,1)
```

Why (3,1)?

- **First dimension:** 3 rows (from the original 3 elements).
- **Second dimension:** 1 column (the new axis we added).

Visualization

```
1. Before ((3,)):
2. [1 2 3] → just a flat row
3. After ((3,1)):
4. [[1]
5. [2]
6. [3]]
```

Think of it like reshaping the array into a column vector.

- These tools allow arrays to **match shapes for broadcasting**, manipulate matrices, and prepare data for AI models.

Mini-Project: Grayscale Image Inversion

Goal: invert a grayscale image (0→255, 255→0) using slicing & broadcasting.

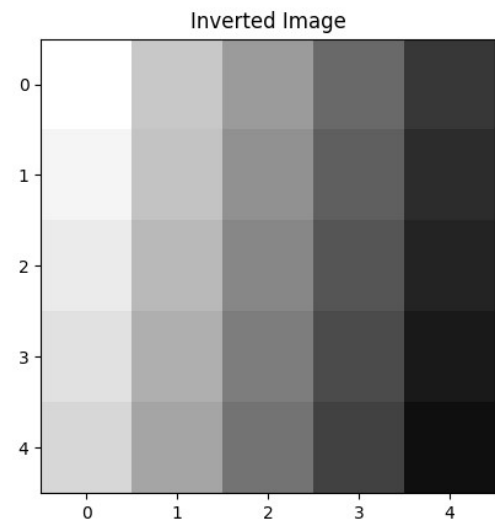

```
1. import numpy as np
2. import matplotlib.pyplot as plt
3.
4. # Simulate a 5x5 grayscale image
5. img = np.array([[0,50,100,150,200],
6.                [10,60,110,160,210],
7.                [20,70,120,170,220],
8.                [30,80,130,180,230],
9.                [40,90,140,190,240]])
10.
11. # Visualize original image
12. plt.imshow(img, cmap='gray', vmin=0, vmax=255)
13. plt.title("Original Image")
14. plt.show()
```

```

15.
16. # Invert image using broadcasting
17. inverted = 255 - img
18.
19. # Visualize inverted image
20. plt.imshow(inverted, cmap='gray', vmin=0, vmax=255)
21. plt.title("Inverted Image")
22. plt.show()

```

Image:



Explanation:

1. The “image” is just numbers (matrix form)

```

1. img = np.array([[0,50,100,150,200],
2.                [10,60,110,160,210],
3.                [20,70,120,170,220],
4.                [30,80,130,180,230],
5.                [40,90,140,190,240]])

```

Here, `img` is a 5×5 **NumPy array**. Each value represents **pixel intensity** in a grayscale image.

- Small numbers (close to 0) → **black/dark pixels**
- Large numbers (close to 255) → **white/bright pixels**
- Numbers in between → shades of gray

👉 Think of it like a chessboard, but instead of black and white squares, each square has a "brightness number."

2. Displaying it with matplotlib

```
1. plt.imshow(img, cmap='gray', vmin=0, vmax=255)
2. plt.title("Original Image")
3. plt.show()
```

- `plt.imshow(img, ...)`: Displays the 2D array as an image.
- `cmap='gray'`: Uses **grayscale colormap** (dark → bright).
- `vmin=0, vmax=255`: Fixes the brightness range so 0 = pure black, 255 = pure white (helps keep consistency).
- `plt.title("Original Image")`: Adds a title.

👉 This lets us **see numbers as shades of light**, turning the matrix into something visually meaningful.

3. Inverting the image using broadcasting

```
1. inverted = 255 - img
```

Here's the magic:

- For each pixel value, we subtract it from 255.
- If a pixel was bright (say 200), it becomes dark (55).
- If a pixel was dark (say 10), it becomes bright (245).

👉 This is like a **photo negative** light becomes dark, dark becomes light.

And NumPy's broadcasting makes it effortless: no loops, it applies to the whole matrix at once.

4. Visualizing inverted

```
1. plt.imshow(inverted, cmap='gray', vmin=0, vmax=255)
2. plt.title("Inverted Image")
3. plt.show()
```

This shows the inverted picture. Now the image looks like a photographic **negative filter**.

- Each pixel value is just a number from 0–255.
- `255 - img` **broadcasts 255** to every pixel → subtraction happens **everywhere at once**.
- This is much faster than looping pixel by pixel.

Visualization:

- You can see the **original image** → **inverted image**, making indexing + broadcasting *real and tangible*.

✅ We'll come to **Matplotlib** later to *visualize how the line fits the data over time*. For now, just know NumPy gives us the tools to **do the math that powers machine learning, data science, and scientific computing**.

Numerical Operations in NumPy

Think of NumPy arrays as **supercharged Excel spreadsheets** where each cell is a number. The difference is that instead of painfully writing formulas cell by cell, NumPy can apply operations **to entire rows, columns, or even the whole sheet at once** fast and memory efficient.

We'll break this into categories:


1. Element-wise Operations

- **Definition:** Operations happen *element by element* across arrays of the same shape.
- **Symbols:** +, -, *, /, **

Example:

