Basic Color Spaces (OpenCV)

- **RGB (Red, Green, Blue)**: is an additive color space where colors are created by combining red, green, and blue light intensities. **OpenCV uses BGR** (Blue, Green, Red) instead of RGB as its default channel order.
  cv2.COLOR_BGR2RGB to convert from BGR to RGB.

- **Grayscale**: A single-channel representation of an image where each pixel represents intensity (brightness) on a scale from black (0) to white (255). Used when color information is not required.
  cv2.COLOR_BGR2GRAY to convert from BGR to grayscale.

- **HSV (Hue, Saturation, Value)**: HSV separates color information (hue) from intensity (value) and saturation. Hue refers to the type of color (e.g., red or blue), saturation measures the purity of the color, and value represents brightness. Use Case: Ideal for tasks like color-based segmentation because it simplifies the representation of colors.
  cv2.COLOR_BGR2HSV to convert from BGR to HSV.

https://en.wikipedia.org/wiki/Color_model#/media/File:RGBCube_a.svg
https://en.wikipedia.org/wiki/RGB_color_model

Color Space RGB

The RGB model works by adding light. When the intensities of red, green, and blue are all at their maximum, the result is white light. When all are at their minimum (zero intensity), the result is black.

Each color channel (R, G, B) typically has a range of values, often from 0 to 255 For example:

(255,0,0)(255,0,0): Pure red
(0,255,0)(0,255,0): Pure green
(0,0,255)(0,0,255): Pure blue
(255,255,0)(255,255,0): Yellow (red + green)

The RGB color space is widely used in electronic displays like computer screens, TVs, and cameras because these devices emit light.

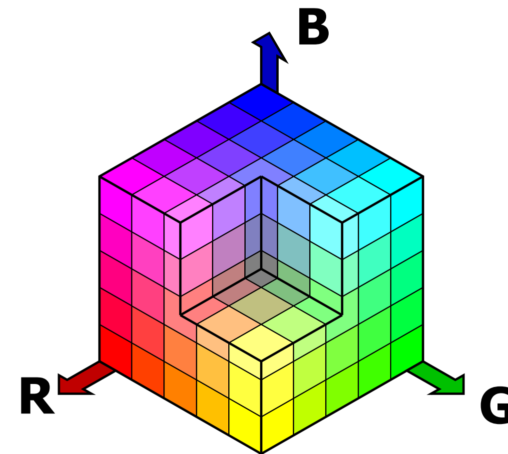

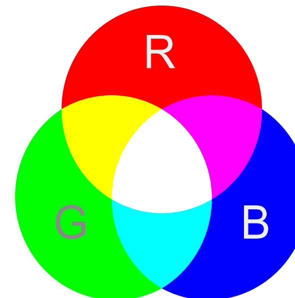

https://en.wikipedia.org/wiki/Color_model#/media/File:RGBCube_a.svg
https://en.wikipedia.org/wiki/RGB_color_model

Color Space RGB

Since colors in the RGB color space are coded using the three channels (in the range *[0, 255]*), it is more difficult to segment an object in the image based on its color.

```
1 import cv2
2 import numpy as np
3 import matplotlib.pyplot as plt
4
5 # Read the image
6 image = cv2.imread("red-tlight.png")
7
8 # Resize the image
9 image = cv2.resize(image, (100, 200))
10
11 # Obtain shape of the image
12 h, w, c = image.shape
13
14 # Split BRG channels
15 b, g, r = cv2.split(image)
```

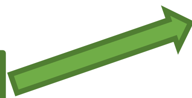

Original Image

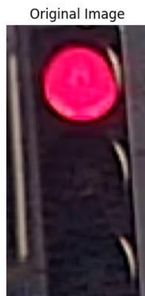

Red Channel

Green Channel

Blue Channel

https://en.wikipedia.org/wiki/Color_model#/media/File:RGBCube_a.svg

Color Space HSV

HSV colorspace

HSV (hue, saturation, value) colorspace is a model to represent the colorspace similar to the RGB color model. Since the hue channel models the color type, it is very useful in image processing tasks that need to segment objects based on its color. Variation of the saturation goes from unsaturated to represent shades of gray and fully saturated (no white component). Value channel describes the brightness or the intensity of the color. Next image shows the HSV cylinder.

For an HSV image, you can use cv.inRange() to isolate colors by specifying lower and upper HSV bounds.
Use Case: Ideal for color segmentation or detecting objects based on specific ranges of pixel intensities.

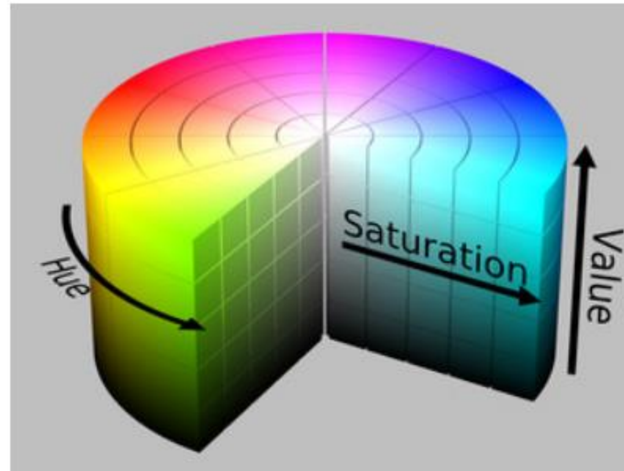

By SharkDderivative work: SharkD [CC BY-SA 3.0 or GFDL], via Wikimedia Commons

https://docs.opencv.org/4.11.0/df/d9d/tutorial_py_colorspaces.html
https://docs.opencv.org/4.11.0/da/d97/tutorial_threshold_inRange.html
https://colorizer.org/

How to find HSV values to track?

This is a common question found in stackoverflow.com. It is very simple and you can use the same function, **cv.cvtColor()**. Instead of passing an image, you just pass the BGR values you want. For example, to find the HSV value of Green, try the following commands in a Python terminal:

```
>>> green = np.uint8([[[0,255,0 ]]])
>>> hsv_green = cv.cvtColor(green,cv.COLOR_BGR2HSV)
>>> print( hsv_green )
[[[ 60 255 255]]]
```

Now you take [H-10, 100,100] and [H+10, 255, 255] as the lower bound and upper bound respectively. Apart from this method, you can use any image editing tools like GIMP or any online converters to find these values, but don't forget to adjust the HSV ranges.

Note

For HSV, hue range is [0,179], saturation range is [0,255], and value range is [0,255]. Different software use different scales. So if you are comparing OpenCV values with them, you need to normalize these ranges.

https://docs.opencv.org/4.11.0/df/d9d/tutorial_py_colorspaces.html
https://docs.opencv.org/4.11.0/da/d97/tutorial_threshold_inRange.html
https://colorizer.org/

Thresholding?

- The simplest segmentation method
- Application example: Separate out regions of an image corresponding to objects which we want to analyze. This separation is based on the variation of intensity between the object pixels and the background pixels.
- To differentiate the pixels we are interested in from the rest (which will eventually be rejected), we perform a comparison of each pixel intensity value with respect to a *threshold* (determined according to the problem to solve).
- Once we have separated properly the important pixels, we can set them with a determined value to identify them (i.e. we can assign them a value of 0 (black), 255 (white) or any value that suits your needs).

https://docs.opencv.org/4.11.0/db/d8e/tutorial_threshold.html

Types of Thresholding

- OpenCV offers the function **cv::threshold** to perform thresholding operations.
- We can effectuate 5 types of Thresholding operations with this function. We will explain them in the following subsections.
- To illustrate how these thresholding processes work, let's consider that we have a source image with pixels with intensity values $src(x, y)$. The plot below depicts this. The horizontal blue line represents the threshold $thresh$ (fixed).

