

# Convolutional Neural Networks

Computer Vision  
CS 543 / ECE 549  
University of Illinois

Jia-Bin Huang

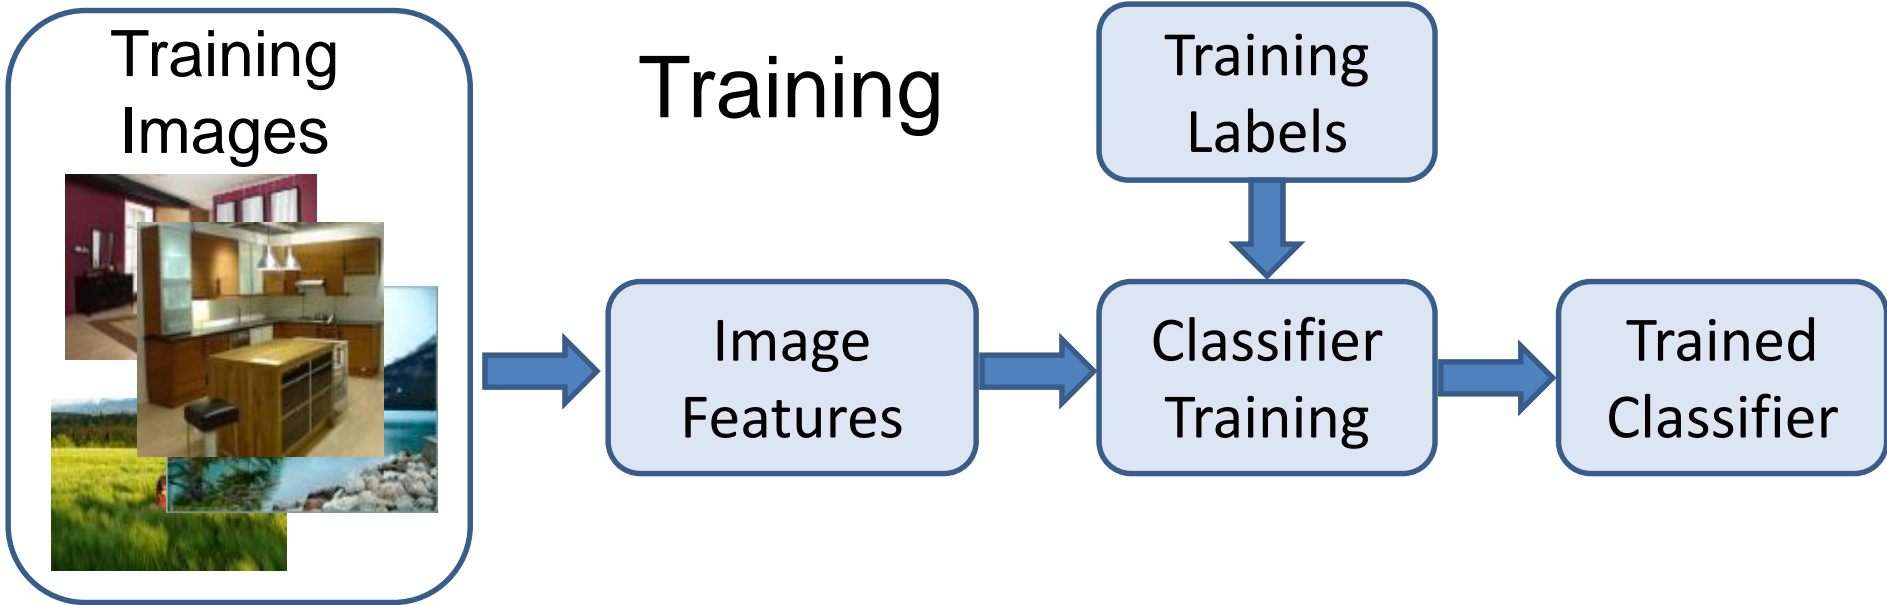
# Reminder: final project

- Posters on **Friday, May 8** at **7pm** in **SC 2405**
  - two rounds, 1.25 hr each
- Papers due **May 11** by email
- Cannot accept late papers/posters due to grading deadlines
- Send Derek an email if you can't present your poster

# Today's class

- Overview
- Convolutional Neural Network (CNN)
- Understanding and Visualizing CNN
- Training CNN
- Probabilistic Interpretation

# Image Categorization: Training phase



# Image Categorization: Testing phase

Training Images



Training

Training Labels

Image Features

Classifier Training

Trained Classifier

Testing

Image Features

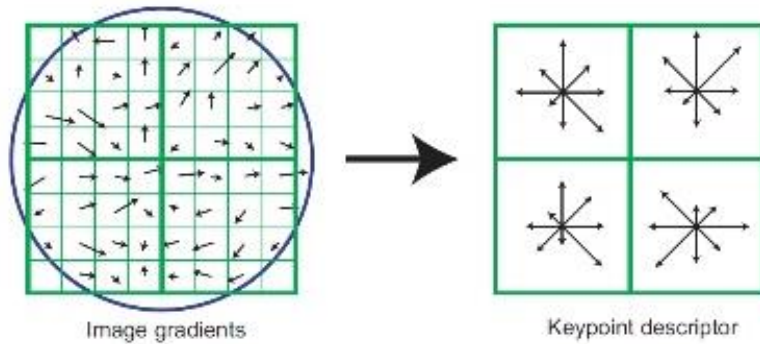
Trained Classifier

Prediction  
**Outdoor**

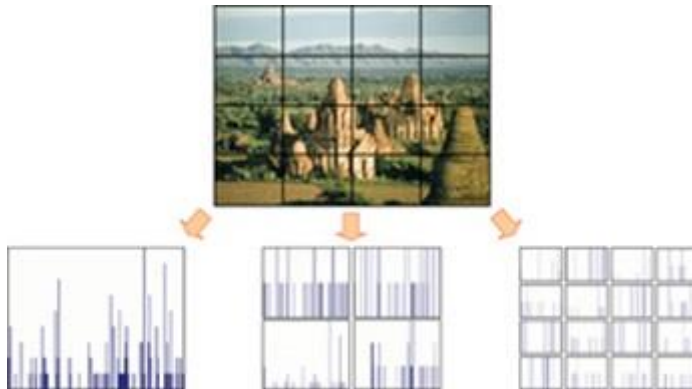


Test Image

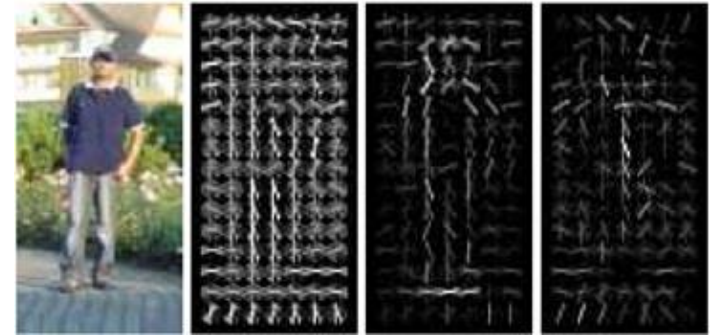
# Features are the Keys



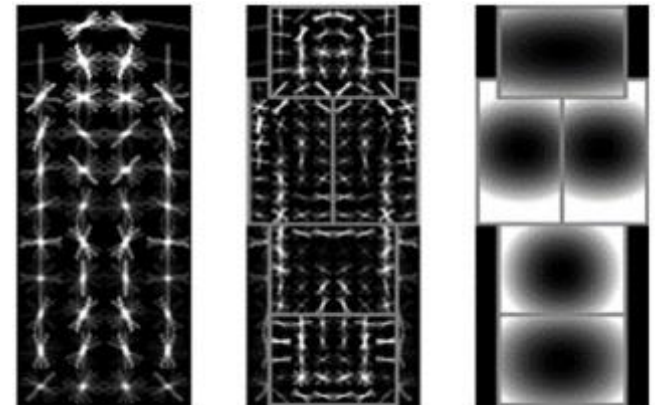
SIFT [Loewe IJCV 04]



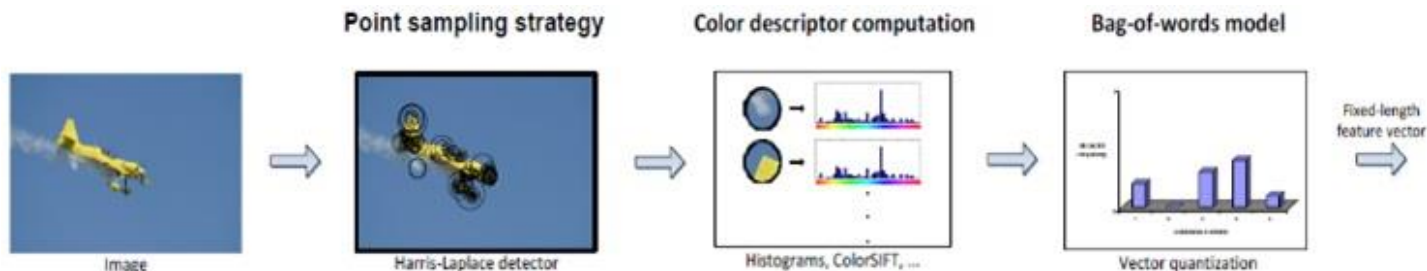
SPM [Lazebnik et al. CVPR 06]



HOG [Dalal and Triggs CVPR 05]



DPM [Felzenszwalb et al. PAMI 10]



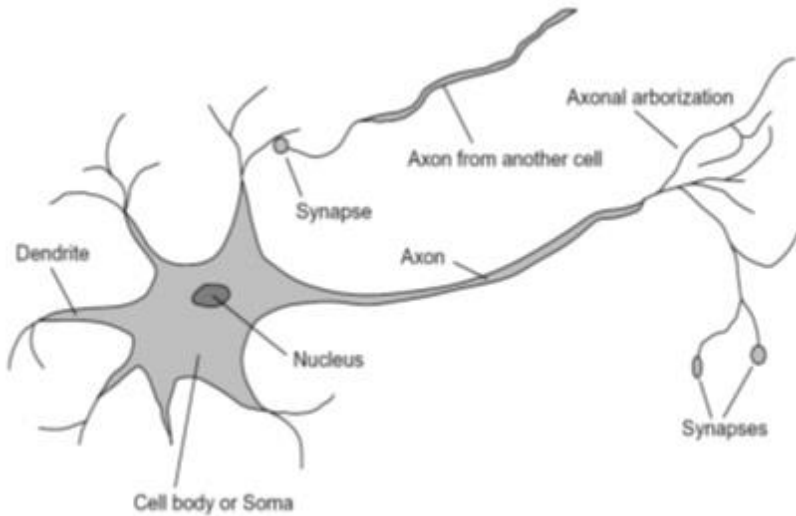
Color Descriptor [Van De Sande et al. PAMI 10]

# Learning a Hierarchy of Feature Extractors

- Each layer of hierarchy extracts features from output of previous layer
- All the way from pixels  $\rightarrow$  classifier
- Layers have the (nearly) same structure

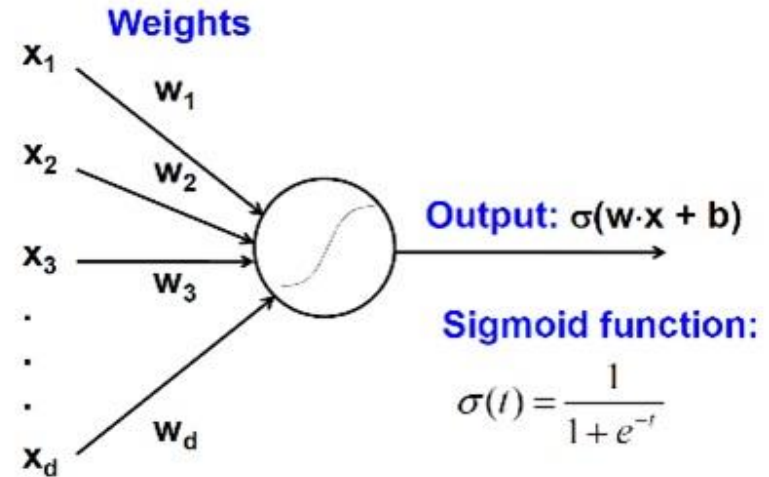


# Biological neuron and Perceptrons



A biological neuron

Input



An artificial neuron (Perceptron)  
- a linear classifier





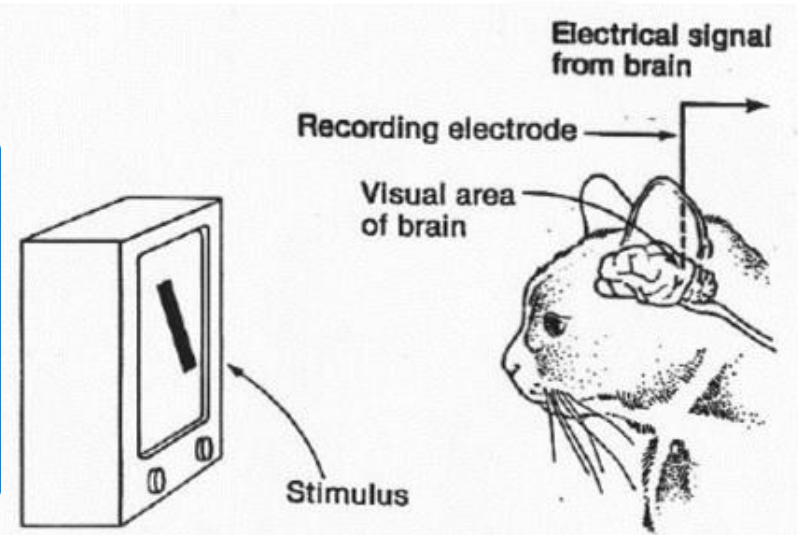
# Simple, Complex and Hypercomplex cells



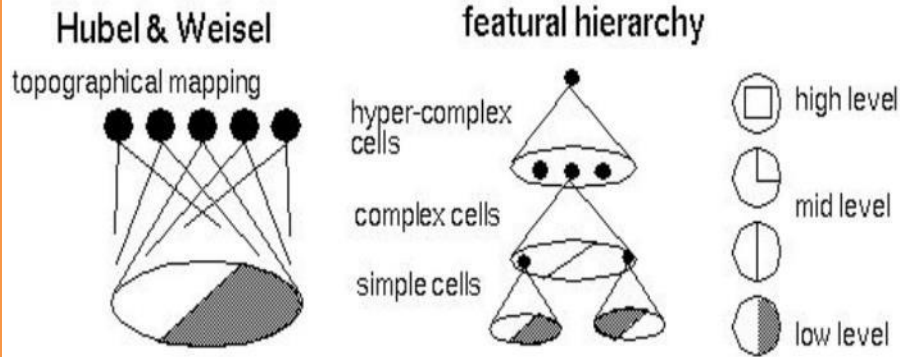
David H. Hubel and Torsten Wiesel

Suggested a **hierarchy of feature detectors** in the visual cortex, with higher level features responding to patterns of activation in lower level cells, and propagating activation upwards to still higher level cells.

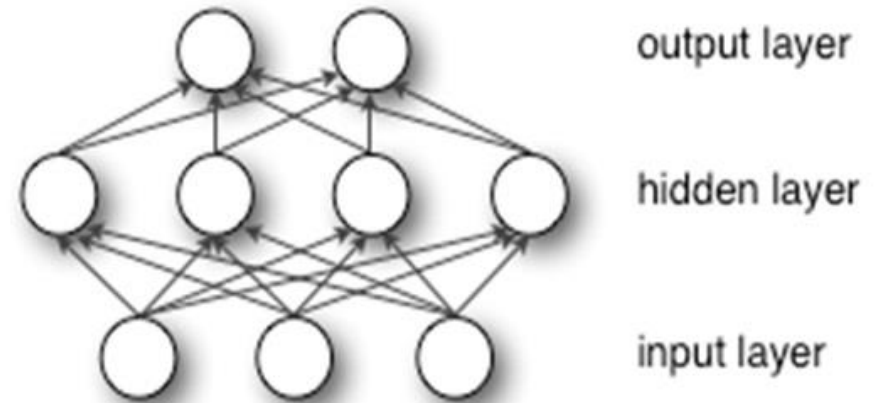
David Hubel's [Eye, Brain, and Vision](#)



# Hubel/Wiesel Architecture and Multi-layer Neural Network



Hubel and Wiesel's architecture



Multi-layer Neural Network  
- *A non-linear classifier*

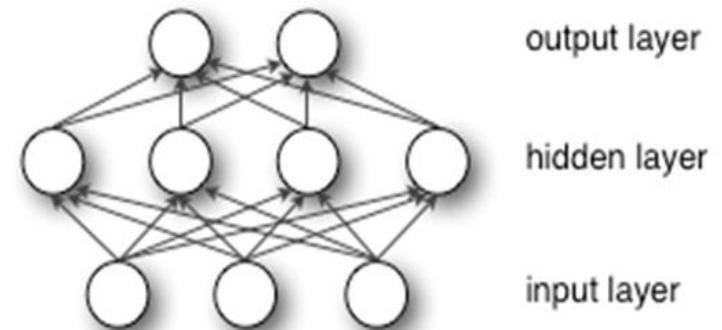


# Multi-layer Neural Network

- A non-linear classifier
- **Training:** find network weights  $\mathbf{w}$  to minimize the error between true training labels  $y_i$  and estimated labels  $f_{\mathbf{w}}(\mathbf{x}_i)$

$$E(\mathbf{w}) = \sum_{i=1}^N (y_i - f_{\mathbf{w}}(\mathbf{x}_i))^2$$

- Minimization can be done by gradient descent provided  $f$  is differentiable
- This training method is called **back-propagation**



# Convolutional Neural Networks (CNN, ConvNet, DCN)

- CNN = a multi-layer neural network with
  - **Local** connectivity
  - **Share** weight parameters across spatial positions
- One activation map (a depth slice), computed with one set of weights

