

Quick start with

Traitement d'images
2^e année SN Multimédia et 3EA IATI

Version: v2026.1.0-rc4
(master @ 4a2de53 2026-01-30)

Contents

1	About this document	1
2	The tutorial code	2
2.1	Building the code with CMake	2
3	OpenCV	3
4	Mat - The Basic Image Container	4
4.1	Creating explicitly a Mat object	5
4.1.1	Access to Mat elements	7
4.2	Print out formatting	9
4.3	Print for other common items	10
5	Load and Display an Image	12
5.1	Goal	12
5.2	Source Code	12
5.3	Explanation	13
5.4	Result	14
6	Load, Modify, and Save an Image	16
6.1	Goals	16
6.2	Code	17
6.3	Explanation	17
6.4	Result	18

1 About this document

This documents collects some of the tutorials that are part of the OpenCV library. These are the most basic yet relevant ones that will help you to get started with the library and complete the TP more easily. Some of them has been adapted, reduced or rewritten with the code you will have to write during the TP in mind. You can always access to the original tutorials and all the other ones from the online documentation that you can find at this [address](#).

2 The tutorial code

The archive containing the tutorials is organized as follows:

- `data` contains some images that can be used with the programs of this tutorial.
- `src` contains the source files for the tp they are organize in directories:
 - `tutorials` contains the code used in these tutorials .

2.1 Building the code with CMake

CMake is a cross-platform free software program that helps managing the build process of software using a compiler-independent method. In simpler words, this means that CMake is an utility that helps to set up the compilation environment for a given source code, independently from the compiler and the building system that is actually used to generate the code, be it Linux's `make`, Apple's Xcode, or Microsoft Visual Studio.

In the case of this TP, CMake will check for all the libraries that are needed to compile our programs and it will automatically generate the corresponding Makefile. Let's see how this work.

From the starting path of the code, create a new directory, let's call it `build`:

```
mkdir build
cd build
```

This directory will contain the results of the compilation, *i.e.* all the executables that you will generate. Now we can tell CMake to configure and create the relevant makefile for us:

```
cmake ..
```

CMake will look for all the necessary libraries and generate the makefile for all the programs. On your screen you should see something like this:

```
simone@simone-XPS-L501X: ~/Documents/ocv/build
File Edit View Search Terminal Help
simone@simone-XPS-L501X:~/Documents/ocv/build$ cd ..; rm -r build; mkdir build; cd build; cmake .. -DOpenCV_DIR~/MOOV3D/dev/code/lbs/opencv/
-- The C compiler identification is GNU
-- The CXX compiler identification is GNU
-- Check for working C compiler: /usr/bin/gcc
-- Check for working CXX compiler: /usr/bin/g++ -- works
-- Detecting C compiler ABI info
-- Detecting CXX compiler ABI info
-- Found GL: /home/simone/Documents/ocv/3rdparty/gln/build/lib/libGL.so
-- Found GLUT
-- Found OpenGL: /usr/lib/x86_64-linux-gnu/libGL.so; /usr/lib/x86_64-linux-gnu/libGL.so
-- Found GLUT
-- Found Doxygen: /usr/bin/doxygen
-- Configuring done
-- Generating done
-- Build files have been written to: /home/simone/Documents/ocv/build
simone@simone-XPS-L501X:~/Documents/ocv/build$
```

This has to be done only once at the beginning, from now on you have just to compile the code you modify in `src` using the usual `make` command. Indeed, in the `build` directory you can see the Makefile we will use to compile the code. The code can be compiled from the `build` directory with the usual

```
make <filename_without_extension>
```

and

```
make clean
```

to delete all the executable and the compilation objects. Try for example to build the code for the first tutorial with

```
make mat_the_basic_image_container
```

The executable will be placed in `build/bin/` and you can run the code with

```
./bin/mat_the_basic_image_container
```

3 OpenCV

OpenCV (Open Source Computer Vision Library) is an open source computer vision and machine learning software library. OpenCV was built to provide a common infrastructure for computer vision applications and to accelerate the use of machine perception in the commercial products.

The library has more than 2500 optimized algorithms, which includes a comprehensive set of both classic and state-of-the-art computer vision and machine learning algorithms. These algorithms can be used to detect and recognize faces, identify objects, classify human actions in videos, track camera movements, track moving objects, extract 3D models of objects, produce 3D point clouds from stereo cameras, stitch images together to produce a high resolution image of an entire scene, find similar images from an image database, remove red eyes from images taken using flash, follow eye movements, recognize scenery and establish markers to overlay it with augmented reality, etc.

