

# Python Face Detection & OpenCV Examples Mini-Guide

This mini-guide contains two practical OpenCV tutorials from the Real Python web site. Learn more about OpenCV and other Python computer vision topics at [realpython.com/blog/categories/opencv/](http://realpython.com/blog/categories/opencv/).

---

## Face Recognition with Python, in Under 25 Lines of Code

**In this post we'll look at a surprisingly simple way to get started with face recognition using Python and the open source library OpenCV.**

Before you ask any questions in the comments section:

1. Do not skip over the blog post and try to run the code. You must understand what the code does not only to run it properly but to troubleshoot it as well.
2. Make sure to use OpenCV v2.
3. You need a working webcam for this script to work properly.
4. Review the other comments/questions as your questions have probably already been addressed.

Thank you.

## OpenCV

OpenCV is the most popular library for computer vision. Originally written in C/C++, it now provides bindings for Python.

OpenCV uses machine learning algorithms to search for faces within a picture. For something as complicated as a face, there isn't one simple test that will tell you if it found a face or not. Instead, there are thousands of small patterns/features that must be matched. The algorithms break the task of identifying the face into thousands of smaller, bite-sized tasks, each of which is easy to solve. These tasks are also called classifiers.

For something like a face, you might have 6,000 or more classifiers, all of which must match for a face to be detected (within error limits, of course). But therein lies the problem: For face detection, the algorithm starts at the top left of a picture and moves down across small blocks of data, looking at each block, constantly asking, "*Is this a face? ... Is this a face? ... Is this a face?*" Since there are 6,000 or more tests per block, you might have millions of calculations to do, which will grind your computer to a halt.

To get around this, OpenCV uses cascades. What's a cascade? The best answer can be found from the dictionary: *A waterfall or series of waterfalls*

Like a series of waterfalls, the OpenCV cascade breaks the problem of detecting faces into multiple stages. For each block, it does a very rough and quick test. If that passes, it does a slightly more detailed test, and so on. The algorithm may have 30-50 of these stages or cascades, and it will only detect a face if all stages pass. The advantage is that the majority of the pictures will return negative during the first few stages, which means the algorithm won't waste time testing all 6,000 features on it. Instead of taking hours, face detection can now be done in real time.

### **Cascades in practice**

Though the theory may sound complicated, in practice it is quite easy. The cascades themselves are just a bunch of XML files that contain OpenCV data used to detect objects. You initialize your code with the cascade you want, and then it does the work for you.

Since face detection is such a common case, OpenCV comes with a number of built-in cascades for detecting everything from faces to eyes to hands and legs. There are even cascades for non-human things. For example, if you run a banana shop and want to track people stealing bananas, this guy has built one for that!

### **Installing OpenCV**

First, you need to find the correct setup file for your operating system

I found that installing OpenCV was the hardest part of the task. If you get strange unexplainable errors, it could be due to library clashes, 32/64 bit differences, etc. I found it easiest to just use a Linux virtual machine and install OpenCV from scratch.

Once installed, you can test whether or not it works by firing up a Python session and typing:

```
$ python
>>> import cv2
>>>
```

If you don't get any errors, you can move on to the next part.

### **Understanding the code**

Let's break down the actual code, which you can download from the repo. Grab the *face\_detect.py* script, the *abba.png* pic, and the *haarcascade\_frontalface\_default.xml*.

```
# Get user supplied values
```

```
imagePath = sys.argv[1]
```

```
cascPath = sys.argv[2]
```

You first pass in the image and cascade names as command-line arguments. We'll use the Abba image as well as the default cascade for detecting faces provided by OpenCV.

```
# Create the haar cascade
```

```
faceCascade = cv2.CascadeClassifier(cascPath)
```

Now we create the cascade and initialize it with our face cascade. This loads the face cascade into memory so it's ready for use. Remember, the cascade is just an XML file that contains the data to detect faces.

```
# Read the image
```

```
image = cv2.imread(imagePath)
```

```
gray = cv2.cvtColor(image, cv2.COLOR_BGR2GRAY)
```

Here we read the image and convert it to grayscale. Many operations in OpenCv are done in grayscale.

```
# Detect faces in the image
```

```
faces = faceCascade.detectMultiScale(  
    gray,  
    scaleFactor=1.1,  
    minNeighbors=5,  
    minSize=(30, 30),  
    flags = cv2.cv.CV_HAAR_SCALE_IMAGE  
)
```

This function detects the actual face - and is the key part of our code, so let's go over the options.