```
1. import numpy as np
2.
3. a = np.array([1,2,3])
4. b = np.array([10,20,30])
5.
6. print(a + b) # [11, 22, 33]
7. print(a * b) # [10, 40, 90]
8. print(a ** 2) # [1, 4, 9]
```

Imagine you and your friend have shopping lists. If you both list the number of apples, bananas, and oranges, then $a+b$ is like adding your lists item by item, the total fruit for each type.

 **Key point:** Python list can't do this directly ($[1,2,3] * [10,20,30]$ would error). NumPy makes it natural.

2. Matrix Multiplication

- **Symbols:**
 - $@$ → Modern operator for matrix multiplication
 - `np.dot()` → Works for both dot product and matrix multiplication
 - `np.matmul()` → Specifically for matrix multiplication

First: Why do we even need "matrix multiplication"?

Imagine you have:

- A list of students (rows)
- Each student's scores in subjects (columns)

Now, you want to compute each student's **weighted average** (say, Math counts 50%, English 30%, Science 20%).

That's exactly what **matrix multiplication** does:

It combines **rows of one matrix** with **columns of another matrix**, following certain rules.

1. The Three Faces of Matrix Multiplication in NumPy

1. @ (modern operator)

This is just a **Python shorthand** for matrix multiplication.

It makes code look **clean and mathematical**.

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

✓ Think of this like **normal algebra**: $A \times B$ in textbooks = $A @ B$ in Python.

2. np.dot()

Historically older.

- If both inputs are **1D arrays**, it does the **dot product**.
(e.g. $[1,2,3] \cdot [4,5,6] = 1*4 + 2*5 + 3*6 = 32$)
- If inputs are **2D arrays**, it does **matrix multiplication**.

So it's a bit **overloaded** (does two things depending on shape).

```
1. a = np.array([1,2,3])
2. b = np.array([4,5,6])
3. print(np.dot(a, b)) # dot product → scalar 32
4.
5. print(np.dot(A, B)) # matrix multiplication
```

3. np.matmul()

This is the **clean, strict one**.

It **only does matrix multiplication**, even for higher dimensions.

Useful in ML when multiplying batches of data (e.g. multiplying many input vectors by weights at once).

```
1. print(np.matmul(A, B)) # same as A @ B
```

Difference

Operation	Shape	Result	Analogy
dot (1D)	$(n,) \cdot (n,)$	scalar	How similar are two directions? (cosine similarity idea)
matmul / @	$(m,n) \times (n,p)$	(m,p)	Combining features with weights (AI, physics, graphics)

Example 1: Weighted Grades

```
1. scores = np.array([[80, 70, 90], # Student 1
2.                 [60, 85, 75]]) # Student 2
3.
4. weights = np.array([[0.5], # Math weight
5.                   [0.3], # English weight
6.                   [0.2]]) # Science weight
7.
8. # Matrix multiplication: scores @ weights
9. final = scores @ weights
10. print(final)
11. Student 1 final = 80*0.5 + 70*0.3 + 90*0.2
12. Student 2 final = 60*0.5 + 85*0.3 + 75*0.2
```

Why Three Options?

- **@** → cleanest, most modern (use this in new code).
- **dot** → legacy, but still common in textbooks and tutorials.
- **matmul** → strict, batch-friendly, used under the hood in deep learning frameworks like TensorFlow/PyTorch.

Takeaway:

Matrix multiplication is about **combining information** rows meeting columns.

- **@** is the clean operator.
- **dot** does double duty (dot product or matrix multiply).
- **matmul** is strict matrix multiplication, good for higher dimensions.

Another Example:

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

```
6. print(A @ B)
7. Output:
8. [[19 22]
9. [43 50]]
```

3. Aggregations (Summaries)

Functions: `sum`, `mean`, `std`, `var`, `min`, `max`, `argmin`, `argmax`

What Are Aggregations?

Aggregation = “squash down” data into a single meaningful number.

Instead of keeping track of 1000 values, you ask: “What’s the average? What’s the biggest? How spread out are they?”

Think of it like **summarizing a whole book into one line**.

The Functions

1. `sum()` → Total

- `arr.sum()` → adds up everything.
- **Analogy:** Imagine passing around a basket to collect everyone’s money in class. At the end, sum is the total inside the basket.

2. `mean()` → Average

- `arr.mean()` → total ÷ number of elements.
- **Analogy:** You and 4 friends go to dinner, spend \$15k in total. Mean = \$3k each (if split fairly).

3. `std()` → Standard Deviation

- Measures how “spread out” the numbers are.
- Small std → values are tightly packed.
- Large std → values are scattered.
- **Analogy:** If your friends are all about the same height (± 2 cm), std is small. If one guy is 7ft tall, std explodes.

4. `var()` → Variance

- Square of std. (Less intuitive, but easier for formulas).
- **Analogy:** Like measuring “spread” but in squared units.

5. `min()` / `max()`

- Smallest and largest values.
- **Analogy:** Shortest vs tallest friend.

6. `argmin()` / `argmax()`

- Index of the min/max element.
- **Analogy:** “Who is the tallest?” Instead of giving you height, it points at the person’s position in the lineup.

Example in Code