It has C++, C, Python, Java and MATLAB interfaces and supports Windows, Linux, Android and Mac OS. OpenCV leans mostly towards real-time vision applications and takes advantage of MMX and SSE instructions when available. A full-featured CUDA and OpenCL interfaces are being actively developed right now. There are over 500 algorithms and about 10 times as many functions that compose or support those algorithms. OpenCV is written natively in C++ and has a templated interface that works seamlessly with STL containers.

4 Mat - The Basic Image Container

`Mat` is basically a class having two data parts: the matrix header (containing information such as the size of the matrix, the method used for storing, at which address is the matrix stored and so on) and a pointer to the matrix containing the pixel values (may take any dimensionality depending on the method chosen for storing). The matrix header size is constant. However, the size of the matrix itself may vary from image to image and usually is larger by order of magnitudes. Therefore, when you're passing on images in your program and at some point you need to create a copy of the image the big price you will need to build is for the matrix itself rather than its header. `OpenCV` is an image processing library. It contains a large collection of image processing functions. To solve a computational challenge most of the time you will end up using multiple functions of the library. Due to this passing on images to functions is a common practice. We should not forget that we are talking about image processing algorithms, which tend to be quite computational heavy. The last thing we want to do is to further decrease the speed of your program by making unnecessary copies of potentially *large* images.

To tackle this issue `OpenCV` uses a reference counting system. The idea is that each `Mat` object has its own header, however the matrix may be shared between two instance of them by having their matrix pointer point to the same address. Moreover, the copy operators **will only copy the headers**, and as also copy the pointer to the large matrix too, however not the matrix itself.

```

1 // creates just the header parts
2 Mat A, C;
3
4 // allocate the matrix from the image file
5 A = imread(argv[1], cv::IMREAD_COLOR);
6
7 // Use the copy constructor
8 Mat B(A);
9
10 // Assignment operator
11 C = A;
```

All the above objects, in the end point to the same single data matrix. Their headers are different, however making any modification using either one of them will affect all the other ones too. In practice the different objects just provide different access method to the same underlying data. Nevertheless, their header parts are different. The real interesting part comes that you can create headers that refer only to a subsection of the full data. For example, to create a region of interest (*ROI*) in an image you just create a new header with the new boundaries:

```

1 // using a rectangle
2 Mat D (A, Rect(10, 10, 100, 100) );
3
4 // extracts A columns, from 1 (inclusive) to 3 (exclusive).
5 // like matlab E = A(:,2:3)
6 // columns and rows start from 0 in C/C++
7 Mat E = A(Range::all(), Range(1,3));
```

Here is another way to select and use columns and rows of the matrix

```

1 // add the 6th row, multiplied by 3 to the 4th row
2 M.row(3) = M.row(3) + M.row(5)*3;
3
4 // now copy the 8th column to the 2nd column
5 // M.col(1) = M.col(7); // this will not work!
6 Mat M1 = M.col(1);
7 M.col(7).copyTo(M1);

```

In order to select more than one column (row) you can use the `Mat` method

```

colRange(int start, int end)
(rowRange(int start, int end)).

```

`start` is the index of first column/row to consider, `end` is the exclusive index of the last column/row, *i.e.* `colRange(1, 5)` will return the columns from the second (index 1) to the fifth (index 4).

Now you may ask if the matrix itself may belong to multiple `Mat` objects who will take responsibility for its cleaning when it's no longer needed. The short answer is: the last object that used it. For this a reference counting mechanism is used. Whenever somebody copies a header of a `Mat` object a counter is increased for the matrix. Whenever a header is cleaned this counter is decreased. When the counter reaches zero the matrix too is freed. Because, sometimes you will still want to copy the matrix itself too, there exists the `clone()` or the `copyTo()` function.

```

1 Mat F = A.clone();
2 Mat G;
3 A.copyTo(G);

```

Now modifying `F` or `G` will not affect the matrix pointed by the `Mat` header. What you need to remember from all this is that:

- Output image allocation for OpenCV functions is automatic (unless specified otherwise).
- No need to think about memory freeing with OpenCV C++ interface.
- The assignment operator and the copy constructor copies only the header.
- Use the `clone()` or the `copyTo()` function to copy the underlying matrix of an image.

