

# Convolutional Neural Networks

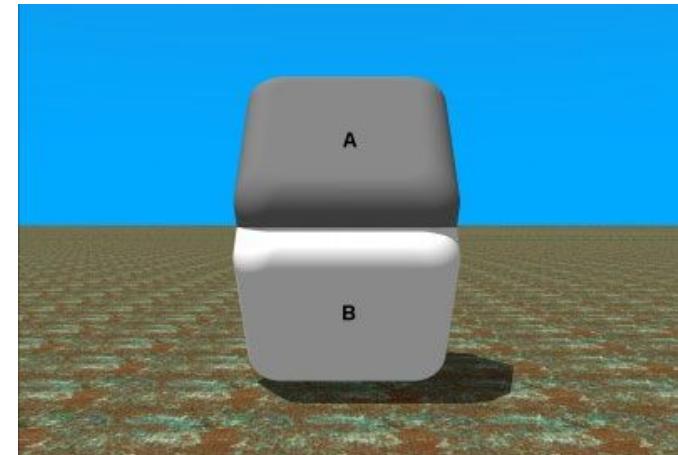
Andrew H. Fagg  
Symbiotic Computing Laboratory  
University of Oklahoma

# Deep Networks for Image Recognition

- Images are composed of large numbers of pixels
- A particular pixel value can vary a lot:
  - Color, illumination
- Objects can vary a lot
  - Size, orientation, perspective

Individual pixels are irrelevant...