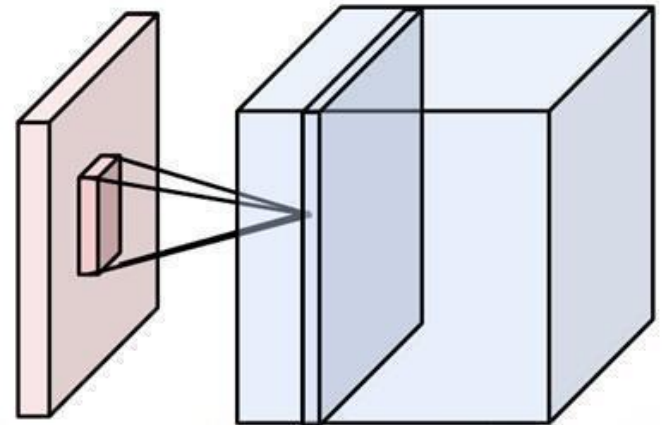


Image credit: A. Karpathy

# Neocognitron [Fukushima, Biological Cybernetics 1980]

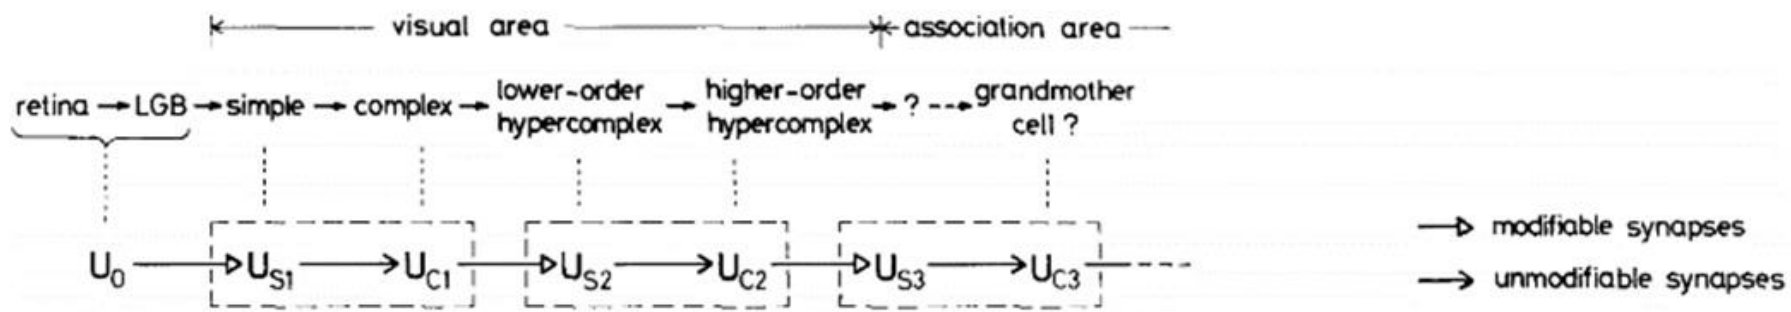
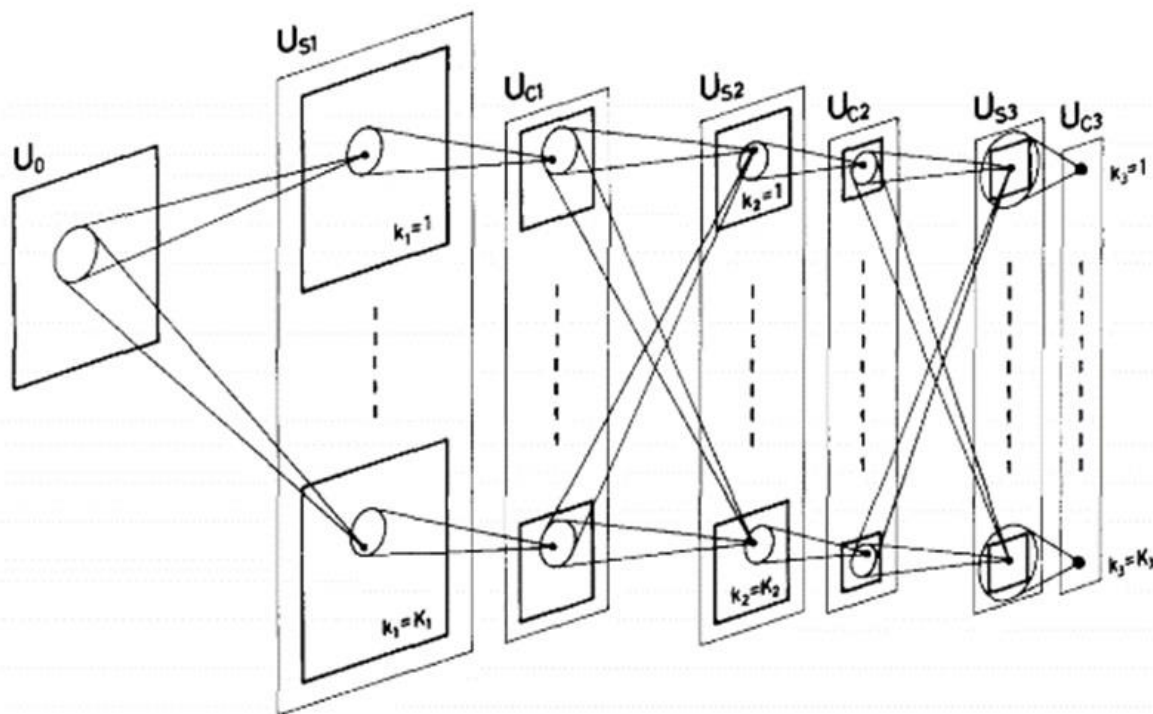


Fig. 1. Correspondence between the hierarchy model by Hubel and Wiesel, and the neural network of the neocognitron

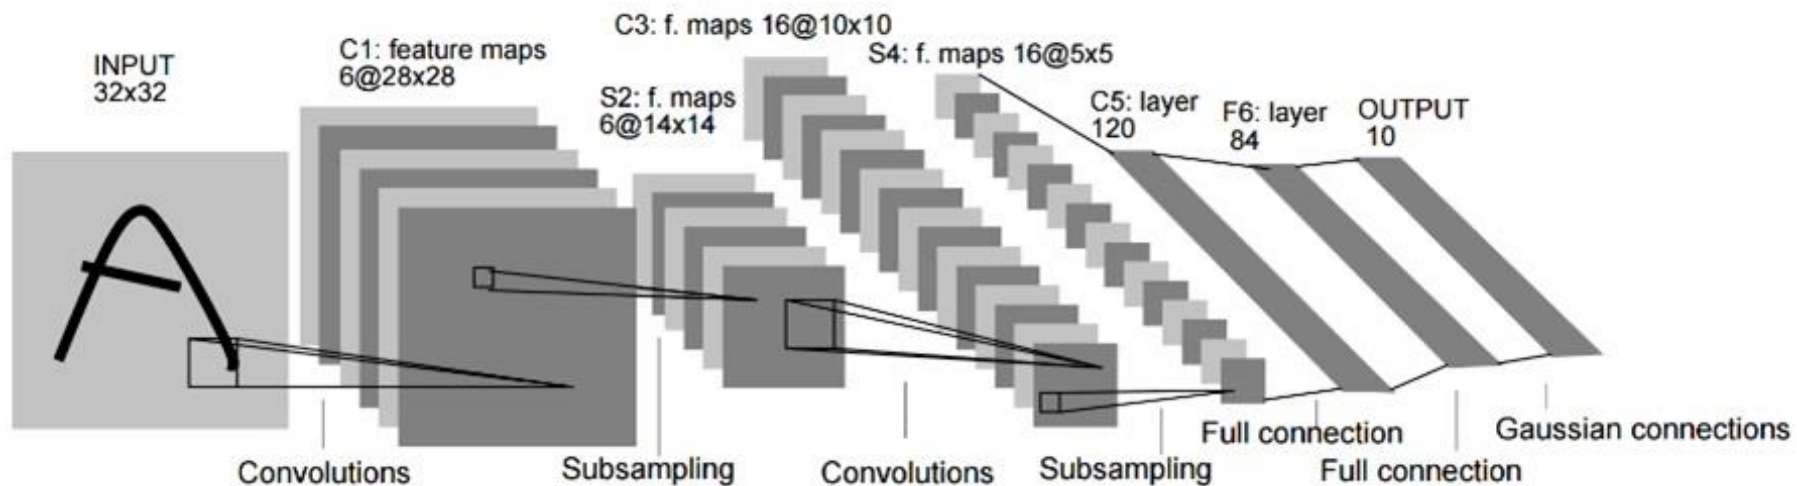


Deformation-Resistant Recognition

S-cells: (simple)  
- extract local features

C-cells: (complex)  
- allow for positional errors

# LeNet [LeCun et al. 1998]



Gradient-based learning applied to document recognition [[LeCun, Bottou, Bengio, Haffner 1998](#)]

LeNet-1 from 1993

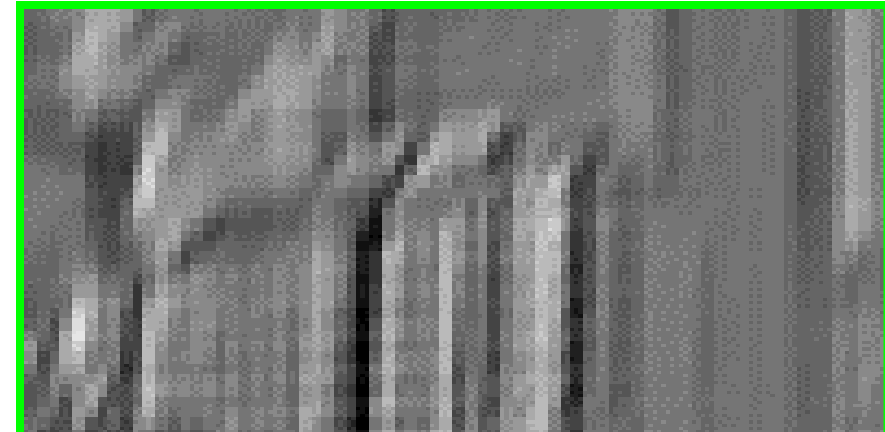


# What is a Convolution?

- Weighted moving sum



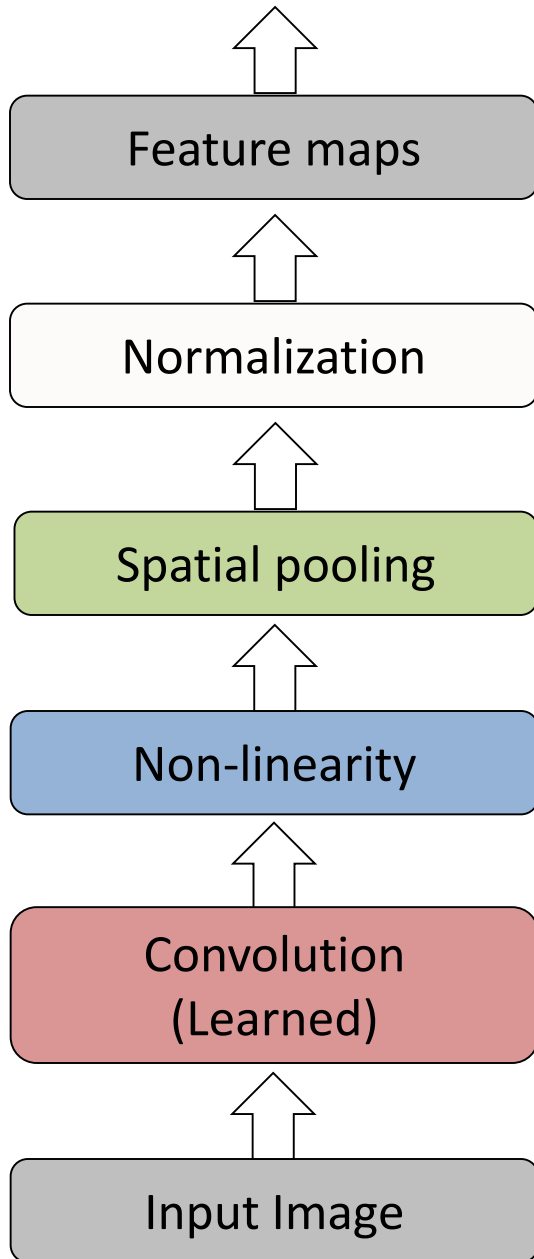
Input



Feature Activation Map

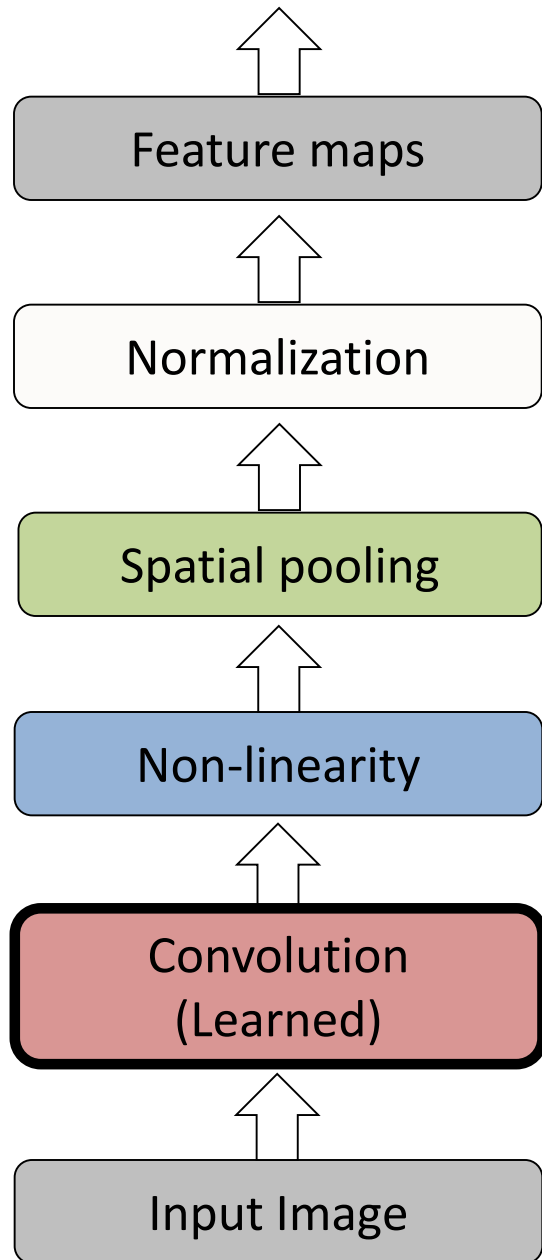
slide credit: S. Lazebnik

# Convolutional Neural Networks





# Convolutional Neural Networks

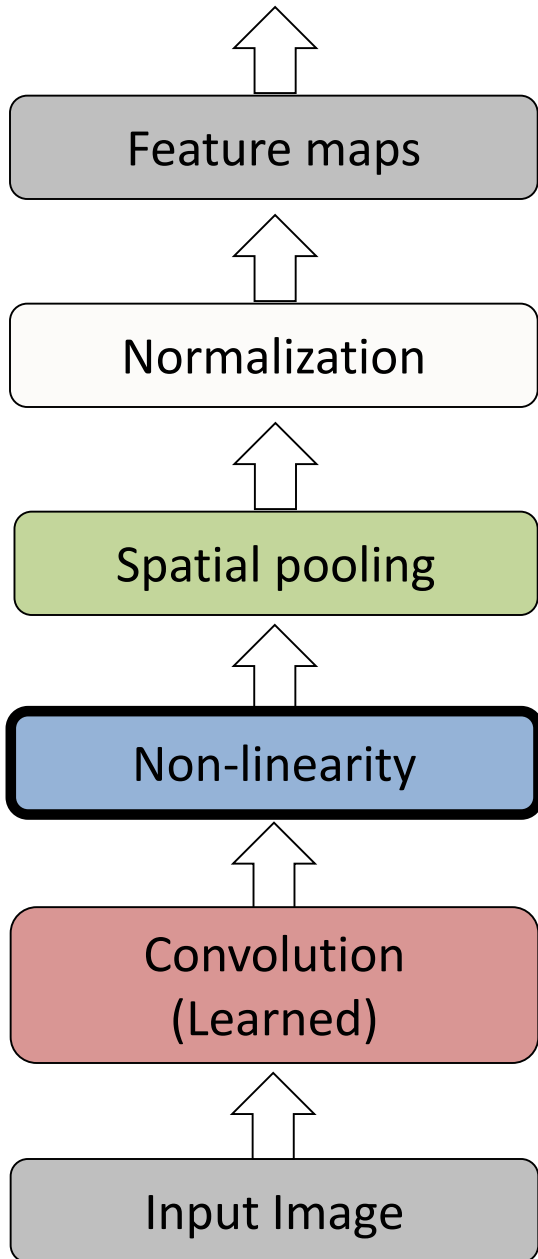


Input

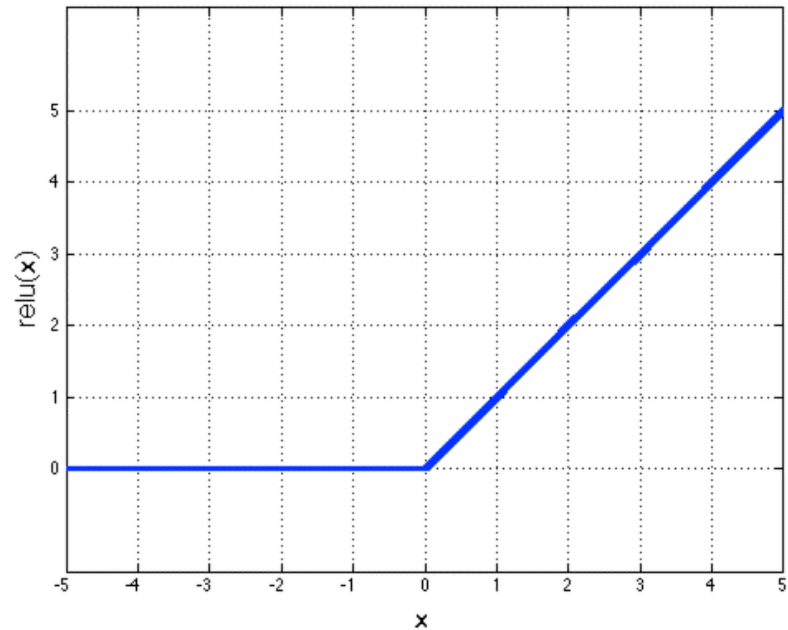


Feature Map

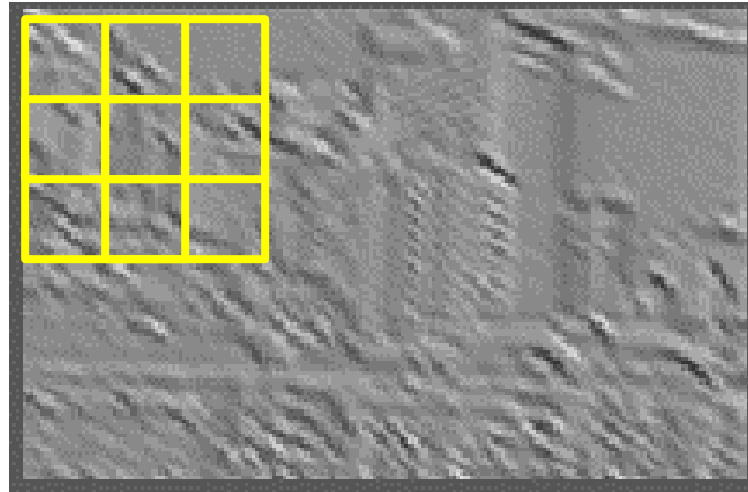
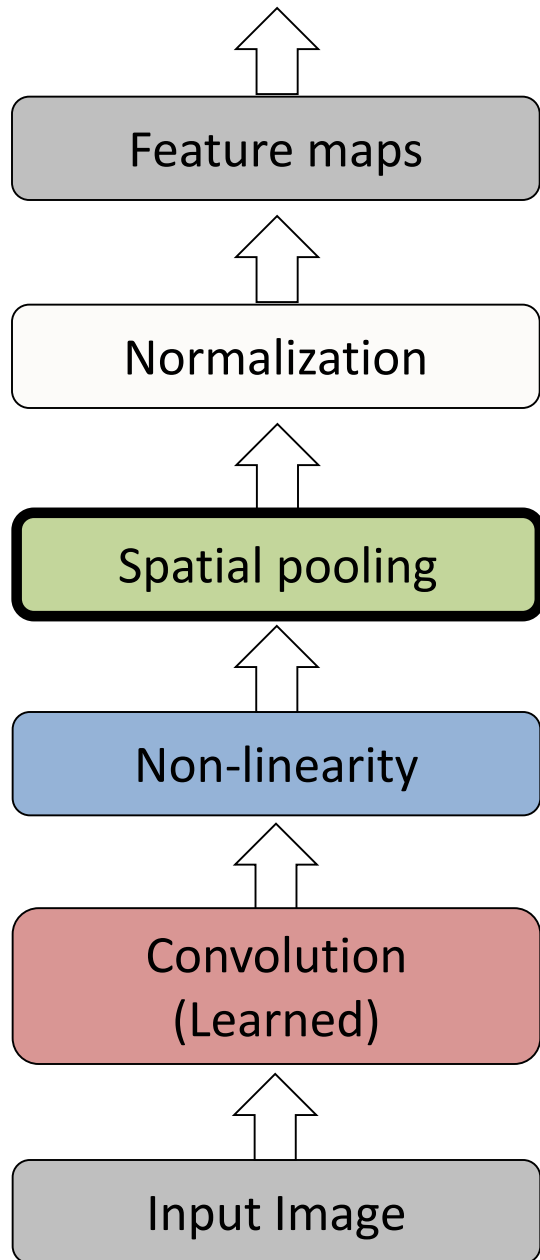
# Convolutional Neural Networks