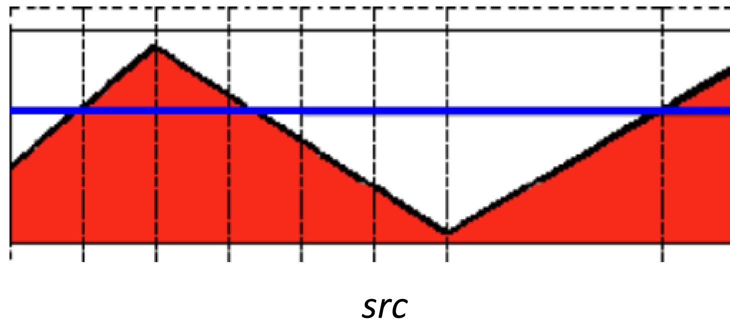

src

https://docs.opencv.org/4.11.0/db/d8e/tutorial_threshold.html

Thresholding

Types of Thresholding

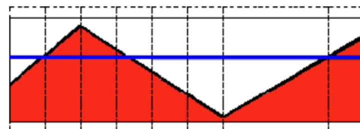

src

Threshold Binary

- This thresholding operation can be expressed as:

$$\texttt{dst}(x,y) = \begin{cases} \texttt{maxVal} & \text{if } \texttt{src}(x,y) > \texttt{thresh} \\ 0 & \text{otherwise} \end{cases}$$

- So, if the intensity of the pixel $src(x,y)$ is higher than $thresh$, then the new pixel intensity is set to a $MaxVal$. Otherwise, the pixels are set to 0.

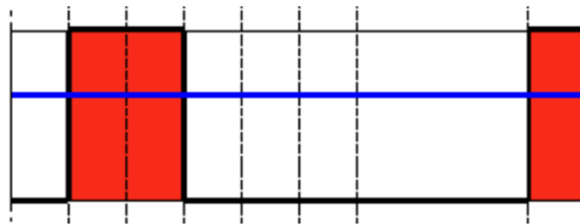

https://docs.opencv.org/4.11.0/db/d8e/tutorial_threshold.html

Thresholding

Types of Thresholding

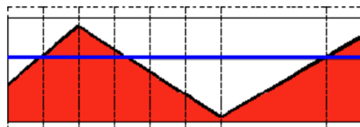

src

Threshold Binary, Inverted

- This thresholding operation can be expressed as:

$$\texttt{dst}(x, y) = \begin{cases} 0 & \text{if } \texttt{src}(x, y) > \texttt{thresh} \\ \texttt{maxVal} & \text{otherwise} \end{cases}$$

- If the intensity of the pixel $src(x, y)$ is higher than $thresh$, then the new pixel intensity is set to a 0. Otherwise, it is set to $MaxVal$.

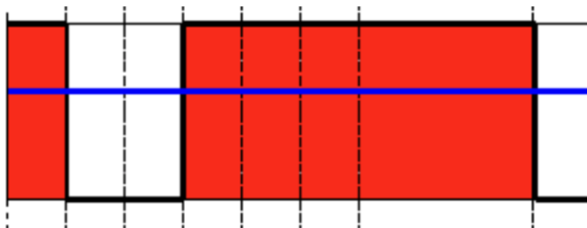

https://docs.opencv.org/4.11.0/db/d8e/tutorial_threshold.html

Thresholding + Morph.

Threshold Binary, Inverted

- This thresholding operation can be expressed as:

$$\texttt{dst}(x,y) = \begin{cases} 0 & \text{if } \texttt{src}(x,y) > \texttt{thresh} \\ \texttt{maxVal} & \text{otherwise} \end{cases}$$

- If the intensity of the pixel $src(x,y)$ is higher than $thresh$, then the new pixel intensity is set to a 0. Otherwise, it is set to $MaxVal$.

th_lst = [50, 100, 150, 200]

cv.threshold(img_input, th_lst[i], 255, cv.THRESH_BINARY_INV)

th: 50 | th: 100 | th: 150 | th: 200

https://docs.opencv.org/4.x/d7/d4d/tutorial_py_thresholding.html https://docs.opencv.org/4.11.0/db/d8e/tutorial_threshold.html https://docs.opencv.org/4.11.0/d3/db4/tutorial_py_watershed.html

Thresholding + Morph.

```
img_input = cv.imread('mon1.png', 0)
kernel = cv.getStructuringElement(cv.MORPH_RECT, (11, 11))
ret, img_th = cv.threshold(img_input, 50, 255, cv.THRESH_BINARY_INV)
img_erode = cv.erode(img_th, kernel, iterations = 1)
```

Morphological Operation

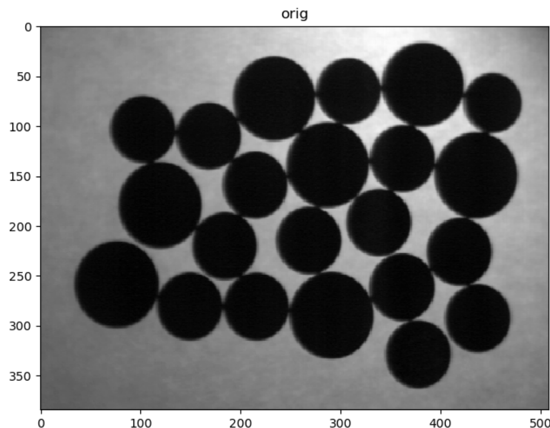

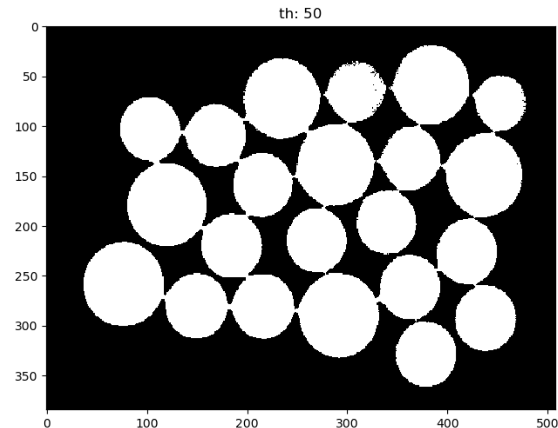

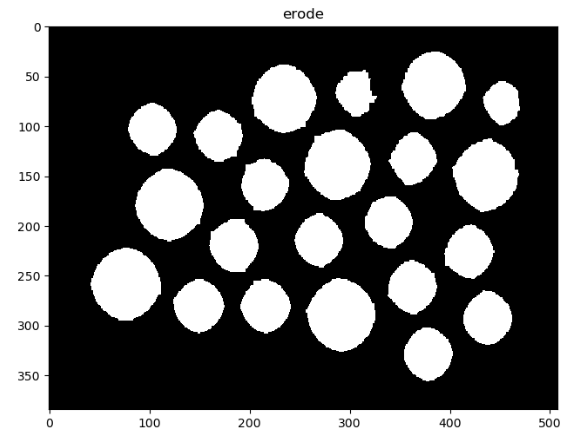

https://docs.opencv.org/4.x/d7/d4d/tutorial_py_thresholding.html https://docs.opencv.org/4.11.0/db/d8e/tutorial_threshold.html https://docs.opencv.org/4.11.0/d3/db4/tutorial_py_watershed.html

