

Title: Machine Learning Lecture

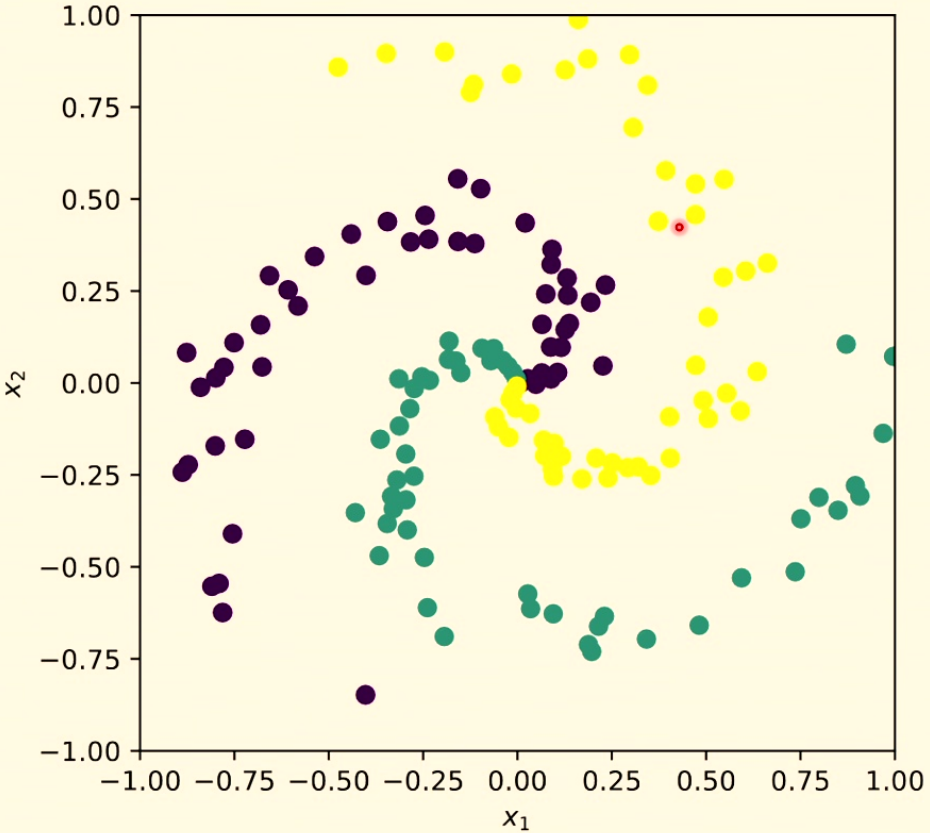
Speakers: Mohamed Hibat Allah

Collection: Machine Learning 2023/24

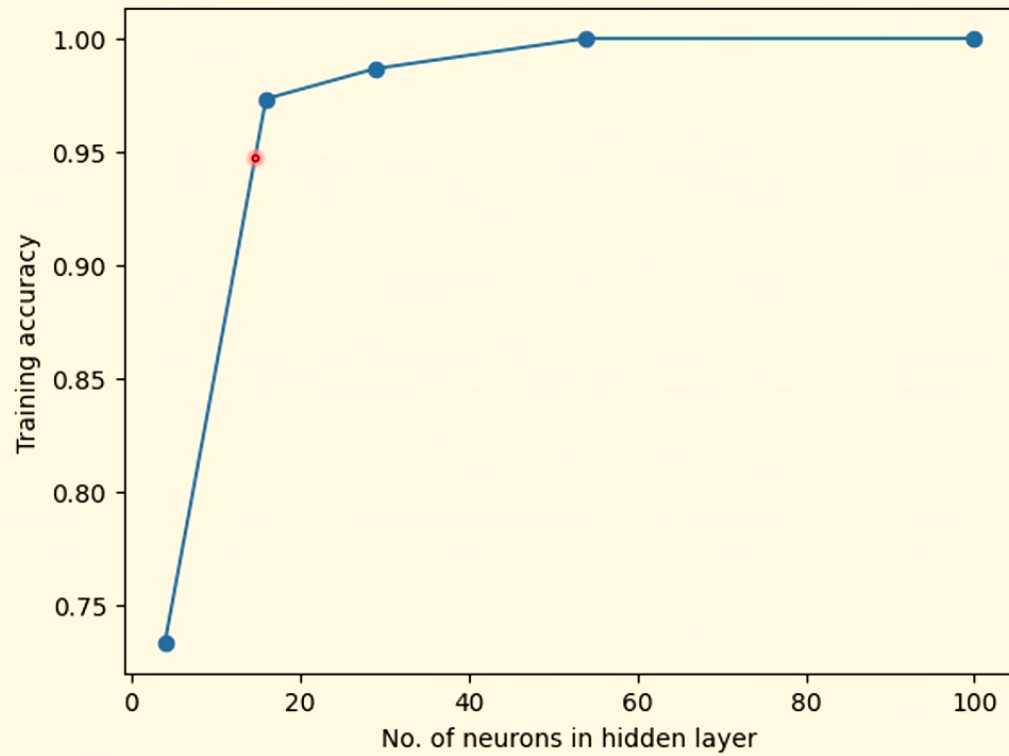
Date: April 18, 2024 - 11:30 AM

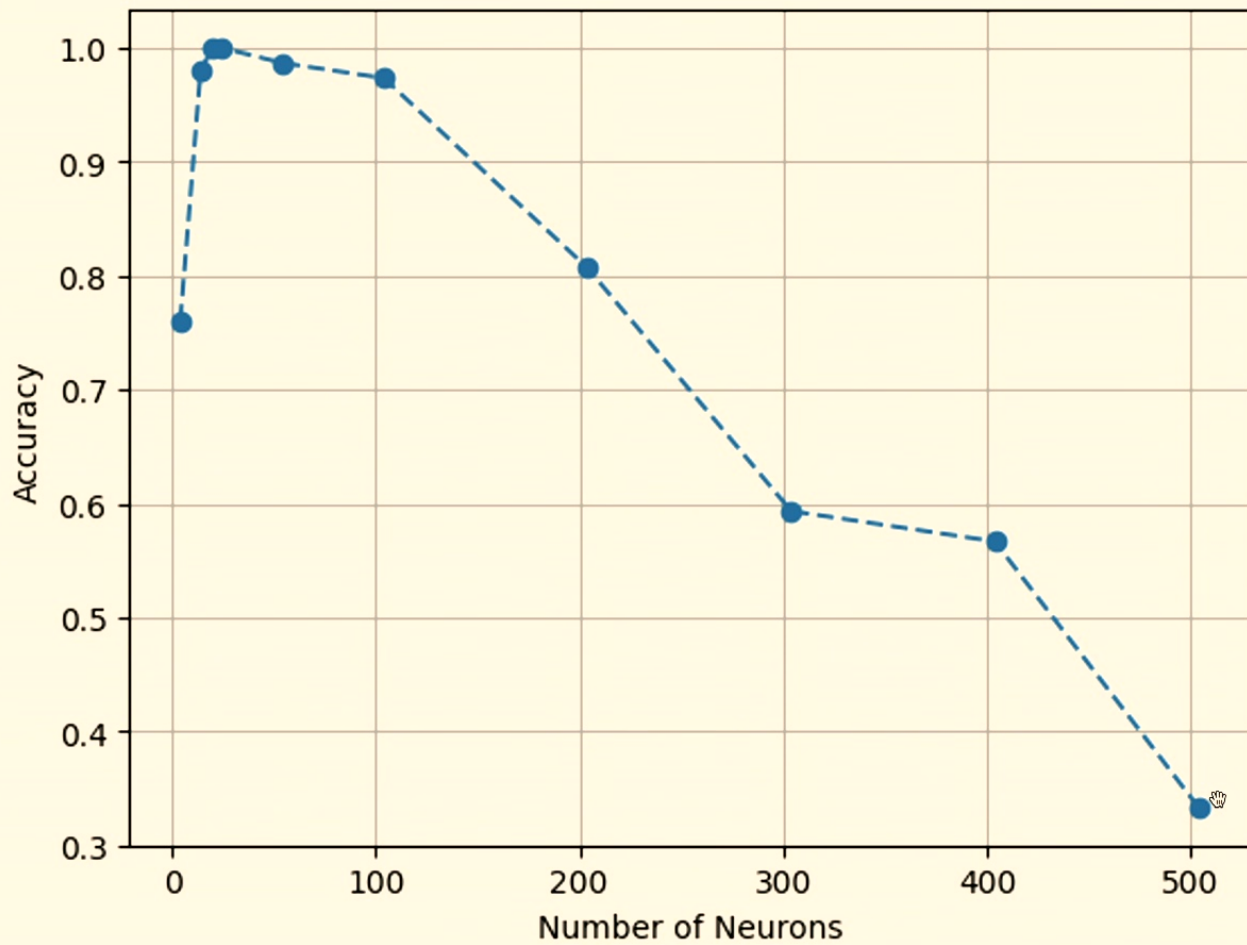
URL: <https://pirsa.org/24040051>

# Tutorial 3

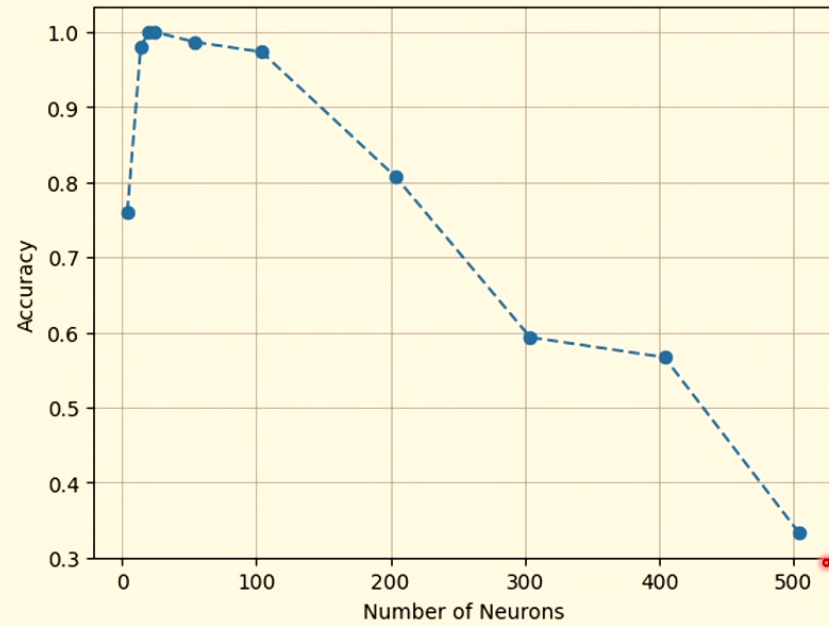


# Tutorial 3 results



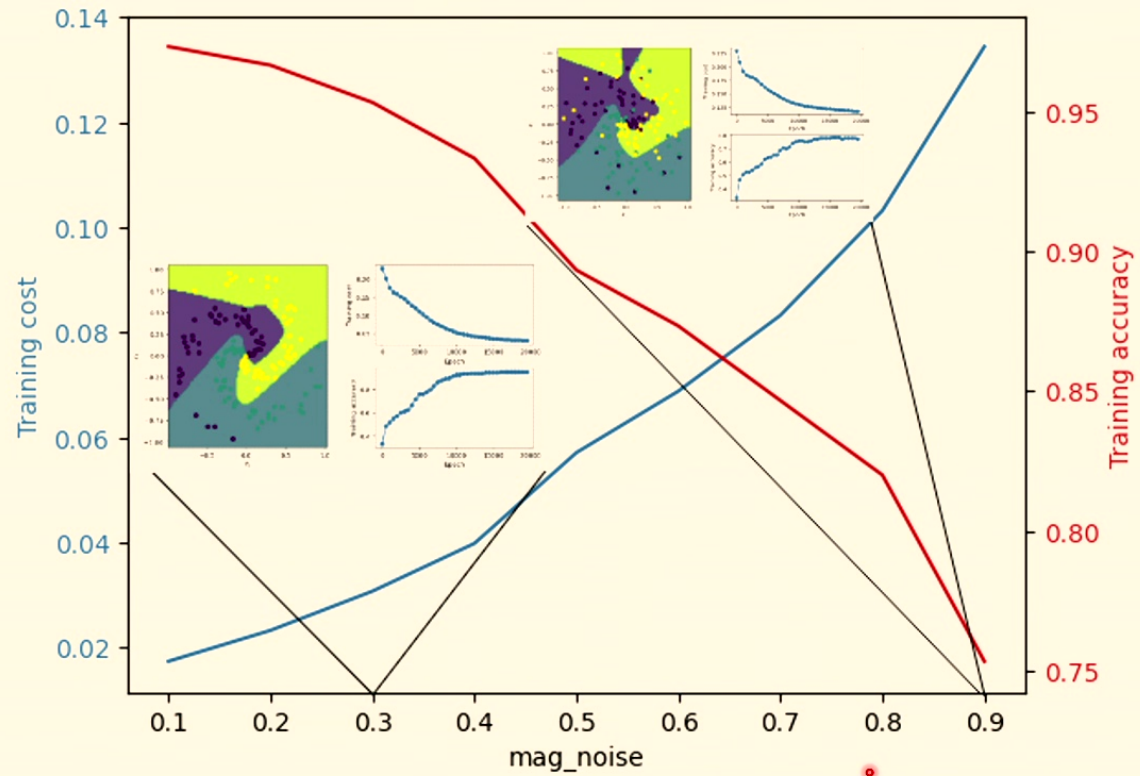


# Tutorial 3 results

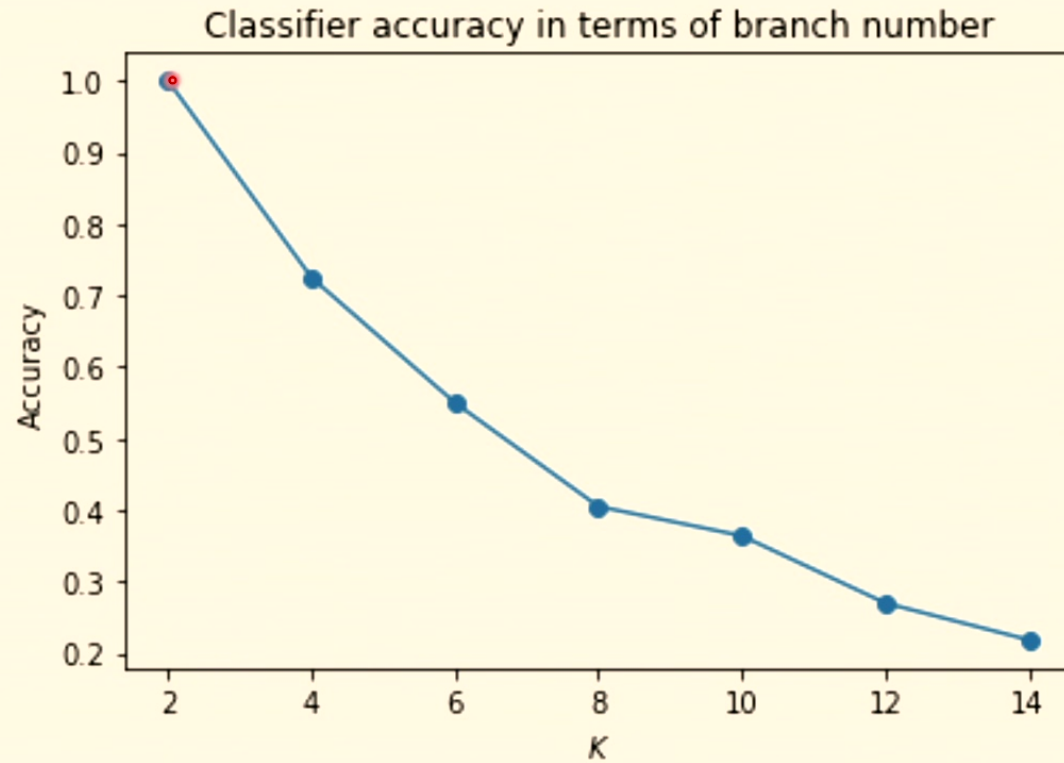


# Tutorial 3 results

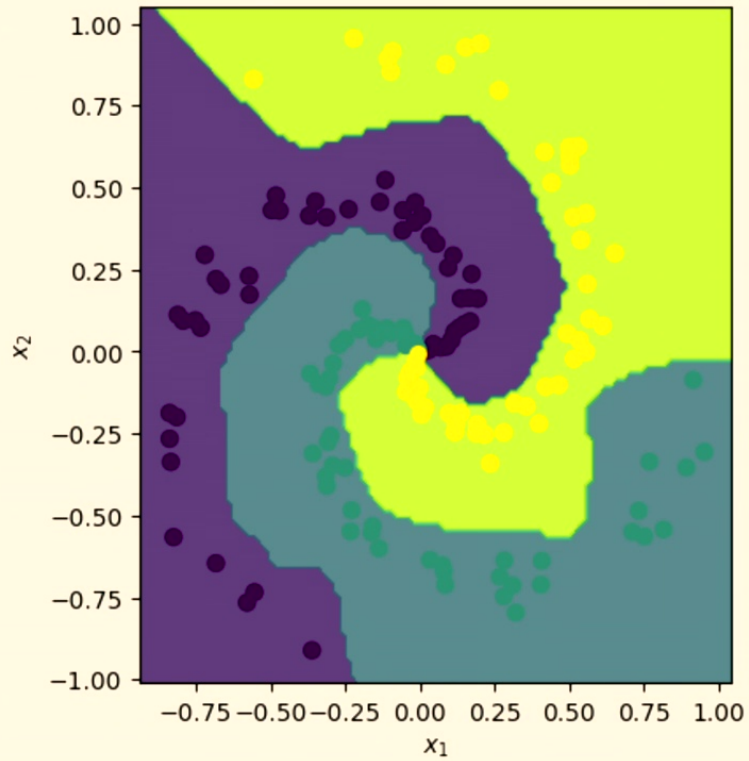
FFNN for 20000 epochs, Adam lr = 0.001



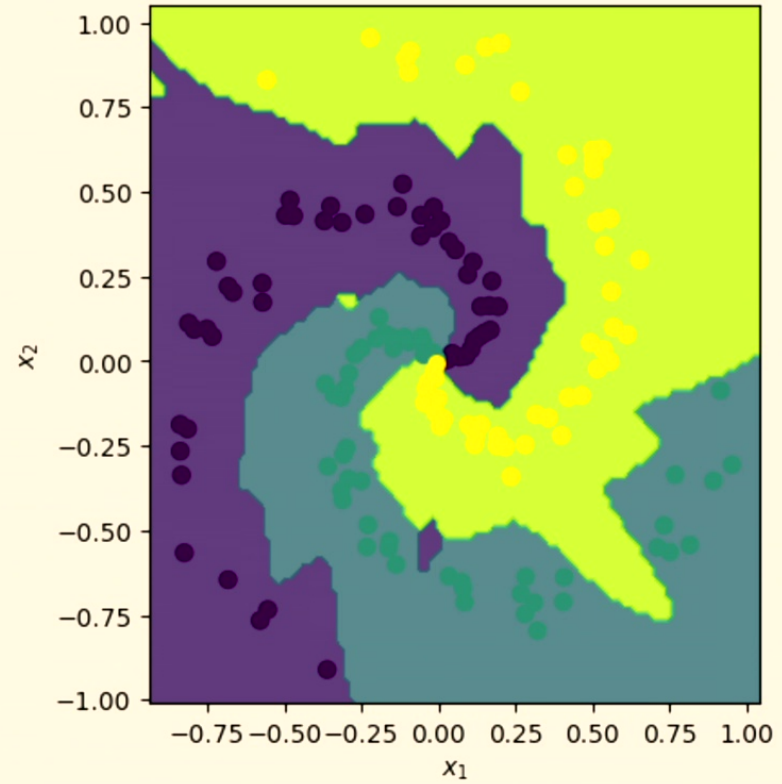
# Tutorial 3 results



# Tutorial 3 results



**Early stopping**





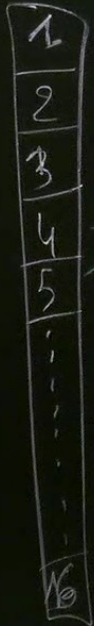
# Lecture 7

Today: Convolutional Neural Networks (CNNs)