it is the groups of pixels that matter

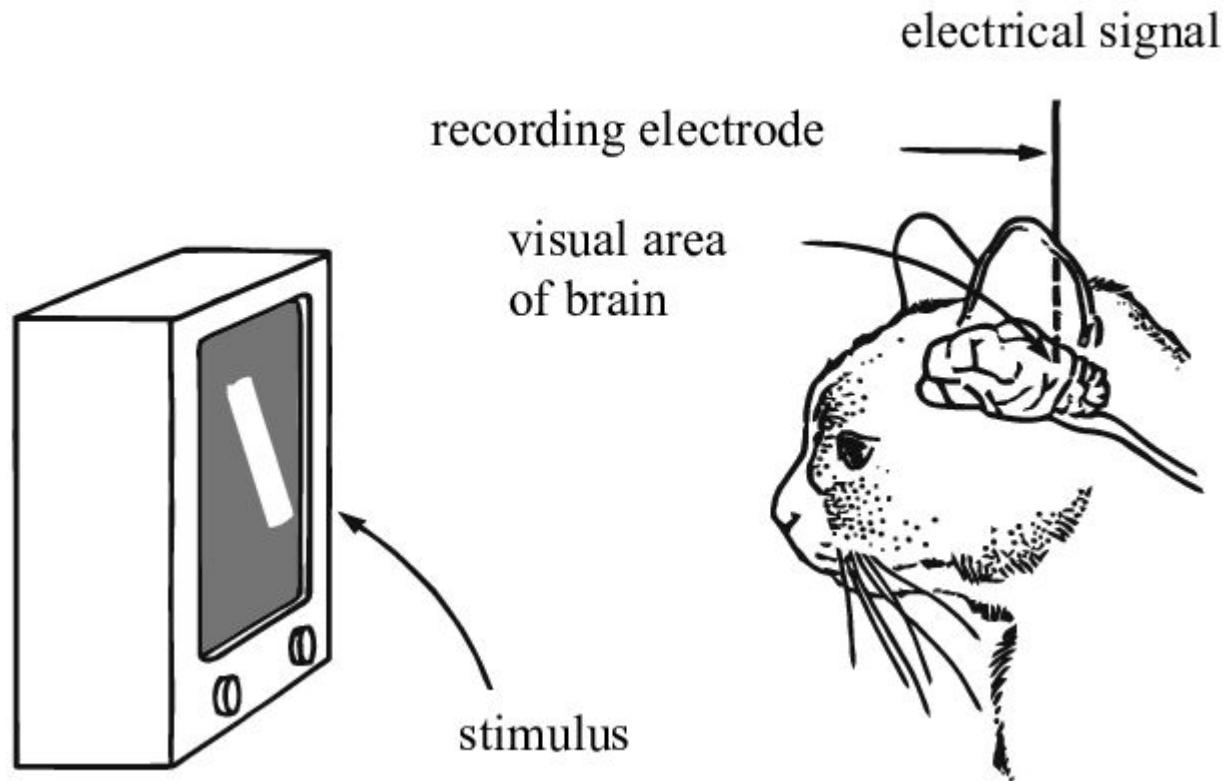


# Deep Networks for Image Recognition

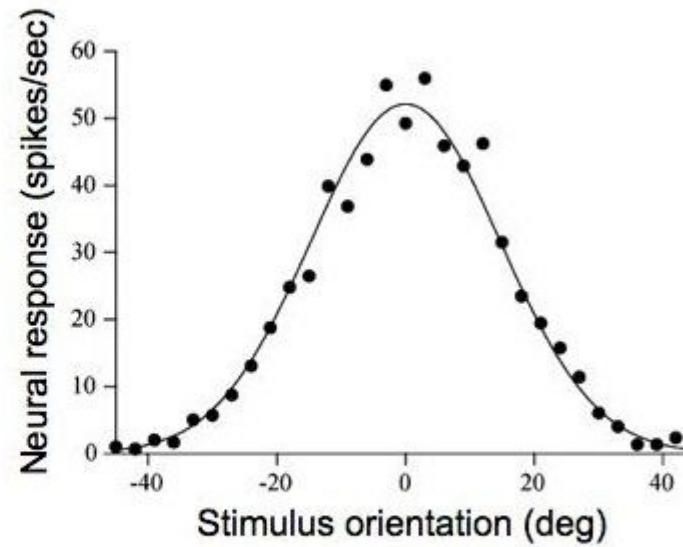
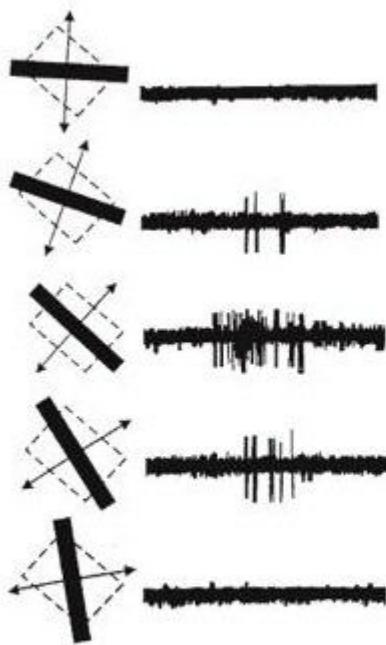
- 1920 (columns) x 1020 (rows) x 3 (channels = RGB) is almost 6 million inputs
- If the next dense layer has 1000 units, then we would have 6 billion parameters!

Need lots of examples and lots of training time. How do we get beyond this?