```
1. import numpy as np
2.
3. arr = np.array([1,2,3,4,5])
4.
5. print(arr.sum()) # 15
6. print(arr.mean()) # 3.0
7. print(arr.std()) # 1.414...
8. print(arr.var()) # 2.0
9. print(arr.min()) # 1
10. print(arr.max()) # 5
11. print(arr.argmax()) # 4 (index of value 5)
12. print(arr.argmin()) # 0 (index of value 1)
```

Why Does This Matter in AI / Data Science?

- **Statistics:** Every dataset starts with summaries (mean income, max score, etc.).
- **AI training:** Models calculate `loss.mean()` (average error).
- **Finance:** Sum of profits, max loss, std = volatility.
- **Sports analytics:** Who’s the top scorer (`argmax`)? What’s the team’s average score (mean)?

Relatable Analogy (Friends’ Heights)

Suppose your friends’ heights = [170, 175, 168, 180, 190]

- **sum** → total combined height = 883 cm
- **mean** → average height = 176.6 cm
- **std** → how varied they are (if you’re all around same height, std small, but if one is a giant, std big).
- **max** → tallest = 190
- **argmax** → index of tallest = 4 (the 5th friend)

It’s like asking:

- “What’s the total height?”
- “What’s the average height?”
- “Who’s the tallest?”
- “How much do we vary?”

4. Cumulative Operations

- **Functions:** `cumsum`, `cumprod`
- They give you a “running total” instead of just the final answer.

Example:

```
1. arr = np.array([1,2,3,4])
2. print(arr.cumsum()) # [1, 3, 6, 10]
3. print(arr.cumprod()) # [1, 2, 6, 24]
```

Relatable analogy:

- **cumsum** = keeping track of your bank balance after every deposit.

Concept: It adds elements one by one **from left to right**, keeping a running total.

Step by step:

- **First element:** 1 → cumulative sum = 1

- **Second element:** 1 + 2 = 3

- **Third element:** 3 + 3 = 6

- **Fourth element:** 6 + 4 = 10

Output: [1, 3, 6, 10]

Analogy: If you save \$1, then \$2, then \$3, then \$4 each day, your cumulative savings are [1, 3, 6, 10].

- **cumprod** = compounding interest or multiplying scores in a game.

Concept: It multiplies elements one by one, keeping a running product.

Step by step:

- **First element:** 1 → cumulative product = 1

- **Second element:** 1 * 2 = 2

- **Third element:** 2 * 3 = 6

- **Fourth element:** $6 * 4 = 24$

Output: [1, 2, 6, 24]

Analogy: Imagine stacking boxes where each box multiplies the total size instead of adding.

5. Axis Operations (Row vs Column)

NumPy arrays are 2D like a table. You can aggregate across:

- **Axis 0** → Column-wise (top to bottom)
- **Axis 1** → Row-wise (left to right)

Example:

```
1. arr = np.array([[1,2,3],
2.               [4,5,6]])
3.
4. print(arr.sum(axis=0)) # [5, 7, 9] (column sums)
5. print(arr.sum(axis=1)) # [6, 15] (row sums)
```

Analogy:

- Axis 0 → “sum by subject” (all students’ scores in Math, English, Science)
- Axis 1 → “sum by student” (total marks per student)

6. Clipping Values

- **Function:** `np.clip(arr, min, max)`
- It forces numbers into a safe range.

👉 **Example:**

```
1. arr = np.array([100, -5, 50, 300])
2. print(np.clip(arr, 0, 255))
3. # [100, 0, 50, 255]
```

Why useful?

- In **images**, clipping ensures pixel values stay between 0 and 255.
- In **ML training**, it prevents exploding gradients (numbers becoming too large and destabilizing learning).

Mini Project: Manual Linear Regression Step

Let’s **simulate one training step of linear regression** using NumPy operations.

Problem: Predict y from x using the formula:

$$y = wx + by = wx + by = wx + b$$

```
3. # Sample data
4. x = np.array([1,2,3,4,5])
5. y = np.array([3,5,7,9,11]) # roughly y = 2x + 1
6.
7. # Initialize weights
8. w = 0.0
9. b = 0.0
10.
11. # Prediction
12. y_pred = w * x + b
13.
14. # Compute error (mean squared error)
15. error = ((y - y_pred) ** 2).mean()
16.
17. # Compute gradients manually
18. grad_w = -2 * (x * (y - y_pred)).mean()
19. grad_b = -2 * (y - y_pred).mean()
20.
21. # Update weights (gradient descent step)
22. w = w - 0.01 * grad_w
23. b = b - 0.01 * grad_b
24.
25. print("Updated w:", w, "Updated b:", b)
```

Step by Step Break Down:

Arrays x and y

```
x = np.array([1,2,3,4,5])
```

```
y = np.array([3,5,7,9,11])
```

- x is a NumPy array with numbers [1,2,3,4,5]
- y is another NumPy array [3,5,7,9,11]
- NumPy arrays let you do **fast math operations** on the whole array at once.

Initial weights w and b