1. The `detectMultiScale` function is a general function that detects objects. Since we are calling it on the face cascade, that's what it detects. The first option is the grayscale image.
2. The second is the `scaleFactor`. Since some faces may be closer to the camera, they would appear bigger than those faces in the back. The scale factor compensates for this.
3. The detection algorithm uses a moving window to detect objects. `minNeighbors` defines how many objects are detected near the current one before it declares the face found. `minSize`, meanwhile, gives the size of each window.

I took commonly used values for these fields. In real life, you would experiment with different values for the window size, scale factor, etc., until you find one that best works for you.

The function returns a list of rectangles where it believes it found a face. Next, we will loop over where it thinks it found something.

```
print "Found {0} faces!".format(len(faces))

# Draw a rectangle around the faces
for (x, y, w, h) in faces:
    cv2.rectangle(image, (x, y), (x+w, y+h), (0, 255, 0), 2)
```

This function returns 4 values: the x and y location of the rectangle, and the rectangle's width and height ( $w$ ,  $h$ ).

We use these values to draw a rectangle using the built-in `rectangle()` function.

```
cv2.imshow("Faces found" ,image)
cv2.waitKey(0)
```

In the end, we display the image, and wait for the user to press a key.

## Checking the results

Let's test against the Abba photo:

```
$ python face_detect.py abba.png haarcascade_frontalface_default.xml
```

That worked. How about another photo:

That ... is not a face. Let's try again. I changed the parameters and found that setting the `scaleFactor` to 1.2 got rid of the *wrong* face.

## What happened?

Well, the first photo was taken fairly close up with a high quality camera. The second one seems to have been taken from afar and possibly from a mobile phone. This is why the `scaleFactor` had to be modified. As I said, you'll have to setup the algorithm on a case by case basis to avoid false positives.

Be warned though that since this is based on machine learning, the results will never be 100% accurate. You will get *good enough* results in most cases, but occasionally the algorithm will identify incorrect objects as faces.

The final code can be found [here](#).

## Extending to a webcam

So what if you want to use a webcam? OpenCV grabs each frame from the webcam and you can then detect faces by processing each frame. You will need a powerful computer, though my five year old laptop seems to cope fine, as long as I don't dance around too much.

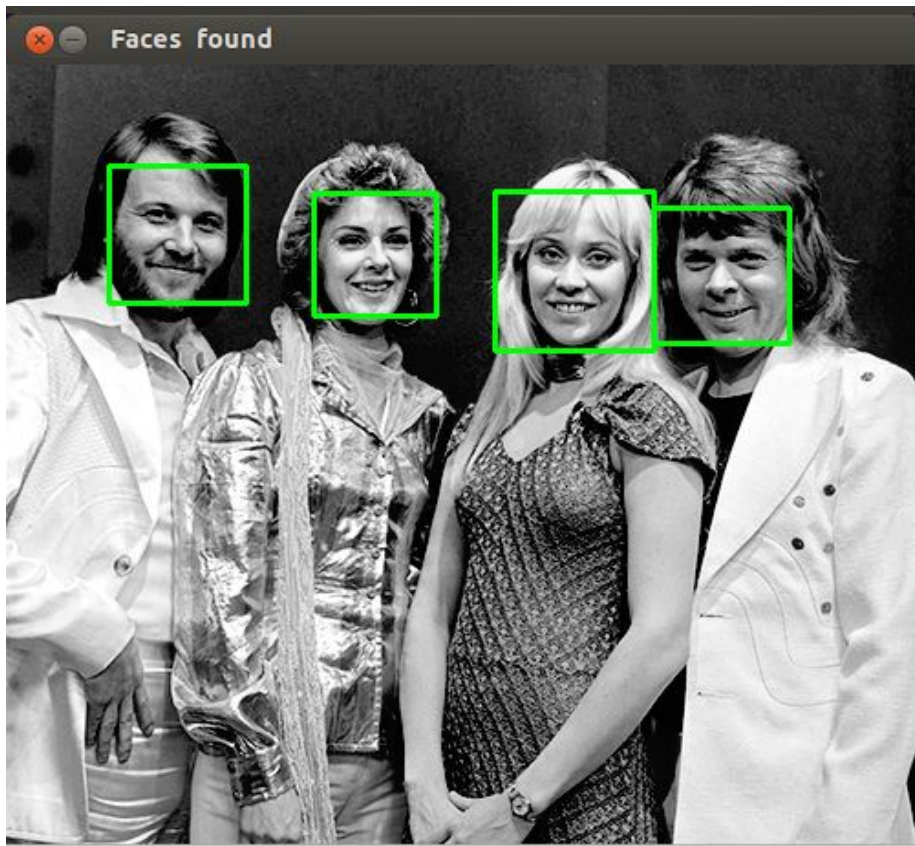


Figure 1:

