ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Image Filtering

Image Processing & Analysis for Life Scientists

Olivier Burri, Romain Guiet & Arne Seitz

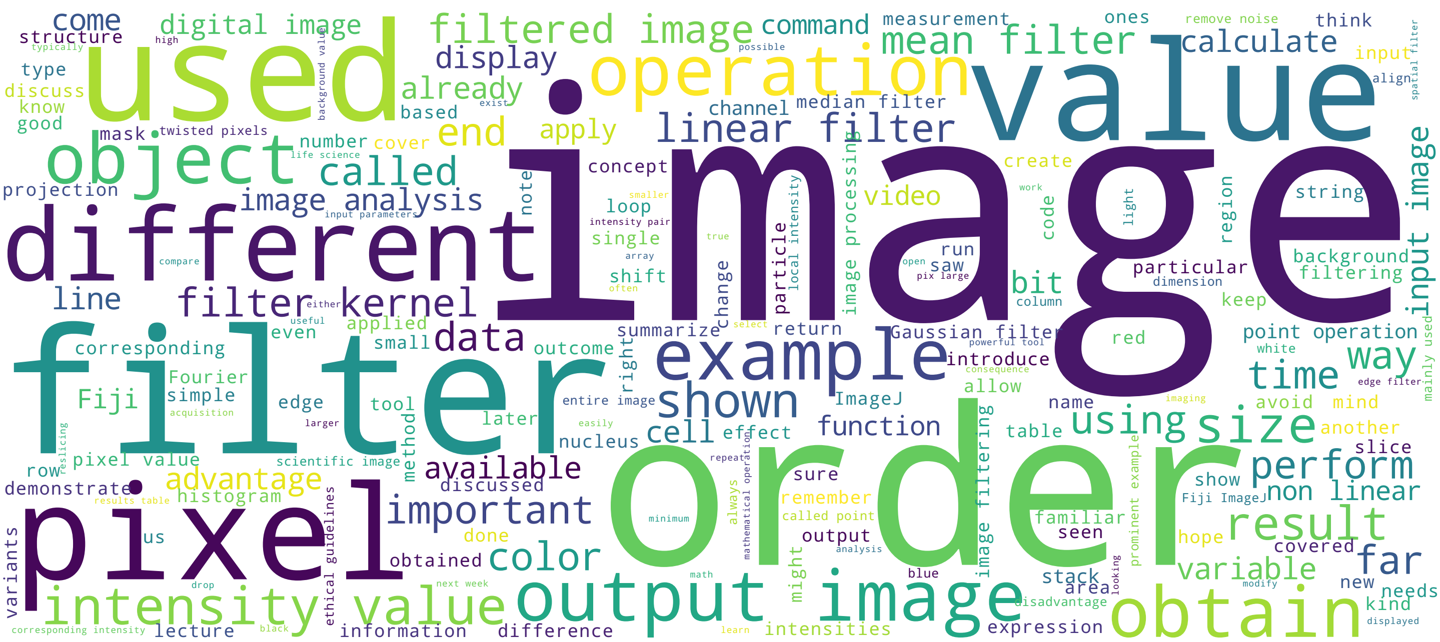

Search MOOC

Video

EPFL

Image Filtering

- What is Filtering
  - Linear Filters
  - Nonlinear Filters
- Uses in Biological Images
  - Artifact/Noise removal
  - Background subtraction
  - Feature enhancement

Welcome to this lecture in which I want to to introduce image filtering. I want to discuss the differences between the point operations we have encountered so far, and the concept of filtering. Moreover, I want to introduce linear and nonlinear filters. In order to avoid to get too theoretical, I will demonstrate the practical use of filters on typical images which are stemming from life sciences. I hope to demonstrate that filters are powerful tools in order to remove noise or similar artifacts from your images. They can also be used for background subtractions; or to enhance particular features within your images.