Morphological Operation

- Simple operations that are very often used on binary images
  - noise elimination
  - isolating individual elements and connecting disparate elements in the image

- To perform morphological operations, we need input image and kernel (structuring element)

- The morphological operations (most common) that we will discuss deeper are:

  - Erosion
  - Dilation
  - Opening
  - Closing
  - Gradient

https://docs.opencv.org/4.10.0/dd/dd7/tutorial_morph_lines_detection.html

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Morphological Operation

- The main principle is that the kernel is moved through to the input image
- The corresponding pixels inside the kernel vs. input image are examined
- Based on the operation (erosion, dilation), the pixels are for example eroded (made to zero)

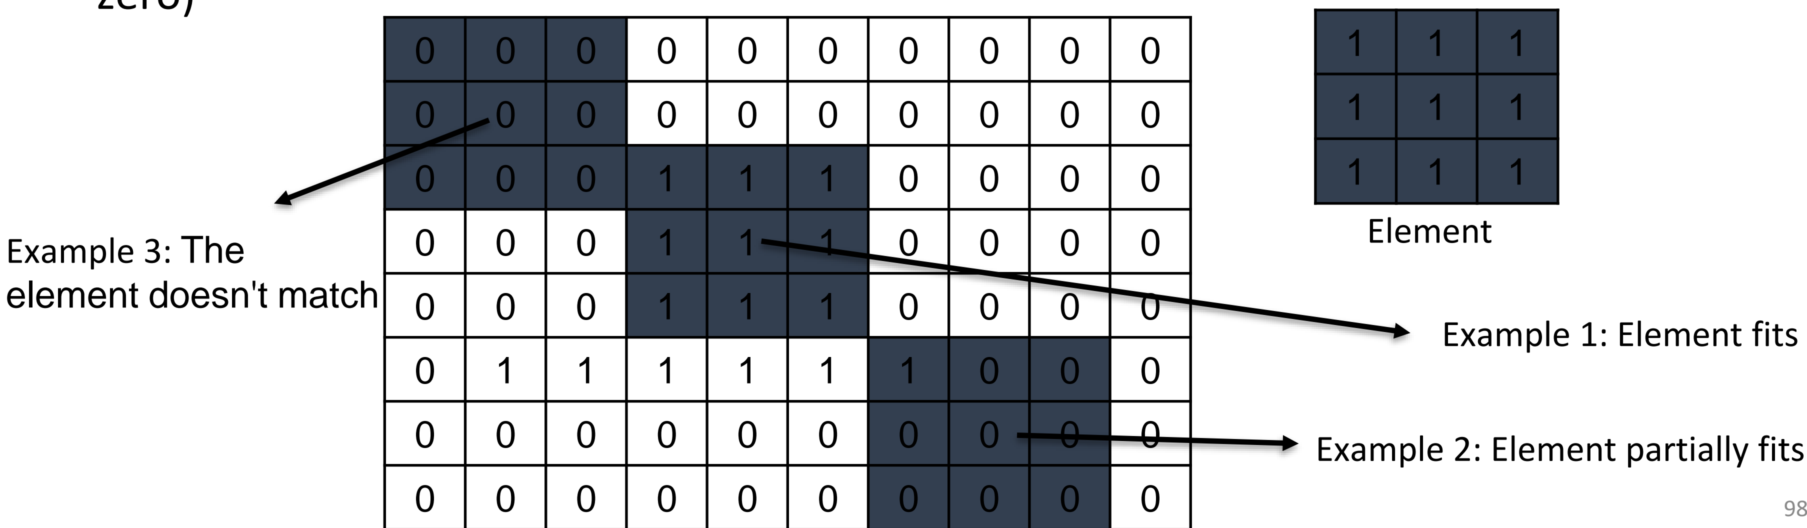

Erosion

The structuring element is placed over the image such that its origin aligns with the pixel being processed.

A pixel in the original image will be considered 1 only if all the pixels under the kernel is 1, otherwise it is eroded (made to zero).

Effects of Erosion and Applications:

- Shrinks objects in an image.
- Removes small noise or isolated bright spots.
- Enlarges gaps or holes within objects.
- Simplifies object boundaries by smoothing sharp edges.
- Object separation.
- Boundary simplification.
- Noise removal.

1	1	1
1	**1**	1
1	1	1

For symmetric structuring elements/kernels (e.g., 3×3 squares, circles), the origin is typically at the center. For example, in a 3×3 square kernel, the center pixel is the origin.

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Erosion

The structuring element is placed over the image such that its origin aligns with the pixel being processed.

A pixel in the original image will be considered 1 only if all the pixels under the kernel is 1, otherwise it is eroded (made to zero).

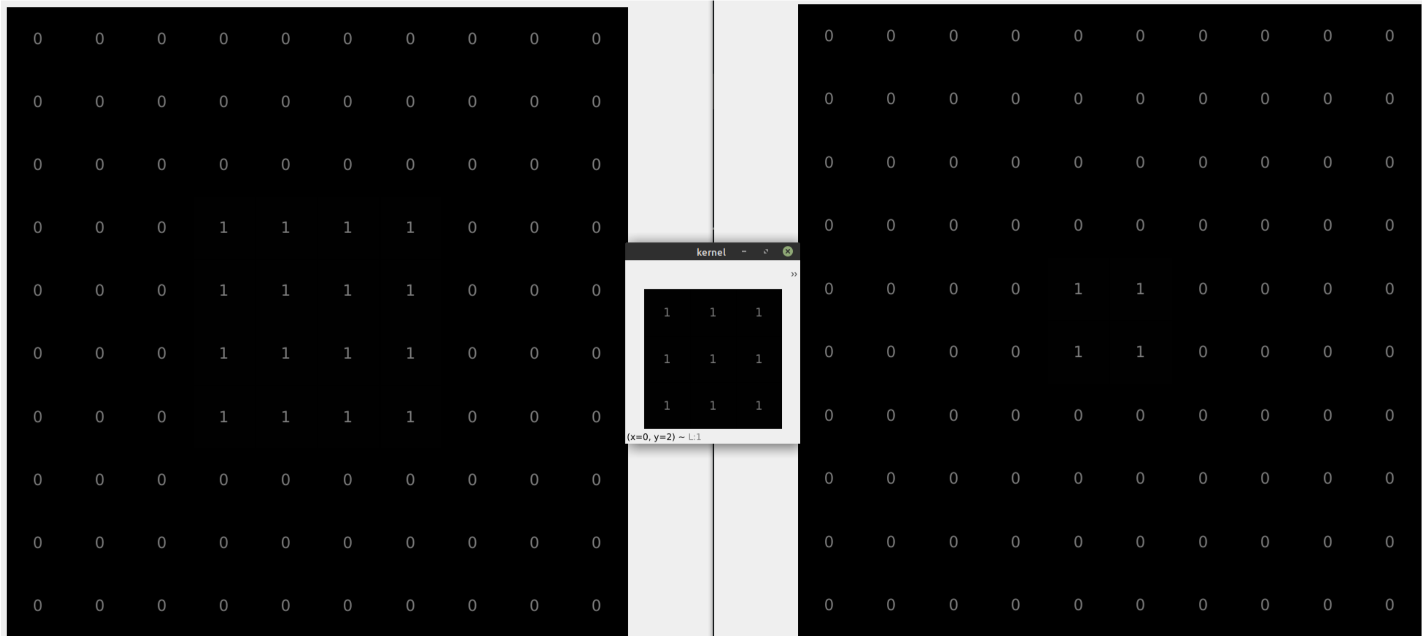

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Erosion

The structuring element is placed over the image such that its origin aligns with the pixel being processed.

