

Finding the Shape of an Image

Name		Date	
Class / Sec		Roll no.	

1 · AIM

To write a Python program that reads an image using **OpenCV** and identifies its **shape** — its height, width and number of colour channels — and calculates how many values are stored in it.

2 · SOFTWARE AND LIBRARIES USED

TOOL	WHAT IT IS FOR
Python 3	The programming language the program is written in
Jupyter Notebook	The environment used to write and run the code (IDLE or Colab also acceptable)
OpenCV (cv2)	The computer vision library. imread() loads an image; .shape reports its dimensions

3 · THEORY

Every image is stored as a grid of numbers. The **shape** tells you the size of that grid — and the order the three numbers come in is the thing students most often get wrong.

POSITION	WHAT IT IS	WHY THAT ORDER
shape[0]	Height (rows)	The array is a list of rows, so rows come first
shape[1]	Width (columns)	Each row holds one entry per column
shape[2]	Channels	3 for a colour image (BGR). Absent for grayscale

Height comes FIRST: We describe screens as **width × height** — 1920 × 1080. NumPy and OpenCV report **height × width**. So a 400-wide, 300-tall image has shape **(300, 400, 3)**, not (400, 300, 3). This reversal is examined constantly.

Grayscale has no third number: A grayscale image stores **one** value per pixel instead of three, so its shape is just **(height, width)** — two numbers, not three. That is why grayscale files are about three times smaller.

4 · PROGRAM

```
import cv2

img = cv2.imread("sample_image.jpg")

if img is None:
    print("Image not found. Check the file name and folder.")
else:
    # The shape gives all three dimensions at once
    print("Shape of the image :", img.shape)


    # Unpack them into three separate variables
    height, width, channels = img.shape
    print("Height   :", height, "pixels")
    print("Width    :", width, "pixels")
    print("Channels  :", channels)
```

```
# How much data is actually stored?
print("Total pixels :", height * width)
print("Total values :", img.size)
```

No window this time: This program only prints numbers, so **imshow()** and **waitKey(0)** are not needed. Nothing is displayed — we are asking about the image, not looking at it.

5 • EXPECTED OUTPUT

The image used here is **400 pixels wide and 300 pixels tall**.



Running the program prints:

```
Shape of the image : (300, 400, 3)
Height   : 300 pixels
Width    : 400 pixels
Channels  : 3
Total pixels : 120000
Total values : 360000
```

[PASTE SCREENSHOT HERE]

Paste a screenshot of YOUR program running, with the code and output visible together

6 • CHECKING THE ARITHMETIC

Every number in the output can be worked out by hand. Show this in your file — it is what earns the observation marks.

QUANTITY	WORKING
Shape	(300, 400, 3) — height 300, width 400, 3 colour channels
Total pixels	height × width = 300 × 400 = 120,000 pixels
Total values stored	120,000 × 3 channels = 360,000 values , one per colour per pixel
If it were grayscale	Shape would be (300, 400) and only 120,000 values — three times smaller
Each value	A whole number from 0 to 255 — 256 possibilities, because $2^8 = 256$

7 • OBSERVATION AND RESULT

Two or three sentences in your own words. Give the dimensions, and say what they mean about how much data the image holds.

Example observation: The image has a shape of (300, 400, 3), meaning it is 300 pixels tall, 400 pixels wide and stores 3 colour channels. That gives 120,000 pixels in total and 360,000 stored values, since each pixel holds one number for blue, green and red. A grayscale version of the same image would need only 120,000 values.

Worth knowing: **img.size** and **img.shape** are different things. **shape** gives the three dimensions; **size** gives the total count of values — the three dimensions multiplied together. Mixing them up is a common slip.

8 · VIVA QUESTIONS FOR THIS PROGRAM

LIKELY QUESTION	SHAPE OF A STRONG ANSWER
Your image is 400 wide and 300 tall. Why does shape say (300, 400, 3)?	NumPy reports height first , then width. Rows come before columns in an array
What does the 3 mean?	Three colour channels — blue, green and red. One value each, per pixel
How many pixels does the image have?	120,000 — that is 300×400 . Not 360,000, which is the number of stored values
What would the shape be if it were grayscale?	(300, 400) — only two numbers, because grayscale stores one value per pixel
What is the difference between shape and size?	shape gives the dimensions; size gives the total number of values — $300 \times 400 \times 3$

9 · HOW THIS ENTRY IS MARKED

FULL MARKS

The image **submitted with** the program · height, width and channels **each printed separately** · the arithmetic checked by hand · an observation giving **total pixels and total values** · you can explain why height comes first

LOSES MARKS

Prints only the shape and stops · no image submitted · says the image is **400 tall and 300 wide** — the wrong way round · confuses total pixels with total values · cannot say what the 3 means

Before you submit: Use **your own** image, and pick one that is clearly **not square** — a landscape photo. If height and width are the same number you cannot prove you understand which one comes first, and the examiner cannot tell either.