Notes

Summary

0m 07s

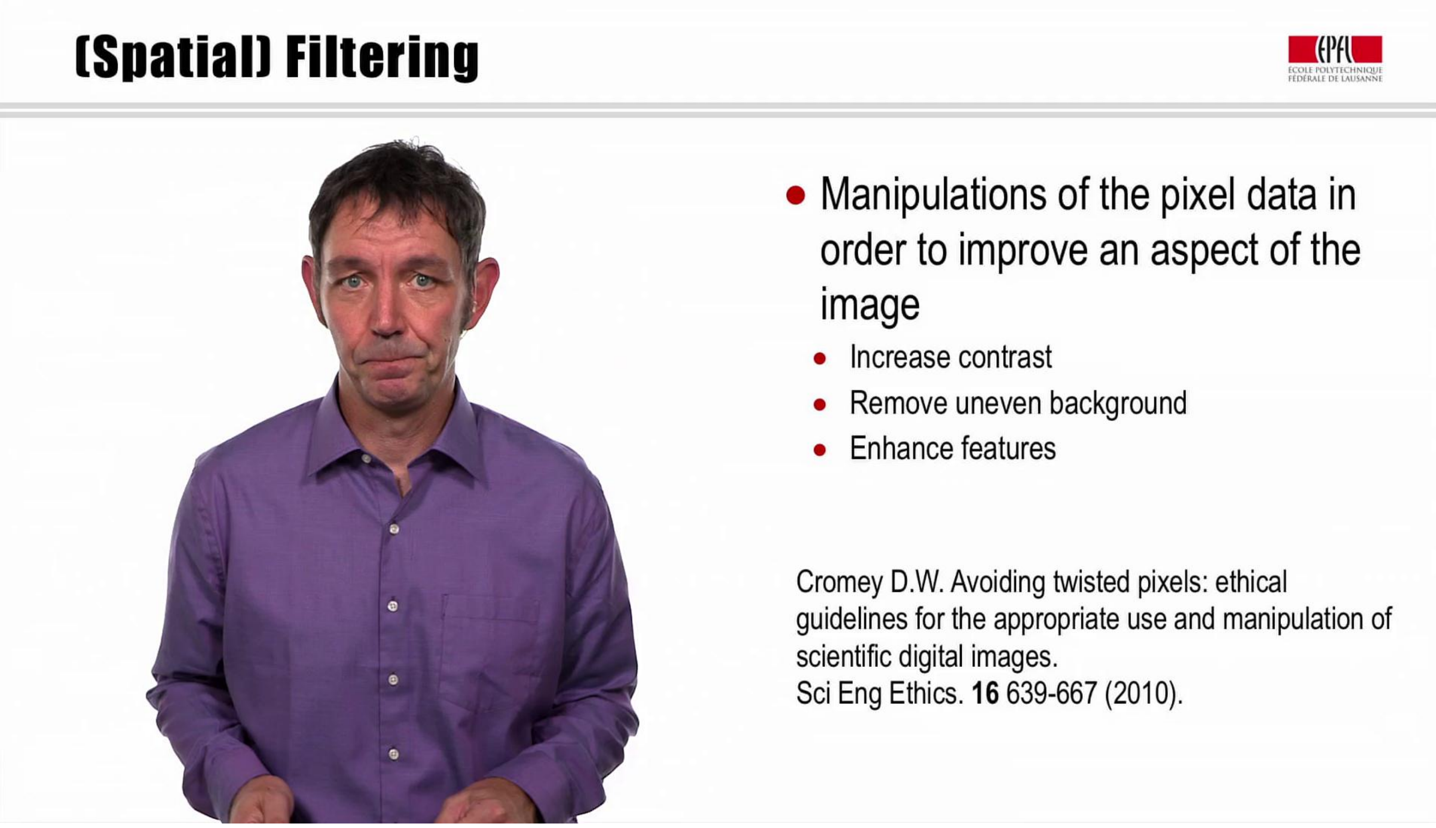

It is important that you keep in mind that filtering is an operation which is manipulating underlying data. Therefore, whenever using filters it needs to be documented. In particular in scientific publications. Also, if features have been enhanced, or some sort of background was removed via a filter, this needs to be mentioned. Otherwise, you're producing so called twisted pixels. If you want to learn more about ethical guidelines regarding digital images manipulations you just have to look a the article Avoiding twisted pixels: ethical guidelines for the appropriate use and manipulation of scientific digital images from Douglas W. Cromey. There, you will find 12 guidelines for proper images handling of scientific images.

