




# Recording Action Potentials





# 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


# On-center Off-surround cells

Receptive Field



Response Profile


Firing Rate



Horizontal Position


# 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


# Orientation Dependence

STIMULUS



$0^\circ$

NEURAL RESPONSE




STIMULUS



$60^\circ$

NEURAL RESPONSE




STIMULUS



$120^\circ$

NEURAL RESPONSE



STIMULUS




$90^\circ$

NEURAL RESPONSE




# Simple Cell: Line Detector








# Complex Cell: 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!)




5 hidden units all looking at the same patch; 5 different masks.

Apply the same 5 masks to each patch = 5 units 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 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

Steel drum

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
- + ReLUs
- + GPUs
- + dropout