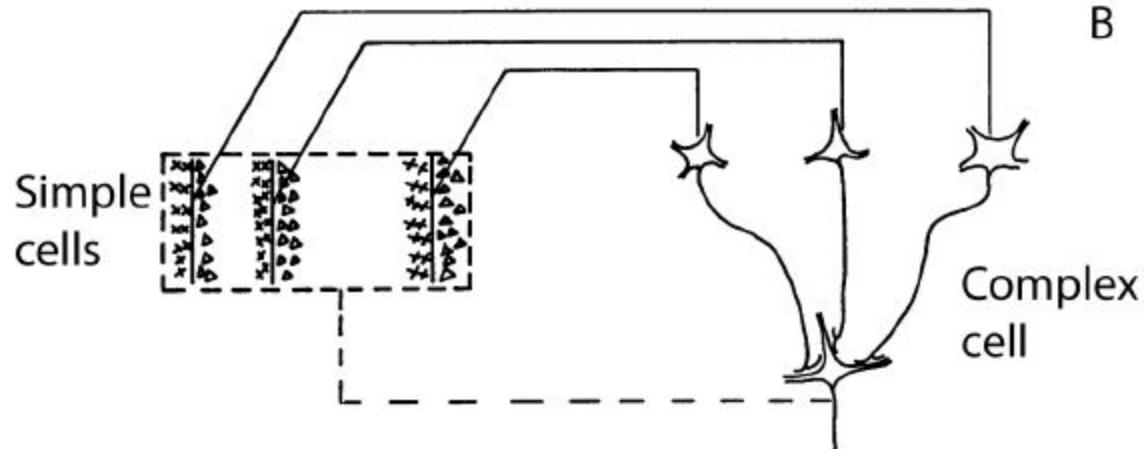
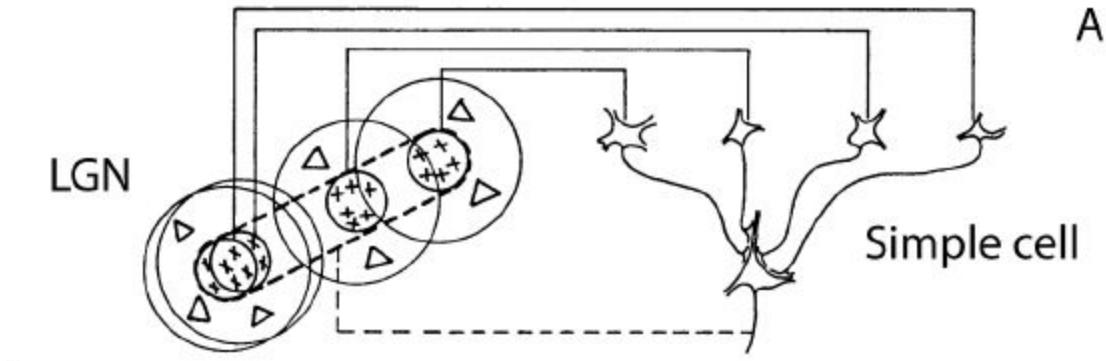
# Hubel and Wiesel (1968)