A pixel in the original image will be considered 1 only if all the pixels under the kernel is 1, otherwise it is eroded (made to zero).

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Dilation

Opposite of erosion: a pixel element is '1' if at least one pixel under the kernel is '1'.

Dilation in image processing expands objects by adding pixels to their boundaries using a structuring element.

Effects of Erosion and Applications:

- Expands object boundaries: Enlarges foreground regions (e.g., white in binary images) by adding pixels around edges.
- Fills small gaps/holes: Bridges minor breaks or cracks within objects, making them more contiguous.
- Merges nearby objects: Connects disjointed regions if they fall within the structuring element’s reach.

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Dilation

Opposite of erosion: a pixel element is '1' if at least one pixel under the kernel is '1'.

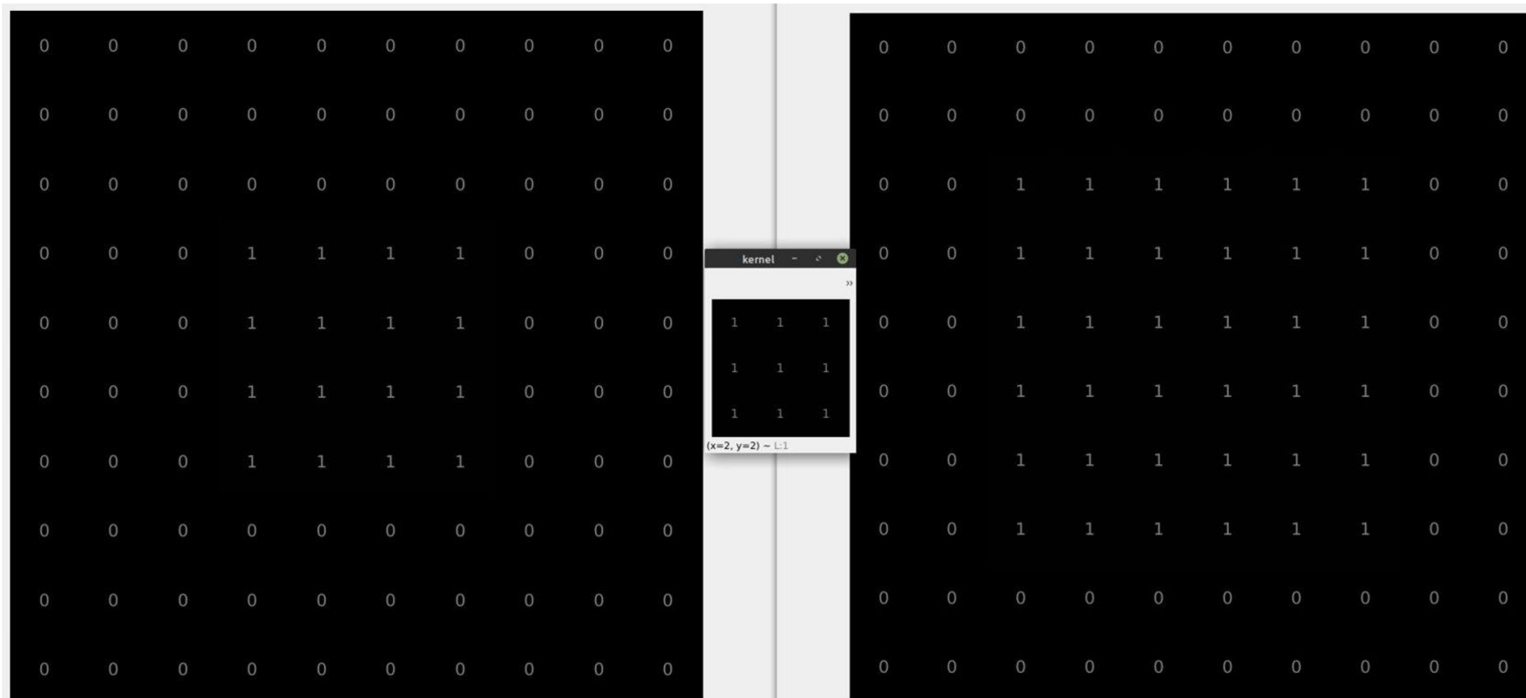

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Dilation

Opposite of erosion: a pixel element is '1' if at least one pixel under the kernel is '1'.

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

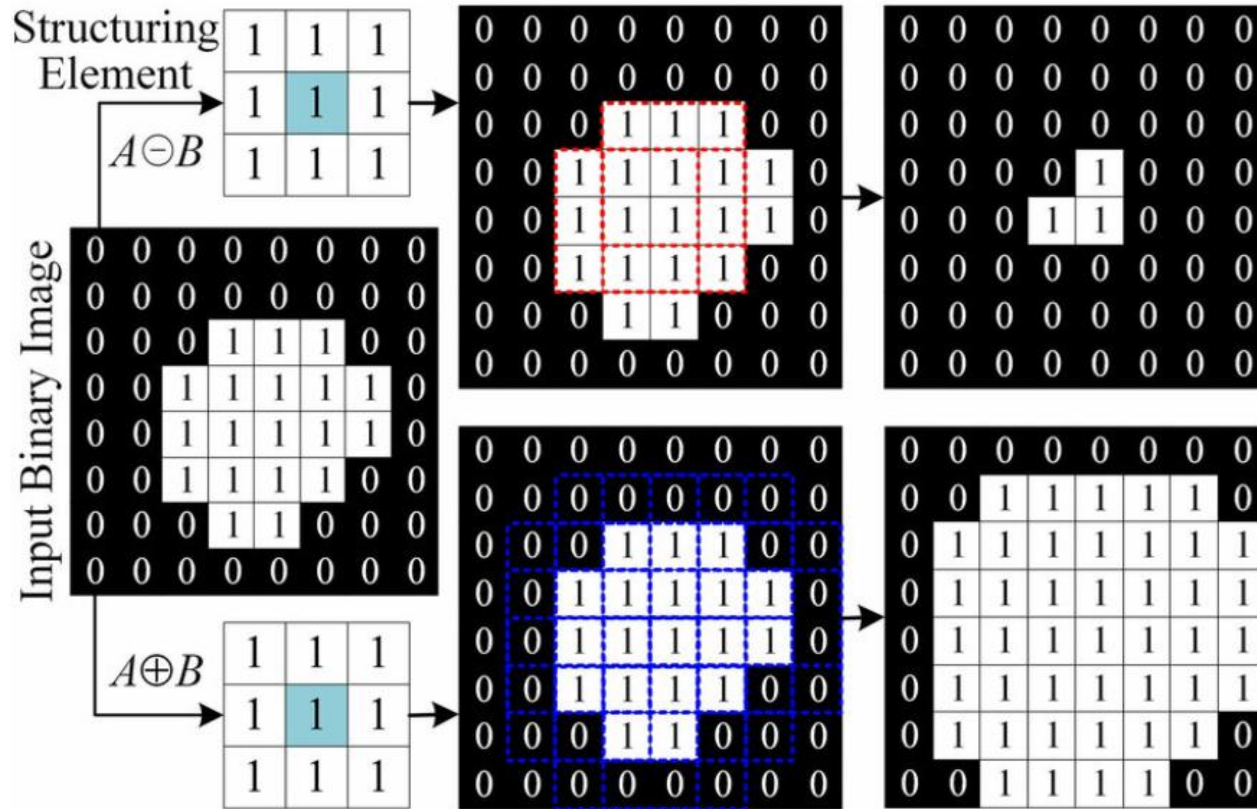

Example of erosion and dilation operation by a 3 × 3 structuring elements.