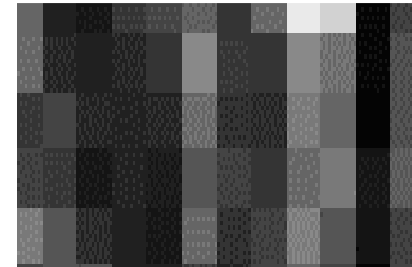
## Rectified Linear Unit (ReLU)



# Convolutional Neural Networks



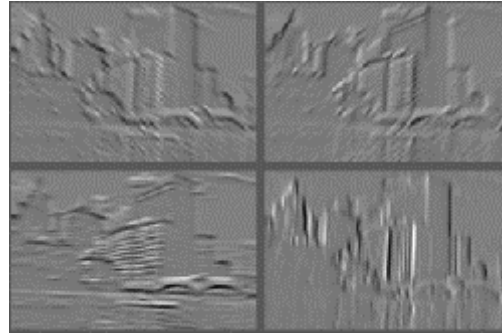
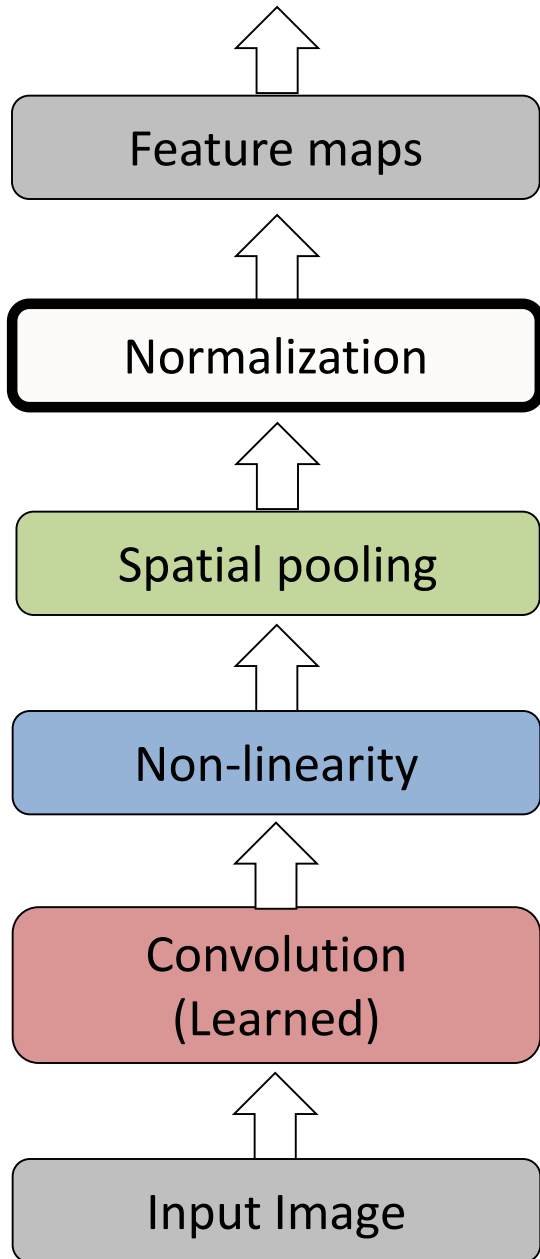
Max pooling



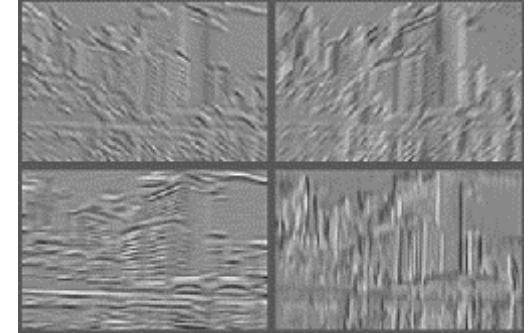
Max-pooling: a non-linear down-sampling

Provide *translation invariance*

# Convolutional Neural Networks

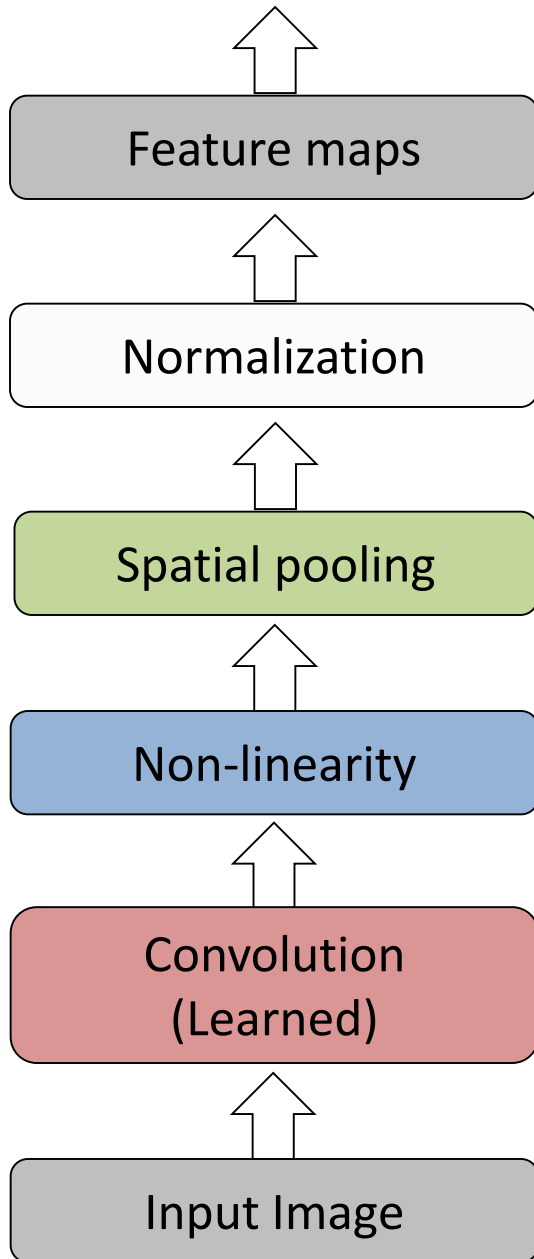


Feature Maps



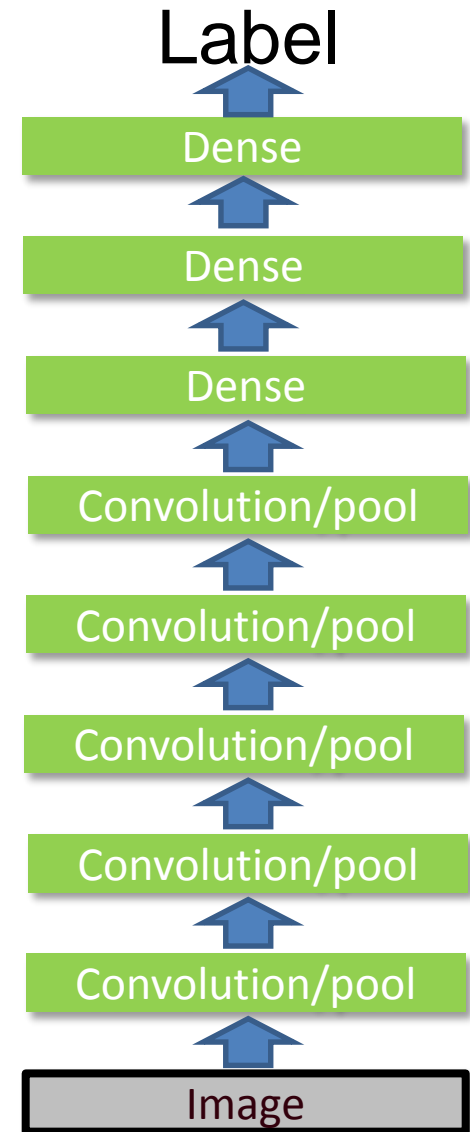
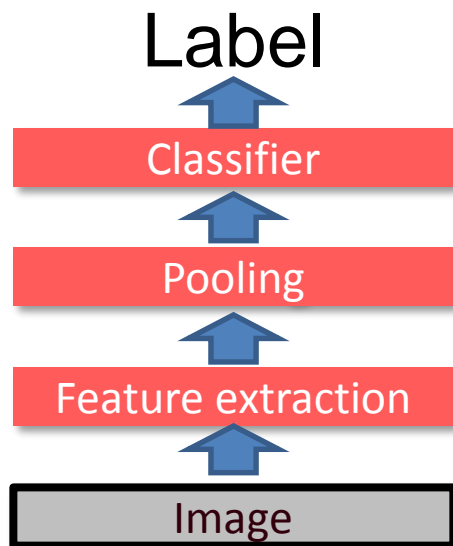
Feature Maps  
After Contrast  
Normalization

# Convolutional Neural Networks



# Engineered vs. learned features

Convolutional filters are trained in a supervised manner by back-propagating classification error

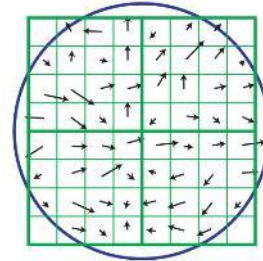
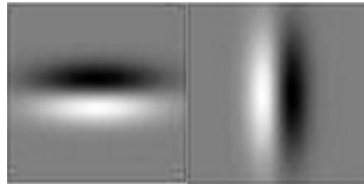


# SIFT Descriptor

Image  
Pixels

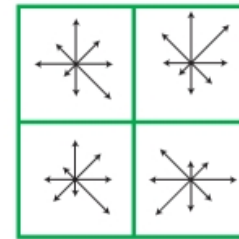
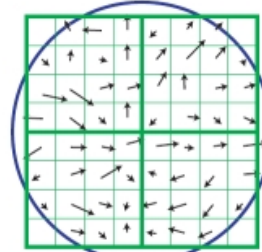


Apply gradient  
filters

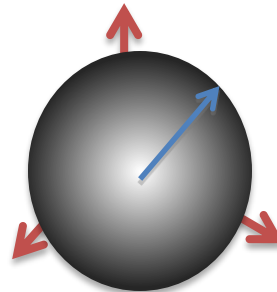


Lowe [IJCV 2004]

Spatial pool  
(Sum)



Normalize to unit  
length



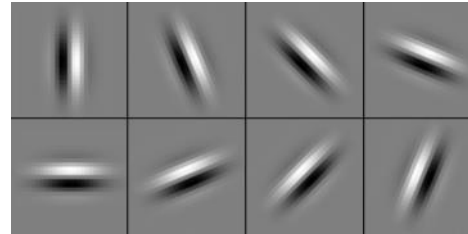
Feature  
Vector

# SIFT Descriptor

Image  
Pixels

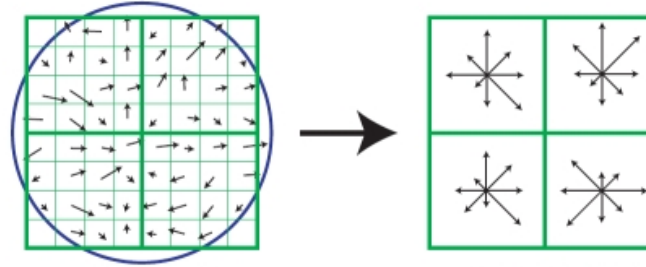


Apply  
oriented filters

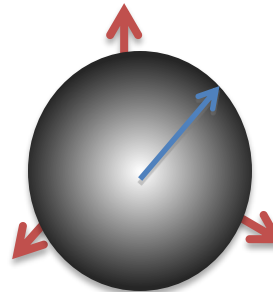


Lowé [IJCV 2004]

Spatial pool  
(Sum)



Normalize to unit  
length



Feature  
Vector



slide credit: R. Fergus



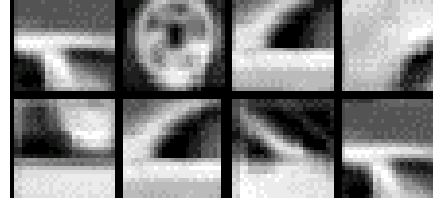
# Spatial Pyramid Matching

Lazebnik,  
Schmid,  
Ponce  
[CVPR 2006]

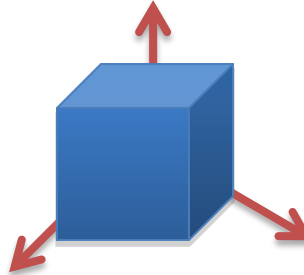
SIFT  
Features



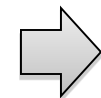
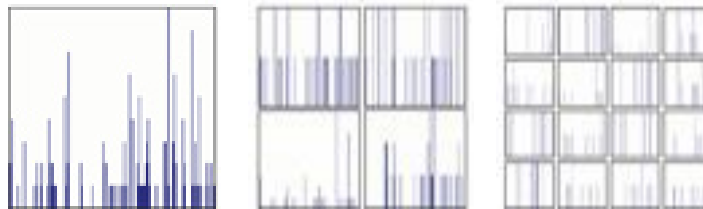
Filter with  
Visual Words



Max



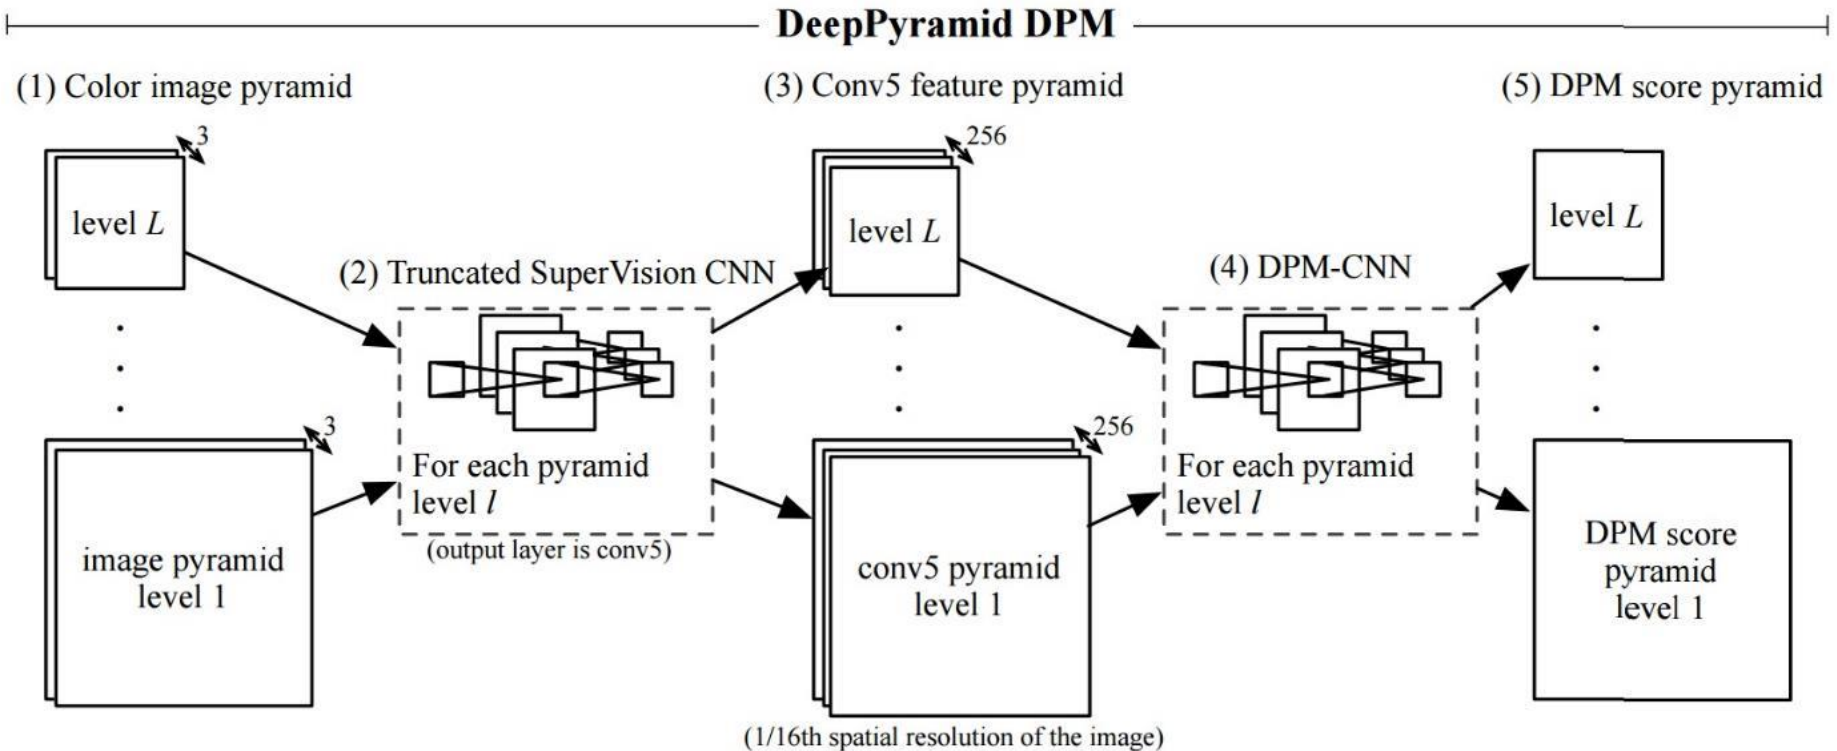
Multi-scale  
spatial pool  
(Sum)



