

From Zero to CNNs to Brains

Andrei Barbu, David Mayo, Colin Conwell



CENTER FOR
Brains
Minds+
Machines



`objectnet.dev`


`0xab.com/papers/conwell-brains-mouse-cnn-2020.pdf`

`bit.ly/31IjH5u`


`bit.ly/3fSAhVh`

Vision

Vision




Vision




Computer vision failures

Computer vision failures



Computer vision failures




NHTSA: Neither Autopilot nor the driver noticed the white side of the tractor-trailer against a brightly lit sky, so the brake was not applied.


Baby machine learning

Classifiers

Classifiers




Classifiers




Linear classifiers


Linear classifiers




Linear classifiers




Linear classifiers




Linear classifiers




Linear classifiers



Linear classifiers




Non-linearities: Sigmoid




Gradient descent

Gradient descent





Linear classifiers and kernels

Linear classifiers and kernels



Linear classifiers and kernels



Intro to PyTorch

Convolution






Image convolution




Gaussian blur kernel



Gaussian blur kernel

0	0	0	5	0	0	0
0	5	18	32	18	5	0
0	18	64	100	64	18	0
5	32	100	100	100	32	5
0	18	64	100	64	18	0
0	5	18	32	18	5	0
0	0	0	5	0	0	0

Gaussian blur





Early edge detection: Sobel


$$\begin{bmatrix} -1 & 0 & +1 \\ -2 & 0 & +2 \\ -1 & 0 & +1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ +1 & +2 & +1 \end{bmatrix}$$


Early edge detection: Sobel




Object detection: HoG



Object detection: DPM



Object detection: Deep learning



Object detection

Object detection

image as input

Object detection

image as input

propose fixed-size regions that are likely to have objects

Object detection

image as input

propose fixed-size regions that are likely to have objects

feed each through a network with *many* layers

Object detection




image as input

propose fixed-size regions that are likely to have objects

feed each through a network with *many* layers

linear classifier at the top

Performance



Vision ...

Vision ...



Vision . . .



A man and a child are sitting on a plane.