https://www.researchgate.net/figure/Example-of-erosion-and-dilation-operation-by-a-3-3-structuring-elements_fig2_372416628

Opening

Opening: **erosion** followed by **dilation**

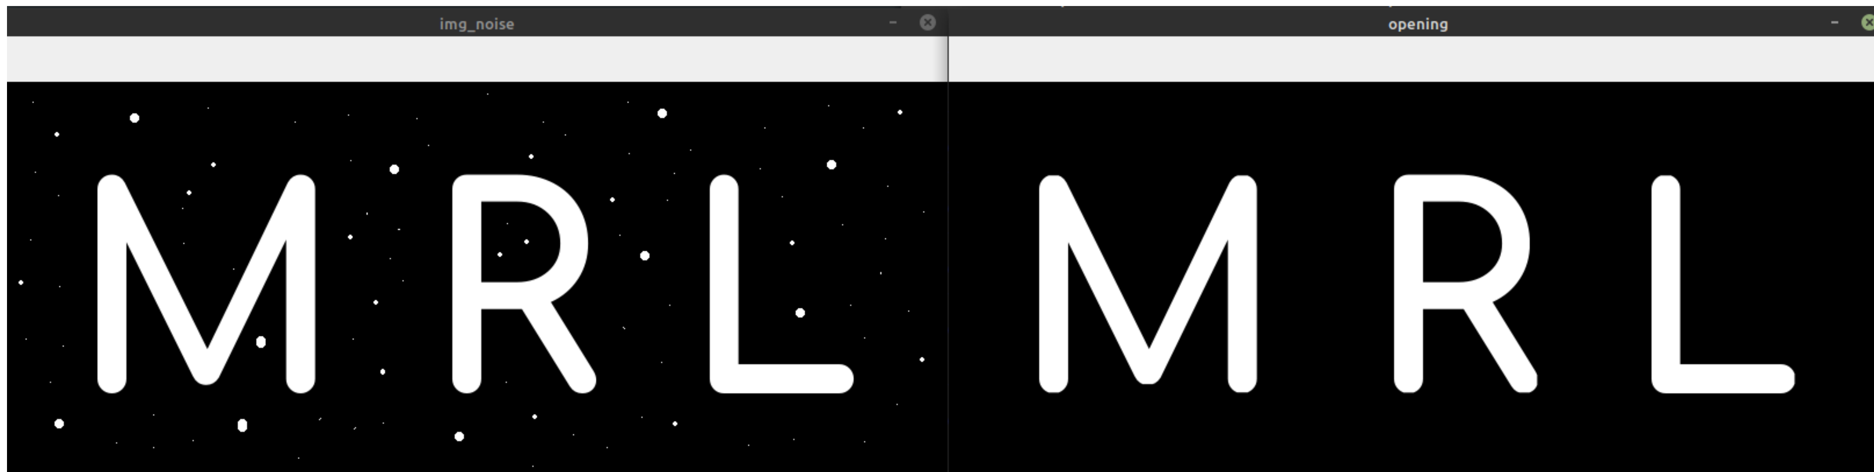

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Closing

Opening: **dilation** followed by **erosion**

Morphological Gradient

Difference between **dilation** and **erosion** of an image

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Structuring Elements

```
kernel = np.ones((5,5), np.uint8)

[[1 1 1 1 1]
 [1 1 1 1 1]
 [1 1 1 1 1]
 [1 1 1 1 1]
 [1 1 1 1 1]]
```

```
kernel = cv.getStructuringElement(cv.MORPH_CROSS, (5,5))

[[0 0 1 0 0]
 [0 0 1 0 0]
 [1 1 1 1 1]
 [0 0 1 0 0]
 [0 0 1 0 0]]
```

```
kernel = cv.getStructuringElement(cv.MORPH_RECT, (5,5))

[[1 1 1 1 1]
 [1 1 1 1 1]
 [1 1 1 1 1]
 [1 1 1 1 1]
 [1 1 1 1 1]]
```

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Structuring Elements

```
kernel = cv.getStructuringElement(cv.MORPH_ELLIPSE, (15,15))

[[0 0 0 0 0 0 0 1 0 0 0 0 0 0 0]
 [0 0 0 1 1 1 1 1 1 1 1 1 0 0 0]
 [0 0 1 1 1 1 1 1 1 1 1 1 1 0 0]
 [0 1 1 1 1 1 1 1 1 1 1 1 1 1 0]
 [0 1 1 1 1 1 1 1 1 1 1 1 1 1 0]
 [1 1 1 1 1 1 1 1 1 1 1 1 1 1 1]
 [1 1 1 1 1 1 1 1 1 1 1 1 1 1 1]
 [1 1 1 1 1 1 1 1 1 1 1 1 1 1 1]
 [1 1 1 1 1 1 1 1 1 1 1 1 1 1 1]
 [1 1 1 1 1 1 1 1 1 1 1 1 1 1 1]
 [0 1 1 1 1 1 1 1 1 1 1 1 1 1 0]
 [0 1 1 1 1 1 1 1 1 1 1 1 1 1 0]
 [0 0 1 1 1 1 1 1 1 1 1 1 1 0 0]
 [0 0 0 1 1 1 1 1 1 1 1 1 0 0 0]
 [0 0 0 0 0 0 0 1 0 0 0 0 0 0 0]]
```


https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Grayscale/Color

Idea of binary morphology can be extended to gray/color images with the use of max (Dilation) and min (Erosion) operation