Classifier

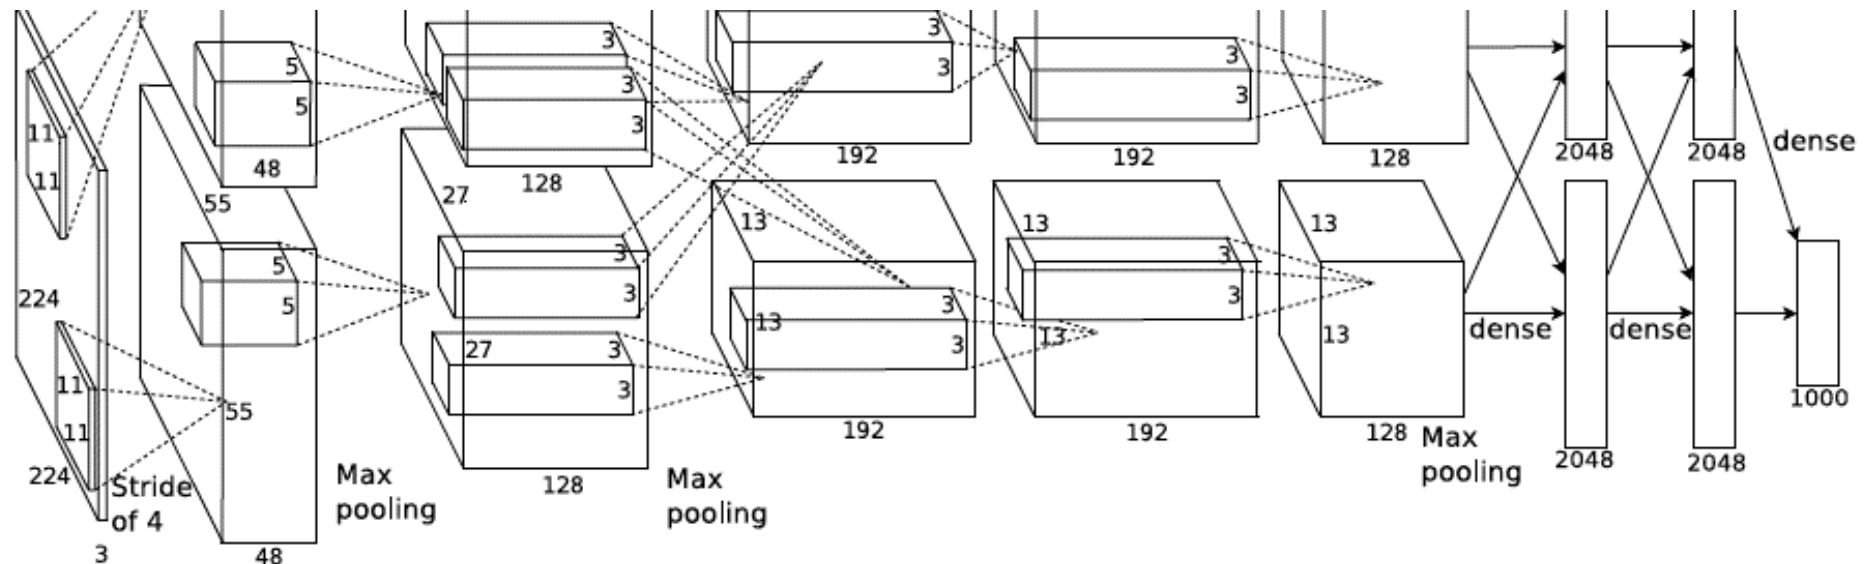
slide credit: R. Fergus

# Deformable Part Model



# AlexNet

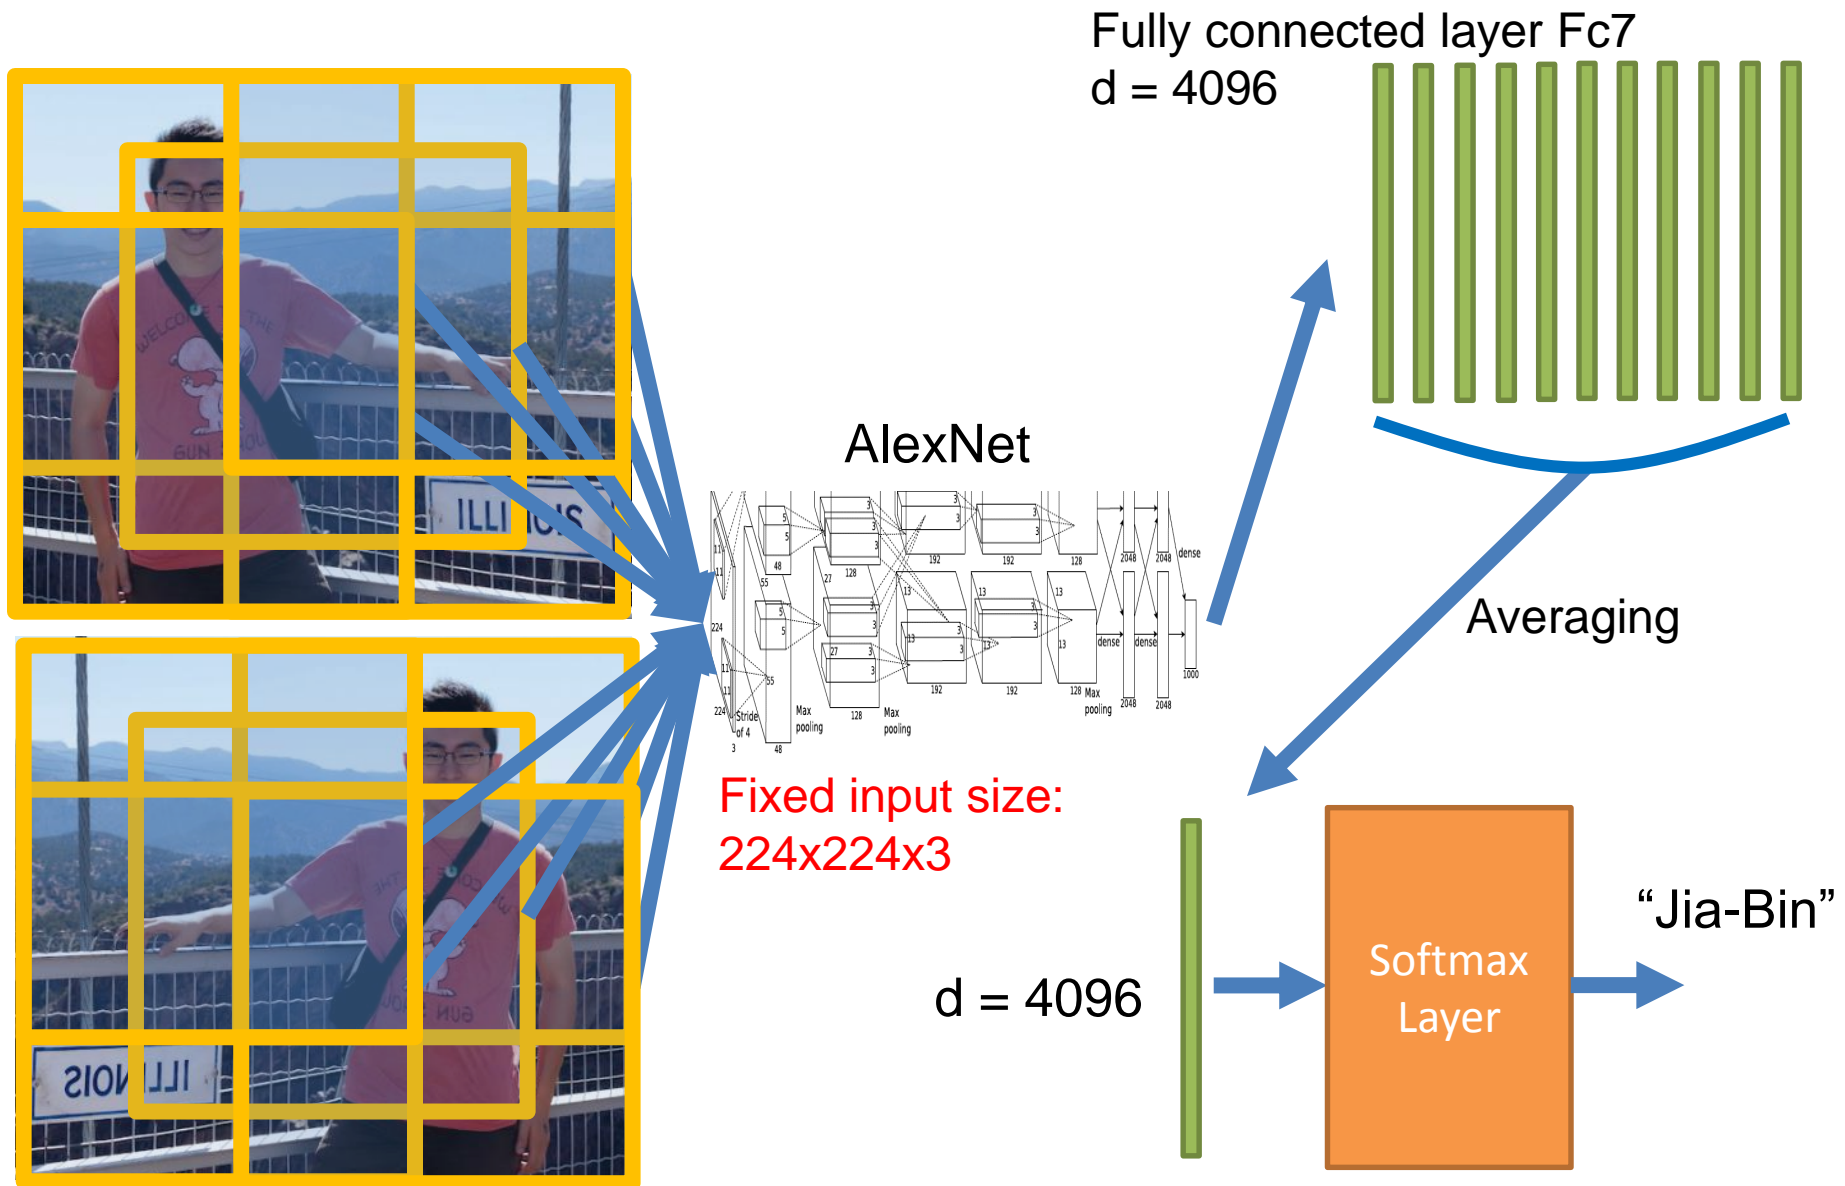
- Similar framework to LeCun'98 but:
  - Bigger model (7 hidden layers, 650,000 units, 60,000,000 params)
  - More data ( $10^6$  vs.  $10^3$  images)
  - GPU implementation (50x speedup over CPU)
    - Trained on two GPUs for a week



A. Krizhevsky, I. Sutskever, and G. Hinton,

[ImageNet Classification with Deep Convolutional Neural Networks](#), NIPS 2012

# Using CNN for Image Classification



# ImageNet Challenge 2012-2014

Best non-convnet in 2012: 26.2%

| Team                             | Year | Place | Error (top-5) | External data |
|----------------------------------|------|-------|---------------|---------------|
| SuperVision – Toronto (7 layers) | 2012 | -     | 16.4%         | no            |
| SuperVision                      | 2012 | 1st   | 15.3%         | ImageNet 22k  |
| Clarifai – NYU (7 layers)        | 2013 | -     | 11.7%         | no            |
| Clarifai                         | 2013 | 1st   | 11.2%         | ImageNet 22k  |
| VGG – Oxford (16 layers)         | 2014 | 2nd   | 7.32%         | no            |
| GoogLeNet (19 layers)            | 2014 | 1st   | 6.67%         | no            |
| <u>Human expert</u> *            |      |       | 5.1%          |               |

| Team                           | Method                                 | Error (top-5) |
|--------------------------------|--|---------------|
| DeepImage - Baidu              | Data augmentation + multi GPU          | 5.33%         |
| PReLU-nets - MSRA              | Parametric ReLU + smart initialization | 4.94%         |
| BN-Inception ensemble - Google | Reducing internal covariate shift      | 4.82%         |

# Beyond classification

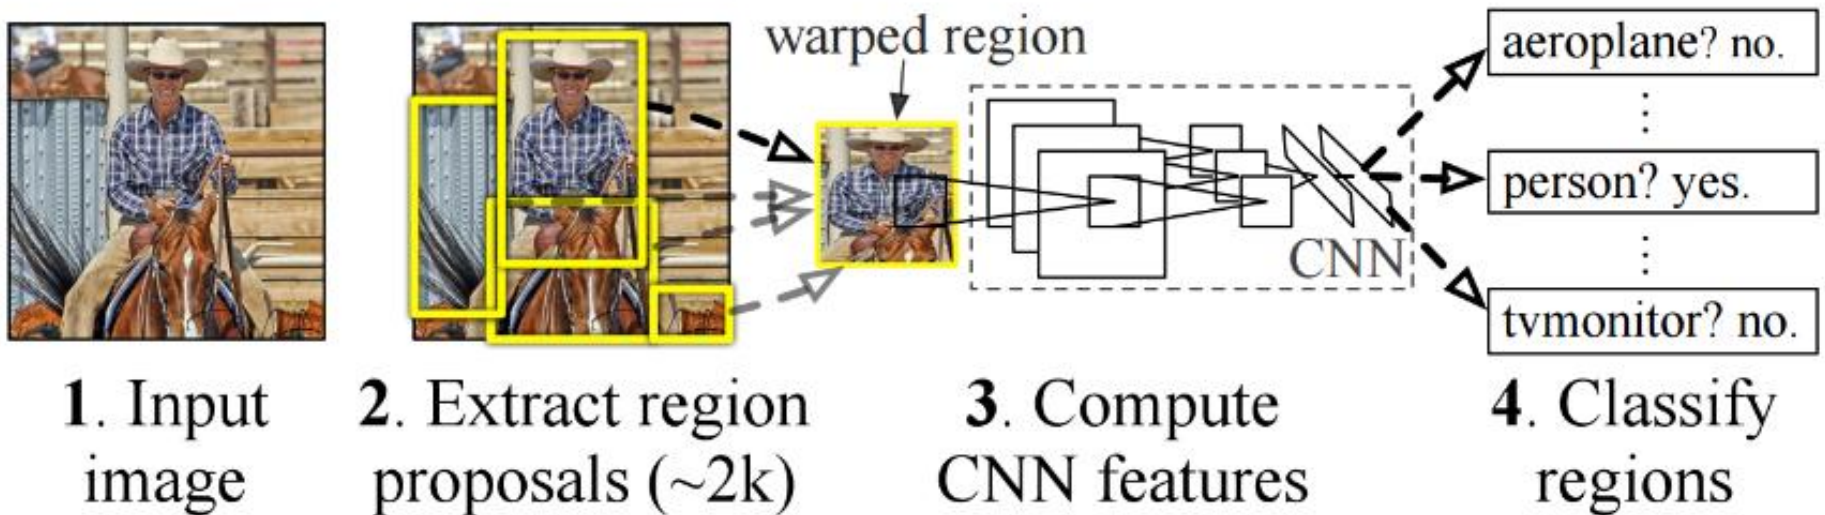
- Detection
- Segmentation
- Regression
- Pose estimation
- Matching patches
- Synthesis

and many more...

# R-CNN: Regions with CNN features

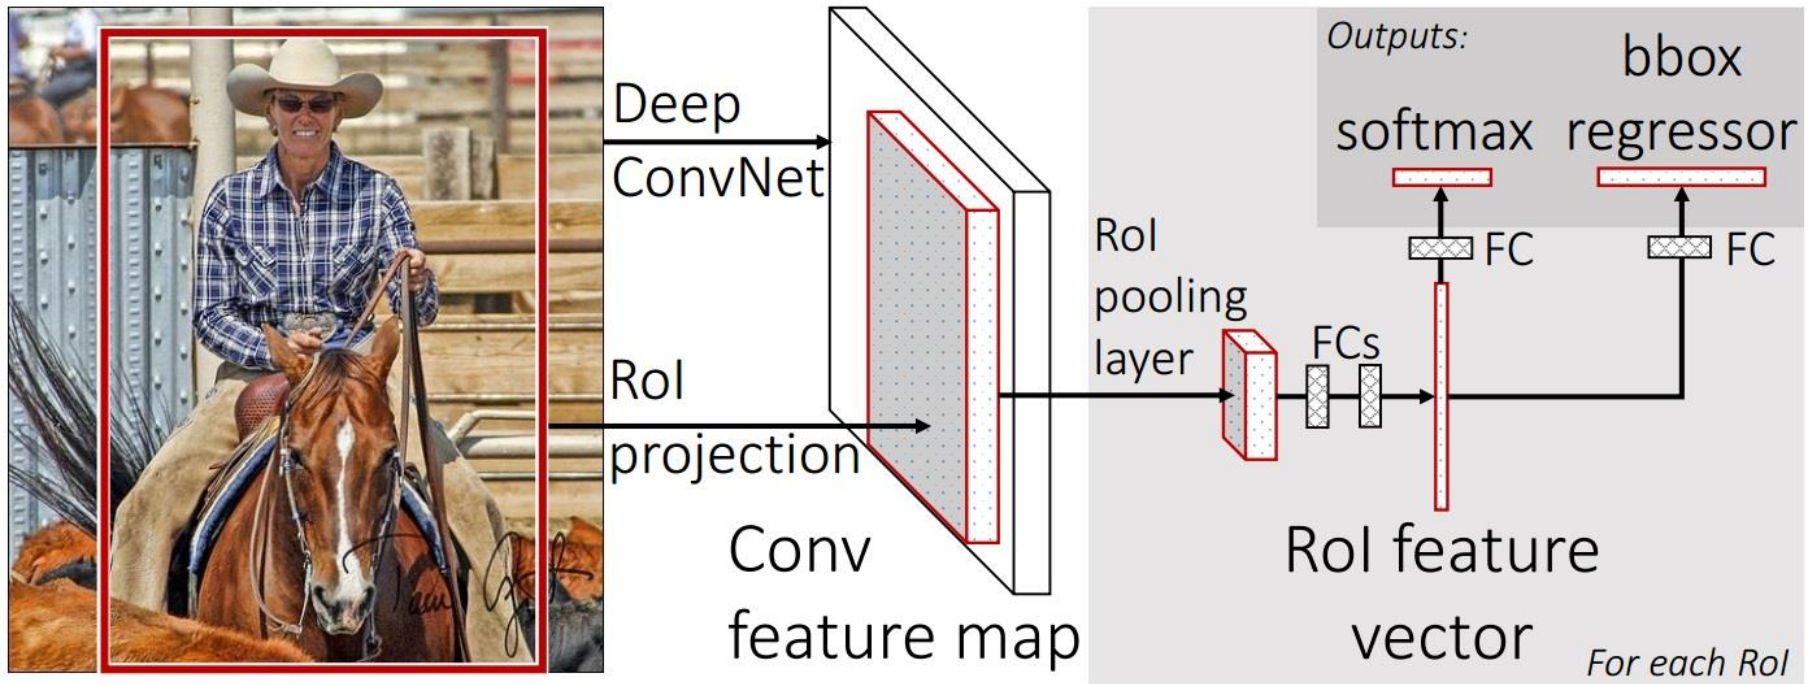
- Trained on ImageNet classification
- Finetune CNN on PASCAL