```
w = 0.0
```

```
b = 0.0
```

- Just numbers for now, think of them as **placeholders**.
- w and b are single values, not arrays.

Prediction: $y_{\text{pred}} = w * x + b$

```
y_pred = w * x + b
```

- NumPy lets you multiply **a number by an array** → this multiplies each element:

- $w * x \rightarrow [0*1, 0*2, 0*3, 0*4, 0*5] = [0,0,0,0,0]$
- Add $b \rightarrow [0+0, 0+0, \dots] = [0,0,0,0,0]$
- **Key idea:** operations on arrays are **element-wise** in NumPy.

Compute error: $error = ((y - y_pred)**2).mean()$

`error = ((y - y_pred)**2).mean()`

- $y - y_pred \rightarrow$ subtracts arrays element by element: $[3-0, 5-0, \dots] = [3, 5, 7, 9, 11]$
- $**2 \rightarrow$ squares each element: $[9, 25, 49, 81, 121]$
- $.mean() \rightarrow$ average: $(9+25+49+81+121)/5 = 57$

Concept: NumPy **lets you do math on arrays without loops.**

Gradients

`grad_w = -2 * (x * (y - y_pred)).mean()`

`grad_b = -2 * (y - y_pred).mean()`

- $grad_w = -2 * \text{mean}(x * (y - y_pred))$
- Compute $(y - y_pred) = [3, 5, 7, 9, 11]$
- Multiply by $x = [1, 2, 3, 4, 5]$ element-wise: $[1*3, 2*5, 3*7, 4*9, 5*11] = [3, 10, 21, 36, 55]$
- Mean = $(3+10+21+36+55)/5 = 125/5 = 25$
- Multiply by $-2 \rightarrow grad_w = -50$

$grad_b = -2 * \text{mean}(y - y_pred)$

- Mean of $[3, 5, 7, 9, 11] = (3+5+7+9+11)/5 = 35/5 = 7$
- Multiply by $-2 \rightarrow grad_b = -14$

Update weights

`w = w - 0.01 * grad_w`

`b = b - 0.01 * grad_b`

- $w = w - 0.01 * grad_w = 0 - 0.01 * (-50) = 0 + 0.5 = \mathbf{0.5}$
- $b = b - 0.01 * grad_b = 0 - 0.01 * (-14) = 0 + 0.14 = \mathbf{0.14}$
- Multiply by 0.01 \rightarrow tiny step

- Subtract from previous value → update the variable
- This is just **regular Python math**, but we used NumPy to calculate `grad_w` and `grad_b` quickly.

This shows how **all those NumPy operations combine into AI**. Under the hood, deep learning is just **millions of these updates, optimized with GPUs**.

Advanced Manipulation

Stacking Arrays

Stacking means **putting arrays together** like LEGO blocks.

- **Horizontal stack (side by side)** → `np.hstack()`
- **Vertical stack (top to bottom)** → `np.vstack()`
- **General stacking** → `np.concatenate()`

Example:

```
1. import numpy as np
2.
3. a = np.array([1,2,3])
4. b = np.array([4,5,6])
5.
6. print(np.hstack([a,b])) # [1 2 3 4 5 6]
7. print(np.vstack([a,b]))
8. # [[1 2 3]
9. # [4 5 6]]
```

Analogy:

- **hstack** = placing books **next to each other on a shelf**.
- **vstack** = stacking books **on top of each other**.

Splitting Arrays

Splitting means **cutting arrays into pieces**.

- `np.split()` → equal splits
- `np.array_split()` → allows unequal splits

Example:

```
1. arr = np.array([1,2,3,4,5,6])
2. print(np.split(arr, 3)) # [[1 2], [3 4], [5 6]]
3. print(np.array_split(arr, 4)) # [[1 2], [3 4], [5 6], [ ]]
```

Analogy: slicing a cake into equal (**split**) or unequal (**array_split**) pieces.

Repeating & Tiling

- **np.repeat()** → repeat each element
- **np.tile()** → repeat the whole array

Example:

```
1. arr = np.array([1,2,3])
2. print(np.repeat(arr, 2)) # [1 1 2 2 3 3]
3. print(np.tile(arr, 2)) # [1 2 3 1 2 3]
```

Analogy:

- **repeat** = saying each word **twice** ("Hi Hi Bye Bye").
- **tile** = replaying the **whole sentence** ("Hi Bye Hi Bye").

Flattening Arrays

Flattening means turning multi-dimensional arrays into **1D**.

- **ravel()** → returns a **view** (shares memory, efficient).
- **flatten()** → returns a **copy** (separate memory).

Example:

```
1. arr = np.array([[1,2],[3,4]])
2. print(arr.ravel()) # [1 2 3 4]
3. print(arr.flatten()) # [1 2 3 4]
```

Copy vs View

This is **super important** for memory efficiency.

- **View:** changes affect the original (they share data).
- **Copy:** changes don't affect the original.

Example:

```
1. arr = np.array([1,2,3])
2. view = arr[:2] # view
```

```
3. copy = arr[:2].copy() # copy
4.
5. view[0] = 99
6. print(arr) # [99 2 3] (changed!)
7.
8. copy[0] = 77
9. print(arr) # [99 2 3] (unchanged!)
```

Analogy:

- **View** = looking at the **same paper** (if you write, it changes the paper).
- **Copy** = photocopying the paper (original stays untouched).

Sorting

- **np.sort()** → sorts values in ascending order
- **np.argsort()** → returns indices of sorted order

Example:

```
1. arr = np.array([30,10,20])
2. print(np.sort(arr)) # [10 20 30]
3. print(np.argsort(arr)) # [1 2 0]
```

Analogy:

- **Sort** = arranging books by **height**.
- **Argsort** = writing **which positions** they should move to.

Unique Values

- **np.unique()** → gives unique elements
- **np.isin(arr, values)** → checks membership
- **np.in1d(arr1, arr2)** → like “is element in other array?”

Example:

```
1. arr = np.array([1,2,2,3,3,3])
2. print(np.unique(arr)) # [1 2 3]
3. print(np.isin([1,4], arr)) # [ True False ]
```

Analogy: like finding **unique students in a class** (ignoring duplicates).

Index Tricks

These are NumPy “hacks” for advanced indexing.

np.ix_ → cross product of indices

```
1. arr = np.arange(1,10).reshape(3,3)
2. print(arr)
3.
4. # pick rows [0,2] and cols [1,2]
5. print(arr[np.ix_([0,2], [1,2])])
```

That prints:

```
1. [[1 2 3]
2.  [4 5 6]
3.  [7 8 9]]
```

Visual grid (Row \ Col):

```
1.   Col→   0 1 2
2. Row 0   1 2 3
3.   1     4 5 6
4.   2     7 8 9
```

Explanation

Imagine a spreadsheet (grid of rows and columns).

- Rows are horizontal lines (0, 1, 2, ...).
- Columns are vertical lines (0, 1, 2, ...).

When you say “give me rows **0 and 2** and columns **1 and 2**”, you want the little rectangle that sits where those rows and columns cross, that’s the **submatrix** or **block**.

np.ix_ is the NumPy helper that **builds the pair of index arrays** you need to select that block.

The concrete array we’ll use

```
1. arr = np.array([
2.  [1, 2, 3],
3.  [4, 5, 6],
4.  [7, 8, 9]
5. ])
```

Think of it with row and column labels:

```
1.   col0 col1 col2
2. row0  1  2  3
3. row1  4  5  6
4. row2  7  8  9
```

You want rows [0, 2] and columns [1, 2].

Step 1 - What `np.ix_([0,2], [1,2])` actually produces

When you call:

```
1. rows = [0, 2]
2. cols = [1, 2]
3. I = np.ix_(rows, cols)
```

I is a **tuple** of two arrays: (I[0], I[1]).

- I[0] is the *row index array*, shaped (2,1) and looks like:

```
[[0],
```

```
 [2]]
```

(two rows, one column)

- I[1] is the *column index array*, shaped (1,2) and looks like:

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

(one row, two columns)

So `I[0].shape == (2,1)` and `I[1].shape == (1,2)`.

Step 2 - Why those shapes? (so they can combine)

NumPy will **broadcast** these two arrays together when used for indexing:

- I[0] will be repeated across columns → becomes a (2,2) array of row indices:

```
[[0, 0],
```

```
 [2, 2]]
```

- I[1] will be repeated down rows → becomes a (2,2) array of column indices:

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

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

Pairing these elementwise gives the full set of (row, column) coordinates:

(0,1), (0,2)

(2,1), (2,2)

That is the **Cartesian product** of rows × columns: every row index paired with every column index.

Step 3 - Using it to select the submatrix

When you do:

```
1. sub = arr[np.ix_(rows, cols)]
```

NumPy uses those broadcasted index arrays and returns the 2x2 block:

```
1. sub =  
2. [[ arr[0,1], arr[0,2] ],  
3. [ arr[2,1], arr[2,2] ]]
```

With numbers:

```
1. sub =  
2. [[2, 3],  
3. [8, 9]]
```

Which matches the crossing of rows 0 & 2 with columns 1 & 2 from the original array.

Step 4 - Contrast with arr[[0,2], [1,2]] (important!)

This other form **does not** take the Cartesian product. Instead it pairs indices element-wise:

- `arr[[0,2], [1,2]]` means: pick the pair (0,1) and the pair (2,2), i.e. exactly two positions:
`[arr[0,1], arr[2,2]] → [2, 9]`

So:

- `np.ix_([0,2],[1,2])` → returns a 2x2 block: `[[2,3],[8,9]]` (cartesian / cross product)
- `arr[[0,2],[1,2]]` → returns a 1-D array with elementwise pairs: `[2, 9]`

Short visual summary

Original:

```
1. [[1 2 3]  
2. [4 5 6]  
3. [7 8 9]]
```

Wanted rows [0,2] and cols [1,2] → submatrix:

```
1. row\col  col1 col2  
2. row0    2   3  
3. row2    8   9
```

That's exactly `arr[np.ix_([0,2],[1,2])] → [[2,3],[8,9]]`.

Final checklist (so you remember which to use)

- Use `np.ix_(rows, cols)` when you want the **full rectangular block** (all combinations of the selected rows with the selected columns).
- Use `arr[[rlist], [clist]]` when you want **elementwise pairs** (the i-th element of rlist paired with the i-th element of clist).

Related helpers (brief)

- **np.meshgrid**: builds coordinate grids from 1-D coordinate arrays, useful for plotting and for building full coordinate pairs; its output shape logic is similar to what `np.ix_` leverages, but **meshgrid** is tailored for geometry/plotting.
- **np.indices**: gives full index grids for all coordinates of an array (handy for fancy coordinate calculations).

Analogy: like selecting specific rows and columns in Excel at the same time.

`np.meshgrid` → make coordinate grids

Super useful in visualization & ML (like plotting decision boundaries).

```
1. x = np.array([1,2,3])
2. y = np.array([10,20])
3. X,Y = np.meshgrid(x, y)
4. print(X)
5. print(Y)
6. #Result:
7. X = [[1 2 3]
8.      [1 2 3]]
9.
10. Y = [[10 10 10]
11.       [20 20 20]]
```

Randomness & Probability

Before diving into math or code, let's start with a question:

Can a computer really be “random”?

Not truly. Computers are machines that follow instructions, so their “randomness” comes from mathematical formulas that **simulate** randomness.

That's what **np.random** in NumPy does, it gives us numbers that *act random enough* for most practical uses.

Let's unpack everything step by step.

1. Random Sampling

Imagine you're rolling dice or picking cards from a deck, that's *random sampling*.

In programming, we simulate this when we want to:

- Create fake data (e.g., random ages, incomes, etc.)
- Shuffle data for training a model
- Add randomness to algorithms (like dropout in neural networks)

NumPy gives us many ways to generate random samples.

■ Uniform Distribution

- “Uniform” means **everything has an equal chance**.
- For example, picking a random number between 0 and 1 — every number in that range is equally likely.

```
1. import numpy as np
2.
3. np.random.uniform(0, 1, 5)
4. # → e.g. [0.25, 0.73, 0.12, 0.94, 0.51]
```

Conceptually:

If you imagine a line between 0 and 1, any point is just as likely to be chosen as any other.

■ Normal Distribution (a.k.a. Bell Curve)

This is *how things often work in real life*:

Most people's heights, test scores, or even noise in data are *clustered around an average*.

```
1. np.random.normal(loc=0, scale=1, size=5)
2. # → e.g. [-0.12, 0.77, -1.45, 0.10, 0.32]
```

Here:

- **loc** = mean (center)
- **scale** = standard deviation (how spread out)
- **size** = how many numbers you want

Think: if **loc=0** and **scale=1**, most values hover near 0, fewer are far away, that's the bell curve.

■ Binomial Distribution

Binomial is about **counting successes** in repeated experiments.

Example: flipping a coin 10 times and counting how many times it lands heads.

```
1. np.random.binomial(n=10, p=0.5, size=5)
2. # → e.g. [6, 4, 5, 7, 3]
```

Meaning:

- **n** = number of trials per experiment (10 flips)
- **p** = probability of success (0.5 = 50%)
- **size** = how many experiments to simulate

🧠 You get results like 6 heads, 4 heads, 5 heads, etc.

🎯 2. np.random.choice - Picking Random Items

Sometimes you want to **pick random items from a list**, like selecting random students or shuffling training data.

```
1. names = ["Alice", "Bob", "Charlie", "Dina"]
2. np.random.choice(names, 2)
3. # → ['Charlie', 'Bob']
```

With **replace=False**, it won't pick the same item twice.

🧠 In AI, this is used for **batch sampling**, selecting small random chunks of data to train on each round (called *mini-batches*).

🌟 3. Seeds & Reproducibility

Randomness is great, but if every run gives *different* results, we can't compare experiments.

A **random seed** "locks" the randomness.

It's like setting the random number generator to start from the same point each time.

```
1. np.random.seed(42)
2. print(np.random.rand(3))
3. # → [0.37454012, 0.95071431, 0.73199394]
```

If you set the seed again to 42, you'll get **the exact same random numbers**, which is crucial for reproducible experiments.

🔄 4. Shuffling & Permutations

Sometimes we need to **mix up data** randomly, for example, before splitting it into training and testing sets.