Notes

Summary

0m 49s

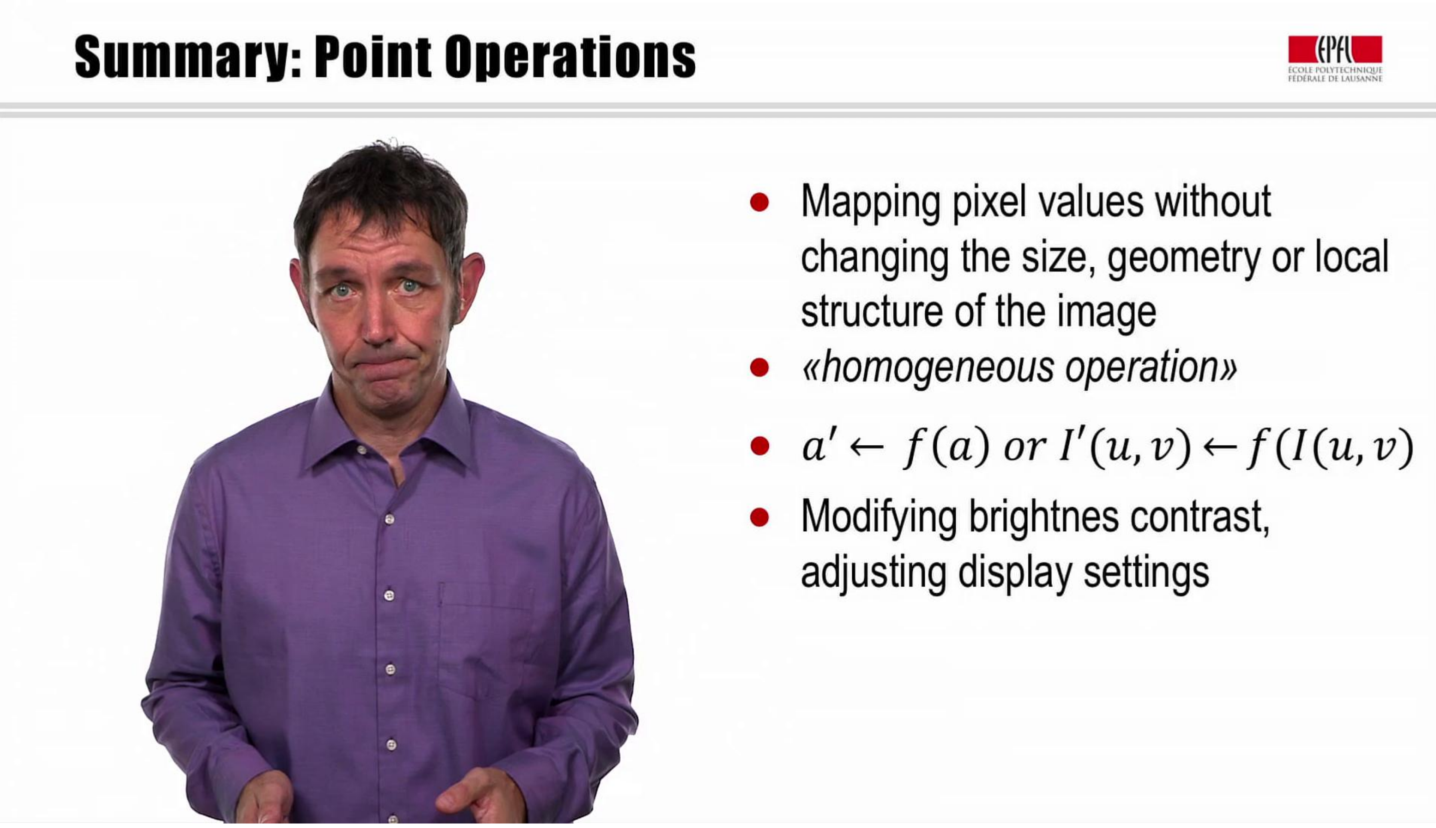

But now, let's come back to our main topic, which is image filtering. So far we have mainly used so called point operations. They are mathematical operations that had been applied to the entire image. So we have changed pixel values without changing the size, geometry, or local structure of the image. The most prominent example were the display settings and Look Up Tables. Mapping functions had been used in order to change the display of the images. The advantage of point operation is that they're completely reversible. The disadvantage is they're are not flexible and powerful enough to perform all needed processing tasks for image analysis.

Notes

Summary

1m 42s

Filtering

- Sharp image: local intensities rise/drop sharply
- Blurred Image: local intensities rise/drop smoothly

Have a look at these 2 images. Without hesitation, you will come to the conclusion that we're looking at the same scene. We have a sharp image where we encounter sharp rises and drops of local intensity. And we have a rather blurred image where the local intensity rises and drops rather smoothly. With the known point operation, we will not manage to convert one image into the other. This is only possible by using so called filters.

Notes

Summary

2m 30s

Linear Filters

1	1	1	1	1	1	1
1	1	1	1	1	1	1
1	1	5	5	5	1	1
1	1	5	10	5	1	1
1	1	5	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

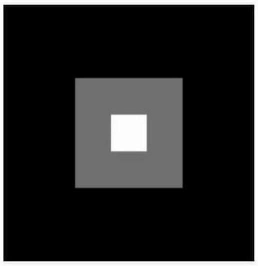

Filter

Output

Indeed, a mean intensity filter has have been used in order to calculate this smooth image. So what is a filter and how does it work? Mathematically speaking a filter is doing a convolution The ones being familiar with that operation might go through the following slides a bit faster. For the other ones, I will explain this operation based on this small image shown at the bottom. The intensity values corresponding to the images are shown above. Now, let's apply this squared filter shown in the middle. Which is, by the way, the previously mentioned mean filter.