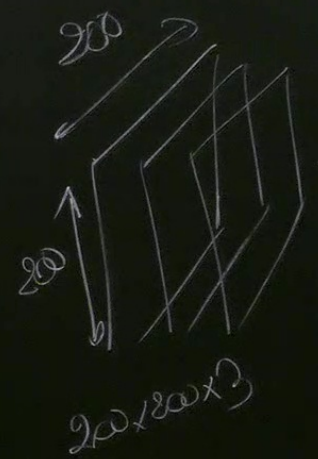
- ↳ Local receptive fields (Filters or kernels)
- ↳ Convolution    ↳ Weight & Biases in CNNs
- ↳ Channels       ↳ Stride & Padding
- ↳ Pooling

# \* Motivation for using CNNs

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	16



Not spatially close  
Locality is lost



$W, b$

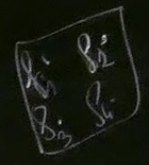
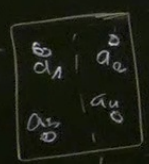


1000 hidden neurons

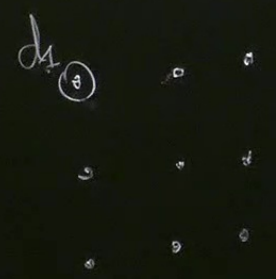
$$[W] = 200 \times 200 \times 3 \times 1000 = 120 \text{ M parameters}$$

Solution: Convolution

Patch  
Local receptive  
Kernel  
filter



$$d_1 = \sum_{i=1}^4 a_i \times b_i$$



Feature map

Firefox File Edit View History Bookmarks Tools Window Help zoom Thu Apr 18 11:48 AM

conv\_arithmetic/README.md a x Image Kernels explained visually x CNN Explorer x Machine Learning for Many-Bo... x PSI Portal

https://github.com/vdumoulin/conv\_arithmetic/blob/master/README.md

master Preview Code Blame 112 lines (91 loc) · 3.22 KB Code 55% faster with GitHub Copilot Raw

bin gif .gitignore arbitrary\_padding\_no\_strides.gif arbitrary\_padding\_no\_strides\_tr... dilation.gif full\_padding\_no\_strides.gif full\_padding\_no\_strides\_transp... no\_padding\_no\_strides.gif no\_padding\_no\_strides\_transpo... no\_padding\_strides.gif no\_padding\_strides\_transposed... padding\_strides.gif padding\_strides\_odd.gif padding\_strides\_odd\_transpose... padding\_strides\_transposed.gif same\_padding\_no\_strides.gif same\_padding\_no\_strides\_trans...