```
1. arr = np.array([1, 2, 3, 4, 5])
2. np.random.shuffle(arr)
3. print(arr)
4. # → [3, 5, 1, 4, 2]
```

.shuffle() changes the array *in place*, it modifies the original one.

If you want a shuffled *copy* instead:

```
1. shuffled = np.random.permutation(arr)
2. print(shuffled)
```



5. Probability Distributions

A **probability distribution** describes *how likely different outcomes are*.

You can think of it like a “map” of randomness.

- **Uniform** → everything equally likely
- **Normal** → most around the middle
- **Binomial** → counts of successes/failures
- **Exponential** → time until something happens (like waiting times)

AI models, especially probabilistic ones depend on understanding these distributions to model *uncertainty* or *noise*.



Why This Matters (Even for Beginners)

Even if you’re just learning NumPy, randomness shows up **everywhere** in AI:

- In **training neural networks**, weights start randomly 🧠
- In **data shuffling**, so models don’t memorize the order
- In **simulations** and **Monte Carlo methods**
- In **reinforcement learning**, where agents explore randomly before improving

You do need to understand this, but don’t worry:

you don’t need heavy math yet, just the **concepts** and **intuition** we covered here.

Practical Applications for AI/ML



1. Image Representation as NumPy Arrays

Concept:

Every image is just a **grid of numbers** (pixels).

- A **black & white** image = 2D array → each value shows brightness (0 = black, 255 = white).
- A **colored (RGB)** image = 3D array → each pixel has **3 values** (red, green, blue).

Example:

```
1. import numpy as np
```

```

2.
3. # A 2x2 grayscale image (black & white)
4. gray_image = np.array([[0, 255],
5.                        [128, 64]])
6.
7. # A 2x2 color image (RGB)
8. color_image = np.array([
9.     [[255, 0, 0], [0, 255, 0]], # red, green
10.    [[0, 0, 255], [255, 255, 255]] # blue, white
11. ])
12.
13. print(color_image.shape) # (2, 2, 3) -> height, width, color_channels

```

Why this matters: AI models *see images as numbers*. So understanding how they're stored helps you process them later for training.

2. Normalization & Standardization

Goal: Make all your data on a similar scale.

Term	Meaning	Formula
Normalization	Scale values between 0 and 1	$x_{norm} = (x - \min) / (\max - \min)$
Standardization	Center values around 0 with std = 1	$x_{std} = (x - \text{mean}) / \text{std}$

Example:

```

1. arr = np.array([10, 20, 30, 40])
2.
3. # Normalization (0 to 1)
4. norm = (arr - arr.min()) / (arr.max() - arr.min())
5.
6. # Standardization (mean = 0, std = 1)
7. std = (arr - arr.mean()) / arr.std()
8.
9. print("Normalized:", norm)
10. print("Standardized:", std)

```

Why this matters: Neural networks learn better when numbers are small and balanced.

3. One-Hot Encoding

Concept:

AI models can't understand words or categories, only numbers.

So we convert categories like:

["cat", "dog", "bird"]

into **binary vectors** (a 1 for the correct label, 0 for others):

```

cat -> [1, 0, 0]
dog -> [0, 1, 0]
bird -> [0, 0, 1]

```

Example:

```
1. import numpy as np
2.
3. labels = np.array([0, 1, 2]) # 0=cat, 1=dog, 2=bird
4. one_hot = np.eye(3)[labels] # identity matrix trick
5. print(one_hot)
```

Why this matters: This encoding lets the model *compare outputs* correctly for classification tasks.

4. Dataset Shuffling & Batching

Concept:

When training AI models, we don't show the data in one go.

We:

1. **Shuffle** → mix the order (avoid bias)
2. **Batch** → take small groups (to train in chunks)

Example:

```
1. data = np.arange(10)
2. np.random.shuffle(data) # shuffle in place
3. print("Shuffled:", data)
4.
5. # Create batches of size 3
6. batch_size = 3
7. for i in range(0, len(data), batch_size):
8.     batch = data[i:i+batch_size]
9.     print("Batch:", batch)
```

Why this matters: Models train faster and more fairly with shuffled batches.

5. Convolution Basics (Before Deep Learning Frameworks)

Concept:

A **convolution** is like sliding a small window (called a *kernel*) over an image to detect patterns (edges, colors, etc.).

Analogy:

Imagine dragging a small magnifying glass over a photo and recording what you see each time.

Example:

```
1. image = np.array([[1, 2, 3],
2.                  [4, 5, 6],
3.                  [7, 8, 9]])
4.
5. kernel = np.array([[1, 0],
6.                   [0, -1]])
7.
8. output = np.zeros((2, 2))
9. for i in range(2):
```

```

10. for j in range(2):
11.     patch = image[i:i+2, j:j+2]
12.     output[i, j] = np.sum(patch * kernel)
13.
14. print(output)

```

Why this matters: This is the **foundation of CNNs (Convolutional Neural Networks)**, the core of computer vision.

6. Implementing Loss Functions (MSE, Cross-Entropy)

Concept:

Loss functions tell the model **how wrong** its predictions are.

Loss	Used For	Formula (simplified)	Intuition
MSE (Mean Squared Error)	Regression (numbers)	$\text{average}((y_{\text{pred}} - y_{\text{true}})^2)$	Measures how far predictions are
Cross-Entropy	Classification (labels)	$-\sum y_{\text{true}} * \log(y_{\text{pred}})$	Punishes wrong confidence

Examples:

