


Convolutional Neural Networks, an intuition...

Tuesday, October 15, 2019 2:01 PM





'80 Fukushima "Neocognitron" 16x16
Late '90 Yann LeCun LeNet-5 32x32 7 layers Digits Read ZipCodes

"Convolutional Network"




2018 Turing Award.



Image Processing.




Convolution layer.



activation function $f = \text{ReLU} = \max(0, z)$




Pooling:



LeNet-5

Pool






1. initializing weights
 2. we feed examples (input images)
 3. we compute error
 4. modify weights to minimize error.
- Repeat 2,3,4 until error is acceptable.

2009. - "CUDA" GPU for computation.
2012. - AlexNet by Alex Krizhevsky (ImageNet Challenge)
2013. - ZFNet by Matthew Zeiler
2014. - GoogLeNet by Alex Krizhevsky & Ilya Sutskever
- ⋮



<http://scs.ryerson.ca/~aharley/vis/conv/flat.html>

Current Research:

- Train The Network
- check what activates each unit.



Sometimes units don't activate





- When do these classifiers fail?
 - How to fool them

Hardware has enabled neural model to produce exciting results!!