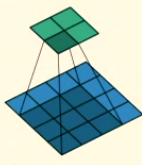
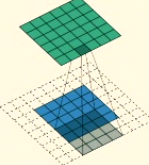
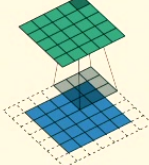
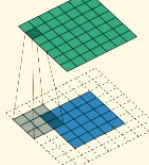
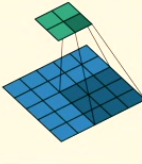
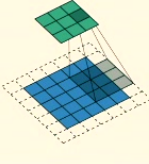
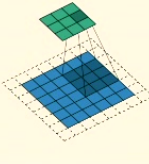
A technical report on convolution arithmetic in the context of deep learning.

The code and the images of this tutorial are free to use as regulated by the licence and subject to proper attribution:

- [1] Vincent Dumoulin, Francesco Visin - [A guide to convolution arithmetic for deep learning \(BibTeX\)](#)

### Convolution animations

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

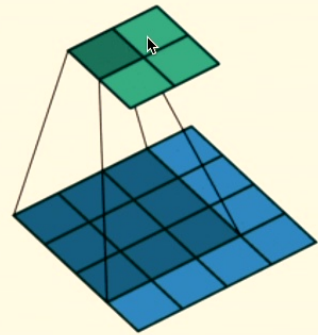
			
No padding, no strides	Arbitrary padding, no strides	Half padding, no strides	Full padding, no strides
			
No padding, strides	Padding, strides	Padding, strides (odd)	

### Transposed convolution animations

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

https://github.com/vdumoulin/conv\_arithmetic/blob/master/gif/no\_padding\_no\_strides.gif

36.9 KB Code 55% faster with GitHub Copilot



Firefox File Edit View History Bookmarks Tools Window Help zoom Thu Apr 18 11:49 AM


https://setosa.io/ev/image-kernels/

Let's walk through applying the following 3x3 **top sobel** kernel to the image of a face from above.

top sobel

$$\begin{pmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{pmatrix}$$

Below, for each 3x3 block of pixels in the image on the left, we multiply each pixel by the corresponding entry of the kernel and then take the sum. That sum becomes a new pixel in the image on the right. Hover over a pixel on either image to see how its value is computed.



input image


$$\begin{pmatrix} 97 & 96 & 104 \\ \times 1 & \times 2 & \times 1 \\ + & 97 & 88 & 73 \\ \times 0 & \times 0 & \times 0 \\ + & ? & ? & ? \\ \times -1 & \times -2 & \times -1 \end{pmatrix}$$

=

?

kernel:

top sobel



output image

One subtlety of this process is what to do along the edges of the image. For example, the top left corner of the input image only has three neighbors. One way to fix this is to extend the edge values out by one in the original image while keeping our new image the

$$\begin{pmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{pmatrix}$$

Below, for each 3x3 block of pixels in the image on the left, we multiply each pixel by the corresponding entry of the kernel and then take the sum. That sum becomes a new pixel in the image on the right. Hover over a pixel on either image to see how its value is computed.

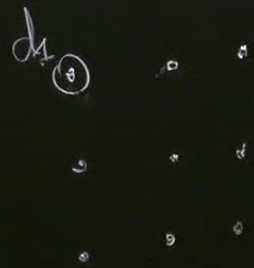
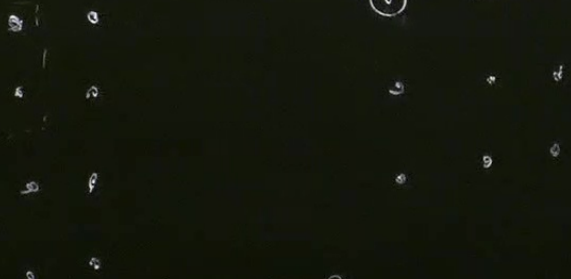
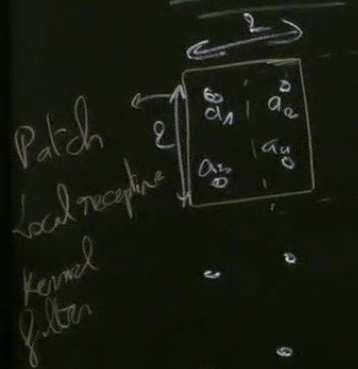
input image

output image

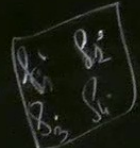
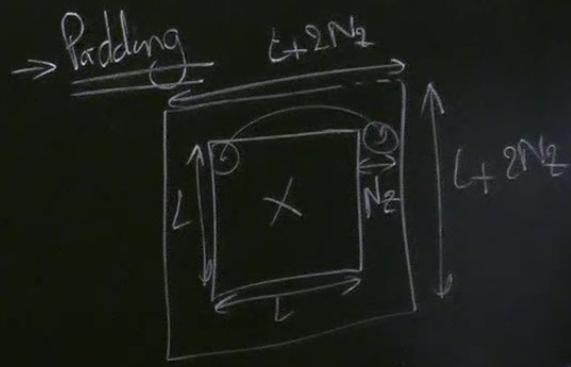
kernel:

One subtlety of this process is what to do along the edges of the image. For example, the top left corner of the input image only has three neighbors. One way to fix this is to extend the edge values out by one in the original image while keeping our new image the same size. In this demo, we've instead ignored those values by making them black.