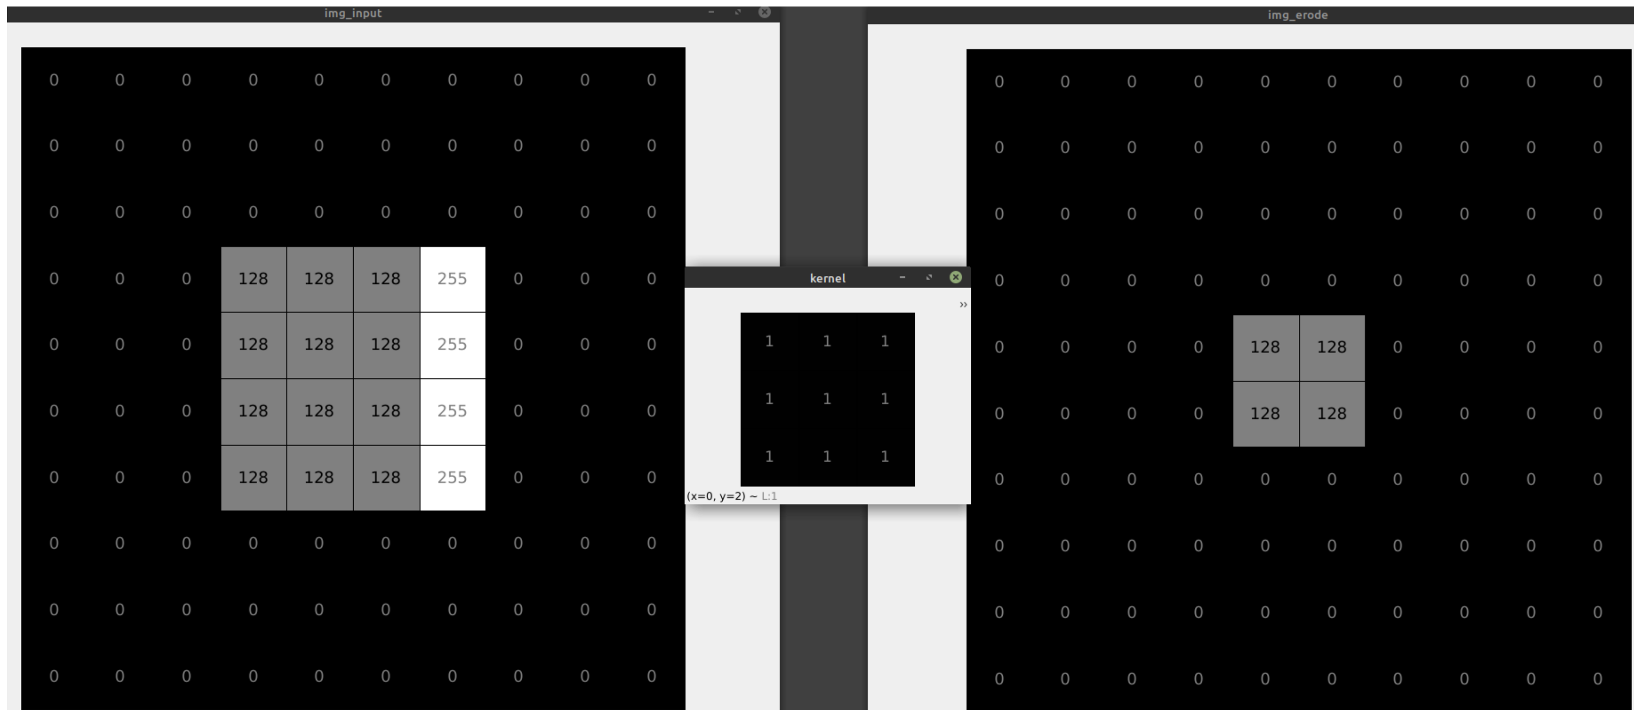

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Grayscale/Color

Idea of binary morphology can be extended to gray/color images with the use of max (Dilation) and min (Erosion) operation

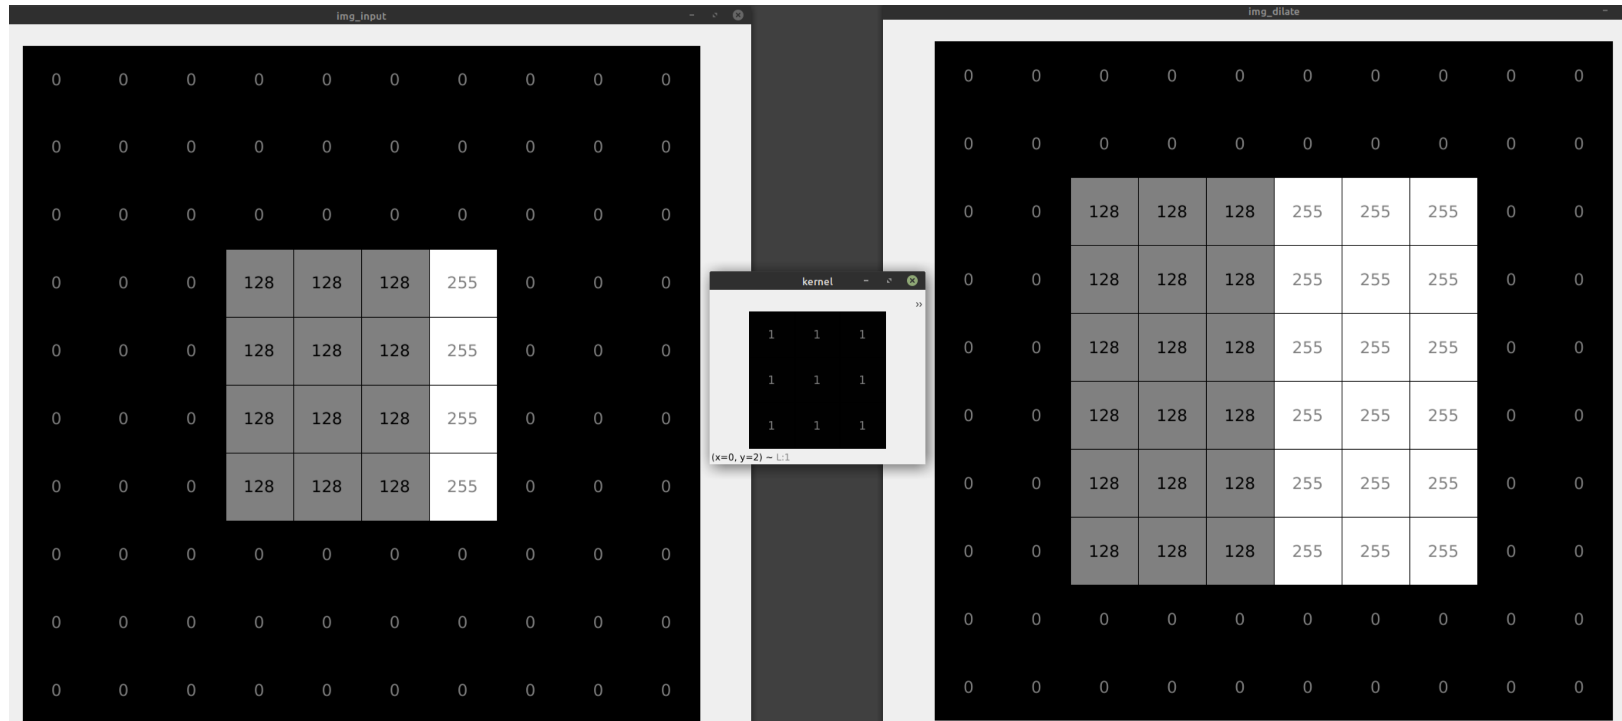

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Grayscale/Color

Idea of binary morphology can be extended to gray/color images with the use of max (dilation) and min (erosion) operation. Dilation and erosion change the brightness.
dilation - increased overall brightnes vs. erosion - dims the image

