

Title: Lecture - Machine Learning, PHYS 777

Speakers: Mohamed Hibat Allah

Collection/Series: Machine Learning (Elective), PHYS 777, February 24 - March 28, 2025

Subject: Condensed Matter, Other

Date: March 11, 2025 - 9:00 AM

URL: <https://pirsa.org/25030038>

Lecture 7: Convolutional Neural Networks

Today:

«CNNs»

↳ Local receptive fields.

↳ Channels

↳ Weight and bias sharing.

↳ Coarse graining through pooling.

Lecture F: Convolutional Neural Networks

Today:

“CNNs”

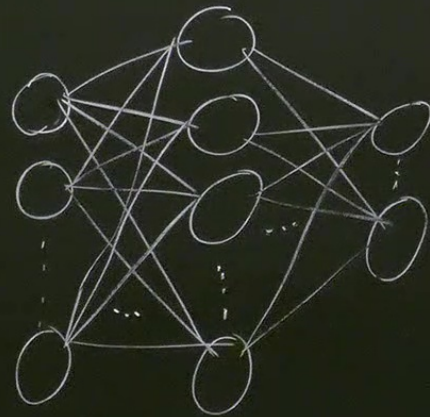
↳ Local receptive fields.

↳ Channels

↳ Weight and bias sharing.

↳ Coarse graining through pooling.

Lectures 3-6. Fully-connected feed-forward neural networks (FFNNs)



Every neuron in layer l is connected
to every neuron in layer $l+1$

Issue: FFNNs do not exploit properties such as locality and translation invariance

↳ A lot of spatial information is lost or must be learned.

1	2	3
4	5	6
7	8	9

Flatten

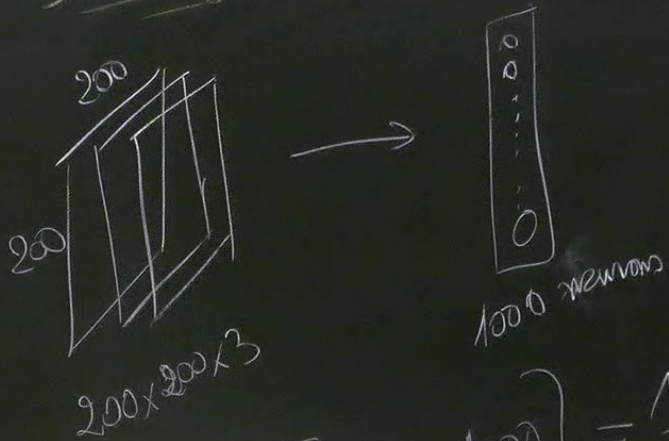
Spatially close

1
2
3
4
5
6
7
8
9

Not spatially close

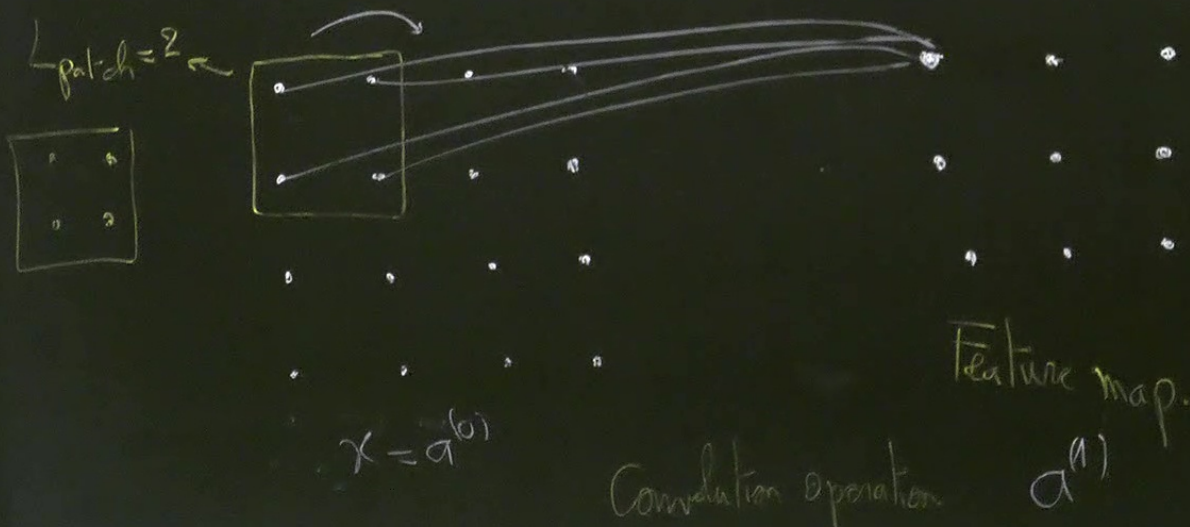
« Locality is lost »