Solution: Convolution



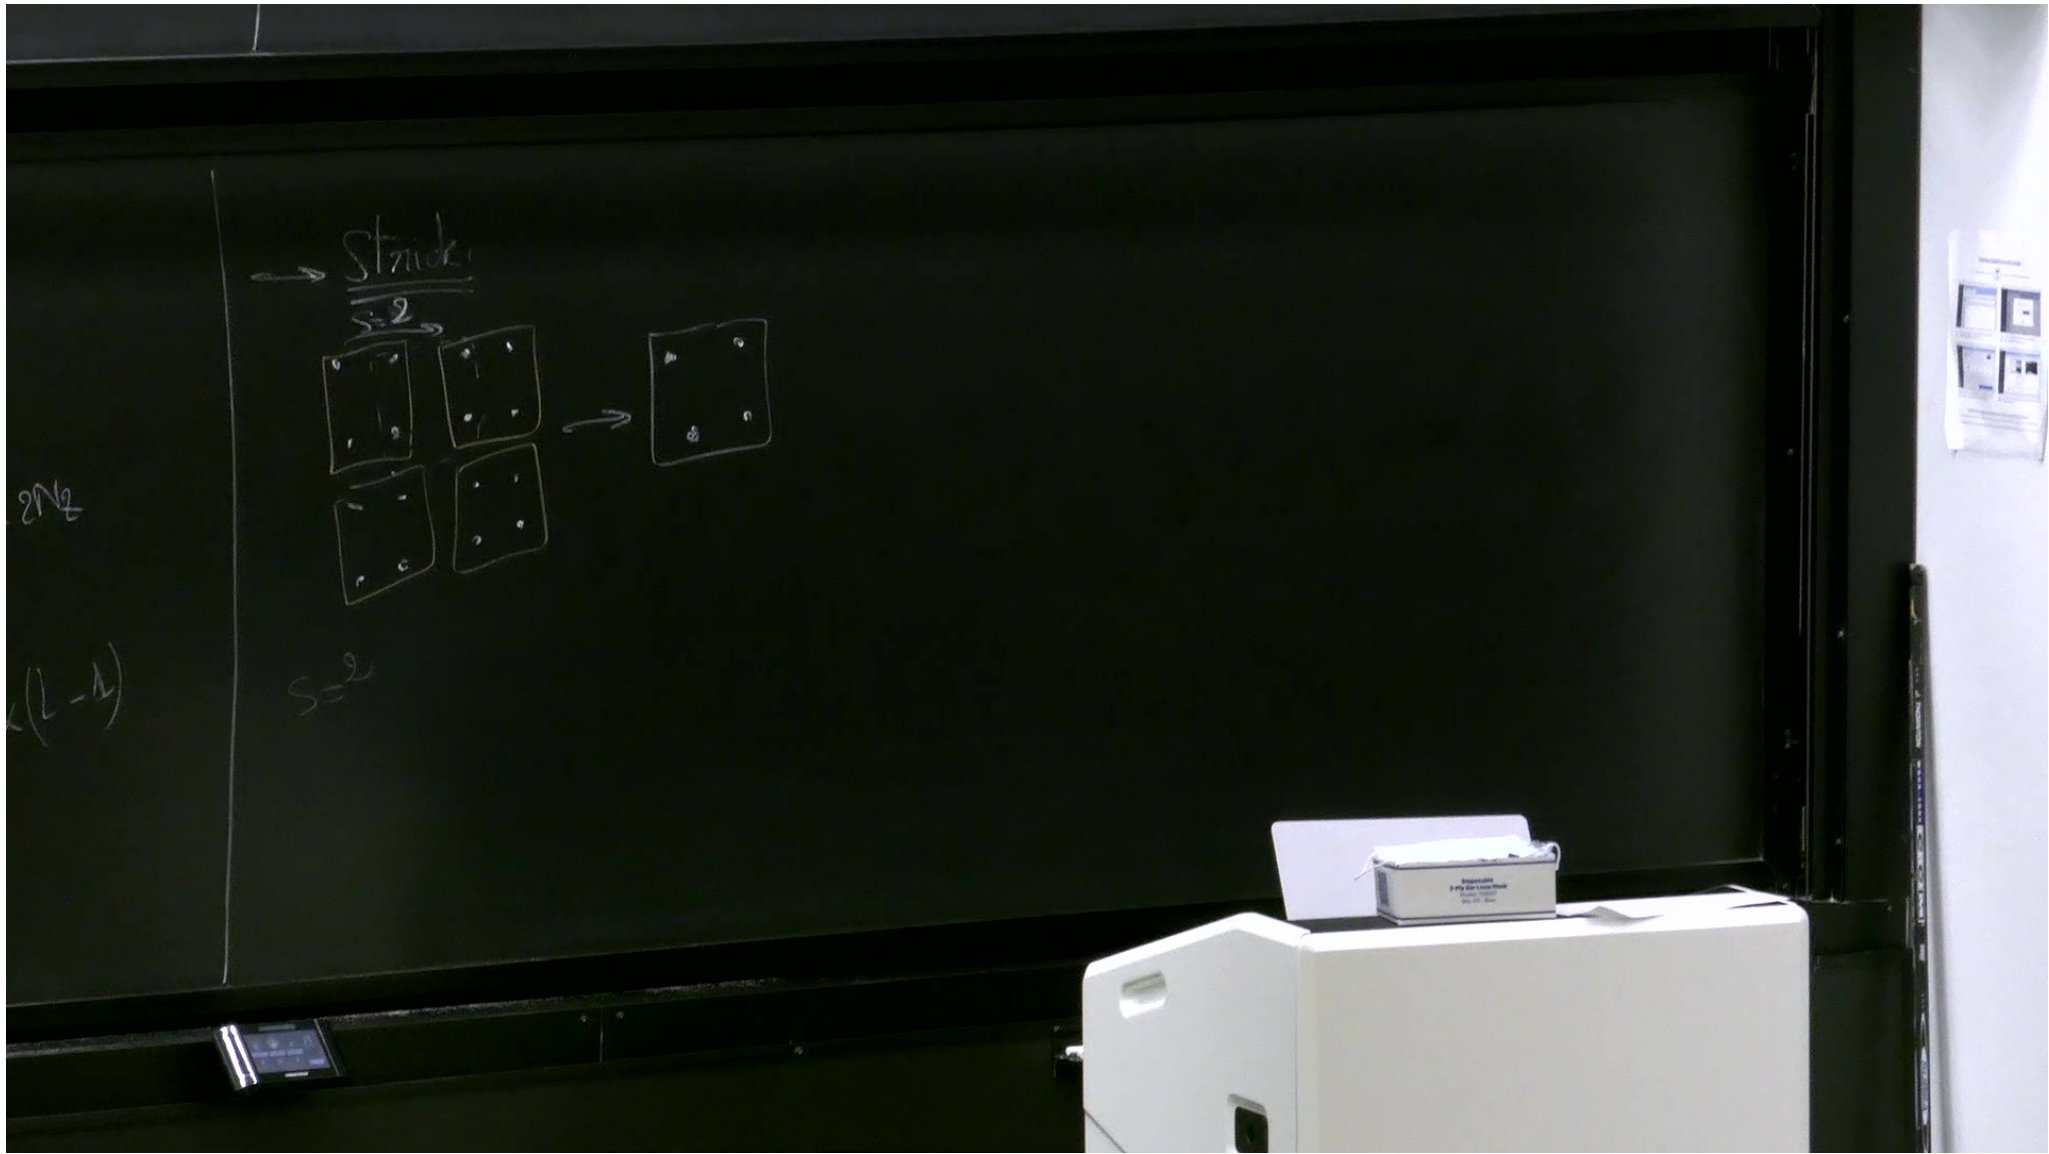
← Feature map



$$d_1 = \sum_{i=1}^4 a_i \times f_i$$

\* Patch size:  $2 \times 2 \rightarrow L$



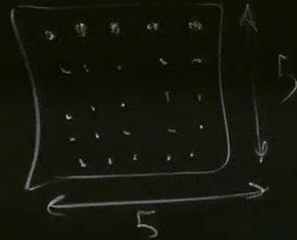
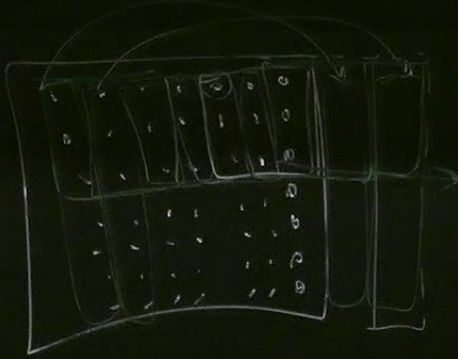


$$x = a^{(b)} \rightarrow L \times L \rightarrow a^{(1)} \rightarrow L_1 \times L_1$$

$$L_1 = \frac{L - L_{\text{patch}} + 2N_z}{\Delta} + 1 \rightarrow \text{has to be an integer.}$$

Using periodic boundary conditions.