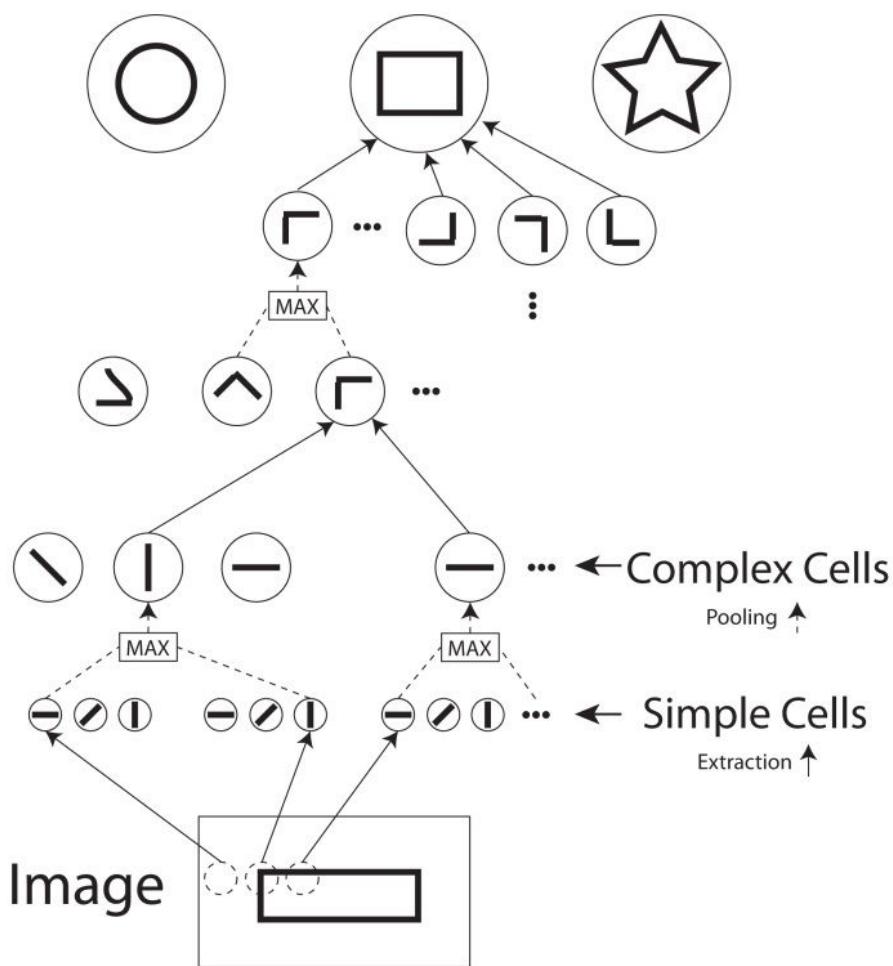


# Orientation Sensitivity



# Complex Features Formed from Simple Ones

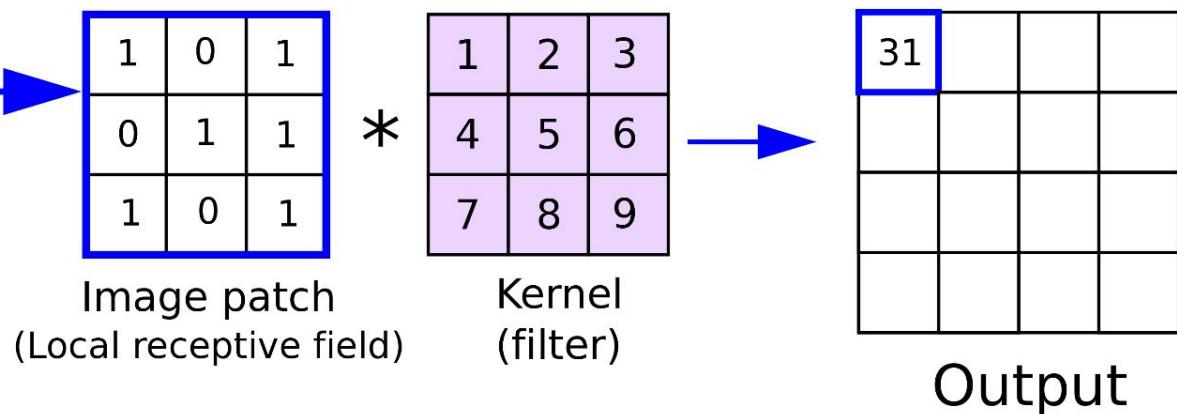




# Convolution

1	0	1	0	1	0
0	1	1	0	1	1
1	0	1	0	1	0