Notes

Summary

2m 59s

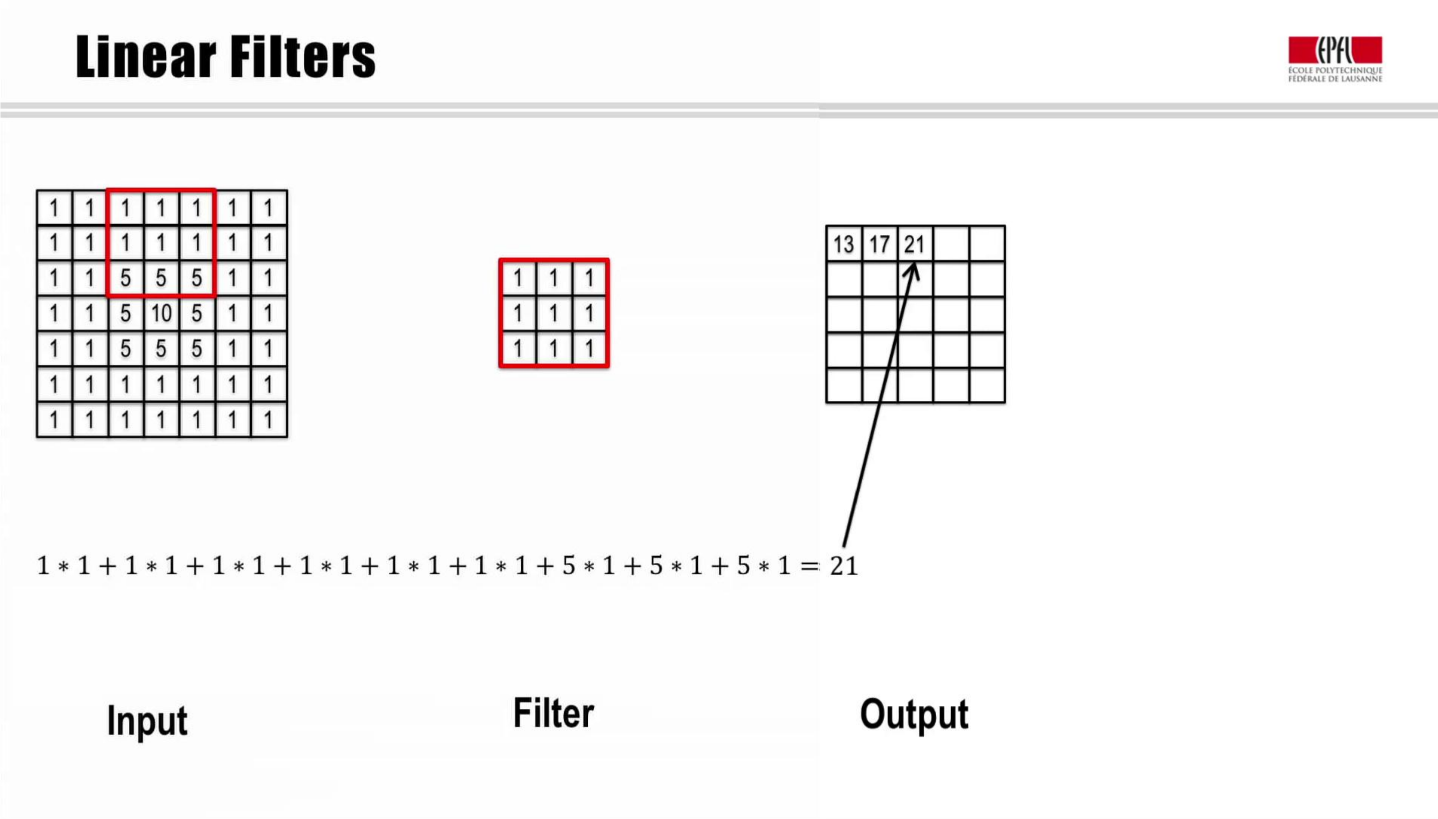

In order to calculate the filtered image we align the filter and the image as shown in the diagram. We multiply the corresponding intensity value of the image with the intensity value of the filter. Now we repeat this procedure for the next intensity pair. And we continue with the other corresponding intensity pair. So we end up with 9 products. And all we have to do is to sum them up. This sum is the 1st intensity value of the filtered image. If you have a closer look to the input and the output image, you see immediately that the latter one is smaller. The input image is 7x7 pix large. The output image is 5x5 pix large. Obviously we lose a row and a column on each side. Let's just note it for the moment. We'll see later on how to compensate this effect. In order to calculate the other values of the output image, we shift the filter as indicated by the red square. Now we perform the same operations again. Multiplying the individual elements and summing them up. This gives the next value of our output image. I think most of you already have a right guess about the continuation. You shift the filter to the next position, and perform the operations again. If you do so for the remaining pixels will end up building the filtered image. So let's have a look at the outcome!

Notes

Summary

3m 37s

Linear Filters

1	1	1	1	1	1	1
1	1	1	1	1	1	1
1	1	5	5	5	1	1
1	1	5	10	5	1	1
1	1	5	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

13	17	21	17	13
17	30	38	30	17
21	38	50	38	21
17	30	38	30	17
13	17	21	17	13

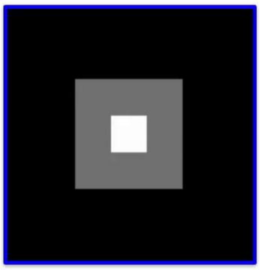

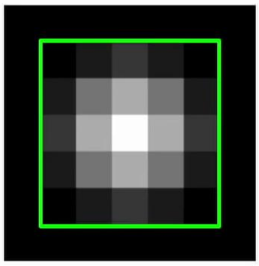

You can see the input image and the image which was filtered with the kernel shown between the images. Kernel is the other name for the 3x3 matrix which was used in order to obtain the output image. The shown image has the same size as the input image. This was done by an operation which is called padding. This means we're just resizing the input image to an 8x8 image. For the intensity values of the new pixels we clearly need to make an assumption. One can either set them to 0 or, what had been done here, copy the values of the nearby pixel. This is clearly arbitrary. But thereby we obtain an image with the same size. This is an advantage of the output image, used even in the image analysis workflows. Let's imagine you want to subtract images which are not identical in size. In order to avoid such situations, padding is a really powerful concept. However, it is important to remember to be careful when interpreting the edges of filtered images.

Notes

Summary

5m 19s

Filtering

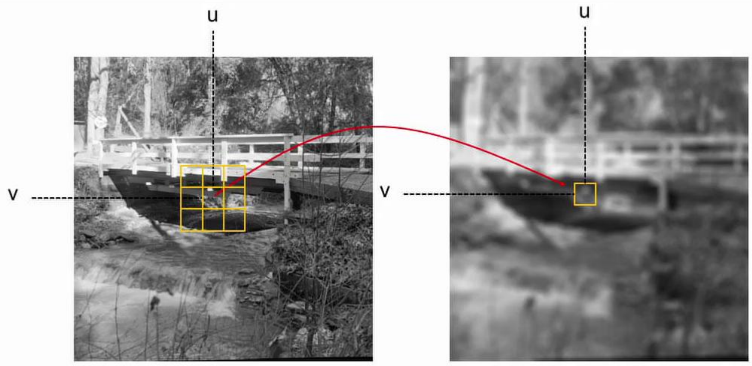

$$I'(u,v) = \frac{p_1 + p_2 + p_3 + p_4 + p_5 + p_6 + p_7 + p_8 + p_9}{9}$$

$$I'(u,v) = \frac{1}{9}\sum_{j=-1}^{1}\sum_{i=-1}^{1} I(u+i, v+j)$$

Filter Matrix:

$$H(i,j) = \frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

It is quite common to use quadratic filter kernels with an odd number of elements. Either 3x3, 5x5, 7x7, and so on. You can for sure also use a 2x2 kernel. But the resulting image is then shifted by a half a pixel. It will then be cumbersome to align the input and the output image. Therefore, the use of odd filter kernel is preferred. Just to summarize the working principle of a spatial filter. The filter matrix is multiplied with intensity value of the image, and the individual products are then summed up. The filter is then moved over the entire image in order to obtain the output image. In order to make sure that the sum of intensity values of the output image is equaled to the sum of the intensities of the input image, we need to introduce a normalization factor. For our mean filter, the factor is 1/9. It is obviously the inverse of the sum of all matrix elements.

Notes

Summary

6m 32s

Image Filtering in Fiji

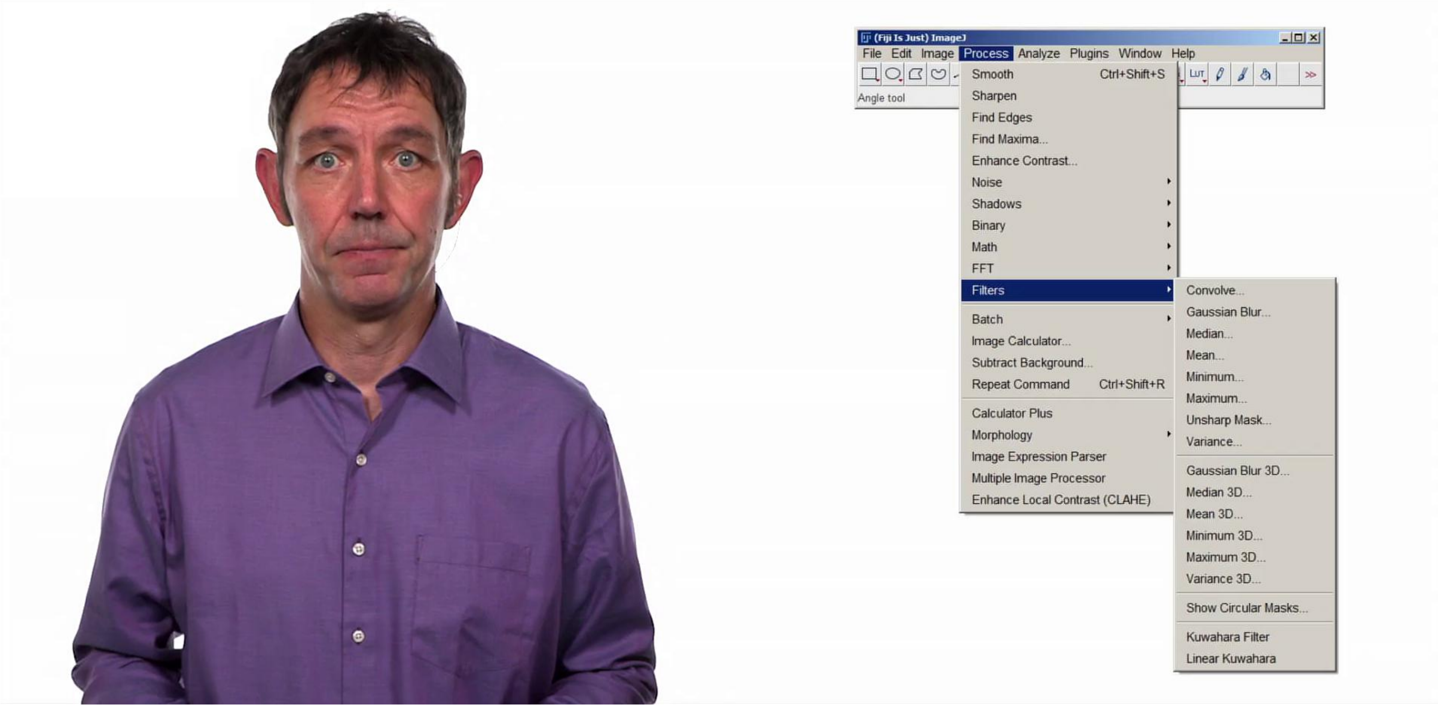

Filters are also available in Fiji/ImageJ. You can find them in the process submenu. So do not hesitate to try the different filters on one of your preferred images.