Scability of FFNNs




$$W. [200^2 \times 3 \times 1000] = 120 \text{ million parameters.}$$

Solution. Convolution with local receptive fields / Filters / Patches...



1d dilation support

55% faster with GitHub Copilot



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

sharpen


$$\begin{pmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{pmatrix}$$

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 kernel: sharpen 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 same size. In this demo, we've instead ignored those values by making them black.

Here's a playground where you can select different kernel matrices and see how they effect the original image or build your own kernel. You can also upload your own image or use live video if your browser supports it.



Files

- master
- Go to file
- 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_transposed...
- padding_strides_transposed.gif
- same_padding_no_strides.gif
- same_padding_no_strides_trans...
- natbib
- pdf
- png
- templates
- .gitignore
- LICENSE
- README.md
- bibliography.bib
- conv_arithmetic.tex

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

Convolution arithmetic

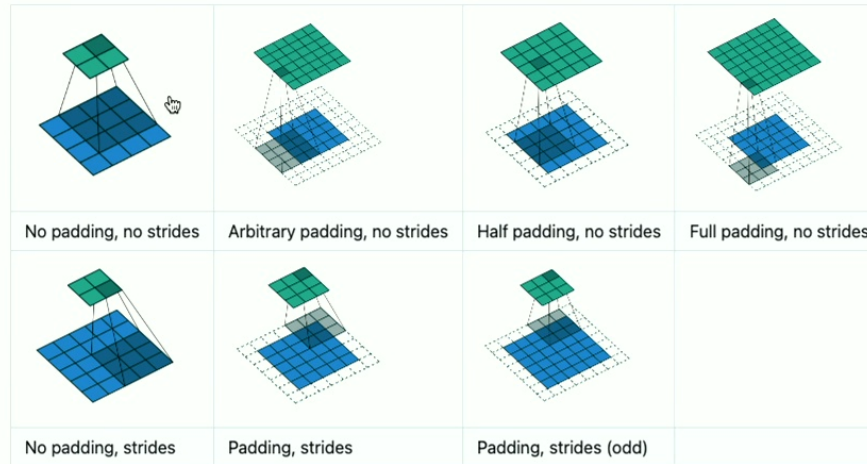
A technical report on convolution arithmetic in the context of deep learning.

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.




Transposed convolution animations

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




Generalizations:


↳ Patch size: can be a square of size $L_{patch} \times L_{patch}$ or a more general shape.



↳ Padding:



↳ Stufe:



$$\begin{bmatrix} 1 & & & \\ & 2 & & \\ & & 3 & \\ & & & 4 \end{bmatrix}$$

$$s=2$$

$$[A^{(0)}] = 4 \times 4$$

$$L_1 = \frac{L - L_{\text{patch}} + 2N_z}{\beta} + 1$$

↓
Size of Feature
map

→ If L_1 is not an integer


↳ $(L_{\text{patch}}, N_z, \beta)$ is invalid.

Channels:

^(c)
 $a_{x,y,c}$
↓ ↓
Spatial Channel
indices index

→ For color images, c can label R, G, B channels.

→ For spin configuration on lattices, c can label sublattices.




○ : Lattice site.

□ : sublattice 1

△ : sublattice 2

$$\{d_i\} = L \times L \times 2$$



$$z_{x,y,K}^{(l)} = \sum_{m=1}^{L_{\text{patch}}} \sum_{n=1}^{L_{\text{patch}}} \sum_{c=1}^{N_c} a_{x+m-1, y+n-1, c}^{(l-1)} W_{mn,K}^{(l)} + b_K^{(l)} \quad 1 \leq K \leq N_f$$

of channels
Convolution

$K \rightarrow$ bias

→ At Layer l , we take N_f filters of size $L_{\text{patch}} \times L_{\text{patch}} \times N_c$.

$$a_{x,y,K}^{(l)} = g \left(z_{x,y,K} \right)$$

↓ Activation function.

