



Additional Parameters :-

① Padding :-



to prevent shrink.



② Stride :-



stride = 1 (default)



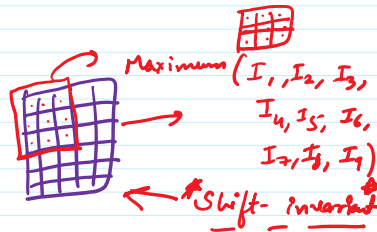
stride = 2.

N.B.: Blue maps are inputs, and cyan maps are outputs.

<p>outputs</p> <p>inputs</p>			
No padding, no strides	Arbitrary padding, no strides	Half padding, no strides	Full padding, no strides
No padding, strides	Padding, strides	Padding, strides (odd)	

Pooling :-

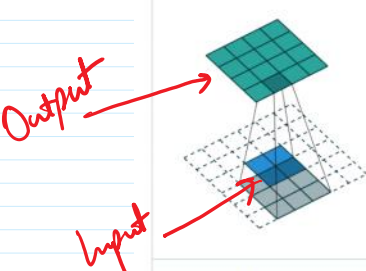
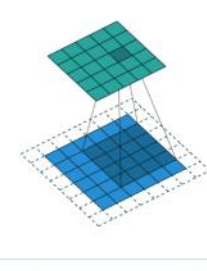
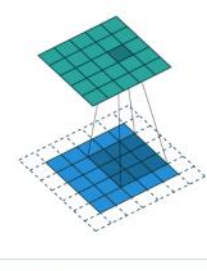
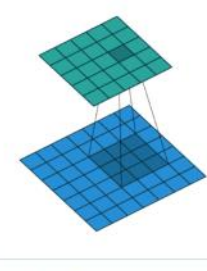
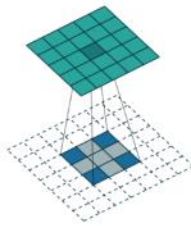
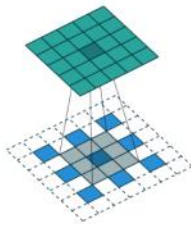
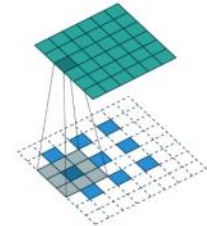
Max Pooling →



Average Pooling → Average Linear Filter.


Unpooling / Transposed Convolution :-

N.B.: Blue maps are inputs, and cyan maps are outputs.

			
No padding, no strides, transposed	Arbitrary padding, no strides, transposed	Half padding, no strides, transposed	Full padding, no strides, transposed
			
No padding, strides, transposed	Padding, strides, transposed	Padding, strides, transposed (odd)	

Downsampling :

Naïve image downsampling



Naïve image downsampling




What is the 1/8 image so pixelated (and do you know what this effect is called)?

aliasing

What is the 1/8 image so pixelated (and do you know what this effect is called)?

Aliasing



Better image downsampling

1	2	1
2	9	2
1	2	1



1/2

Apply a smoothing filter first, then throw away half the rows and columns

Gaussian filter
delete even rows
delete even columns



1/4

Gaussian filter
delete even rows
delete even columns



1/8

Better image downsampling



1/2




1/4 (2x zoom)



1/8 (4x zoom)

Naïve image downsampling



Naïve image downsampling



1/2



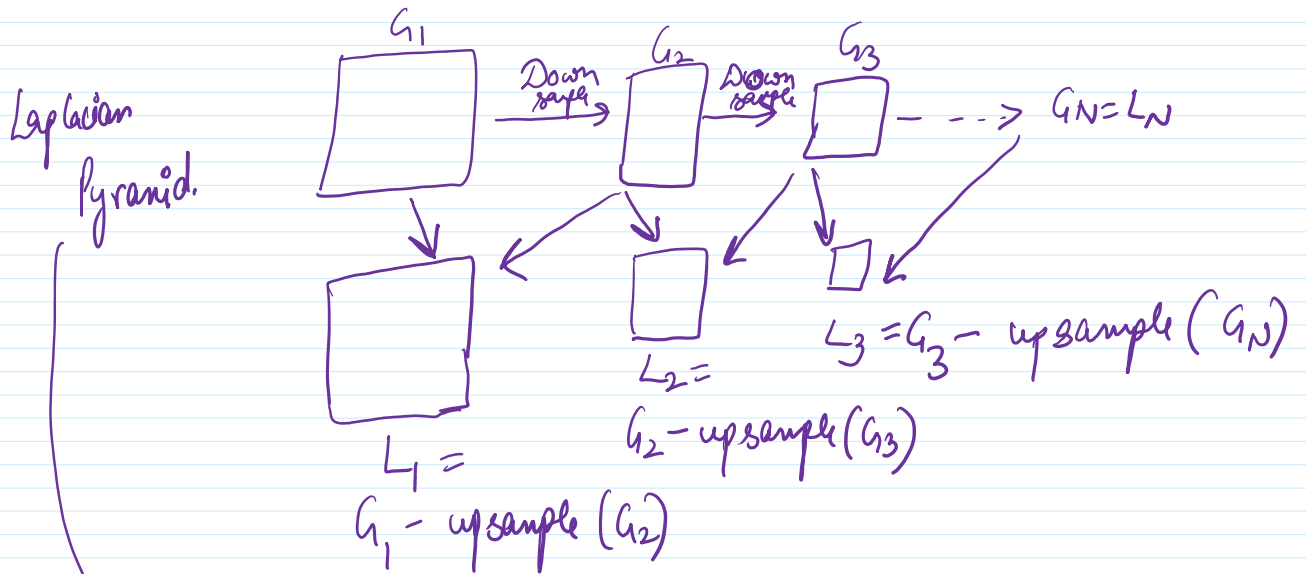
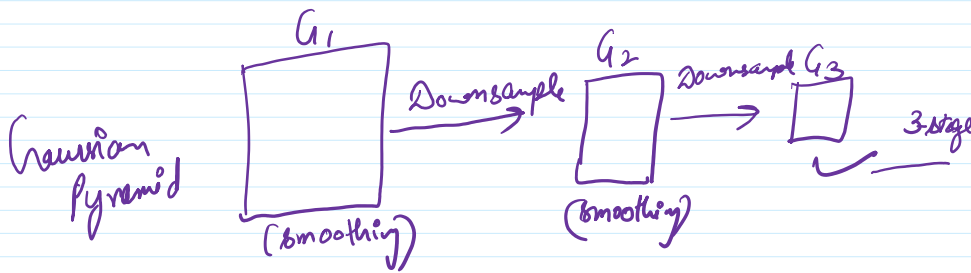
1/4 (2x zoom)



1/8 (4x zoom)

What is the 1/8 image so pixelated (and do you know what this effect is called)?

Image Pyramids :-



Why?

Laplacian can be written as a difference of two Gaussians

0	1	0
1	4	1
0	1	0

0	-1	0
-1	4	-1
0	-1	0

1	4	1
0	1	0

4	4	1
0	-1	0

u

v