Notes

Summary

7m 38s

Linear Filters - Examples

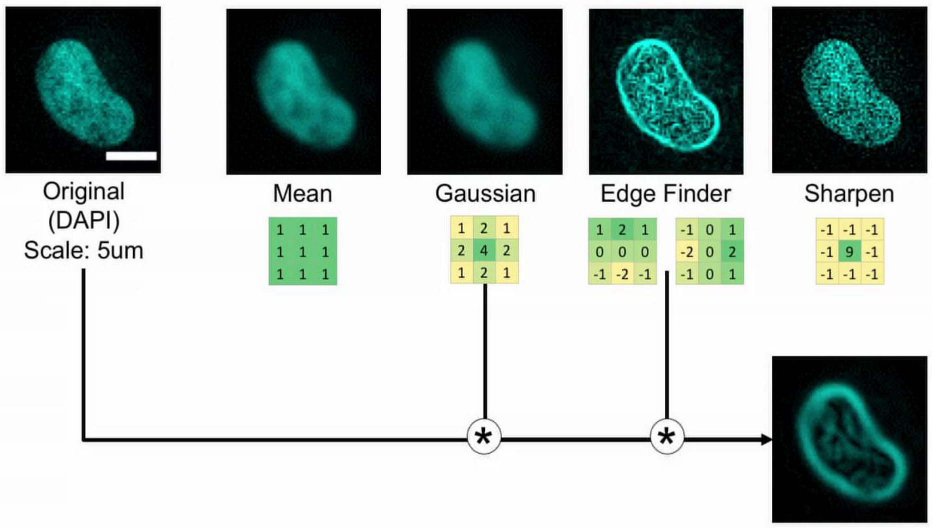

Here you can see the effect of different linear filters on the image of a nucleus which had been stained with DAPI. Mean filter should already be familiar to you. It is good to smooth a structure. In particular if you have a low signal to noise ratio. A kind of similar effect is obtained when a Gaussian filter is used. If you compare the 2 filter kernel, you see that the central pixel has much higher weight in the Gaussian filter. Also the direct neighbors have a higher width compared to the one in the mean filter. As a consequence, you obtain an image which is more blurred than the one obtain with a mean filter. A filter called edge filter is differing in certain aspects. It is the first filter kernel which contains negative elements. And it comes in 2 variants. If you apply that filter to a homogeneous area it returns a pixel value with the intensity 0. If you're skeptical, you can now do the math on your own. At the transition between an area with high an low intensity, this filter emphasizes the edge. There are 2 variants of the filter. In order to detect edges which are oriented parallel to the rows of the zeros in the kernel. Last, but not least, the effect of the sharpen filter. By combining the output of the Gaussian filter and the edge filter, one obtains the edge which is defining the nucleus.

Notes

Summary

7m 51s

Non-Linear Filters

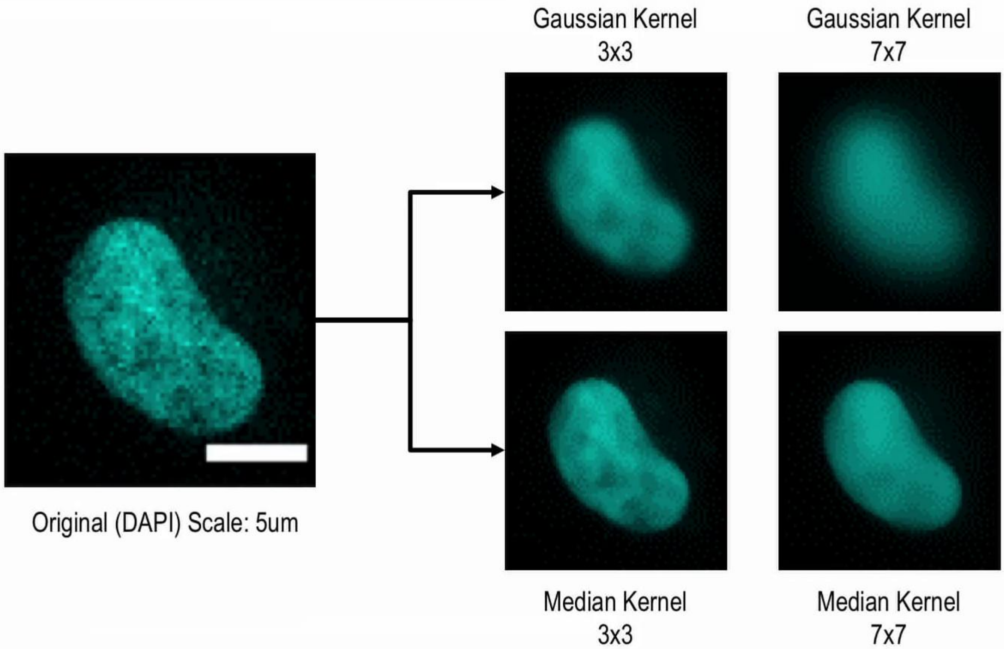

1	1	1	1	1	1	1
1	1	1	1	1	1	1
1	1	5	5	5	1	1
1	1	5	10	5	1	1
1	1	5	5	5	1	1
1	1	1	1	1	1	1
1	1	1	1	1	1	1

Sum: 13
Mean: 13/9
Median:1
Max: 5
Min:1

So far we have looked at linear filters. However we can also use non linear operations in order to come to the value of the filtered image. A prominent example of a non linear filter is the median filter. The advantage of a median filter is that it's preserving edges much better than the Gaussianor a Mean filter. Other examples of non linear filter are minimum or maximum filter. In particular, the 1st one can be used in order to calculate an image which is representing the background. You just have to pay attention that the filter kernel is larger than the object within your image.