4.1 Creating explicitly a `Mat` object

In the [Load, Modify, and Save an Image](#) tutorial you will see how to write a matrix to an image file by using the `imwrite` function. However, for debugging purposes it's much more convenient to see the actual values. You can achieve this via the `<<` operator of `Mat`. However, be aware that this only works for two dimensional matrices.

Although `Mat` is a great class as image container it is also a general matrix class. Therefore, it is possible to create and manipulate multidimensional matrices. You can create a `Mat` object in multiple ways:

- `Mat()` Constructor

```

1 Mat M(2,2, CV_8UC3, Scalar(0,0,255));
2 cout << "M = " << endl << " " << M << endl << endl;

```

```
M =
[0, 0, 255, 0, 0, 255;
 0, 0, 255, 0, 0, 255]
```

For two dimensional and multichannel images we first define their size: row and column count wise.

Then we need to specify the **data type** to use for storing the elements and the number of channels per matrix point. To do this we have multiple definitions made according to the following convention:

```
1 CV_[number of bits per item][Type Prefix]C[Channels]
```

For instance, `CV_8UC3` means we use unsigned char types that are 8 bit long and each pixel has three items of this to form the three channels. This are predefined for up to four channel numbers. The **Scalar** is four element short vector. Specify this and you can initialize all matrix points with a custom value. However if you need more you can create the type with the upper macro and putting the channel number in parenthesis as you can see below.

- Use C\C++ arrays and initialize via constructor

```
1 int sz[3] = {2,2,2};
2 Mat L(3,sz, CV_8UC(1), Scalar::all(0));
```

The upper example shows how to create a matrix with more than two dimensions.

- Specify its dimension, then pass a pointer containing the size for each dimension and the rest remains the same.
- **Create()** function:

```
1 M.create(4,4, CV_8UC(2));
2 cout << "M = " << endl << " " << M << endl << endl;
```

```
M =
[205, 205, 205, 205, 205, 205, 205, 205;
 205, 205, 205, 205, 205, 205, 205, 205;
 205, 205, 205, 205, 205, 205, 205, 205;
 205, 205, 205, 205, 205, 205, 205, 205]
```

You cannot initialize the matrix values with this construction. It will only reallocate its matrix data memory if the new size will not fit into the old one.

- MATLAB style initializer: `zeros()`, `ones()`, `eye()`.
- Specify size and data type to use:

```

1  Mat E = Mat::eye(4, 4, CV_64F);
2  cout << "E = " << endl << " " << E << endl << endl;
3
4  Mat O = Mat::ones(2, 2, CV_32F);
5  cout << "O = " << endl << " " << O << endl << endl;
6
7  Mat Z = Mat::zeros(3,3, CV_8UC1);
8  cout << "Z = " << endl << " " << Z << endl << endl;

```

```

E =
[1, 0, 0, 0;
 0, 1, 0, 0;
 0, 0, 1, 0;
 0, 0, 0, 1]

Z =
[0, 0, 0;
 0, 0, 0;
 0, 0, 0]

O =
[1, 1;
 1, 1]

```

- For small matrices you may use comma separated initializers:

```

1  Mat C = (Mat_<double>(3,3) << 0, -1, 0, -1, 5, -1, 0, -1, 0);
2  cout << "C = " << endl << " " << C << endl << endl;

```

```

C =
[0, -1, 0;
 -1, 5, -1;
 0, -1, 0]

```

- Create a new header for an existing Mat object and `clone()` or `copyTo()` it.

```

1  Mat RowClone = C.row(1).clone();
2  cout << "RowClone = " << endl << " " << RowClone << endl << endl;

```

```

RowClone =
[-1, 5, -1]

```

Note: You can fill out a matrix with random values using the `randu()` function. You need to give the lower and upper value between what you want the random values:

```

1  Mat R = Mat(3, 2, CV_8UC3);
2  randu(R, Scalar::all(0), Scalar::all(255));

```

4.1.1 Access to Mat elements

In order to access to a single element of a matrix you can use its template method `Mat::at()`:

```
1 template<typename T> T& Mat::at(int i, int j)
```

where the type `T` is the type of the matrix we are accessing, so it can be `float`, `double`, `int` etc. The indices `i` and `j` are the indices of the element (always 0-based). For example:

```
1 // a 3x4 matrix of floats
2 Mat A = Mat(3, 4, CV_32FC1);
3
4 for(int i = 0; i < A.rows; ++i)
5 {
6     for(int j = 0; j < A.cols; ++j)
7     {
8         // set the value of the element
9         A.at<float>(i,j) = i*A.cols + j;
10        // get the value of the element value
11        cout << A.at<float>(i,j) << endl;
12    }
13 }
```

Access to elements in multichannel matrices When dealing with multichannel matrices (for example color images), each matrix element is no longer a single scalar value, but a *small fixed-size vector*, containing one value per channel. In this case, the template parameter of `Mat::at<>` must be a vector type.