→ Weights and Biases sharing
↳ Translation invariance
↳ Reduction in terms of # parameters.

Firefox File Edit View History Bookmarks Tools Window Help

conv_arithmetic/gif/no_padding x Image Kernels explained visually x CNN Explorer x NN SVG x +

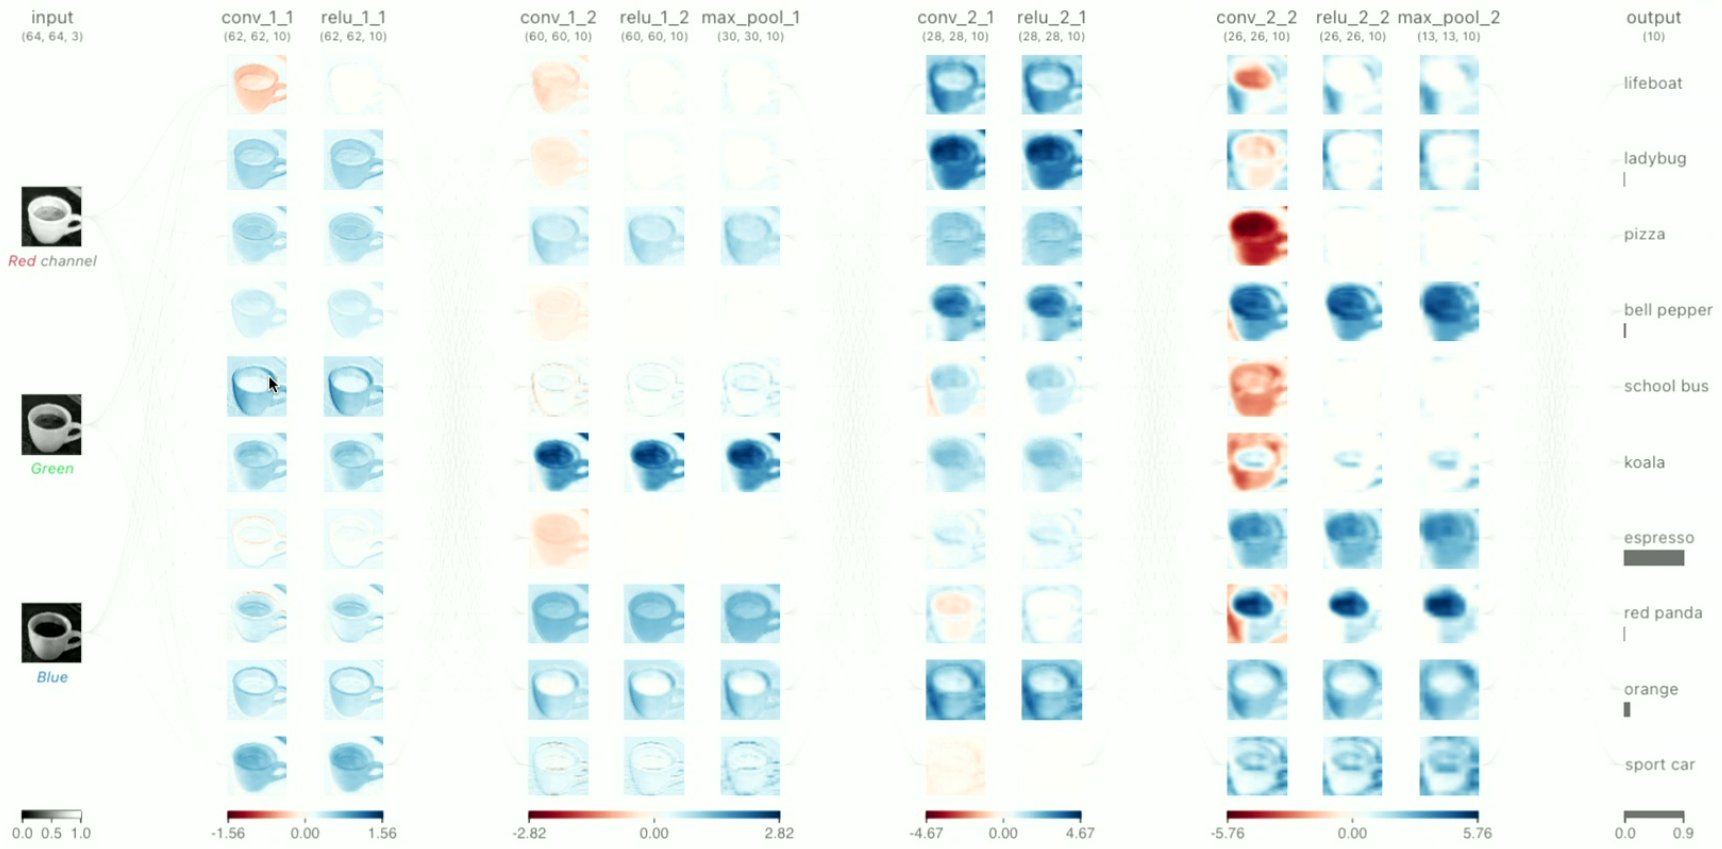
https://poloclub.github.io/cnn-explainer/ 120% ☆