## R-CNN: *Regions with CNN features*





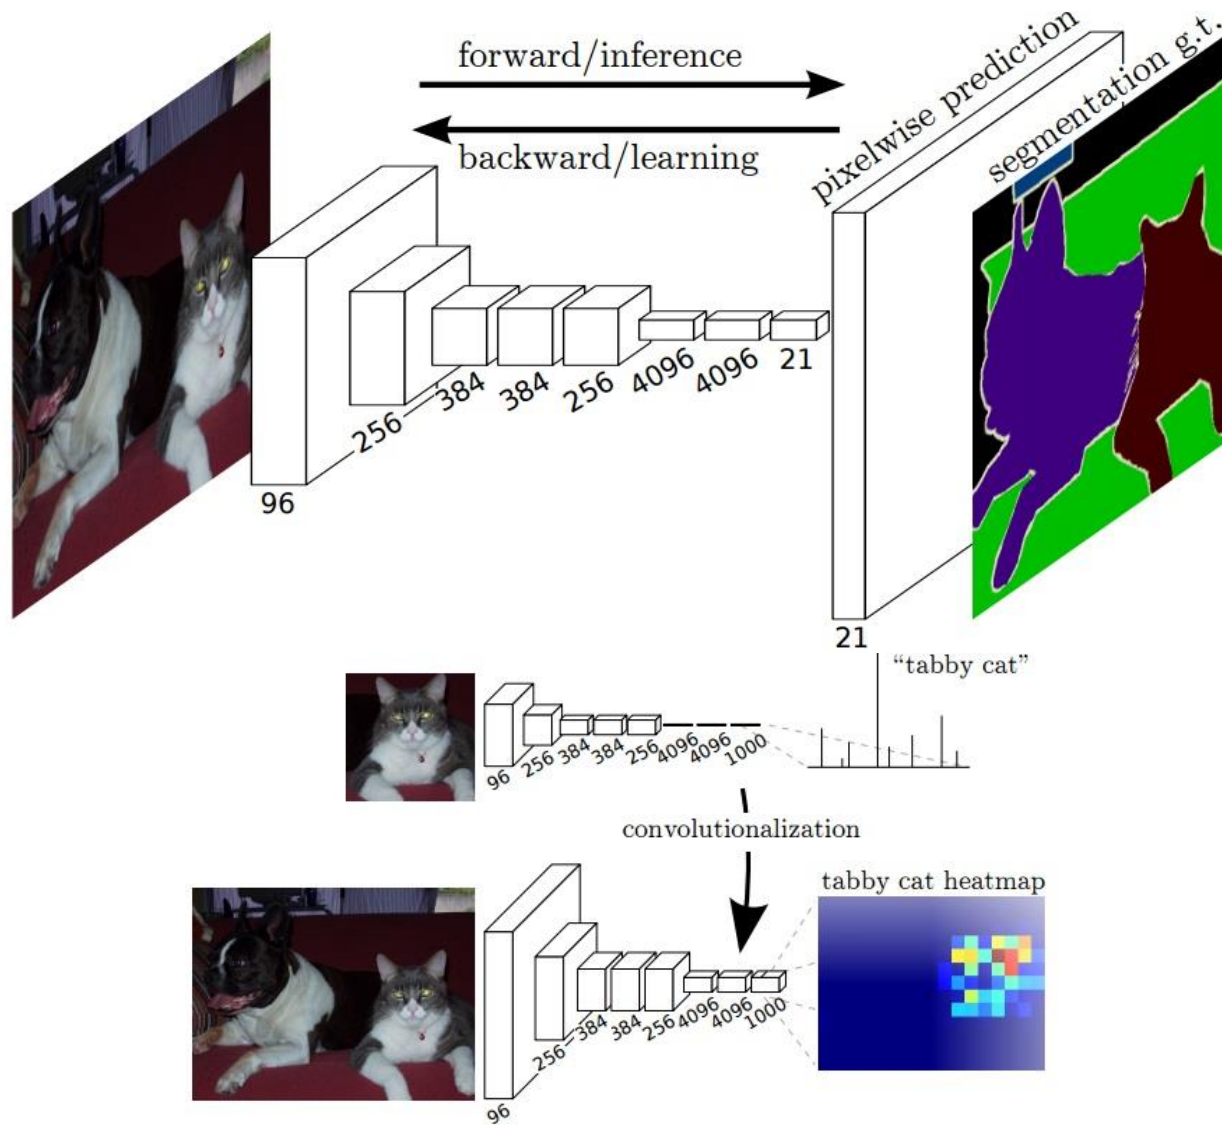
# Fast R-CNN



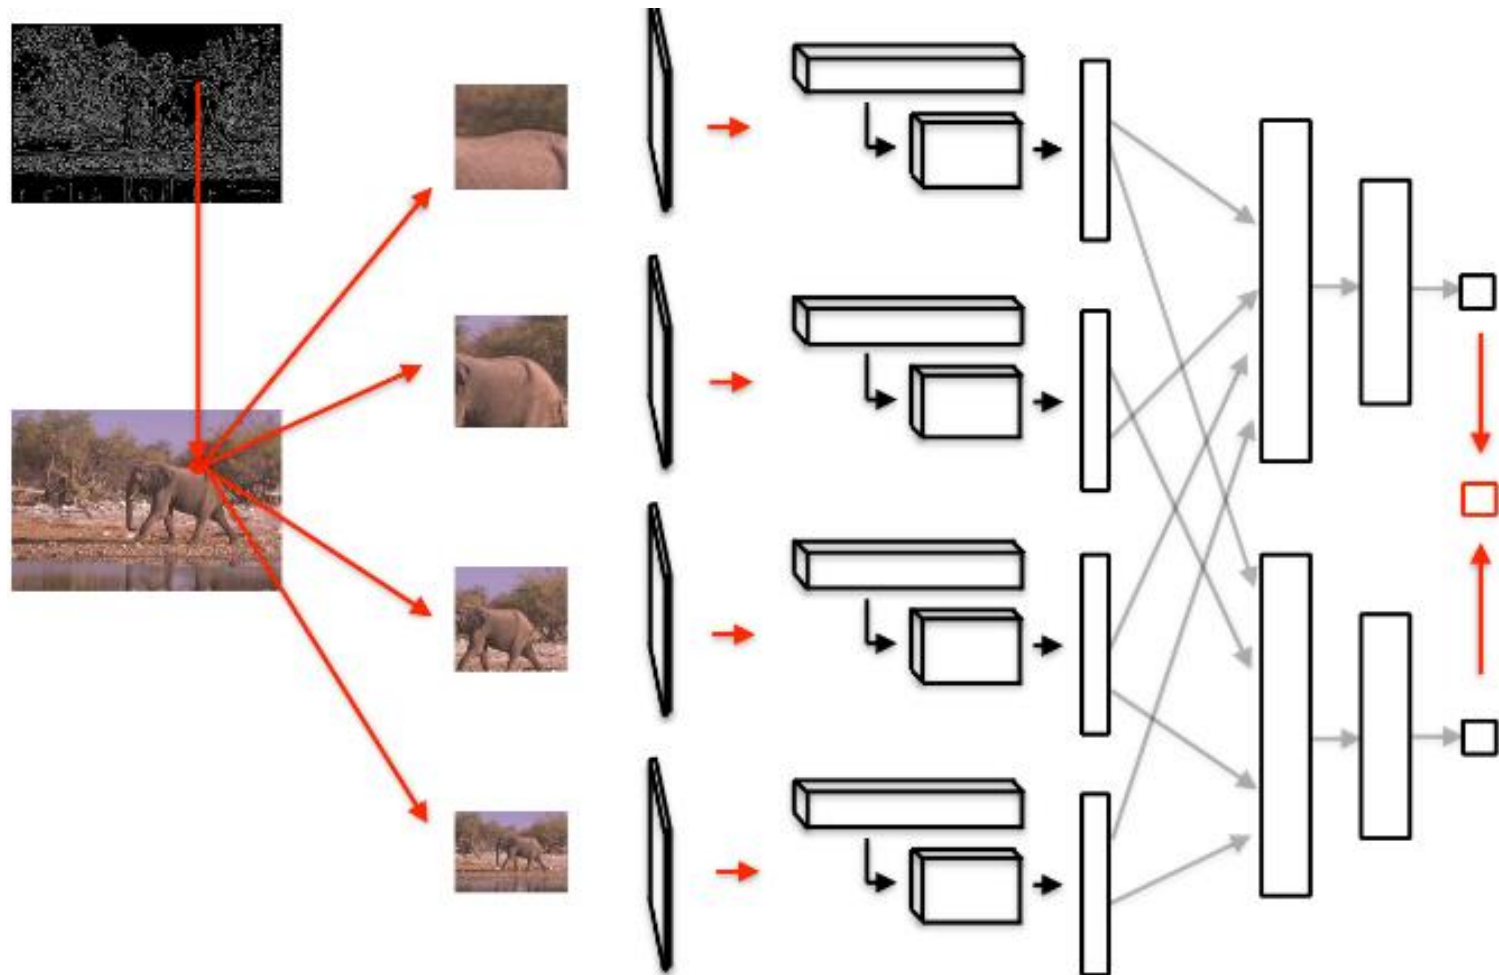
Fast RCNN [[Girshick, R 2015](https://arxiv.org/abs/1504.08083)]  
<https://github.com/rbgirshick/fast-rcnn>



# Labeling Pixels: Semantic Labels



# Labeling Pixels: Edge Detection

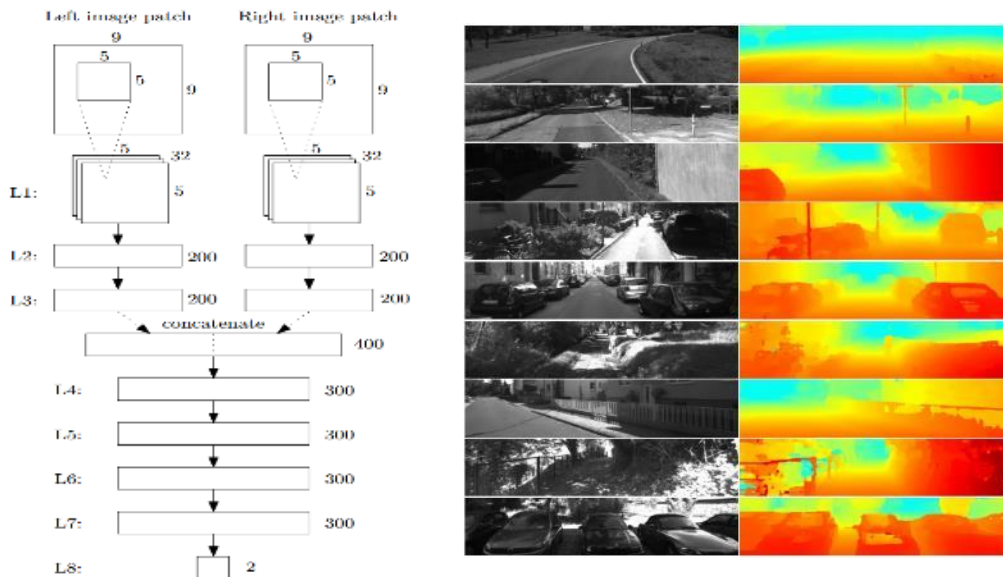


DeepEdge: A Multi-Scale Bifurcated Deep Network for Top-Down Contour Detection  
[Bertasius et al. CVPR 2015]

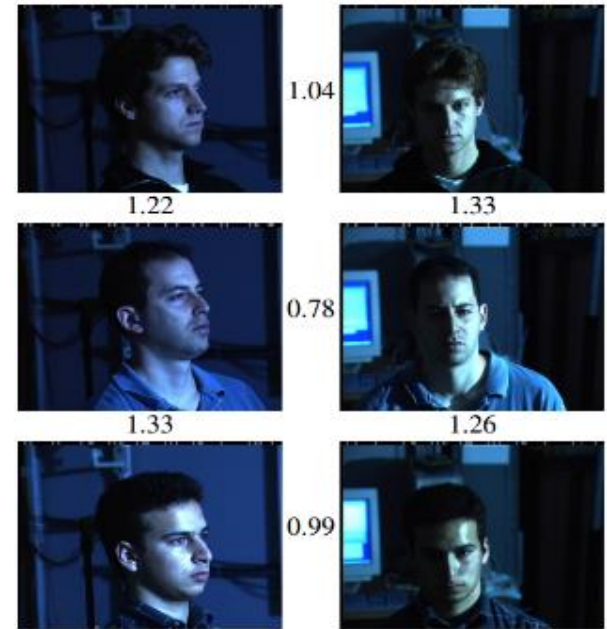




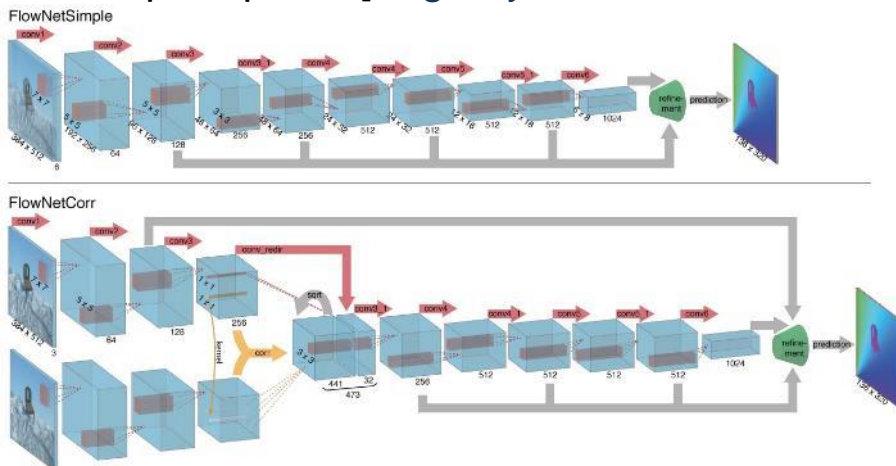
# CNN as a Similarity Measure for Matching



Stereo matching [Zbontar and LeCun CVPR 2015]  
 Compare patch [Zagoruyko and Komodakis 2015]



FaceNet [Schroff et al. 2015]

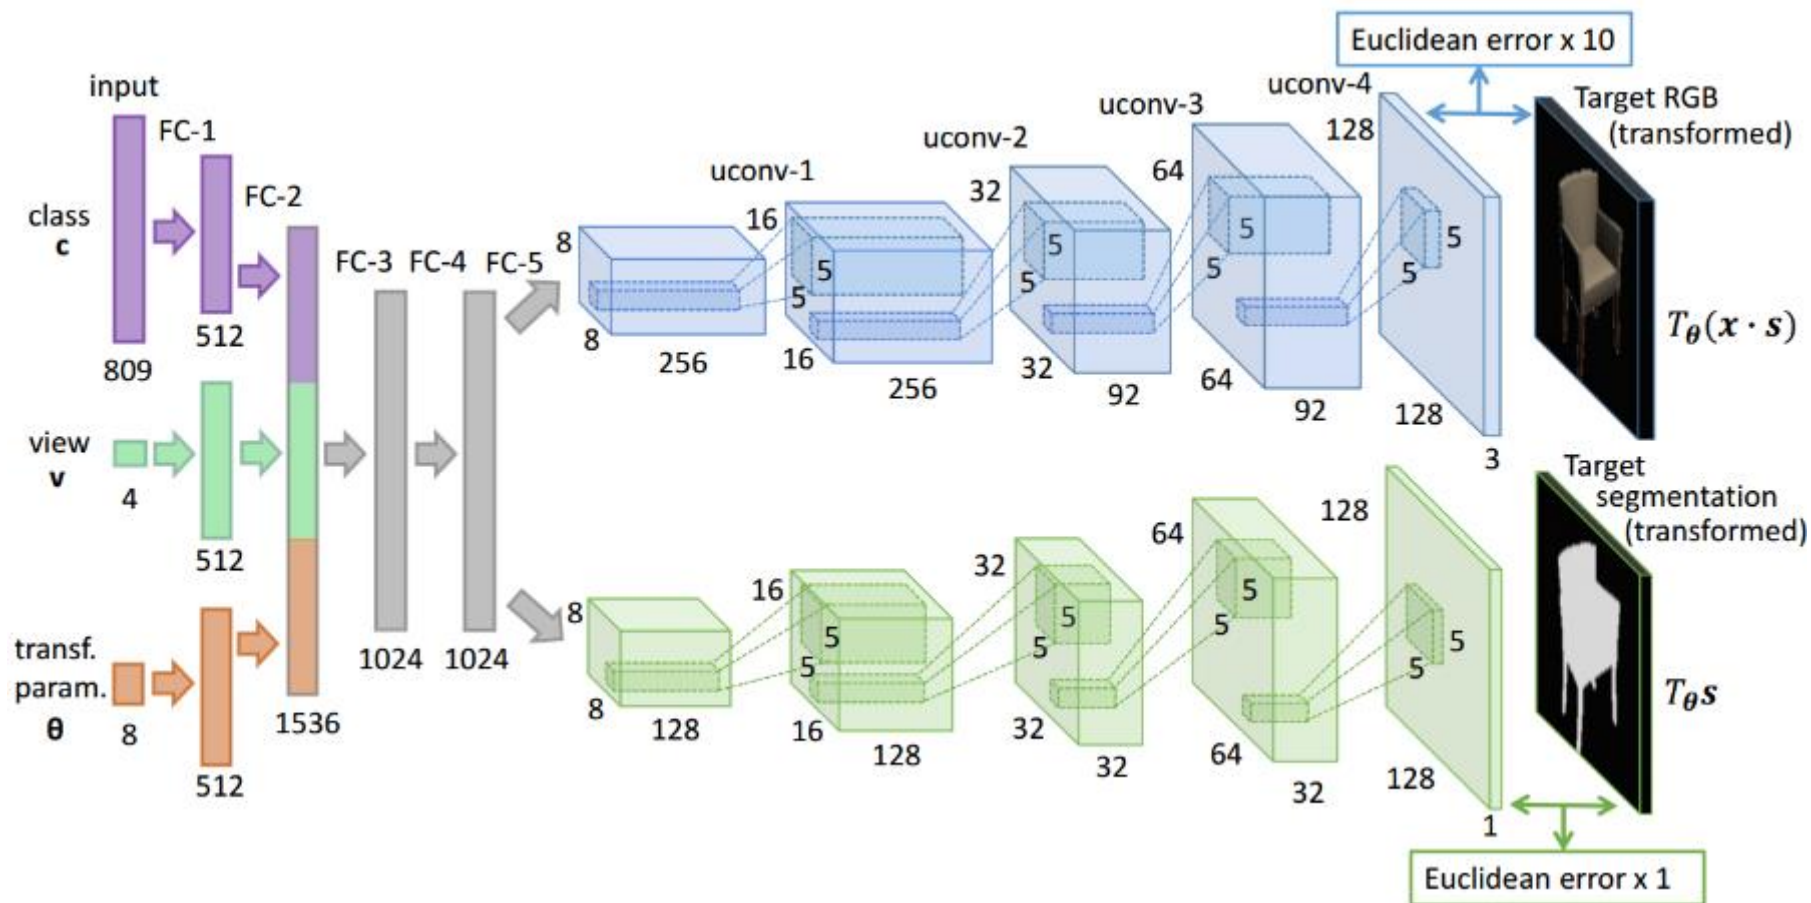


FlowNet [Fischer et al 2015]