OpenCV provides the `Vec` class to represent such fixed-size vectors. It is defined as:

```
1 template<typename _Tp, int cn> class Vec;
```

where `_Tp` is the data type of each channel and `cn` is the number of channels. For convenience, OpenCV defines several commonly used aliases:

- `Vec3b`: 3 channels of `uchar` (8-bit unsigned)
- `Vec4b`: 4 channels of `uchar` (e.g. BGR images)
- `Vec3f`: 3 channels of `float`
- `Vec4f`: 4 channels of `float`
- `Vec2i`: 2 channels of `int`

Here, the suffix `b` stands for `uchar` (`b` stands for “byte”, because `uchar` and `uchar` are coded in 1 byte), `f` for `float`, `i` for `int`, etc.

For example, a standard color image has type `CV_8UC3`, i.e. an 8-bit unsigned matrix with 3 channels. To access a pixel at coordinates `(i,j)`, we must therefore use `Vec3b`:

```
1 Vec3b pixel = img.at<Vec3b>(i, j);
```

Each channel can then be accessed using the `[]` operator:

```

1 uchar blue  = pixel[0];
2 uchar green = pixel[1];
3 uchar red   = pixel[2];

```

By default, OpenCV stores color images in *BGR* order, not RGB. You can modify individual channels directly:

```

1 img.at<Vec3b>(i, j)[0] = 0; // Blue
2 img.at<Vec3b>(i, j)[1] = 255; // Green
3 img.at<Vec3b>(i, j)[2] = 0; // Red

```

or assign all channels at once:

```

1 img.at<Vec3b>(i, j) = Vec3b(0, 255, 0);

```

The vector type used with `Mat::at<>` must always match the matrix type. For example:

- CV_32FC3 → Vec3f
- CV_8UC4 → Vec4b
- CV_64FC1 → double

As a general rule, for a matrix of type:

```
CV_<depth>C<X>
```

you should use:

```
Vec<X><depth-type>
```

This explicit typing ensures type safety and makes pixel access in multichannel images clear and unambiguous.

4.2 Print out formatting

In the above examples you could see the default formatting option. Nevertheless, OpenCV allows you to format your matrix output format to fit the rules of:

- Default

```

1 cout << "R (default) = " << endl << R << endl << endl;

```

```

R (default) =
[91, 2, 79, 179, 52, 205;
 236, 8, 181, 239, 26, 248;
 207, 218, 45, 183, 158, 101]

```

- Python

```

1 cout << "R (python) = " << endl << R << endl << endl;

```

```
R (python) =
[[[91, 2, 79], [179, 52, 205]],
 [[236, 8, 181], [239, 26, 248]],
 [[207, 218, 45], [183, 158, 101]]]
```

- Comma separated values (CSV)

```
1 cout << "R (csv) = " << endl << R << endl << endl;
```

```
R (csv) =
91, 2, 79, 179, 52, 205
236, 8, 181, 239, 26, 248
207, 218, 45, 183, 158, 101
```

- Numpy

```
1 cout << "R (numpy) = " << endl << R << endl << endl;
```

```
R (numpy) =
array([[[91, 2, 79], [179, 52, 205]],
 [[236, 8, 181], [239, 26, 248]],
 [[207, 218, 45], [183, 158, 101]]], type='uint8')
```

- C

```
1 cout << "R (c) = " << endl << R << endl << endl;
```

```
R (c) =
<91, 2, 79, 179, 52, 205,
236, 8, 181, 239, 26, 248,
207, 218, 45, 183, 158, 101>
```

4.3 Print for other common items

OpenCV offers support for print of other common OpenCV data structures too via the << operator like:

- 2D Point

```
1 Point2f P(5, 1);
2 cout << "Point (2D) = " << P << endl << endl;
```

```
Point (2D) = [5, 1]
```

- 3D Point

```

1 Point3f P3f(2, 6, 7);
2 cout << "Point (3D) = " << P3f << endl << endl;