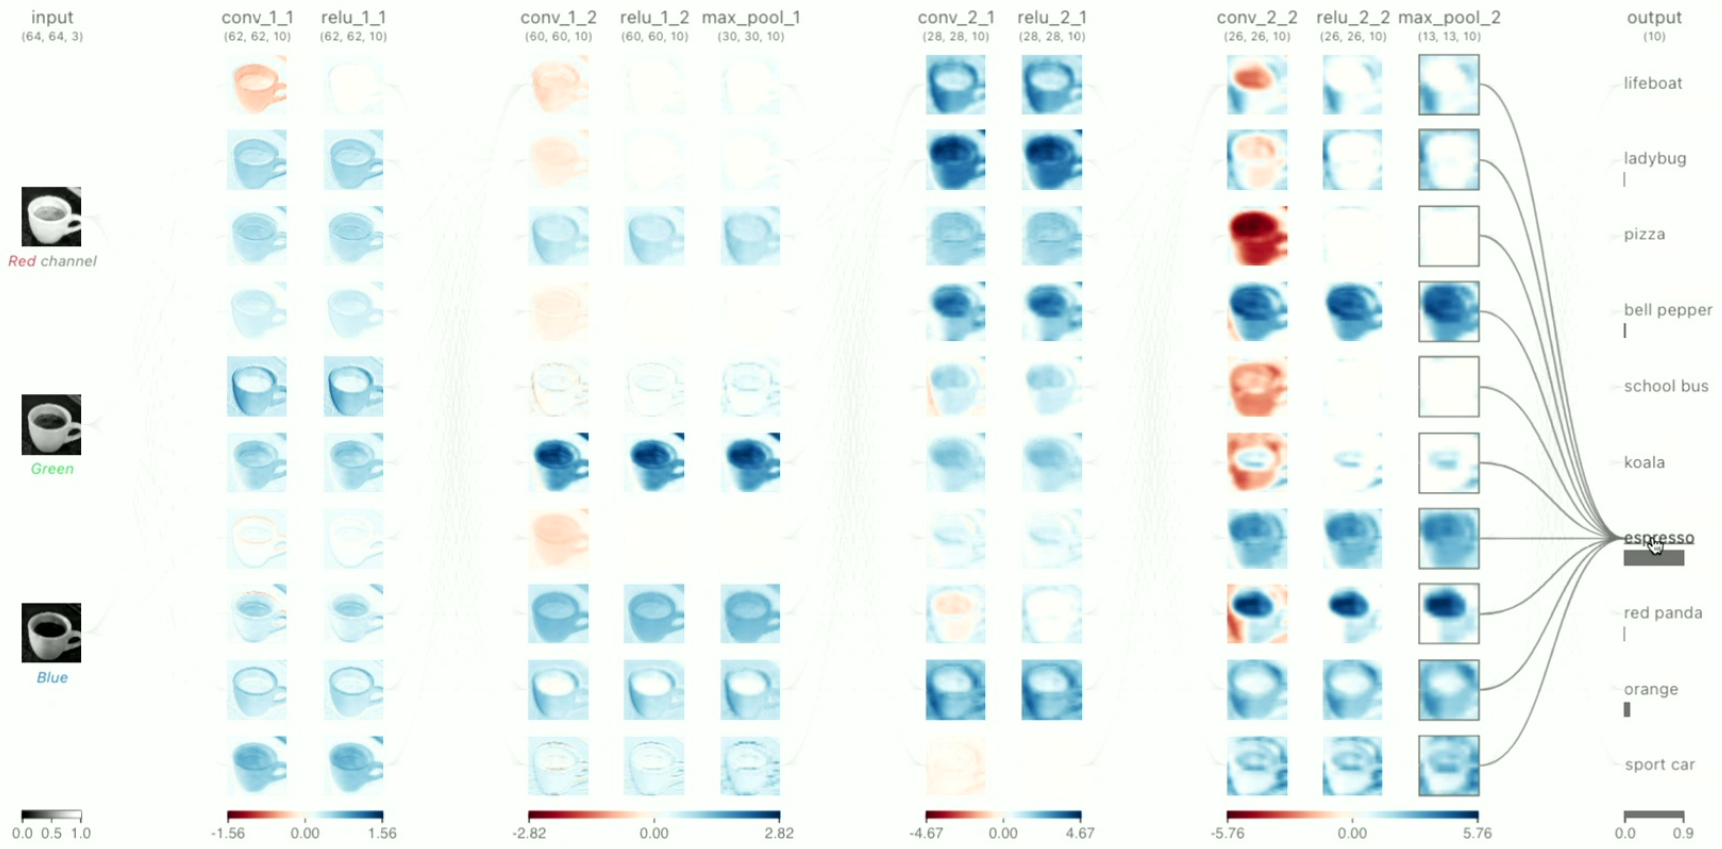
Learn Convolutional Neural Network (CNN) in your browser.