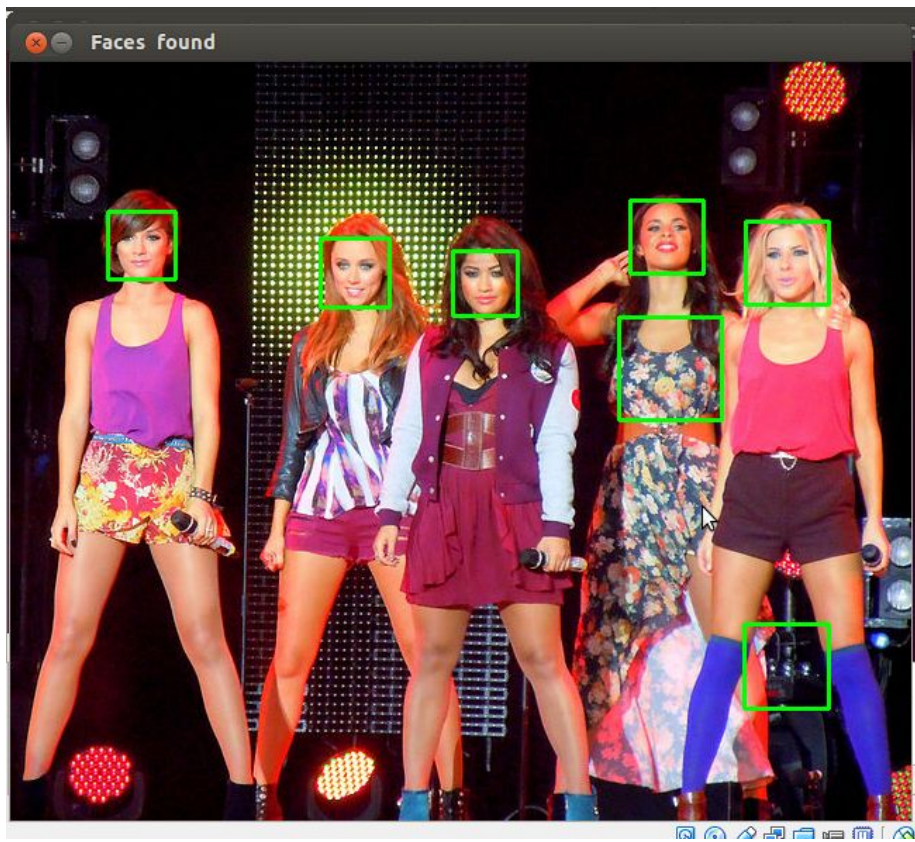


Figure 2:




Figure 3:



**UPDATED** The next blog post is live: Face Detection in Python Using a Webcam. Check it out!

---

## Face Detection in Python Using a Webcam

*This tutorial is a follow-up to **Face Recognition in Python**, so make sure you've gone through that first post.*

**As mentioned in the first post, it's quite easy to move from detecting faces in images to detecting them in video via a webcam - which is exactly what we will detail in this post.**

Before you ask any questions in the comments section:

1. Do not skip over the blog post and try to run the code. You must understand what the code does not only to run it properly but to troubleshoot it as well.
2. Make sure to use OpenCV v2.
3. You need a working webcam for this script to work properly.
4. Review the other comments/questions as your questions have probably already been addressed.

Thank you.