```

```
Point (3D) = [2, 6, 7]
```

- `std::vector` via `cv::Mat`

```

1 vector<float> v;
2 v.push_back( (float)CV_PI);
3 v.push_back(2);
4 v.push_back(3.01f);
5
6 cout << "Vector of floats via Mat = " << Mat(v) << endl << endl;

```

```
Vector of floats via Mat = [3.1415927; 2; 3.01]
```

- `std::vector` of points

```

1 vector<Point2f> vPoints(20);
2 for (size_t i = 0; i < vPoints.size(); ++i)
3     vPoints[i] = Point2f((float)(i * 5), (float)(i % 7));
4
5 cout << "A vector of 2D Points = " << vPoints << endl << endl;

```

```
A vector of 2D Points = [0, 0; 5, 1; 10, 2; 15, 3; 20, 4; 25, 5; 30, 6; 35, 0; 40, 1; 45, 2; 50, 3; 55, 4; 60, 5; 65, 6; 70, 0; 75, 1; 80, 2; 85, 3; 90, 4; 95, 5]
```

Most of the samples here have been included into a small console application. You can download it from [here](#) or in the core section of the cpp samples.

5 Load and Display an Image

5.1 Goal

In this tutorial you will learn how to:

- Load an image (using `imread`)
- Create a named OpenCV window (using `namedWindow`)
- Display an image in an OpenCV window (using `imshow`)

5.2 Source Code

The source code is in `src/tutorials/display_image.cpp`. The program takes as input an image file and it displays it in a window

```
1  # include <opencv2 / core / core.hpp>
2  # include <opencv2 / highgui / highgui.hpp>
3  # include <iostream>
4  using namespace cv;
5  using namespace std;
6  int main( int argc, char** argv )
7  {
8      if( argc != 2)
9      {
10         cout << " Usage: display_image ImageToLoadAndDisplay" << endl;
11         return -1;
12     }
13     Mat image;
14     // Read the file
15     image = imread(argv[1], cv::IMREAD_COLOR);
16     if(! image.data )
17     // Check for invalid input
18     {
19         cout << "Could not open or find the image" << std::endl ;
20         return EXIT_FAILURE;
21     }
22     // Create a window for display.
23     namedWindow( "Display window", cv::WINDOW_AUTOSIZE );
24     // Show our image inside it.
25     imshow( "Display window", image );
26     // Wait for a keystroke in the wind
27     waitKey(0);
28     return EXIT_SUCCESS;
29 }
```

5.3 Explanation

In OpenCV 4 we have multiple modules. Each one takes care of a different area or approach towards image processing. You could already observe this in the structure of the user guide of these tutorials itself. Before you use any of them you first need to include the header files where the content of each individual module is declared.

You'll almost always end up using the:

- *core* section, as here are defined the basic building blocks of the library
- *highgui* module, as this contains the functions for input and output operations

```
1 # include <iostream> // for standard I/O
2 # include <string> // for strings
```

We also include the *iostream* to facilitate console line output and input. To avoid data structure and function name conflicts with other libraries, OpenCV has its own namespace: *cv*. To avoid the need appending prior each of these the *cv::* keyword you can import the namespace in the whole file by using the lines:

```
1 using namespace cv;
2 using namespace std;
```

This is true for the STL library too (used for console I/O). Now, let's analyze the *main* function. We start up assuring that we acquire a valid image name argument from the command line.

```
1 if( argc != 2)
2 {
3     cout <<" Usage: display_image ImageToLoadAndDisplay" << endl;
4     return EXIT_FAILURE;
5 }
```

Then create a *Mat* object that will store the data of the loaded image.

```
1 Mat image;
```

Now we call the *imread* function which loads the image name specified by the first argument (*argv[1]*). The second argument specifies the format in what we want the image. This may be:

- *cv::IMREAD_UNCHANGED* (<0) loads the image as is (including the alpha channel if present)
- *cv::IMREAD_GRAYSCALE* (0) loads the image as an intensity one
- *cv::IMREAD_COLOR* (>0) loads the image in the RGB format

```
1 // Read the file
2 image = imread(argv[1], cv::IMREAD_COLOR);
```

After checking that the image data was loaded correctly, we want to display our image, so we create an OpenCV window using the `namedWindow` function. These are automatically managed by OpenCV once you create them. For this you need to specify its name and how it should handle the change of the image it contains from a size point of view. It may be:

- `cv::WINDOW_AUTOSIZE` is the only supported one if you do not use the Qt backend.
- In this case the window size will take up the size of the image it shows. No resize permitted!
- `cv::WINDOW_NORMAL` on Qt you may use this to allow window resize.
- The image will resize itself according to the current window size.
- By using the `|` operator you also need to specify if you would like the image to keep its aspect ratio (`cv::WINDOW_KEEPRATIO`) or not (`cv::WINDOW_FREERATIO`).

```
1 // Create a window for display.
2 namedWindow( "Display window", cv::WINDOW_AUTOSIZE );
```

Finally, to update the content of the OpenCV window with a new image use the `imshow` function. Specify the OpenCV window name to update and the image to use during this operation:

```
1 // Show our image inside it.
2 imshow( "Display window", image );
```

Because we want our window to be displayed until the user presses a key (otherwise the program would end far too quickly), we use the `waitKey` function whose only parameter is just how long should it wait for a user input (measured in milliseconds). Zero means to wait forever.


```
1 //Wait for a keystroke in the window
2 waitKey(0);
```

5.4 Result

- Compile your code and then run the executable giving an image path as argument.
- For exemple from the build directory:

```
./bin/display_image ../data/images/HappyFish.png
```

- You should get a nice window as the one shown below:



6 Load, Modify, and Save an Image

Note: We assume that by now you know how to load an image using `imread` and to display it in a window (using `imshow`). Read the [Load and Display an Image](#) tutorial otherwise.

6.1 Goals

In this tutorial you will learn how to:

- Load an image using `imread`
- Transform an image from BGR to Grayscale format by using `cvtColor`
- Save your transformed image in a file on disk (using `imwrite`)

6.2 Code

The source code is in `src/tutorials/load_modify_image.cpp`. The program takes as input an image file, it converts it into a greyscale image, it displays the two images inside two separate windows and save the greyscale version into a file.

Here it is:

```
1 # include <cv.h>
2 # include <highgui.h>
3 using namespace cv;
4 int main( int argc, char** argv )
5 {
6     char* imageName = argv[1];
7     Mat image;
8     image = imread( imageName, 1 );
9     if( argc != 2 || !image.data )
10    {
11        printf( " No image data \n " );
12        return EXIT_FAILURE;
13    }
14    Mat gray_image;
15    cvtColor( image, gray_image, cv::COLOR_BGR2GRAY );
16    imwrite( "../data/images/Gray_Image.jpg", gray_image );
17    namedWindow( imageName, cv::WINDOW_AUTOSIZE );
18    namedWindow( "Gray image", cv::WINDOW_AUTOSIZE );
19    imshow( imageName, image );
20    imshow( "Gray image", gray_image );
21    waitKey(0);
22    return EXIT_SUCCESS;
23 }
```

6.3 Explanation

1. We begin by:

- Creating a Mat object to store the image information
- Load an image using `imread`, located in the path given by `imageName`.
- For this example, assume you are loading a RGB image.

2. Now we are going to convert our image from BGR to Grayscale format. OpenCV has a really nice function to do this kind of transformations:

```
1 cvtColor( image, gray_image, cv::COLOR_BGR2GRAY );
```

As you can see, `cvtColor` takes as arguments:

- a source image (`image`)

- a destination image (`gray_image`), in which we will save the converted image.
 - an additional parameter that indicates what kind of transformation will be performed.
 - In this case we use `cv::COLOR_BGR2GRAY` (because of `imread` has BGR default channel order in case of color images).
3. So now we have our new `gray_image` and want to save it on disk (otherwise it will get lost after the program ends).
 4. To save it, we will use a function analogous to `imread`: `imwrite`

```
1 imwrite( "../../images/Gray_Image.jpg", gray_image );
```

Which will save our `gray_image` as `Gray_Image.jpg` in the folder `images` located two levels up of my current location.


5. Finally, let's check out the images. We create two windows and use them to show the original image as well as the new one:

```
1 namedWindow( imageName, cv::WINDOW_AUTOSIZE );
2 namedWindow( "Gray image", cv::WINDOW_AUTOSIZE );
3 imshow( imageName, image );
4 imshow( "Gray image", gray_image );
```

6. Add the `waitKey(0)` function call for the program to wait forever for a user key press.

6.4 Result

When you run your program you should get something like this:



And if you check in your folder (in my case `images`), you should have a newly `.jpg` file named `Gray_Image.jpg`:



Congratulations, you are done with this tutorial!