Match ground and aerial images  
 [Lin et al. CVPR 2015]

# CNN for Image Generation

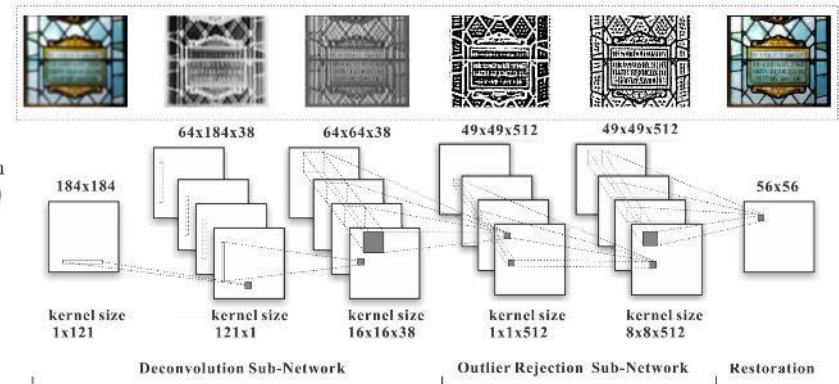
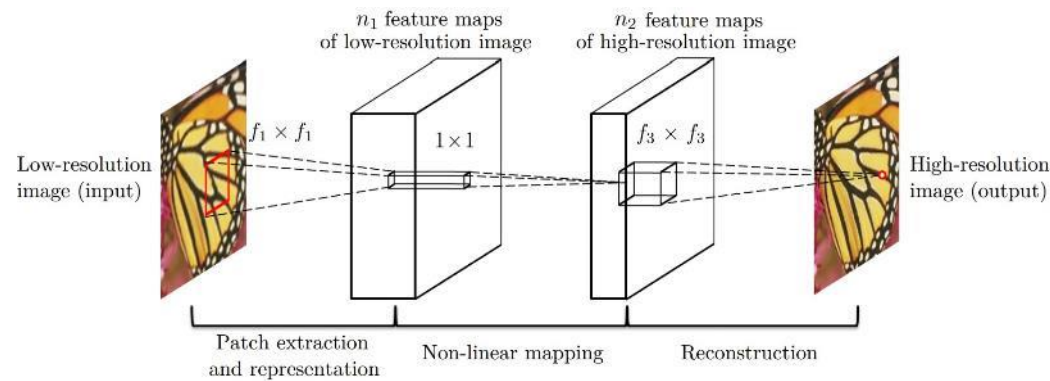


# Chair Morphing

1

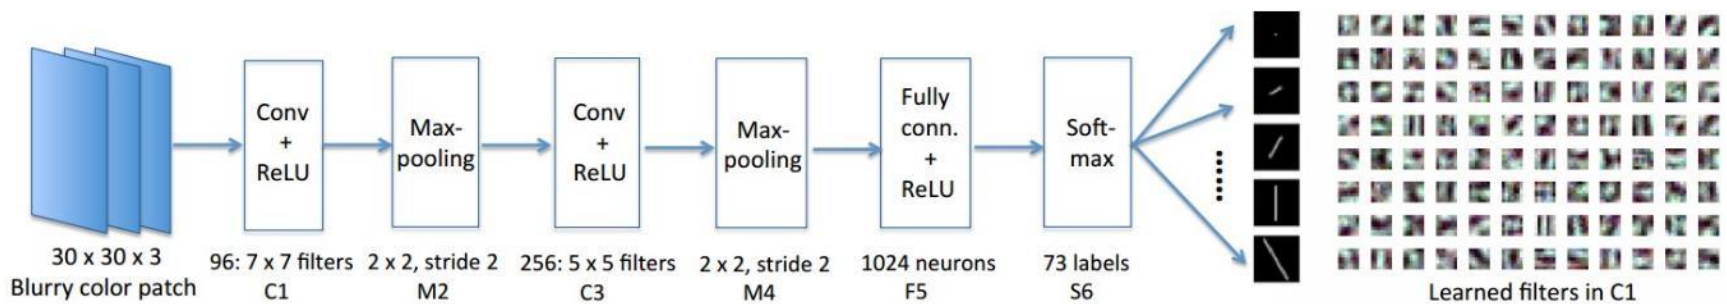


# CNN for Image Restoration/Enhancement



Super-resolution  
[Dong et al. ECCV 2014]

Non-blind deconvolution  
[Xu et al. NIPS 2014]



Non-uniform blur estimation  
[Sun et al. CVPR 2015]



# Take a break...

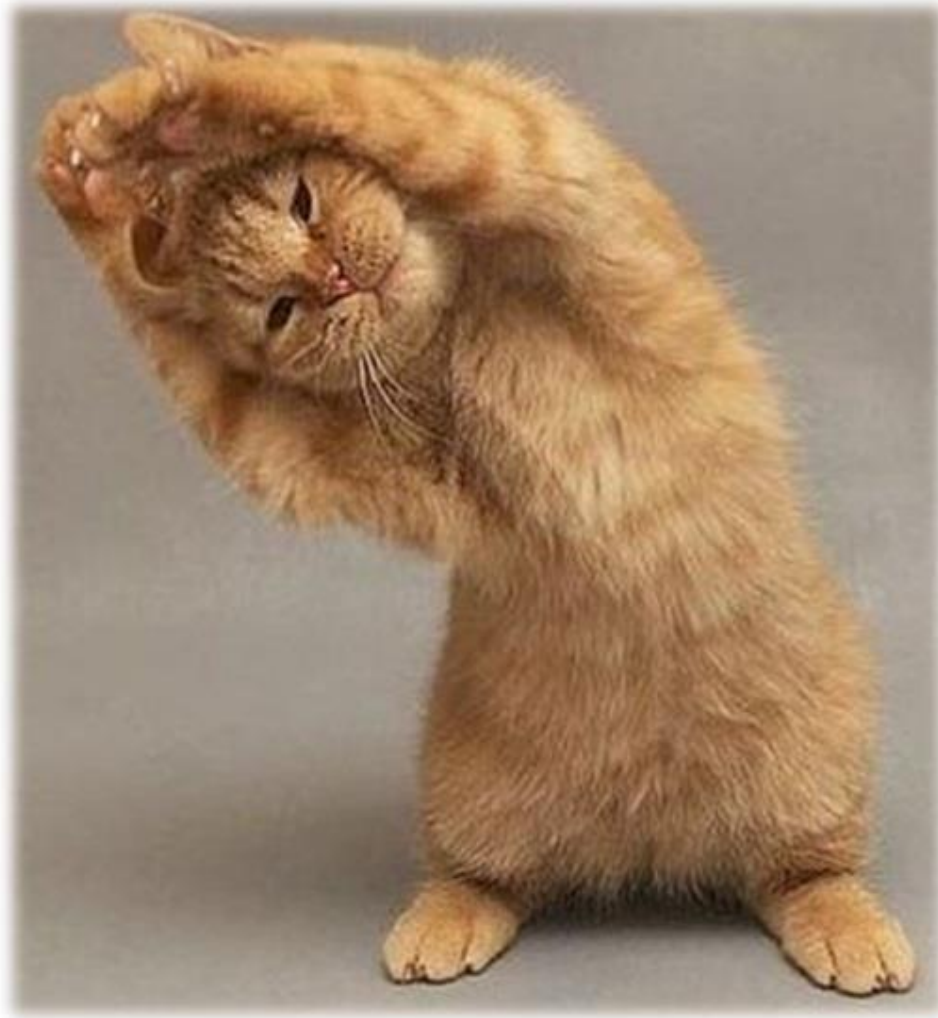


Image source: <http://mehimandthecats.com/feline-care-guide/>



# Understanding and Visualizing CNN

- Find images that maximize some class scores
- Individual neuron activation
- Visualize input pattern using deconvnet
- Invert CNN features
- Breaking CNNs

# Find images that maximize some class scores



dumbbell



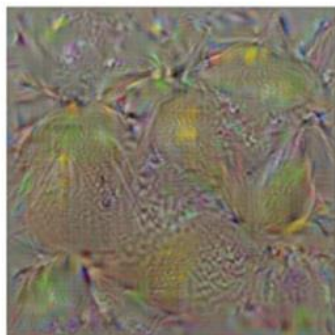
cup



dalmatian



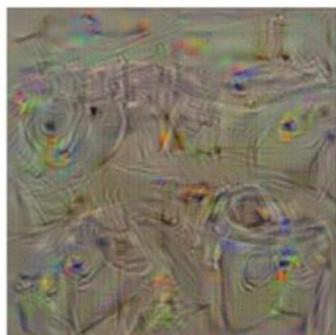
bell pepper



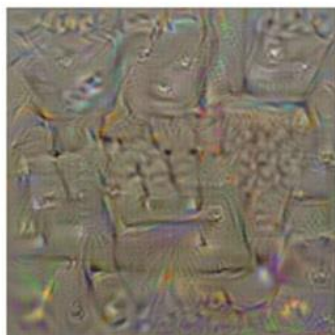
lemon



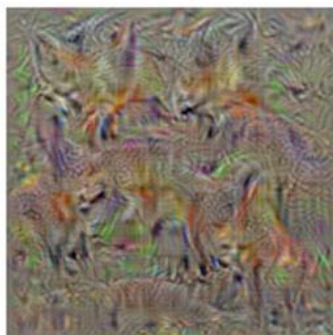
husky



washing machine



computer keyboard



kit fox



person: HOG template

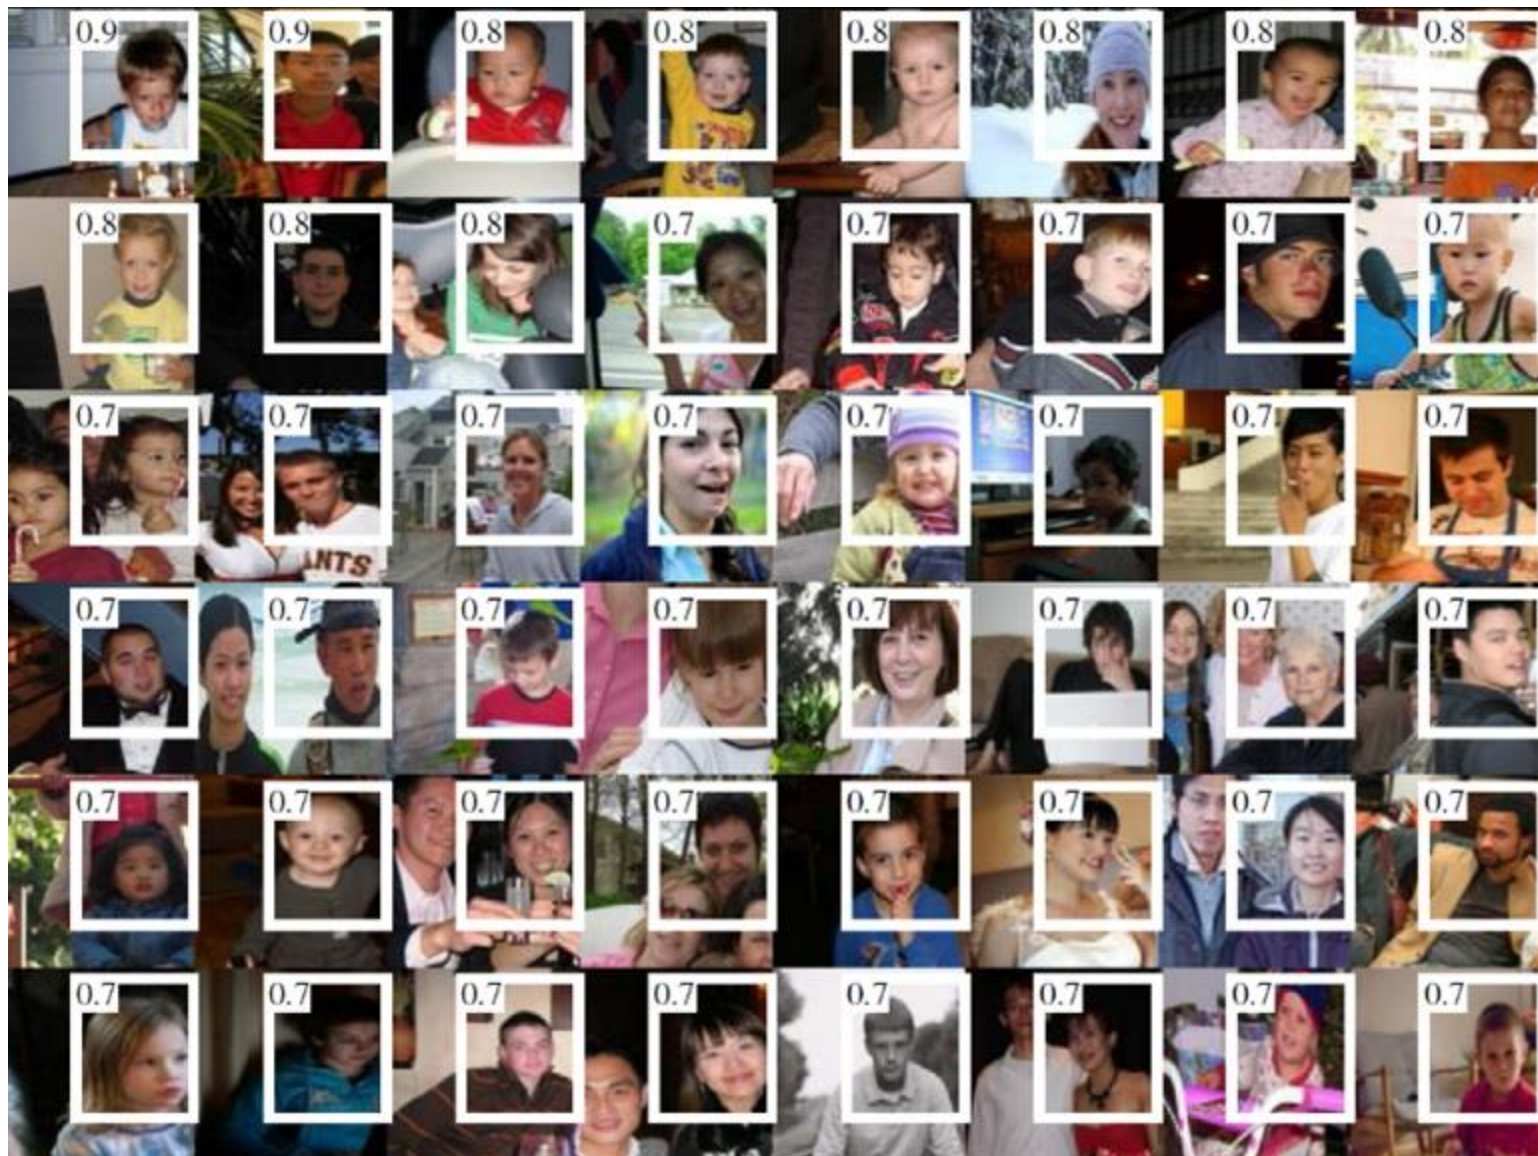
# Individual Neuron Activation



RCNN [Girshick et al. CVPR 2014]