### Pre-requisites

1. OpenCV installed (see the previous blog post for details)
2. A working webcam

### The Code

Let's dive straight into the code, taken from this repository.

```
import cv2
import sys

cascPath = sys.argv[1]
faceCascade = cv2.CascadeClassifier(cascPath)

video_capture = cv2.VideoCapture(0)

while True:
    # Capture frame-by-frame
    ret, frame = video_capture.read()
```



```

gray = cv2.cvtColor(frame, cv2.COLOR_BGR2GRAY)

faces = faceCascade.detectMultiScale(
    gray,
    scaleFactor=1.1,
    minNeighbors=5,
    minSize=(30, 30),
    flags=cv2.cv.CV_HAAR_SCALE_IMAGE
)

# Draw a rectangle around the faces
for (x, y, w, h) in faces:
    cv2.rectangle(frame, (x, y), (x+w, y+h), (0, 255, 0), 2)

# Display the resulting frame
cv2.imshow('Video', frame)

if cv2.waitKey(1) & 0xFF == ord('q'):
    break

# When everything is done, release the capture
video_capture.release()
cv2.destroyAllWindows()

```

Now let's break it down...

```

import cv2
import sys

cascPath = sys.argv[1]
faceCascade = cv2.CascadeClassifier(cascPath)

```

This should be familiar to you. We are creating a face cascade, as we did in the image example.

```
video_capture = cv2.VideoCapture(0)
```

This line sets the video source to the default webcam, which OpenCV can easily capture.

**NOTE:** You can also provide a filename here, and Python will read in the video file. However, you need to have ffmpeg installed for that since OpenCV itself cannot decode compressed video. Ffmpeg acts as the front end for OpenCV, and, ideally, it should be compiled directly into OpenCV. This is not easy to do, especially on Windows.

```

while True:
    # Capture frame-by-frame

```

```
ret, frame = video_capture.read()
```

Here, we capture the video. The `read()` function reads one frame from the video source, which in this example is the webcam. This returns:

1. The actual video frame read (one frame on each loop)
2. A return code

The return code tells us if we have run out of frames, which will happen if we are reading from a file. This doesn't matter when reading from the webcam, since we can record forever, so we will ignore it.

```
# Capture frame-by-frame
ret, frame = video_capture.read()

gray = cv2.cvtColor(frame, cv2.COLOR_BGR2GRAY)

faces = faceCascade.detectMultiScale(
    gray,
    scaleFactor=1.1,
    minNeighbors=5,
    minSize=(30, 30),
    flags=cv2.cv.CV_HAAR_SCALE_IMAGE
)

# Draw a rectangle around the faces
for (x, y, w, h) in faces:
    cv2.rectangle(frame, (x, y), (x+w, y+h), (0, 255, 0), 2)

# Display the resulting frame
cv2.imshow('Video', frame)
```

Again, this code should be familiar. We are merely searching for the face in our captured frame.

```
if cv2.waitKey(1) & 0xFF == ord('q'):
    break
```

We wait for the 'q' key to be pressed. If it is, we exit the script.

```
# When everything is done, release the capture
video_capture.release()
cv2.destroyAllWindows()
```

Here, we are just cleaning up.

**Test!**

[Click here to watch the example video](#)

So, that's me with my driver's license in my hand. And you can see that the algorithm tracks both the real me and the photo me. Note that when I move slowly, the algorithm can keep up. When I move my hand up to my face a bit faster, though, it gets confused and mistakes my wrist for a face.

Like I said in the last post, machine learning based algorithms are rarely 100% accurate. We aren't at the stage where Robocop driving his motorcycle at 100 mph can track criminals using low quality CCTV cameras... yet.

The code searches for the face frame by frame, so it will take a fair amount of processing power. For example, on my five year old laptop, it took almost 90% of the CPU.

### **Next Steps**

Okay, so you know how to detect faces. But what if you want to detect your own object, like your car or your TV or your favorite toy?

OpenCV allows you to create your own cascades, but the process isn't well documented. Here is a blog post that shows you how to train your own cascade to detect a banana.

If you want to take it one step further and recognize individual faces - perhaps to detect and recognize your face amongst many strangers - the task is surprisingly difficult. This is mainly due to the large amount of image pre-processing involved. But if you are willing to tackle the challenge, it is possible by using machine learning algorithms as described here.

### **Want to know more?**

This will be covered in much greater detail along with a number of computational science and machine learning topics and more in my upcoming course. The course is based on a highly successful Kickstarter.

The Kickstarter is over, but you can still order the course at Python for Engineers. Visit to find out more.

Also, post links to your videos below to get direct feedback from me. Comment if you have questions.

Oh—and next time we'll be tackling some motion detection. Stay tuned!