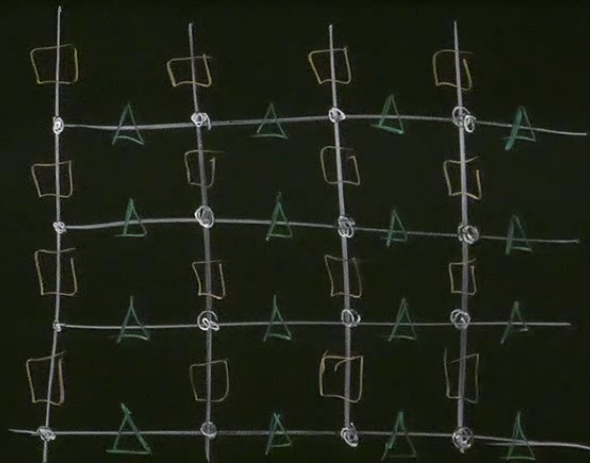
$$L_1 = \frac{L}{\Delta}$$



$$N_2 = 0$$

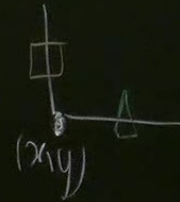
$$S = 1$$

$$L = \frac{7 - 3 + 2 \times 0}{1} + 1 = 5$$

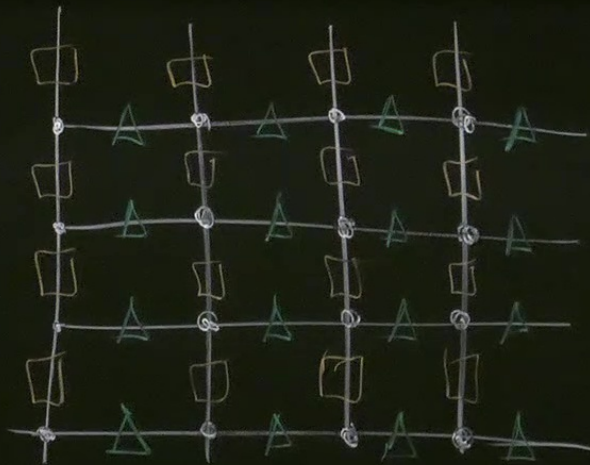


$\Delta$ : Sublattice 1

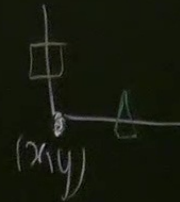
$\square$ : Sublattice 2



$a_{x,y,c}$   
 channel (Sublattice index)  
 spatial information.



$\Delta$ : Sublattice 1  $\rightarrow N_c = 2$   
 $\square$ : Sublattice 2



$a_{x,y,c}$   
 channel (Sublattice index)  
 spatial information

$$a_{xyik}^{(l)} = \phi \left( \sum_{m=1}^{L_{patch}} \sum_{n=1}^{L_{patch}} \sum_{c=1}^{N_c} a_{x+m, y+n, c}^{(l-1)} W_{mn}^{(l)} + b_k^{(l)} \right)$$

Activation function.

$x$  index of the filter  
 $y$  index of the filter  
 $c$  channel or sublattice index  
 $k$  Label for different filters  
 $N_f$

$$a_{xyik}^{(l)} = \phi \left( \sum_{m=1}^{L_{patch}} \sum_{n=1}^{L_{patch}} \sum_{c=1}^{N_c} a_{x+m, y+n-1, c}^{(l-1)} W_{mm}^{(l)} + b_K^{(l)} \right)$$

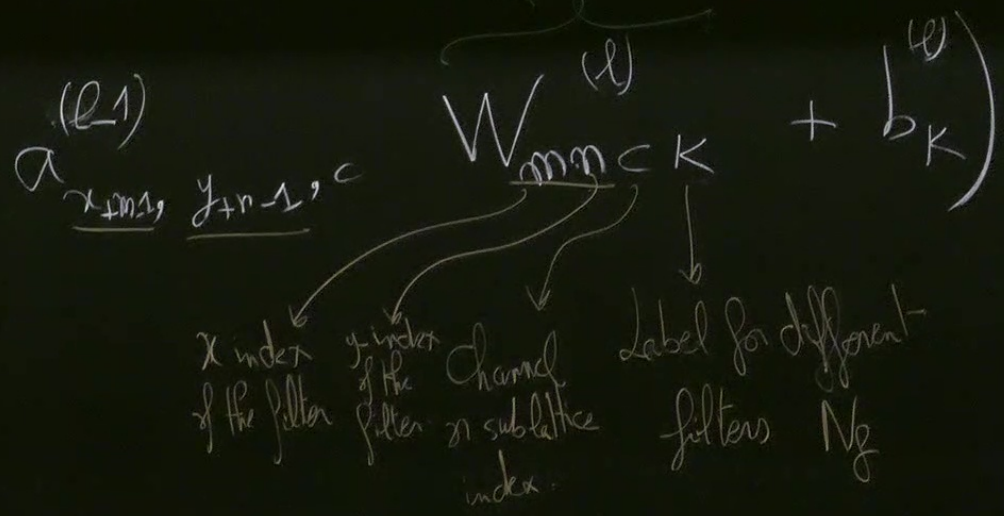
Activation function.

$x$  index of the filter  
 $y$  index of the filter  
 $c$  channel or sublattice index.  
 $K$  Label for different filters  $N_f$

$$a_{xyik}^{(l)} = \sigma \left( \sum_{m=1}^{L_{\text{patch}}} \sum_{n=1}^{L_{\text{patch}}} \sum_{c=1}^{N_c} a_{x+m, y+n, c}^{(l-1)} W_{mn} + b_k^{(l)} \right)$$

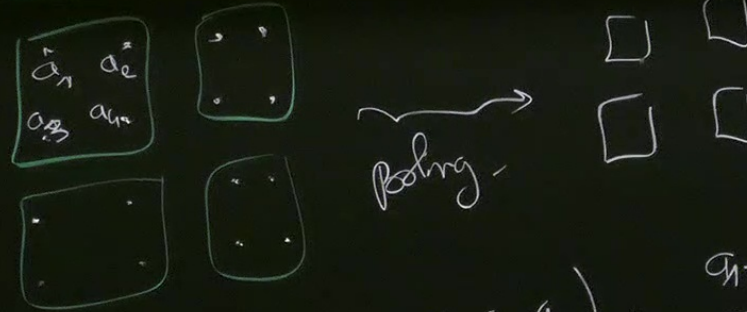
Activation function.

No dependence on  $x$  and  $y$  (translation invariance)





Pooling



$$\square = \max(a_1, a_2, a_3, a_4), \quad \frac{a_1 + a_2 + a_3 + a_4}{4}$$

Firefox File Edit View History Bookmarks Tools Window Help zoom Thu Apr 18 12:28 PM

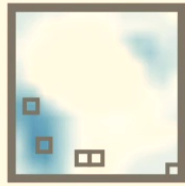
https://poloclub.github.io/cnn-explainer/#article-relu