```

1. # MSE
2. y_true = np.array([2.5, 0.0, 2.1])
3. y_pred = np.array([3.0, -0.5, 2.0])
4. mse = np.mean((y_pred - y_true)**2)
5. print("MSE:", mse)
6.
7. # Cross-Entropy (simple version)
8. y_true = np.array([1, 0, 0]) # true = class 0
9. y_pred = np.array([0.7, 0.2, 0.1]) # model probabilities
10. cross_entropy = -np.sum(y_true * np.log(y_pred))
11. print("Cross-Entropy:", cross_entropy)

```

👉 **Why this matters:** The smaller the loss, the smarter the model gets.

🧩 Summary

Concept	What it Teaches	Why It Matters
Image as Array	Numbers = pixels	Foundation for computer vision
Normalization	Scaling values	Helps model learn
One-Hot Encoding	Convert categories	Needed for classification
Shuffling/Batching	Fair learning	Prevents bias
Convolution	Pattern detection	Base of CNNs
Loss Functions	Model feedback	Core of training




Connecting to AI Frameworks (Beginner Breakdown)

1. How NumPy Underlies TensorFlow & PyTorch

What's going on:

- You already know **NumPy** → it handles **arrays and math**.
- AI frameworks like **TensorFlow** and **PyTorch** do **almost the same thing**, but faster (and often on GPUs instead of CPUs).

Think of it like this:

Level	Tool	Job
 Beginner	NumPy	Teaches you basic math with arrays (CPU only)
 Advanced	PyTorch, TensorFlow	Use arrays called <i>tensors</i> (can run on GPU for big AI models)
 So the core idea is:		

Tensors = Fancy NumPy arrays that can live on GPU and track gradients for training.

2. Converting Between NumPy Arrays and Tensors

Let's look at how the two talk to each other.

Using PyTorch

```
1. import numpy as np
2. import torch
3.
4. # Create NumPy array
5. arr = np.array([[1, 2], [3, 4]])
6.
7. # Convert NumPy → Tensor
8. tensor = torch.from_numpy(arr)
9. print(tensor)
10.
11. # Convert Tensor → NumPy
12. back_to_numpy = tensor.numpy()
13. print(back_to_numpy)
```




Both share the same memory, changing one changes the other!

Using TensorFlow

```
1. import tensorflow as tf
2. import numpy as np
3.
4. arr = np.array([1, 2, 3])
5. tensor = tf.convert_to_tensor(arr)
6.
7. print(tensor) # TensorFlow tensor
8. print(tensor.numpy()) # Back to NumPy
```

3. When to Use NumPy vs Pandas vs SciPy

These three are like siblings, they work together but have different personalities:

Library	Use For	Example
 NumPy	Math, arrays, and linear algebra	Image pixels, ML math
 Pandas	Data handling (tables like Excel)	Datasets, CSV files
 SciPy	Advanced math, scientific tools	Statistics, optimization

Example:

```
1. import numpy as np
2. import pandas as pd
3. from scipy import stats
4.
5. # NumPy: raw math
6. arr = np.array([1, 2, 3, 4])
7. print(np.mean(arr))
8.
```

```
9. # Pandas: data table
10. df = pd.DataFrame({"Age": [10, 20, 30]})
11. print(df["Age"].mean())
12.
13. # SciPy: scientific functions
14. print(stats.zscore(arr)) # Standardize data
```

Rule of thumb:

- **NumPy** = for math
- **Pandas** = for data
- **SciPy** = for advanced math

4. Why NumPy Operations Look Like Tensor Operations

You might've noticed something:

Matrix multiplication, addition, broadcasting, it all looks the same.

That's not an accident!

TensorFlow and PyTorch were *inspired* by NumPy, so you don't have to relearn everything later.

Example:

```
1. import numpy as np
2. import torch
3.
4. # NumPy
5. a = np.array([[1, 2], [3, 4]])
6. b = np.array([[5, 6], [7, 8]])
7. print(np.dot(a, b)) # matrix multiply
8.
9. # PyTorch
10. a_t = torch.tensor([[1, 2], [3, 4]])
11. b_t = torch.tensor([[5, 6], [7, 8]])
12. print(torch.matmul(a_t, b_t)) # same idea!
```

Key idea:

If you understand NumPy's rules, how shapes, broadcasting, and indexing work, you already understand 80% of PyTorch and TensorFlow.

5. Bonus (Optional but Cool): What Are Tensors Really?

Concept:

A *tensor* is just a fancy name for "multi-dimensional array".

Dimension	Example	Name
0D	5	Scalar
1D	[1, 2, 3]	Vector
2D	[[1, 2], [3, 4]]	Matrix
3D	A bunch of matrices stacked	Tensor

So a 4D tensor could be a batch of images: **height** × **width** × **color** × **batch_size**.

Summary Table

Concept	Meaning	Example
NumPy → Base library	Handles arrays and math	np.array, np.dot
TensorFlow / PyTorch	Built on top of NumPy ideas	tf.tensor, torch.tensor
Converting between them	Easy and instant	torch.from_numpy()
Pandas	Works with tabular data	CSV or DataFrames
SciPy	Adds scientific math tools	Stats, optimization
Tensors	Generalized arrays	Used in AI models