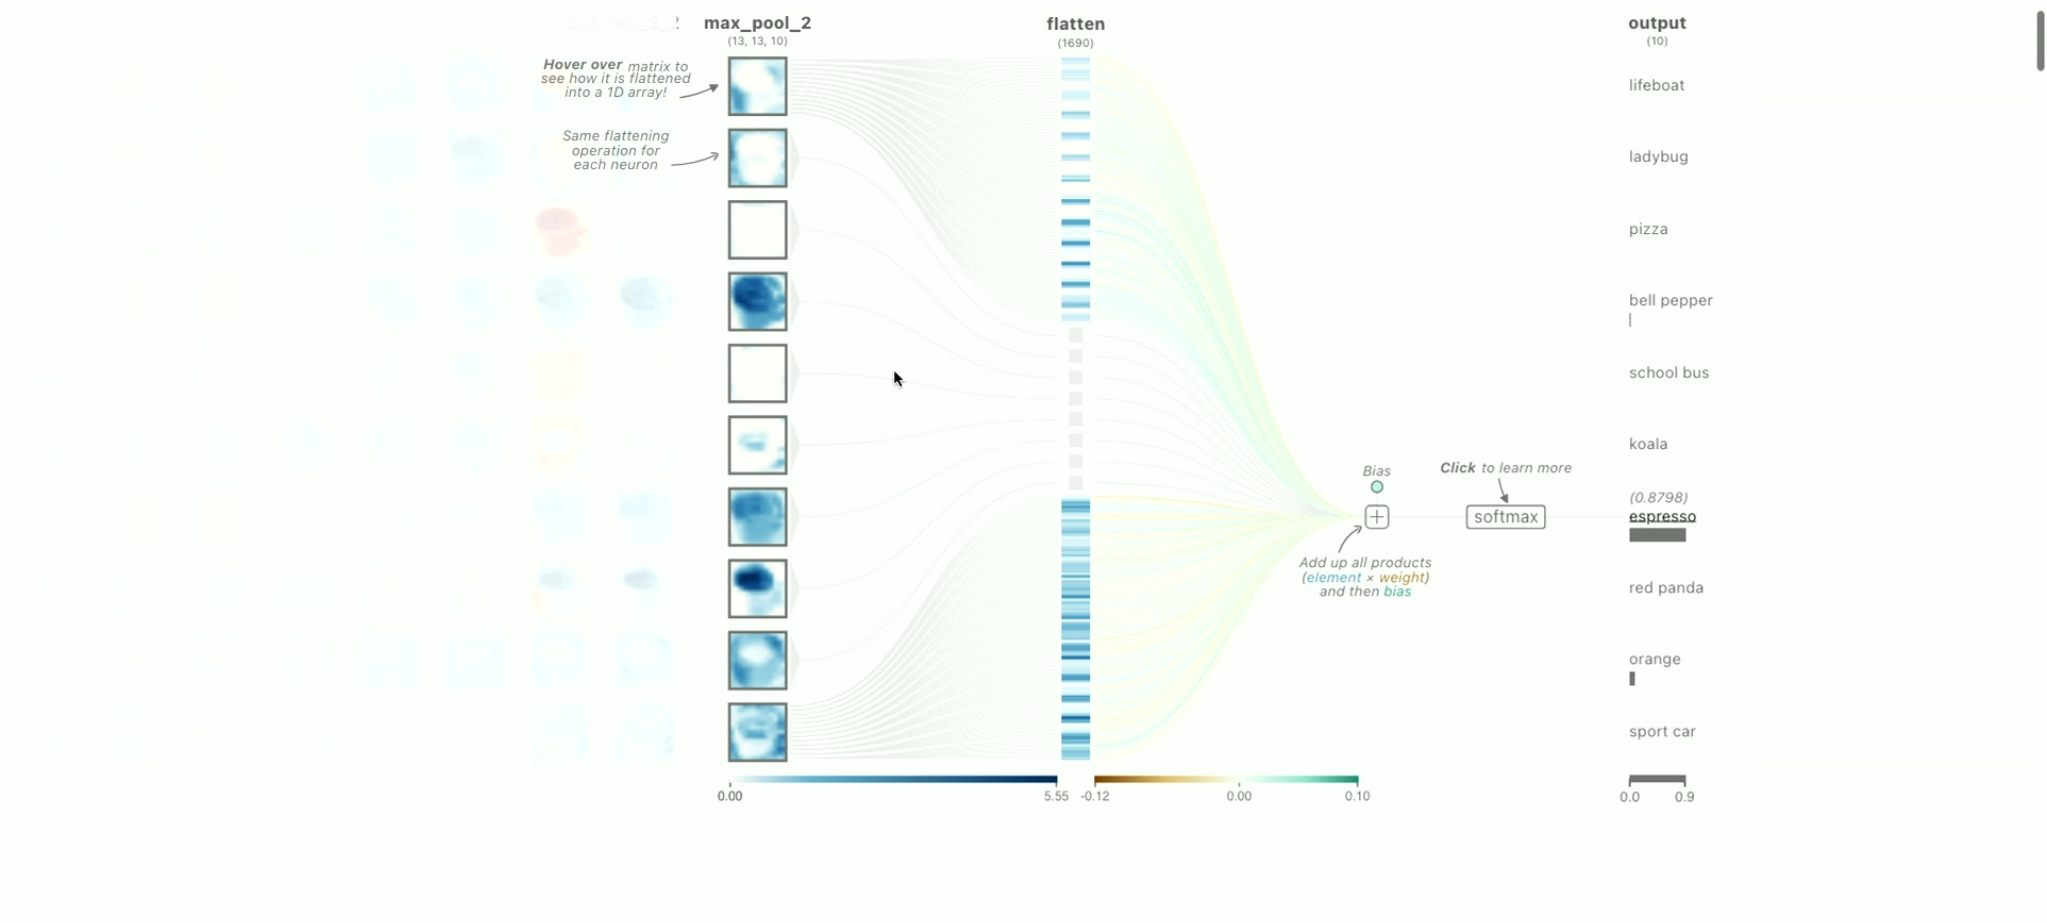
input (64, 64, 3) conv_1_1 (62, 62, 10) relu_1_1 (62, 62, 10) conv_1_2 (60, 60, 10) relu_1_2 (60, 60, 10) max_pool_1 (30, 30, 10) conv_2_1 (28, 28, 10) relu_2_1 (28, 28, 10) conv_2_2 (26, 26, 10) relu_2_2 (26, 26, 10) max_pool_2 (13, 13, 10) output (10)