Red channel  
Green  
Blue

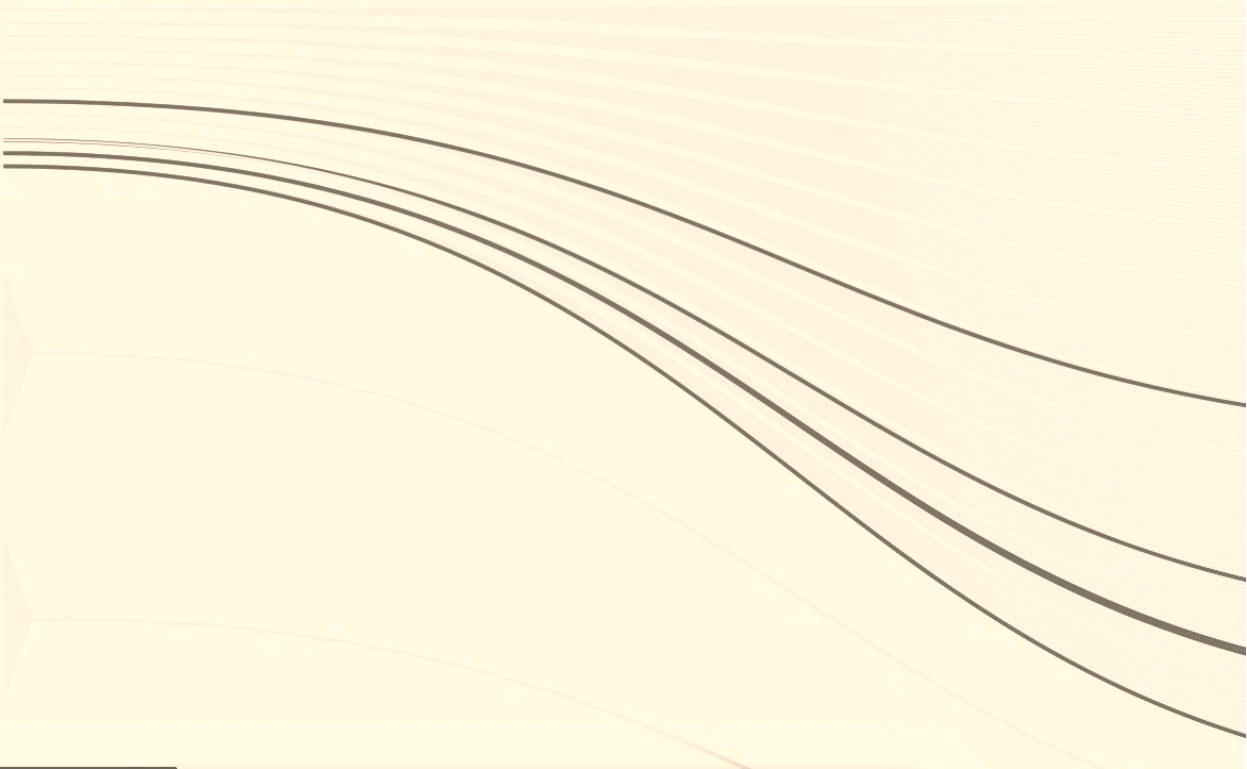
0.0 0.5 1.0  
-1.56 0.00 1.56  
-2.82 0.00 2.82  
-4.67 0.00 4.67  
-5.76

conv\_2\_2\_relu\_2\_2 max\_pool\_2  
(16, 16, 10) (26, 26, 10) (13, 13, 10)

*Hover over matrix to see how it is flattened into a 1D array!*



*Same flattening operation for each neuron*



Firefox File Edit View History Bookmarks Tools Window Help zoom Thu Apr 18 12:31 PM

https://poloclub.github.io/cnn-explainer/#article-relu

# CNN EXPLAINER Learn Convolutional Neural Network (CNN) in your browser!

max\_pool\_2 (13, 13, 10)

flatten (1690)

output (10)