1	0	1	1	1	0
0	1	1	0	1	1
1	0	1	0	1	0

Input



<https://anhreynolds.com/blogs/cnn.html>

# Convolution: Edge Detector

10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0
10	10	10	10	0	0	0	0	0

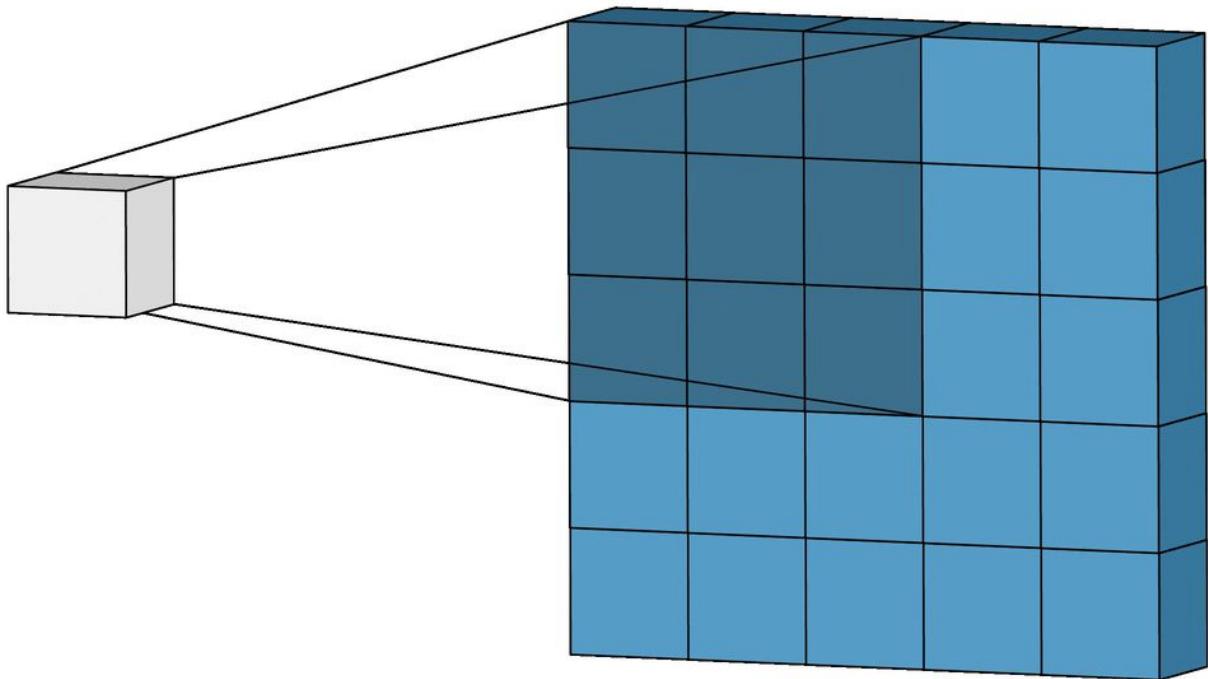
\*

1	0	-1
1	0	-1
1	0	-1

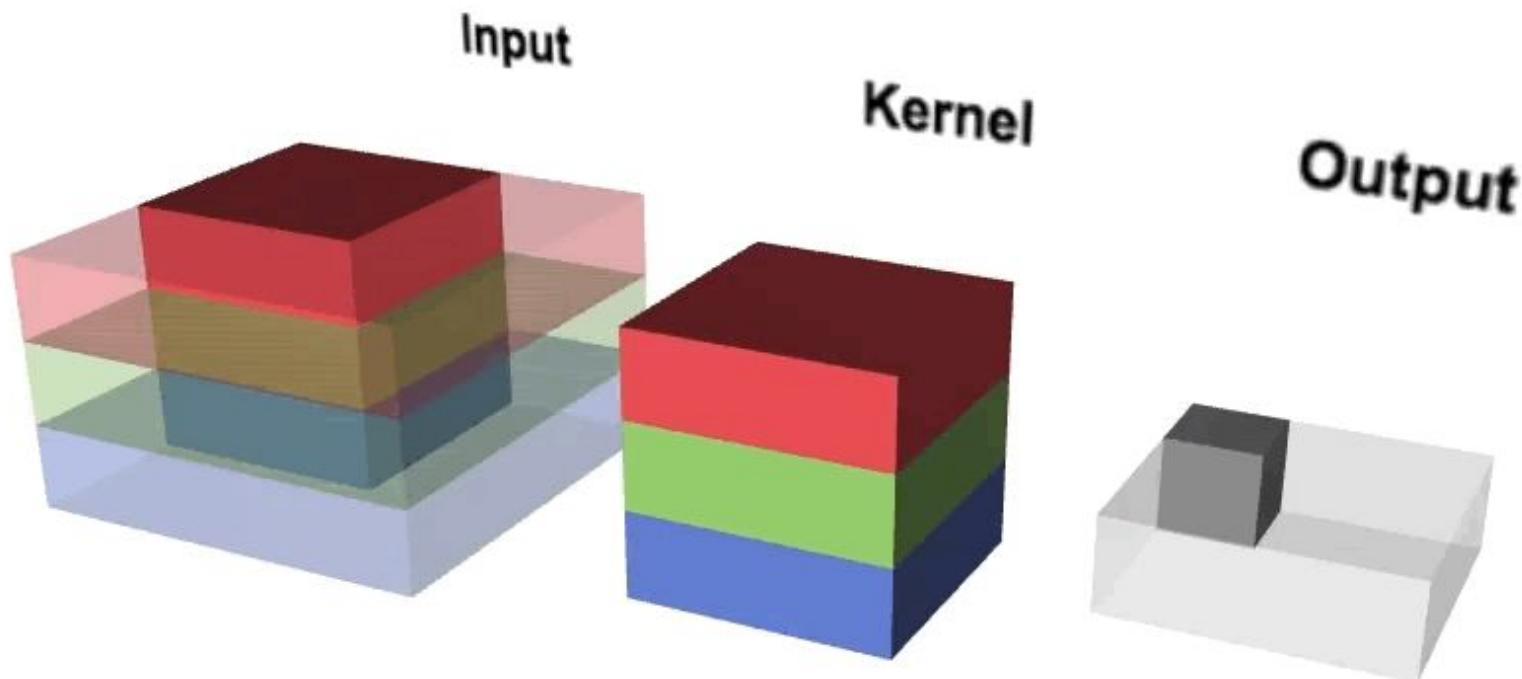
Vertical

=

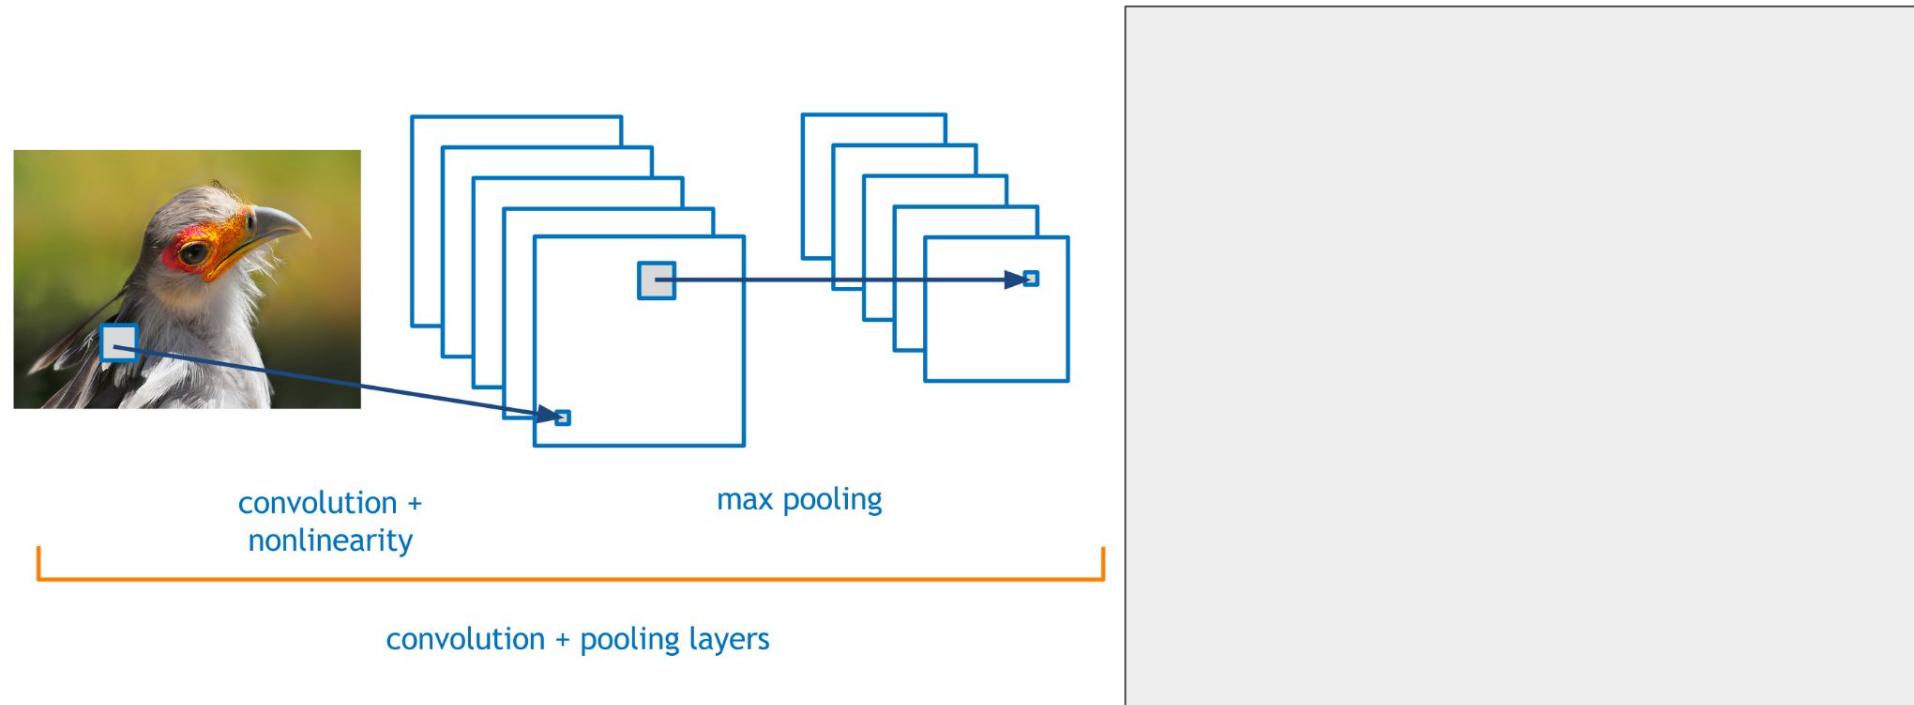
0	0	30	30	0	0
0	0	30	30	0	0
0	0	30	30	0	0
0	0	30	30	0	0
0	0	30	30	0	0