Notes

Summary

9m 30s

Non-Linear Filters

Original (DAPI) Scale: 5um | Kuwahara | Anisotropic Diffusion

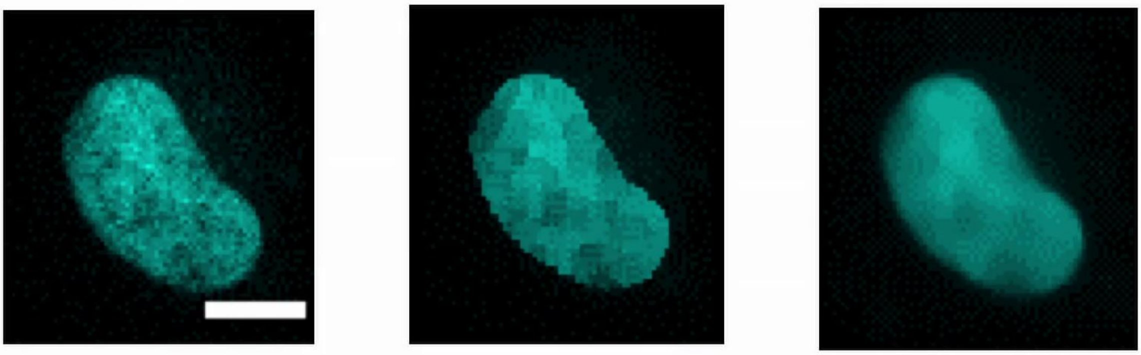

But more settings

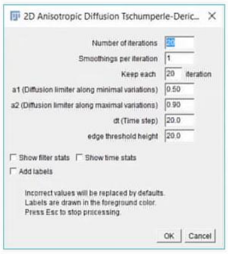

In that case the minimum filter will always return the value which is corresponding to the background value. There might be situations where simple filters as discussed so far are not sufficient anymore. In that case we can use a dedicated algorithm in order to come to a filtered image. Here you see the outcome of 2 more complicated filters: kuwahara and anisotropic diffusion.

