




Single Cell Recording





Receptive field structure in ganglion cells: On-center Off-surround




Stimulus




Electrical response

Receptive field structure in ganglion cells: On-center Off-surround




Stimulus condition




Electrical response

Receptive field structure in ganglion cells: On-center Off-surround




Stimulus condition



Electrical response

Receptive field structure in ganglion cells: On-center Off-surround




Stimulus condition



Electrical response

Receptive field structure in ganglion cells: On-center Off-surround




Stimulus condition



Electrical response

Receptive field structure in ganglion cells: On-center Off-surround




Stimulus condition




Electrical response

On-center Off-surround cells

Receptive Field




Response Profile




Off-center On-surround cells

Receptive Field




Response Profile


Firing Rate



Horizontal Position



Cortical Area V1 aka primary visual cortex



Single-cell recording from visual cortex



David Hubel & Thorston Wiesel


2 Classes of Cells in V1

Simple cells


Complex cells

Simple Cells: Line Detectors

A. Light Line Detector




Firing Rate




Horizontal Position

B. Dark Line Detector




Firing Rate



Horizontal Position

Simple Cells: Edge Detectors

C. Dark-to-light Edge Detector




Firing Rate




Horizontal Position

D. Light-to-dark Edge Detector



Firing Rate



Horizontal Position


Complex Cells: Location Independent

STIMULUS



0°

NEURAL RESPONSE



STIMULUS



60°

NEURAL RESPONSE




STIMULUS



120°

NEURAL RESPONSE



STIMULUS




90°


NEURAL RESPONSE





Neural Line Detector




Location-Independent Line Detector



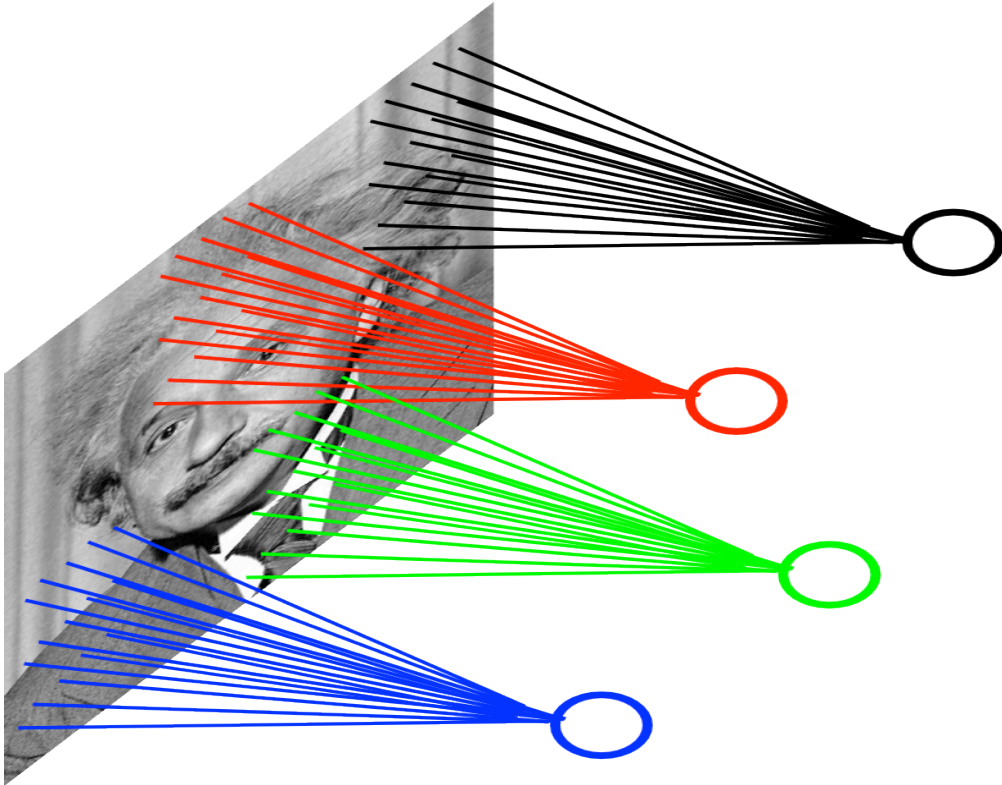
“Retinatopic” map from Retina to V1




Convolutional neural nets (CNNs), LeCun, 1989. LeNet 5 classifier for handwritten digits.




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer



Convolutional Layer



Convolutional Layer




Convolutional of Two Signals




$$* \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix} =$$

mask




(vector, NOT a matrix!)



Hidden layer of “depth” 5:
five neurons all looking at
the same patch; five
different masks.


Apply the same 5 masks to
each patch. Five neurons
per patch.




Hidden layer of “depth” 5:
five neurons all looking at
the same patch; five
different masks.

Apply the same 5 masks to
each patch. Five neurons
per patch.

Convolutional Layer



learned masks






Activations:




convolving the input with the first mask gives the first output slice


Activations:



Convolution layers are often followed by pooling/
subsampling layers that shrink the image. Each
pooled unit is the maximum of a 4-unit block.
Hardwired, no weights to train.



Convolutional neural nets (CNNs), LeCun, 1989. LeNet 5 classifier for handwritten digits.



IMAGENET Large Scale Visual Recognition Challenge

The Image Classification Challenge:
1,000 object classes
1,431,167 images




Output:
Scale
T-shirt
Steel drum
Drumstick
Mud turtle



Output:
Scale
T-shirt
Giant panda
Drumstick
Mud turtle



Convolutional neural nets, Krizevsky et al., 2012



AlexNet: A competition-winning classifier for recognizing images in objects. The ImageNet Large Scale Visual Recognition Challenge, 2012.

- + millions of images
- + GPUs
- + ReLUs
- + dropout