lifeboat

ladybug

pizza

bell pepper

school bus

koala

(0.8798) espresso

red panda

orange

sport car

Bias +

Add up all products (element × weight) and then bias

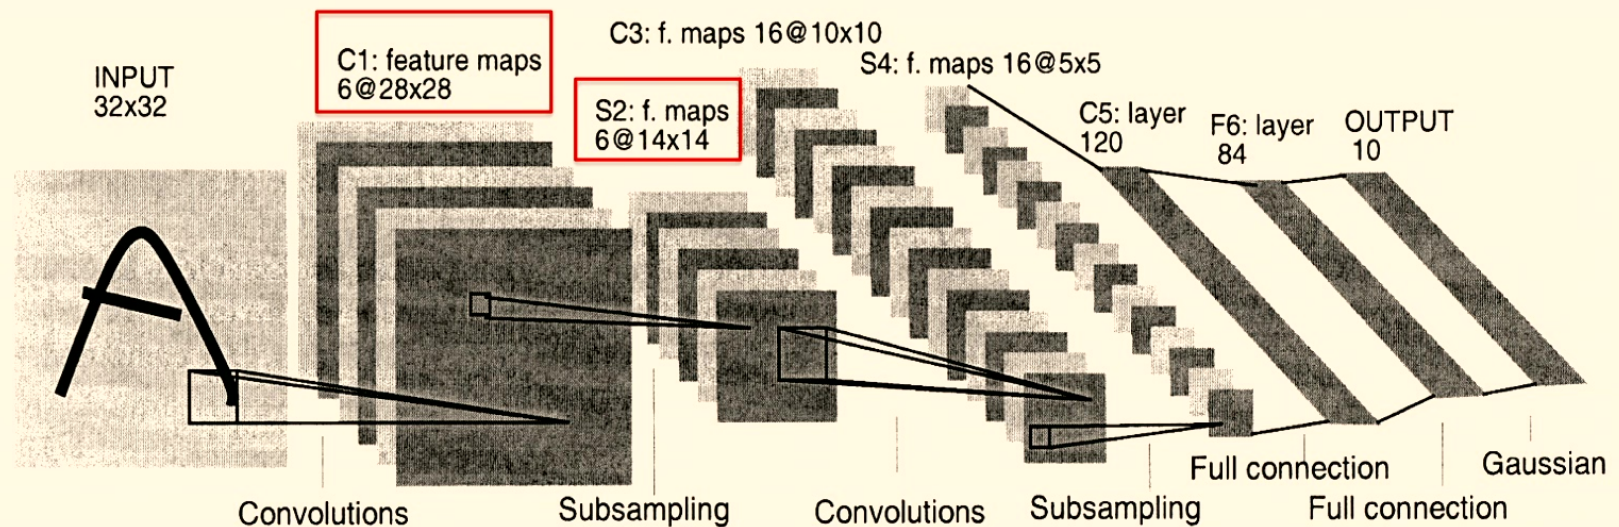
Click to learn more

softmax

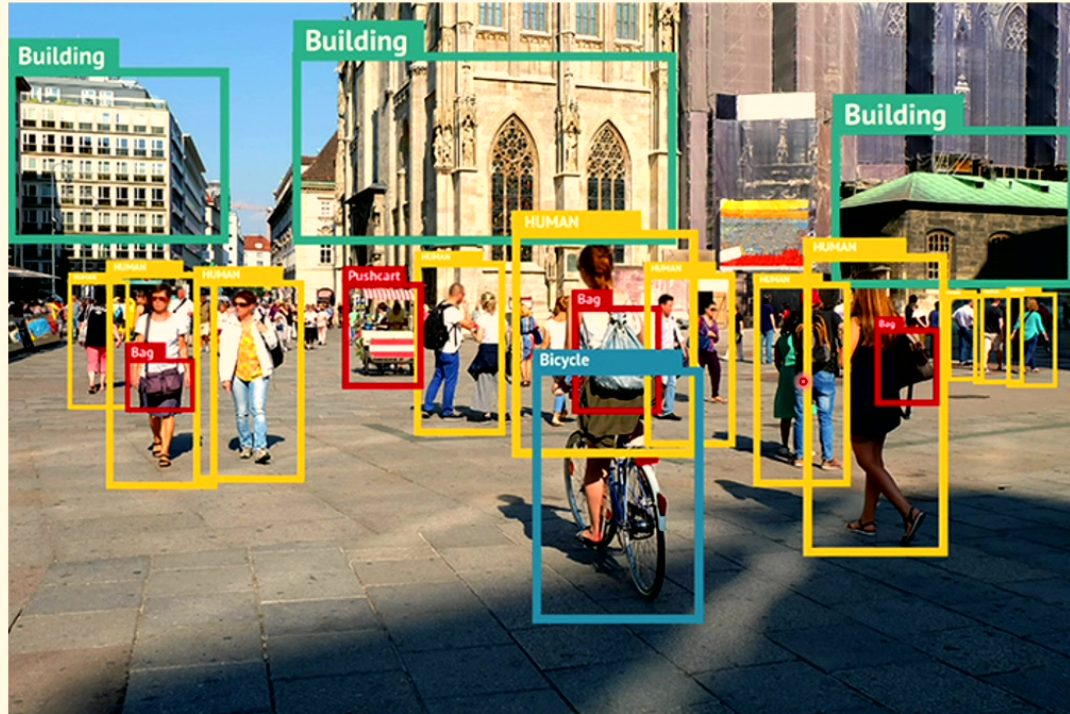
0.00 5.55 -0.12 0.00 0.10 0.0 0.9

# LeNet

- Here's the [LeNet](#) architecture, which was applied to handwritten digit recognition on MNIST in 1998:



# Applications of CNNs



Credit: [www.datasciencecentral.com](http://www.datasciencecentral.com)

# Applications of CNNs

## Two-dimensional frustrated $J_1$ - $J_2$ model studied with neural network quantum states

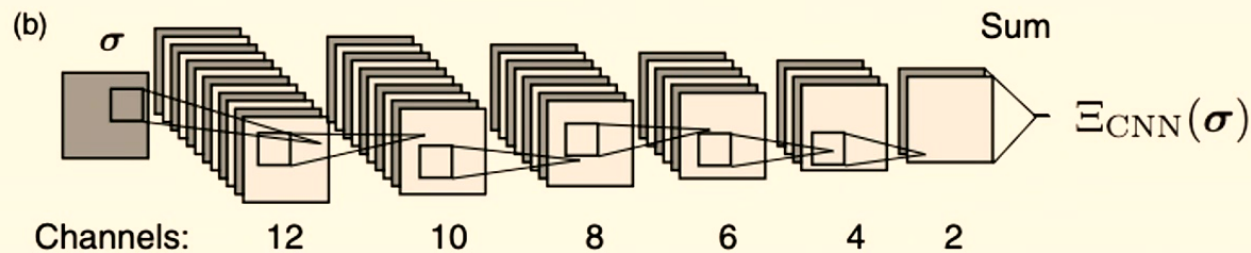
Kenny Choo<sup>1</sup>, Titus Neupert<sup>1</sup> and Giuseppe Carleo<sup>2</sup>

<sup>1</sup>Department of Physics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

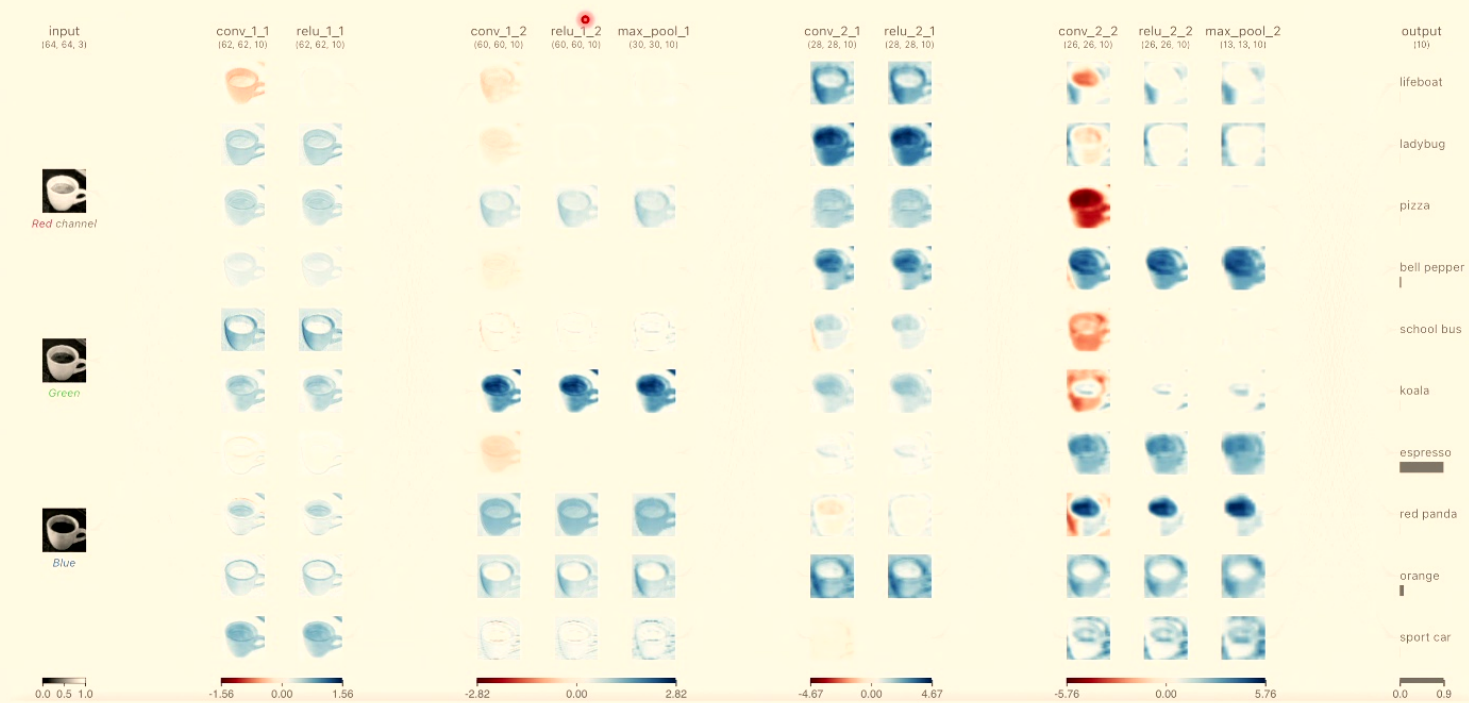
<sup>2</sup>Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, New York 10010, USA

(Received 23 March 2019; revised manuscript received 17 July 2019; published 11 September 2019)

The use of artificial neural networks to represent quantum wave functions has recently attracted interest as a way to solve complex many-body problems. The potential of these variational parametrizations has been supported by analytical and numerical evidence in controlled benchmarks. While approaching the end of the early research phase in this field, it becomes increasingly important to show how neural-network states perform for models and physical problems that constitute a clear open challenge for other many-body computational methods. In this paper, we start addressing this aspect, concentrating on a presently unsolved model describing two-dimensional frustrated magnets. Using a fully convolutional neural network model as a variational ansatz, we study the frustrated spin-1/2  $J_1$ - $J_2$  Heisenberg model on the square lattice. We demonstrate that the resulting predictions for both ground-state energies and properties are competitive with, and often improve upon, existing state-of-the-art methods. In a relatively small region in the parameter space, corresponding to the maximally frustrated regime, our ansatz exhibits comparatively good but not the best performance. The gap between the complexity of the models adopted here and those routinely adopted in deep-learning applications is, however, still substantial, such that further improvements in future generations of neural-network quantum states are likely to be expected.



# Demonstration of CNNs



<https://poloclub.github.io/cnn-explainer/>