Red channel
Green
Blue


lifeboat
ladybug
pizza
bell pepper
school bus
koala
espresso
red panda
orange
sport car







What is a Convolutional Neural Network?



Lecture 7: Convolutional Neural Networks

Today:

“CNNs”

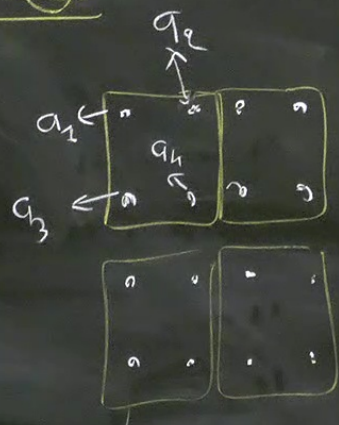
↳ Local receptive fields

↳ Channels


↳ Weight and bias sharing.

↳ Coarse graining through pooling.

Pooling



pooling




where \square can be:

$$\rightarrow \max(a_1, a_2, a_3, a_n)$$

$$\rightarrow \frac{a_1 + a_2 + a_3 + a_n}{4}$$

Majority rule




CNN architecture

Many CNNs implement one or many layers of convolution (+ pooling)
followed by a fully-connected layer(s).

LeNet

- Here's the **LeNet** architecture, which was applied to handwritten digit recognition on MNIST in 1998:




Classification

- ImageNet results over the years. There are 1000 classes. Note that errors are top-5 errors (the network gets to make 5 guesses).

Year	Model	Top-5 error
2010	Hand-designed descriptors + SVM	28.2%
2011	Compressed Fisher Vectors + SVM	25.8%
2012	AlexNet	16.4%
2013	a variant of AlexNet	11.7%
2014	GoogLeNet	6.6%
2015	deep residual nets	4.5%

- Human-level performance is around 5.1%.
- They stopped running the object recognition competition because the performance is already so good.

Applications of CNNs






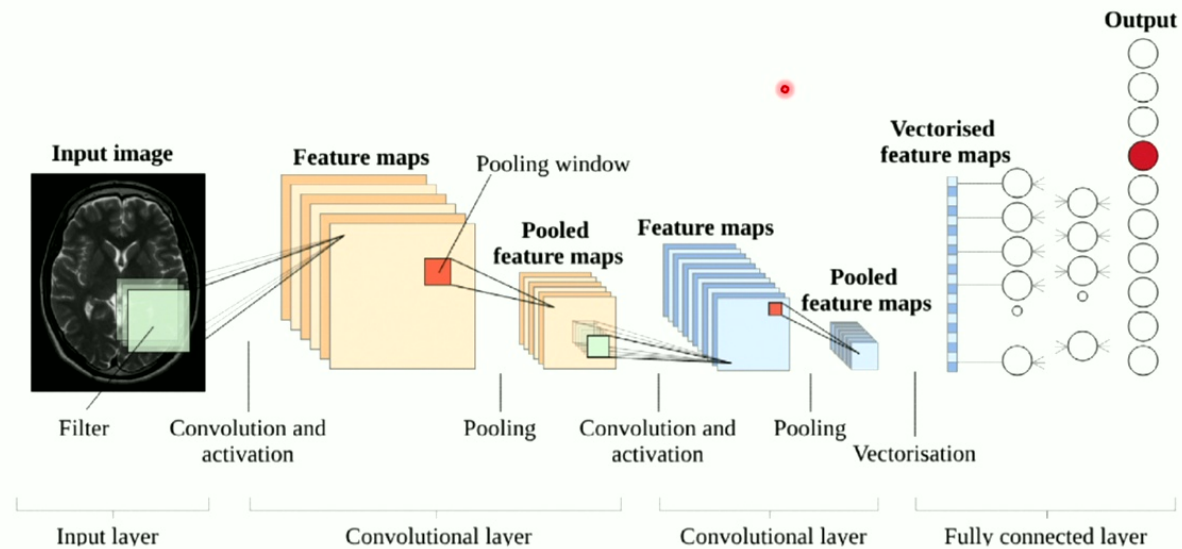
Credit: www.datasciencecentral.com

Applications of CNNs


Review

An overview of deep learning in medical imaging focusing on MRI

Alexander Selvikvåg Lundervold ^{a b}  , Arvid Lundervold ^{a c d} 



Applications of CNNs



Astronomia ex machina: a
history, primer, and outlook on
neural networks in astronomy

Michael J. Smith and James E. Geach

Department of Physics, Astronomy & Mathematics,
School of Physics, Engineering and Computer Science,
University of Hertfordshire, Hatfield, AL10 9AB

Applications of CNNs

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

Kenny Choo¹, Titus Neupert¹ and Giuseppe Carleo²

¹Department of Physics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

²Center 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.