<https://towardsdatascience.com/intuitively-understanding-convolutions-for-deep-learning-1f6f42faee1>



# Local Operators



# Operator Types

- Convolution: Feature detection - recognize some pattern over a small grid of inputs
  - At a given layer, have many different patterns that we are looking for in parallel
- Max Pooling: does there exist some pattern within a small grid of inputs?
- Scaling: Allows simple feature detection and pooling to apply at multiple visual scales

# Local Operators

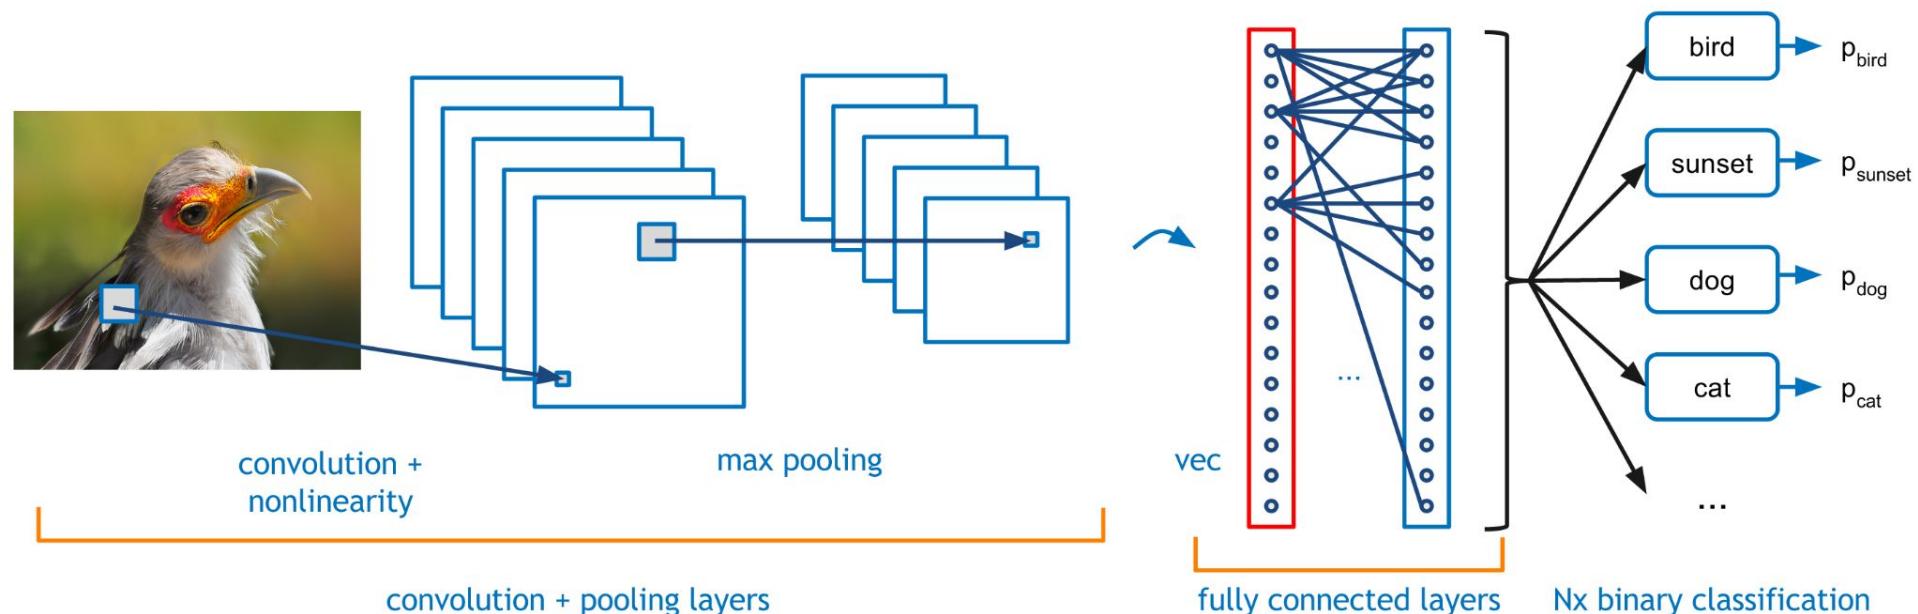
- Multiple stacked modules consisting of pattern recognition (convolution), pooling (max) and scaling (striding)
- With each module, our representation becomes more and more abstract
  - Ultimately: feathers, eyes, beaks ...
  - All have specific visual patterns, though there may be many variations of each