Notes

Summary

10m 12s

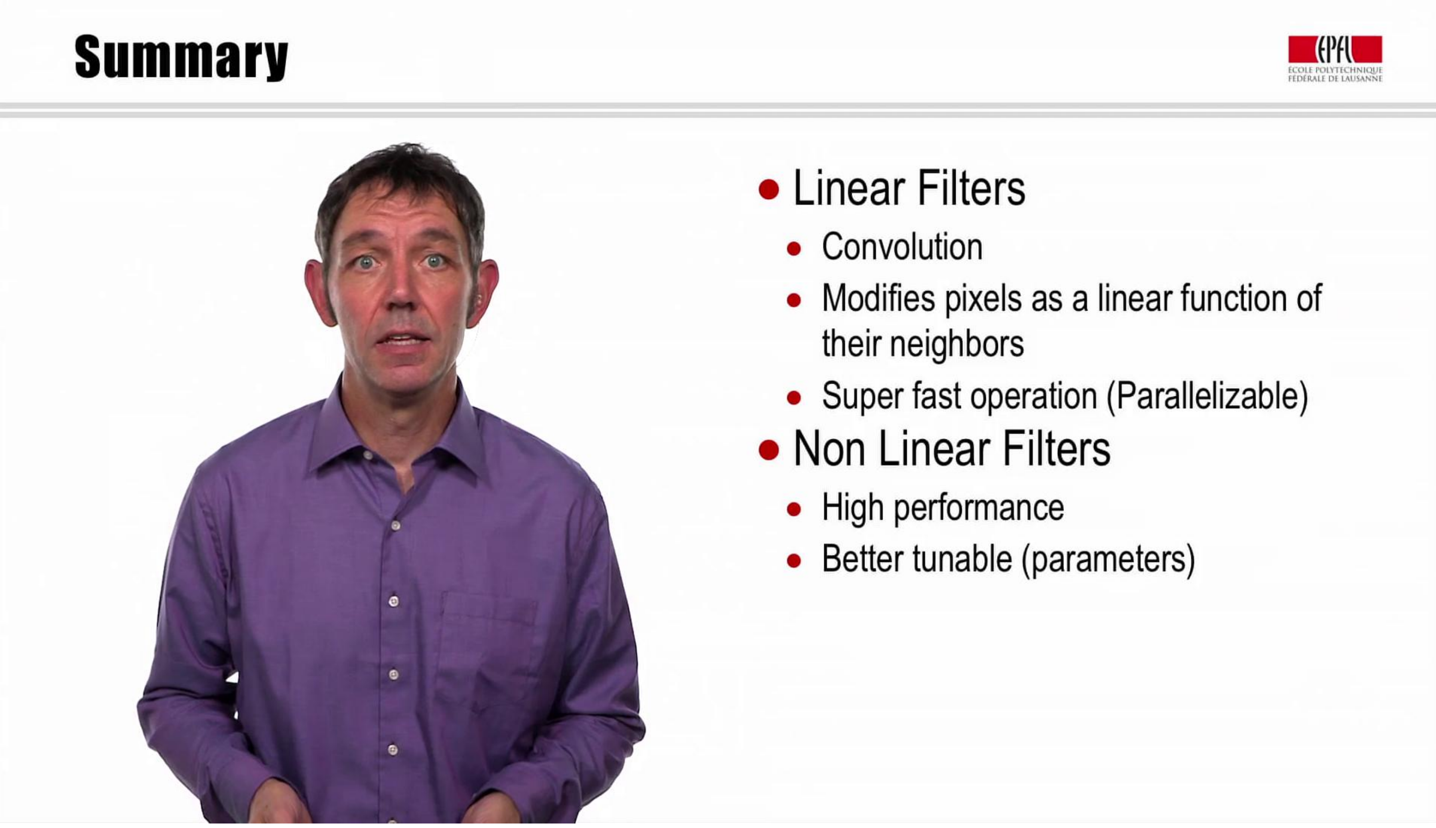

We are not going to do the exact math of these filters. But what is important to remember is that that such filter exists and that they are extremely powerful. However, typically it is required to parameterise your filter. So the outcome depends on a variety of input values. This can be seen as an advantage as it allows to obtain a better representation of objects. But at the same time, it is also a disadvantage. As the result depends on the input parameters. So let's summarize! We have introduced filters. They modify the pixels by taking the values of the neighboring pixels into account Mathematically, we're speaking of the so called convolution. By the way, this is the exact operation which lens is performing. There, the filter kernel is the so called point spread function. In image processing, filters are mainly used in order to remove noise from images; or to calculate an image which corresponds to the background value of an image. Filters can be implemented in parallel. Therefore the output image can be obtained in very short delay. Besides linear filters, also non linear filters are available. These can be extremely powerfull filters which can be tuned via input parameters in order to obtain the filtered image.

Notes

Summary

10m 40s

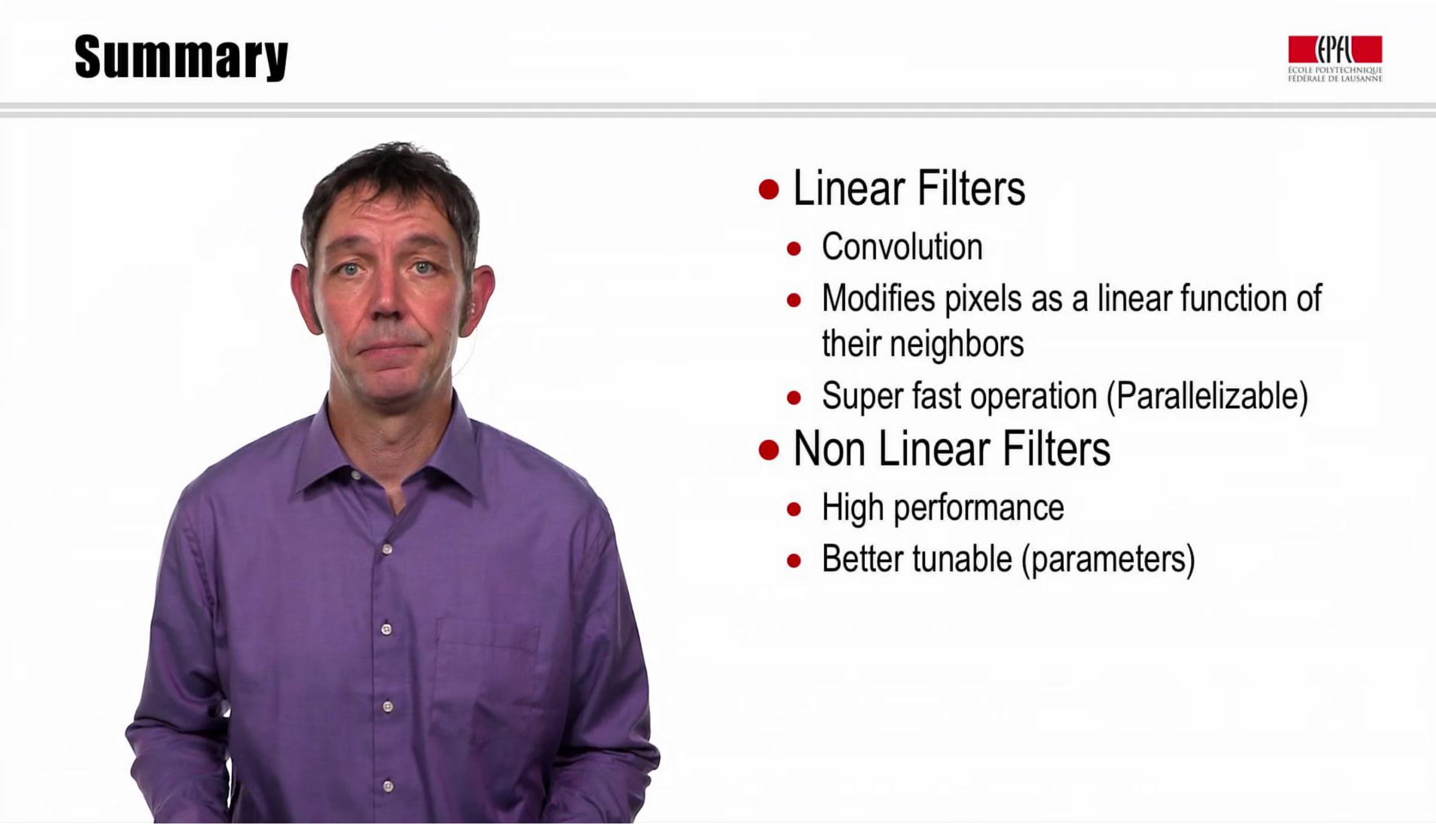

So far, we have only demonstrated spatial filters. However, please keep in mind that filters can also be applied to time series or 3D data. With this one, we're finishing the lecture. Hope to see you next week. And take care!

Notes

Summary

12m 08s