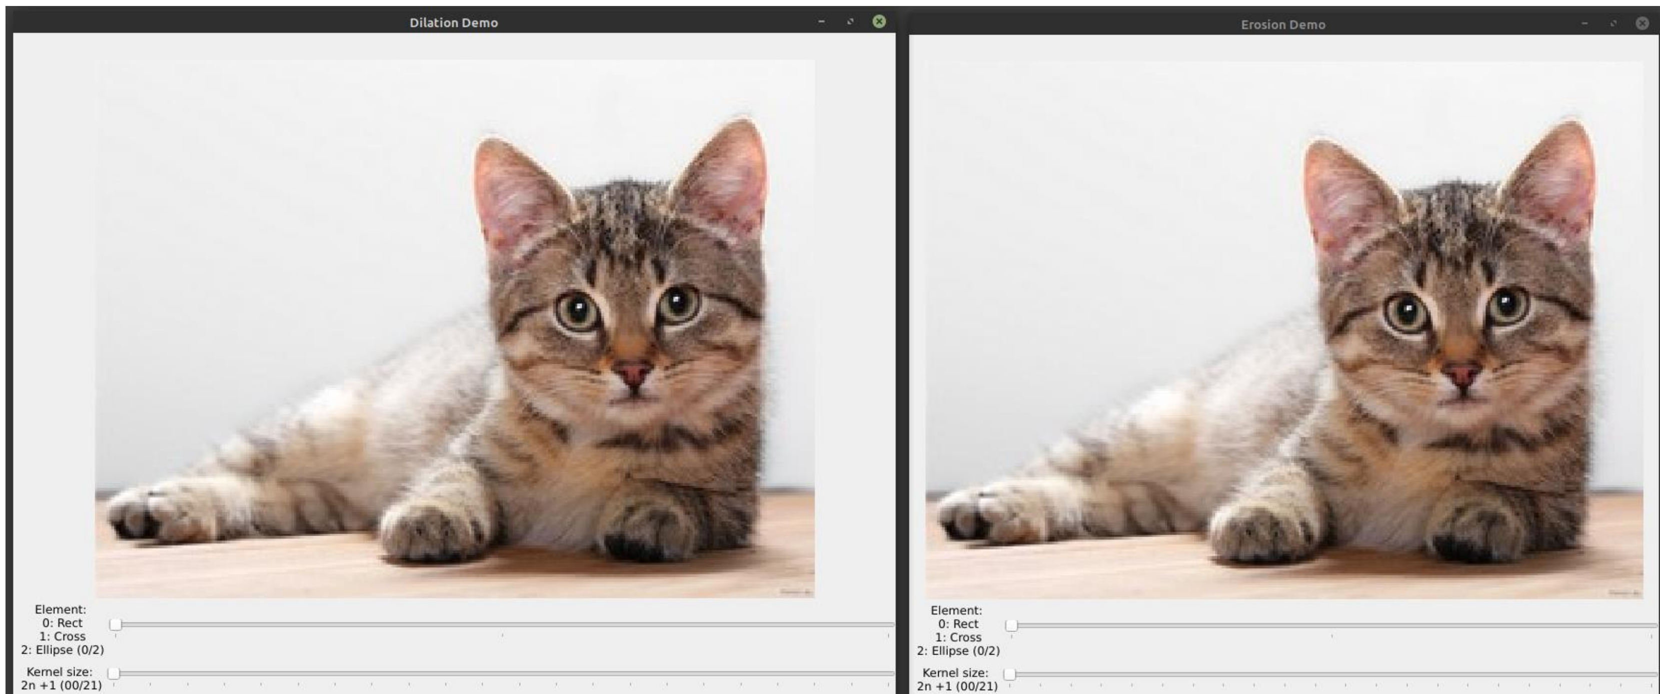

https://docs.opencv.org/4.11.0/d9/d61/tutorial_py_morphological_ops.html

https://docs.opencv.org/4.11.0/db/df6/tutorial_erosion_dilatation.html

Grayscale/Color

Idea of binary morphology can be extended to gray/color images with the use of max (dilation) and min (erosion) operation. Dilation and erosion change the brightness.
dilation - increased overall brightnes vs. erosion - dims the image

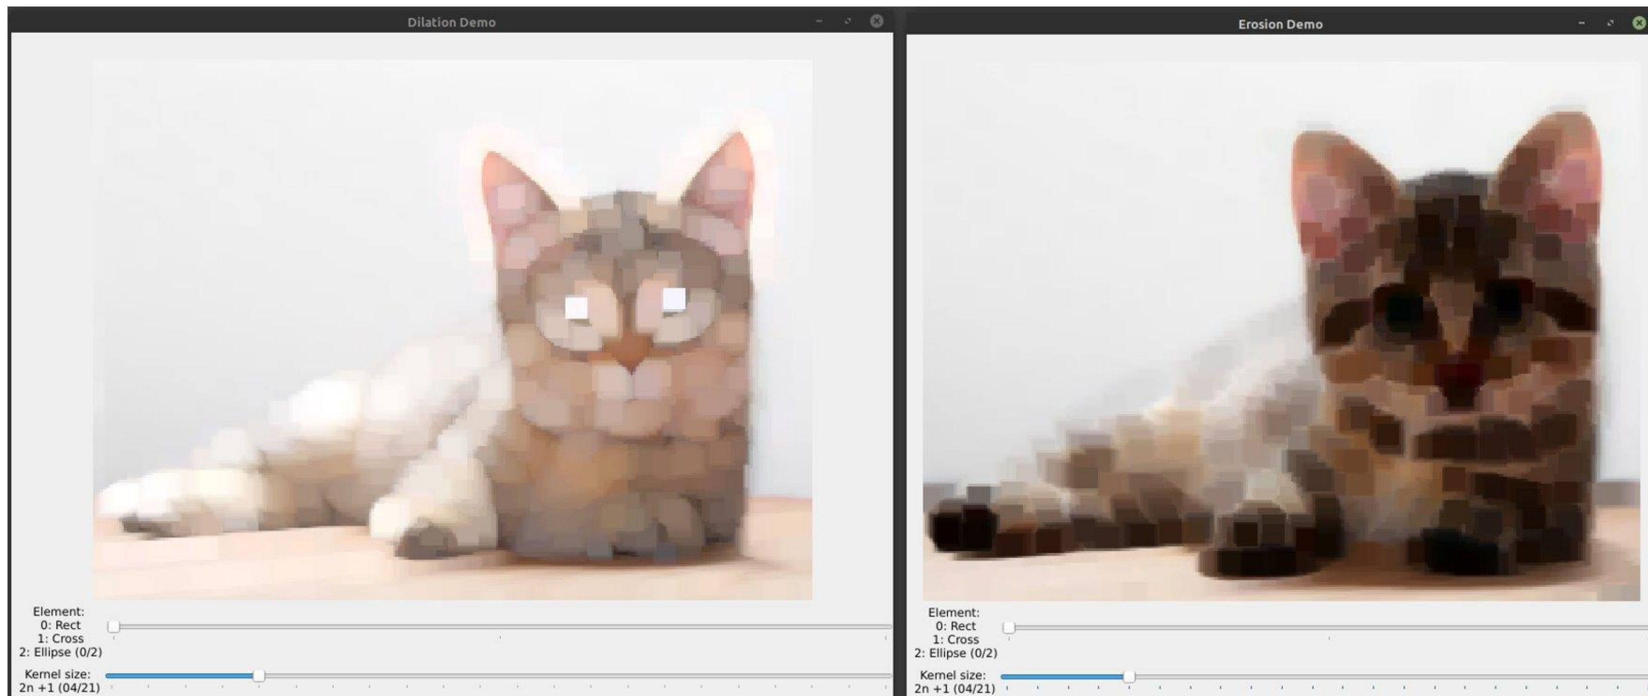

Grayscale/Color

Idea of binary morphology can be extended to gray/color images with the use of max (dilation) and min (erosion) operation. Dilation and erosion change the brightness.
dilation - increased overall brightnes vs. erosion - dims the image