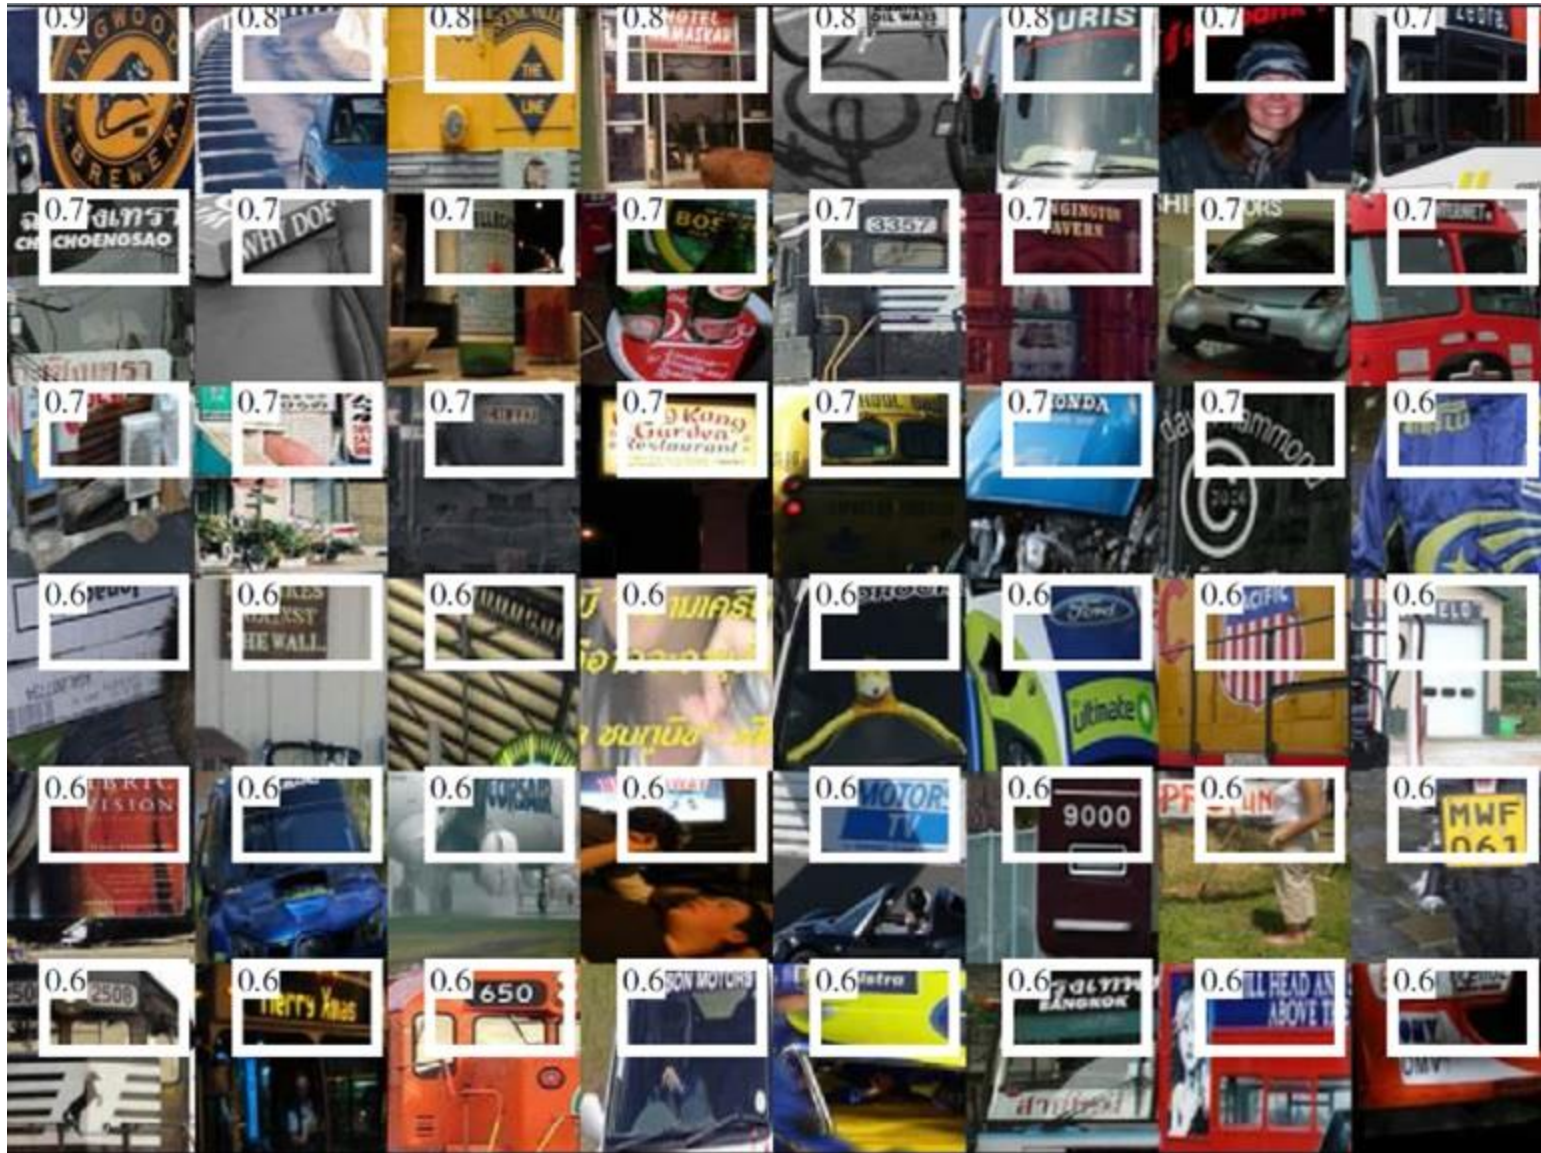


# Individual Neuron Activation



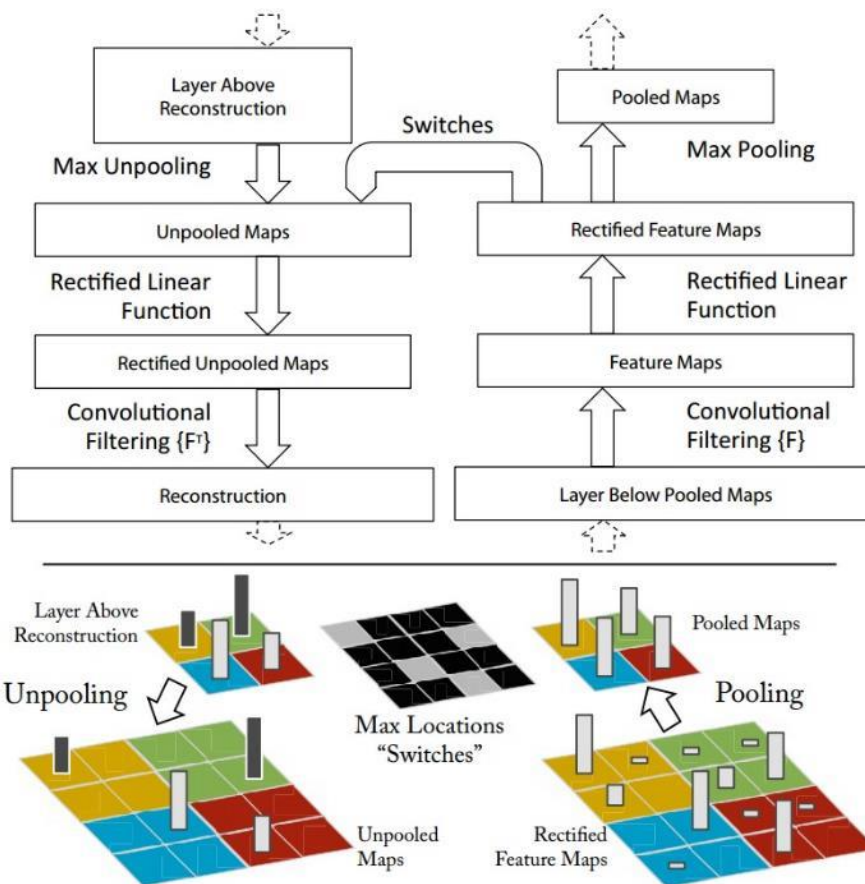


# Individual Neuron Activation



# Map activation back to the input pixel space

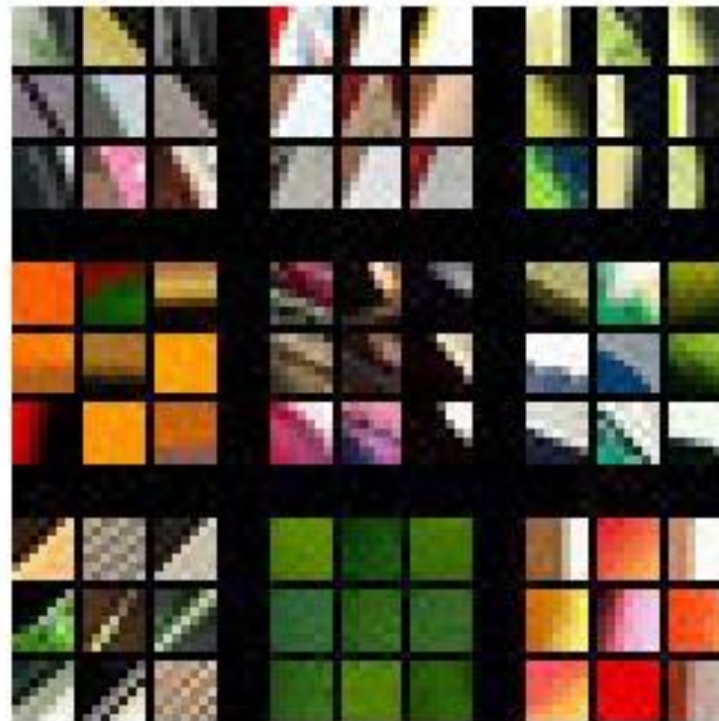
- What input pattern originally caused a given activation in the feature maps?



# Layer 1

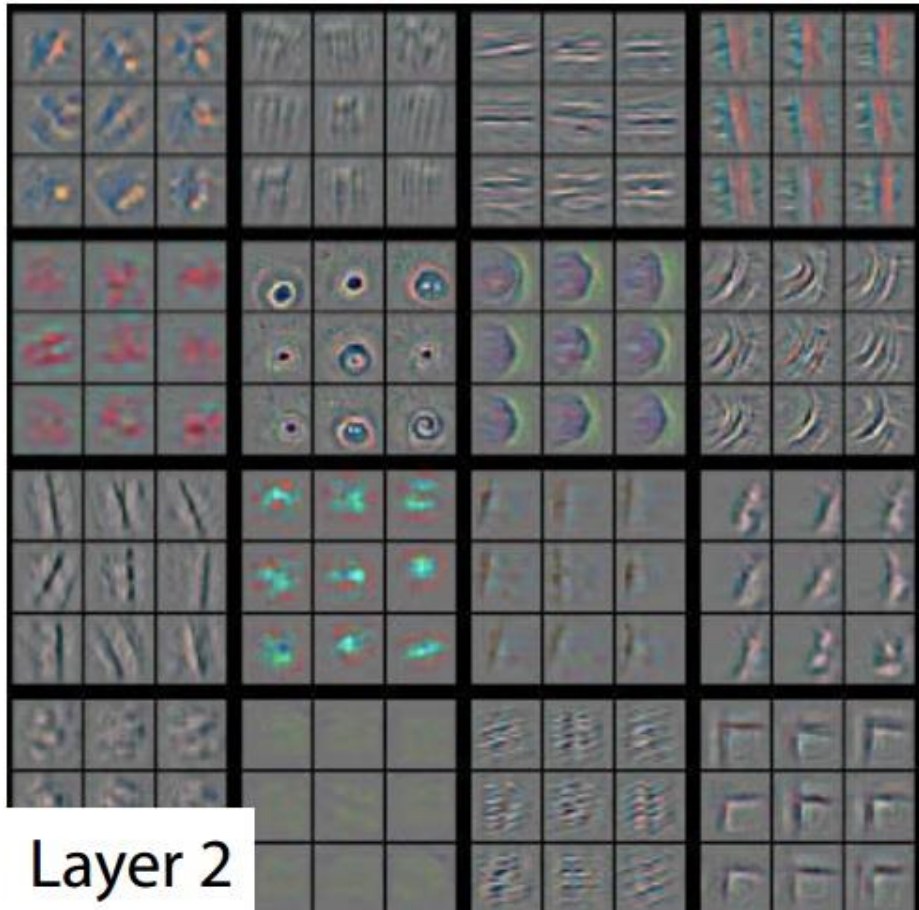


Layer 1



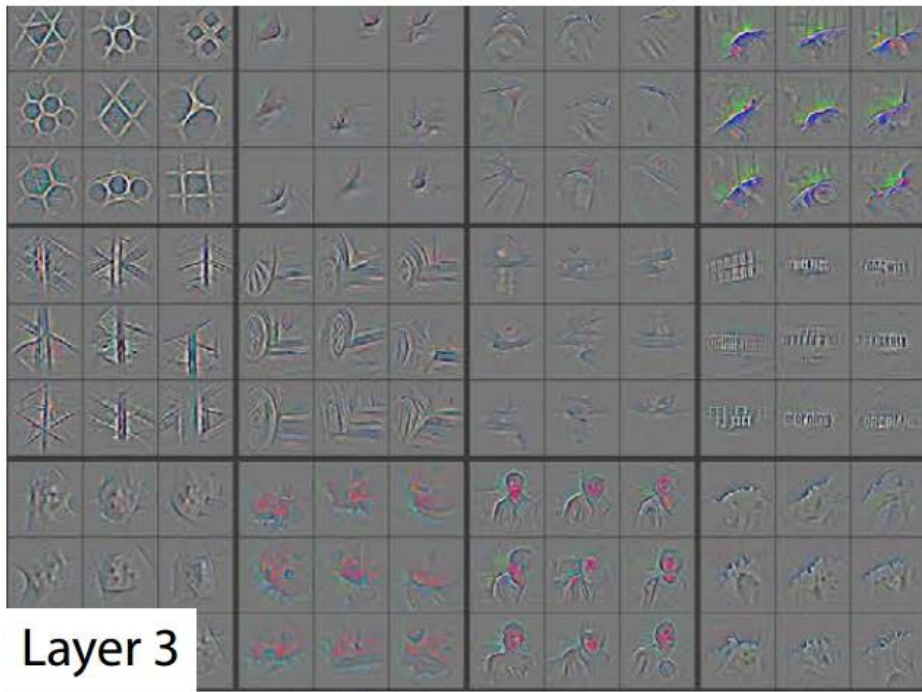


# Layer 2





# Layer 3

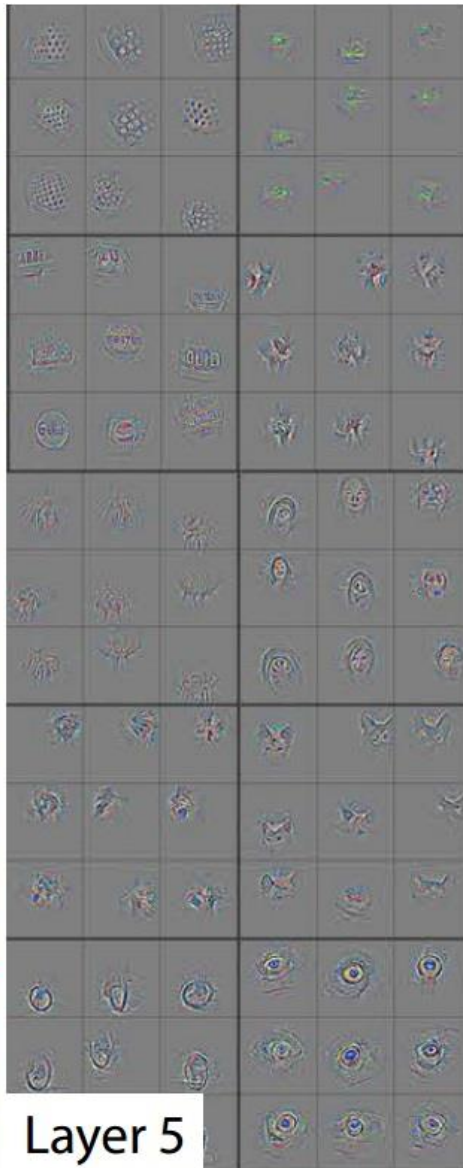




# Layer 4 and 5



Layer 4

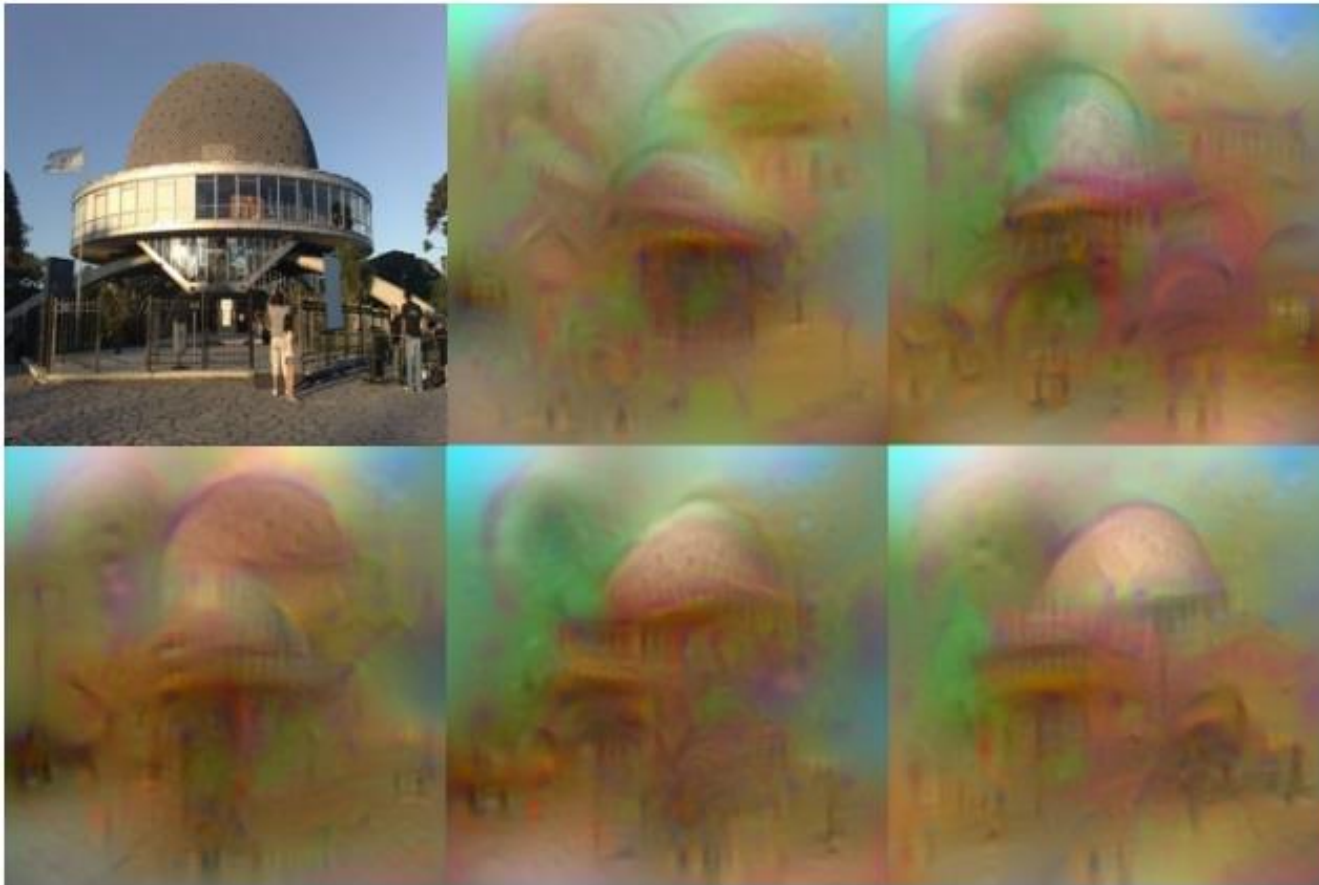


Layer 5



# Invert CNN features

- Reconstruct an image from CNN features



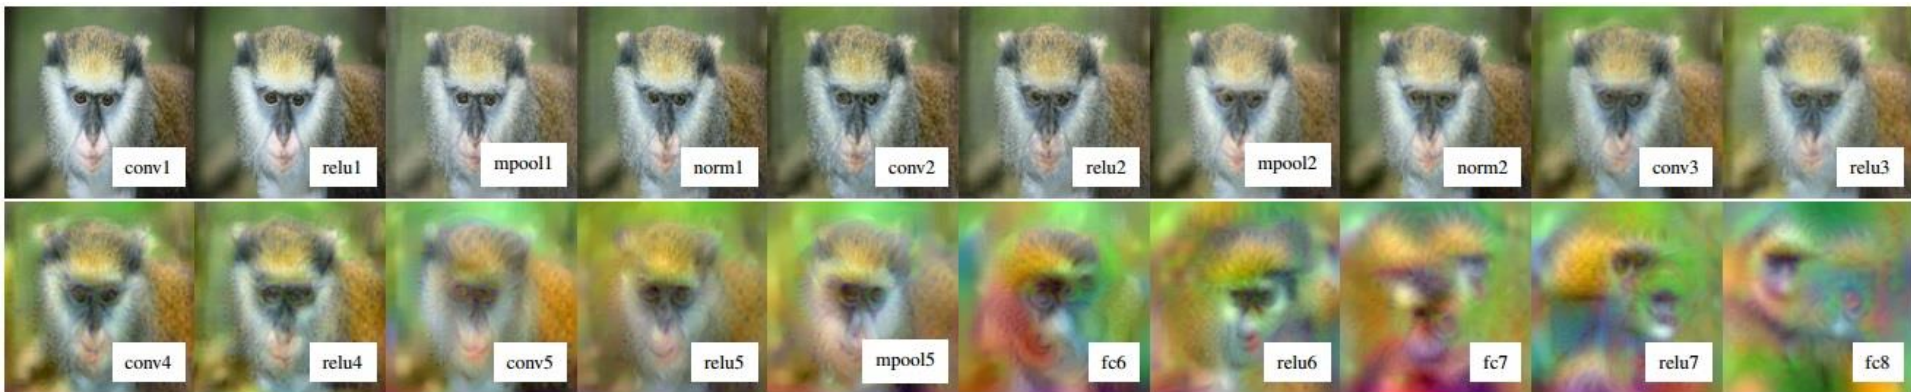
Understanding deep image representations by inverting them

[Mahendran and Vedaldi CVPR 2015]