# Beyond the Primitives

How should the primitives be combined to form more of a semantic representation (dog, cat, grandma, etc.)?

- After computing the primitives in the first layers of our deep network, employ dense layers to allow for arbitrary combinations of the primitives

# Combining Local Operators to Recognize Global Patterns



# Applications of CNNs

- Image classification
- Image recoding: deblurring, colorization, semantic segmentation
- Image generation

1D and 3D data are possible, too



# CNN Details: Convolution

```
from tensorflow.keras.layers import Conv2D
#####
model = Sequential()
model.add(InputLayer(input_shape=(image_size[0],
    image_size[1], nchannels), name='input'))
# Input shape: (rows, cols, chans)
model.add(Conv2D(filters=10,
    kernel_size=3,
    strides=1,
    padding='valid',
    use_bias=True,
    name='C0',
    activation='elu'))
# Output shape: (rows-2, cols-2, 10)
```

# Convolution2D

Conv2d other key properties:

- kernel\_initializer
- bias\_initializer
- kernel\_regularizer
- bias\_regularizer
- activity\_regularizer

# Pooling

```
from tensorflow.keras.layers import MaxPooling2D

#####
# Input shape: (rows, cols, chans)

model.add(MaxPooling2D(pool_size=2,
                      strides=2,
                      padding='same',
                      use_bias=True,
                      name='MP0'))

# Output shape: (rows//2, cols//2, chans)
```

# Global Max Pooling

```
from tensorflow.keras.layers import GlobalMaxPooling2D

#####
# Input shape: (rows, cols, chans)
model.add(GlobalMaxPooling2D())
# Output shape: (chans,)
```

# Dropout

Drop entire channel at once

```
from tensorflow.keras.layers import SpatialDropout2D

#####
# Input shape: (rows, cols, chans)
model.add(SpatialDropout2D(p))
# Output shape: (rows, cols, chans)
```

# CNN Notes

- 1D, 2D, and 3D versions built into Keras/TF
- Can use BatchNormalization() as usual
  - Applies individually to every element in the (rows, cols, chans) Tensor

# CNN Modules

Sequence of layers:

- $k \times \text{Conv2D}$
- $\text{MaxPooling2D}$
- $\text{SpatialDropout2D}$
- $\text{BatchNormalization}$

# CNN for Image Classification

- $n \times$  CNN Module
  - Decreasing rows & cols while increasing filters  
(product should decrease)
- GlobalMaxPooling2D
- $m \times$  Dense
  - Decreasing number of hidden units
- Dense( $n_{\text{classes}}$ , activation='softmax')
  - Classes are exclusive

# Different N-Class Network Configs

All: N output units

	<b>Exclusive Classes</b>	<b>Multi-Class (any combination of classes)</b>
<b>Nonlinearity</b>	softmax	sigmoid
<b>Desired output: 1-hot encoding of class</b>	categorical_crossentropy categorical_accuracy	binary_crossentropy binary_accuracy
<b>Desired output: 1 integer (class number)</b>	sparse_categorical_crossentropy sparse_categorical_accuracy	X