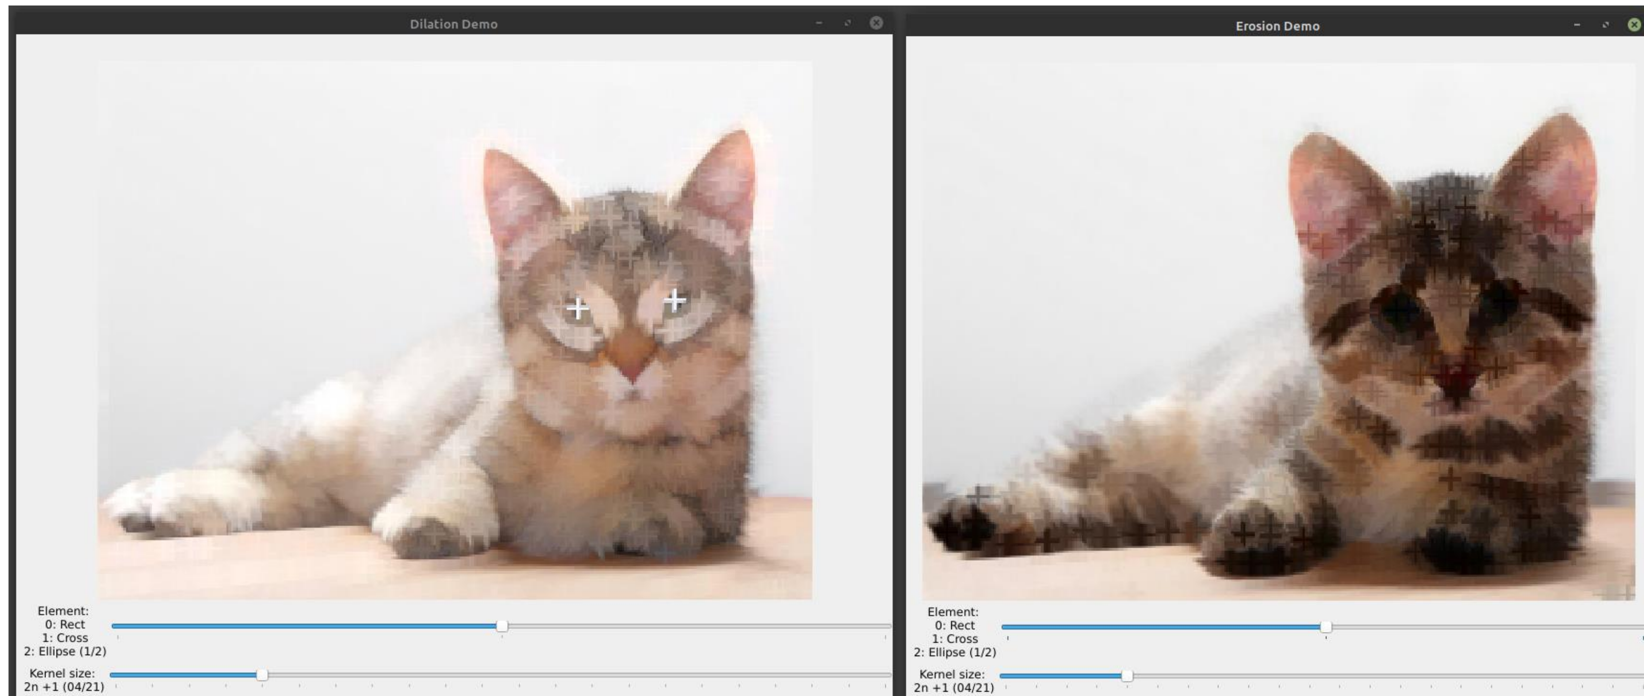

Grayscale/Color

Idea of binary morphology can be extended to gray/color images with the use of max (dilation) and min (erosion) operation. Dilation and erosion change the brightness.
dilation - increased overall brightnes vs. erosion - dims the image

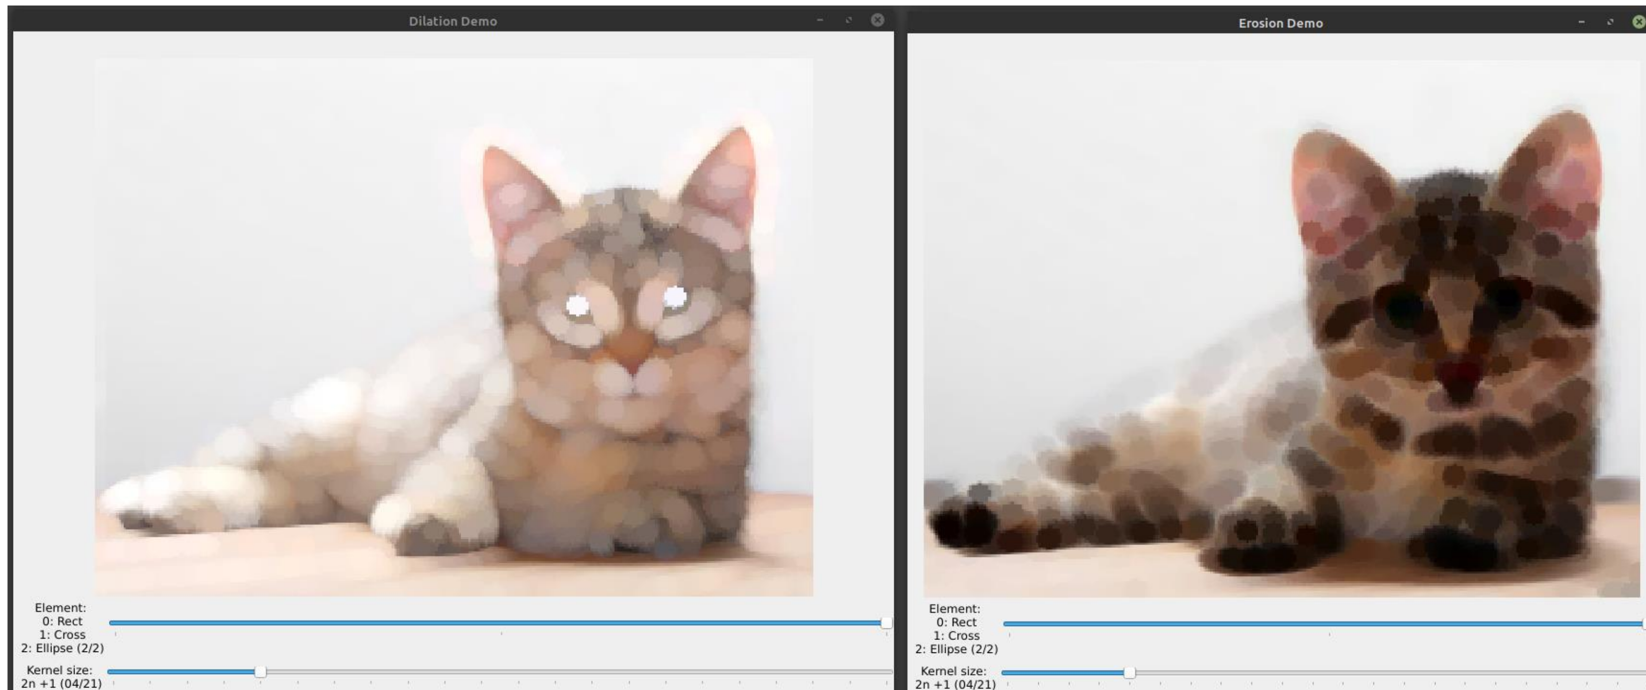

Grayscale/Color

Analogous way: a morphological gradient is the difference between a dilation and an erosion