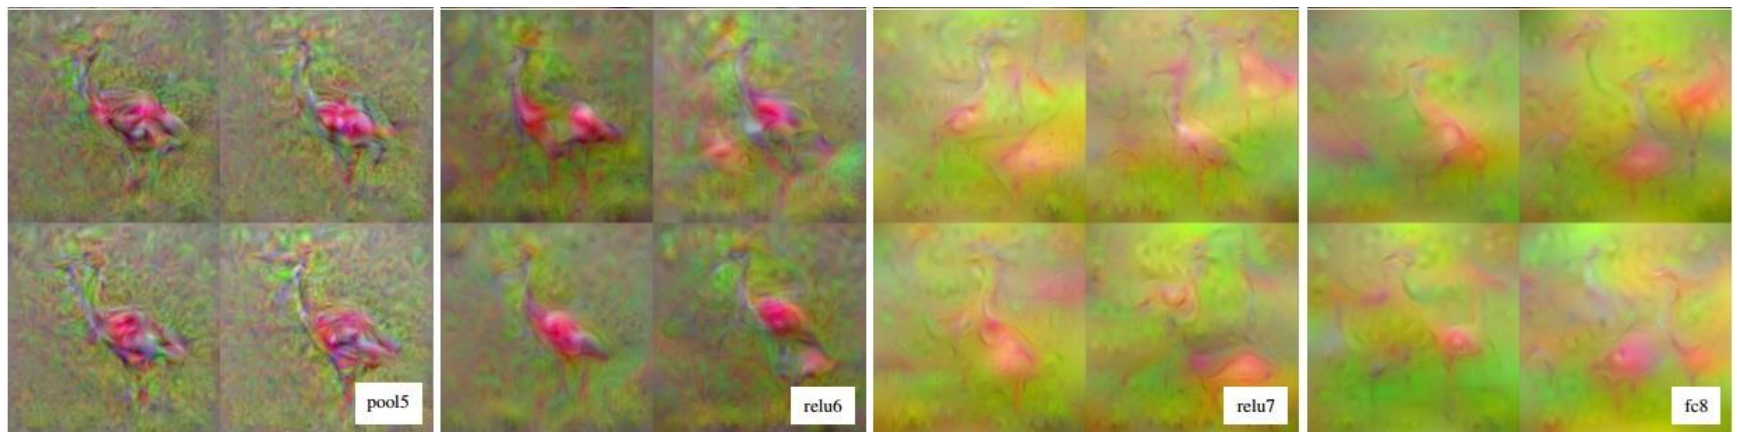


# CNN Reconstruction

## Reconstruction from different layers



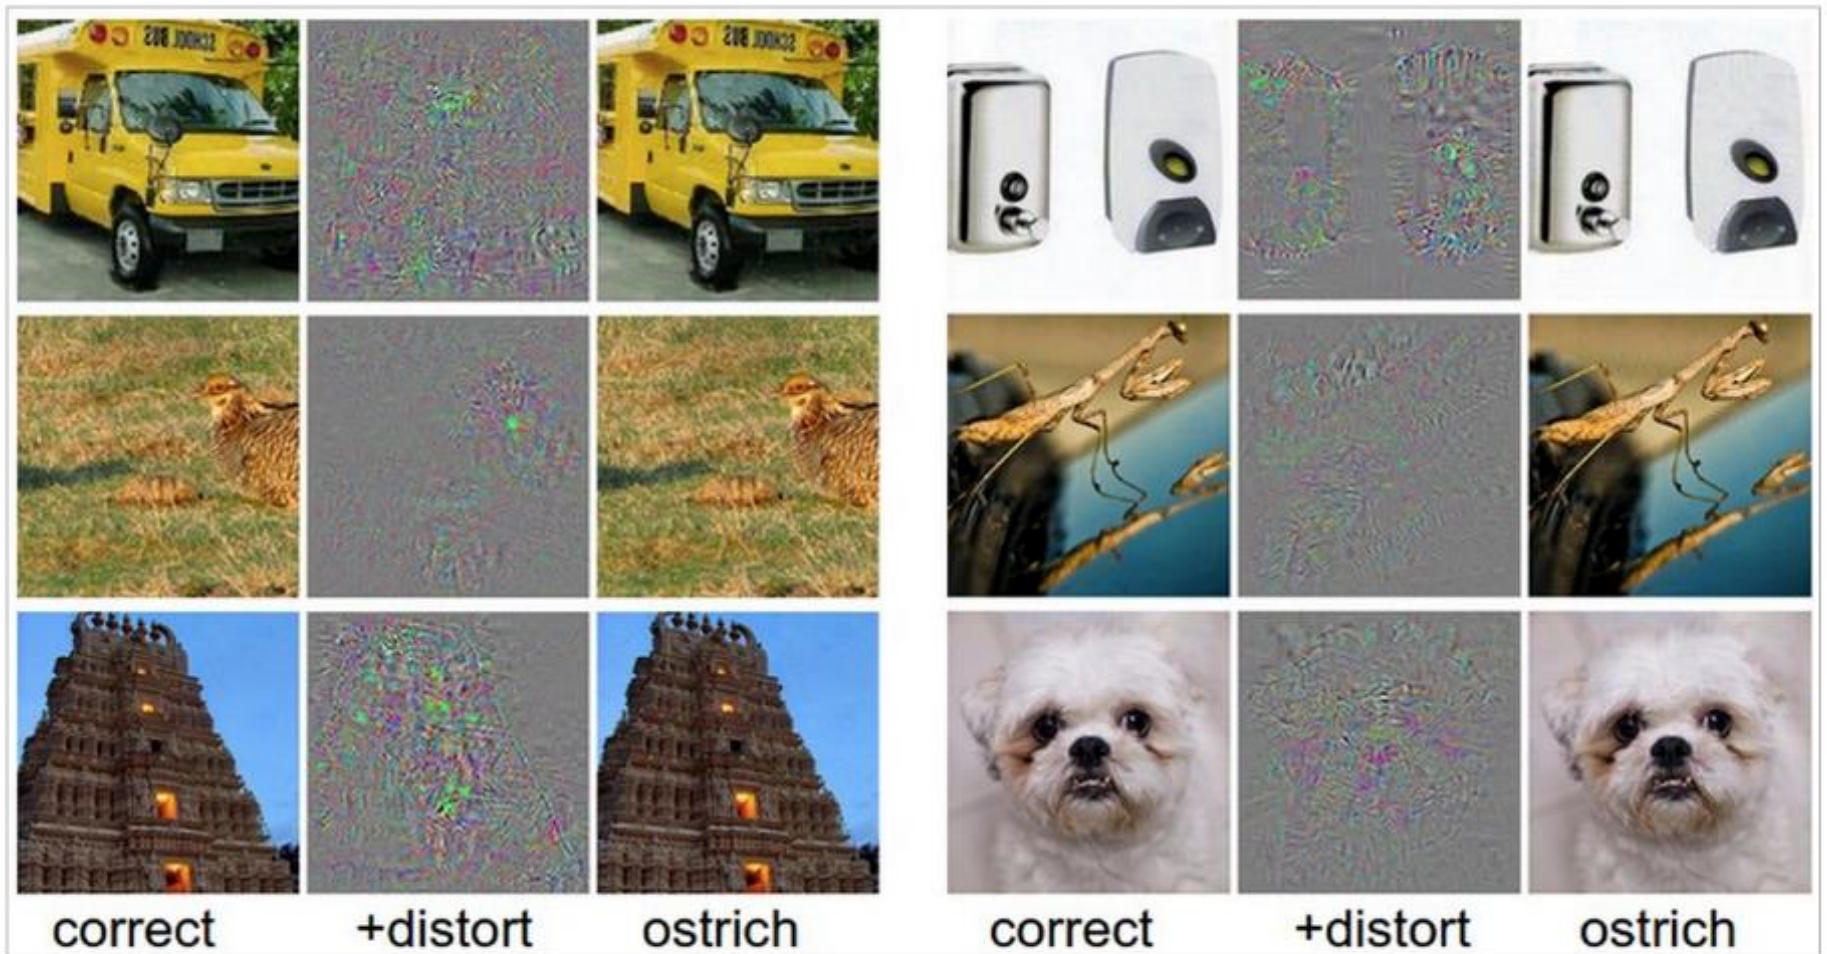
## Multiple reconstructions



Understanding deep image representations by inverting them

[Mahendran and Vedaldi CVPR 2015]

# Breaking CNNs



Take a correctly classified image (left image in both columns), and add a tiny distortion (middle) to fool the ConvNet with the resulting image (right).

Intriguing properties of neural networks [Szegedy ICLR 2014]



# What is going on?

“panda”

57.7% confidence



$x$

+ .007 ×

“nematode”

8.2% confidence



$\frac{\|E\|}{\|x\|}$

=

“gibbon”

99.3 % confidence



$x \rightarrow x + a \frac{\|E\|}{\|x\|}$

Explaining and Harnessing Adversarial Examples [Goodfellow ICLR 2015]

<http://karpathy.github.io/2015/03/30/breaking-convnets/>

# What is going on?

- Recall gradient descent training: modify the weights to reduce classifier error

$$\mathbf{w} \leftarrow \mathbf{w} - \alpha \frac{\partial E}{\partial \mathbf{w}}$$

- Adversarial examples: modify the *image* to *increase* classifier error

$$\mathbf{x} \leftarrow \mathbf{x} + a \frac{\nabla_{\mathbf{x}} E}{\|\nabla_{\mathbf{x}} E\|}$$

# Fooling a linear classifier

- Perceptron weight update: add a small multiple of the example to the weight vector:

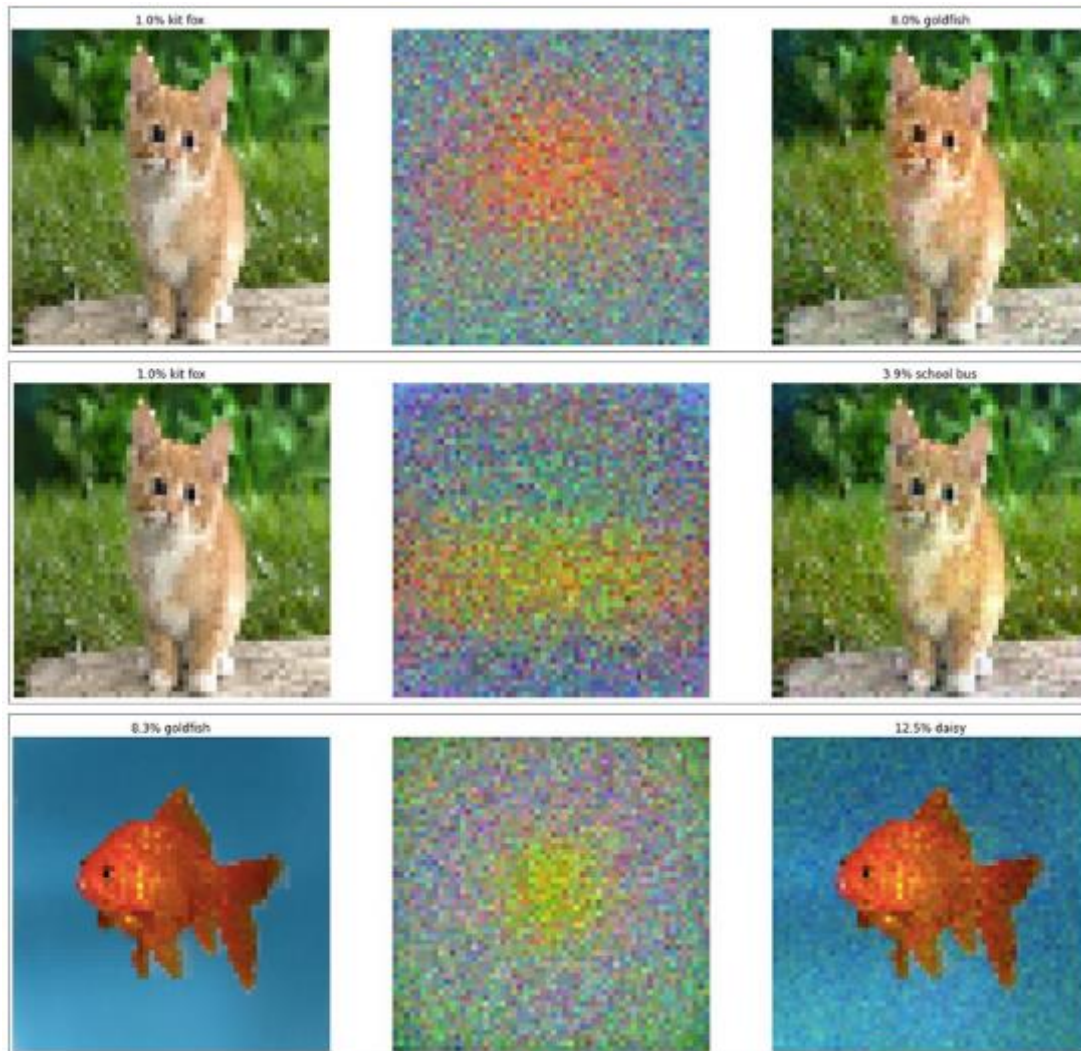
$$\mathbf{w} \leftarrow \mathbf{w} + \alpha \mathbf{x}$$

- To fool a linear classifier, add a small multiple of the weight vector to the training example:

$$\mathbf{x} \leftarrow \mathbf{x} + \alpha \mathbf{w}$$



# Fooling a linear classifier



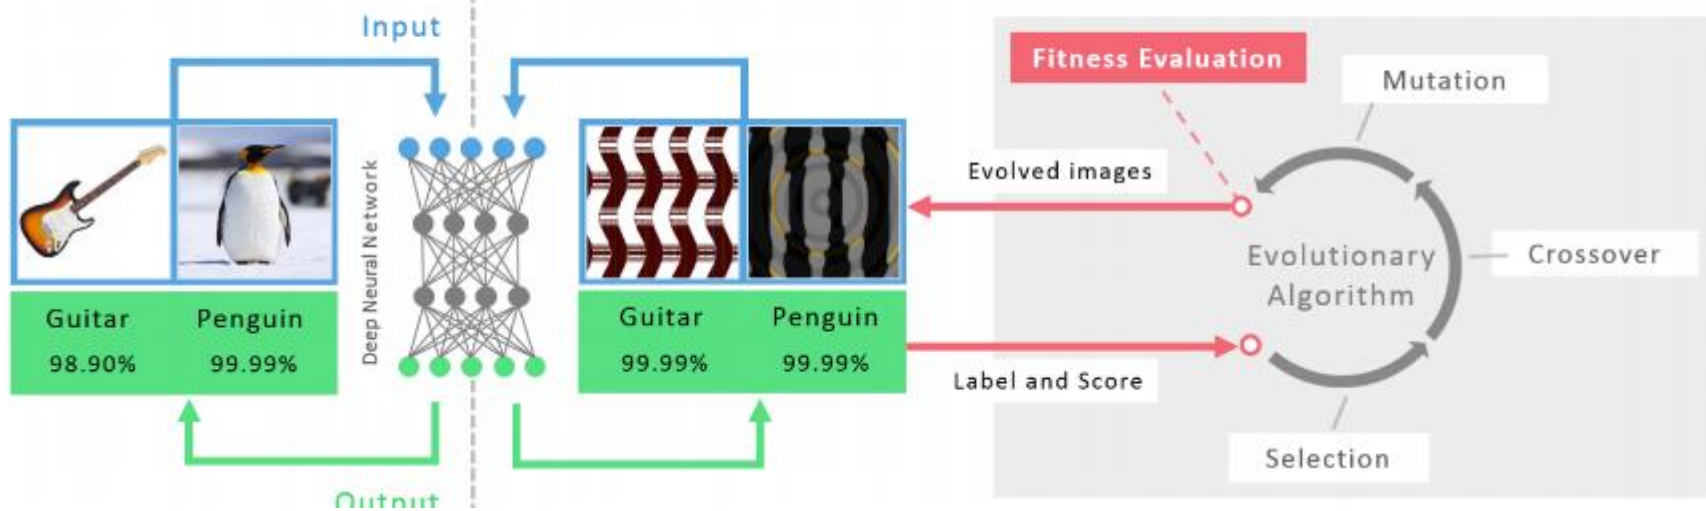
Fooled linear classifier: The starting image (left) is classified as a kit fox. That's incorrect, but then what can you expect from a linear classifier? However, if we add a small amount 'goldfish' weights to the image (top row, middle), suddenly the classifier is convinced that it's looking at one with high confidence. We can distort it with the school bus template instead if we wanted to.

<http://karpathy.github.io/2015/03/30/breaking-convnets/>

# Breaking CNNs

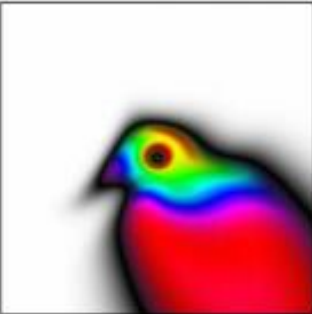



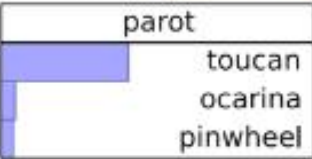
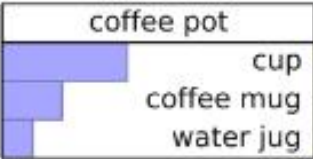
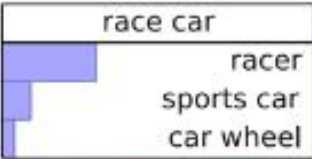
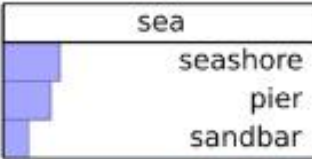

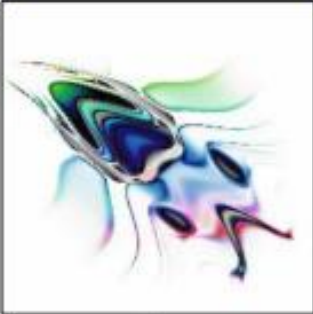
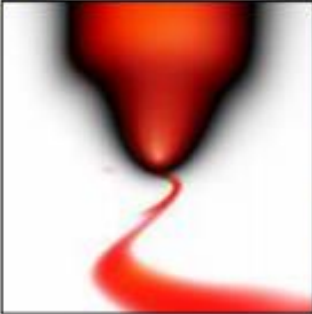
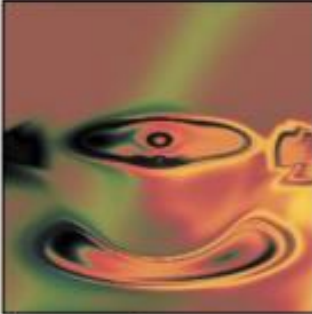
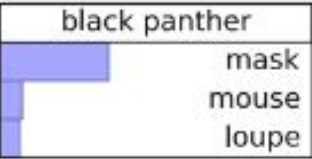

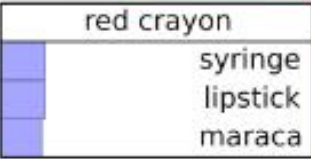
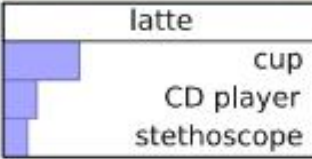
1 State-of-the-art DNNs can recognize real images with high confidence

2 But DNNs are also easily fooled: images can be produced that are unrecognizable to humans, but DNNs believe with 99.99% certainty are natural objects



Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images [Nguyen et al. CVPR 2015]

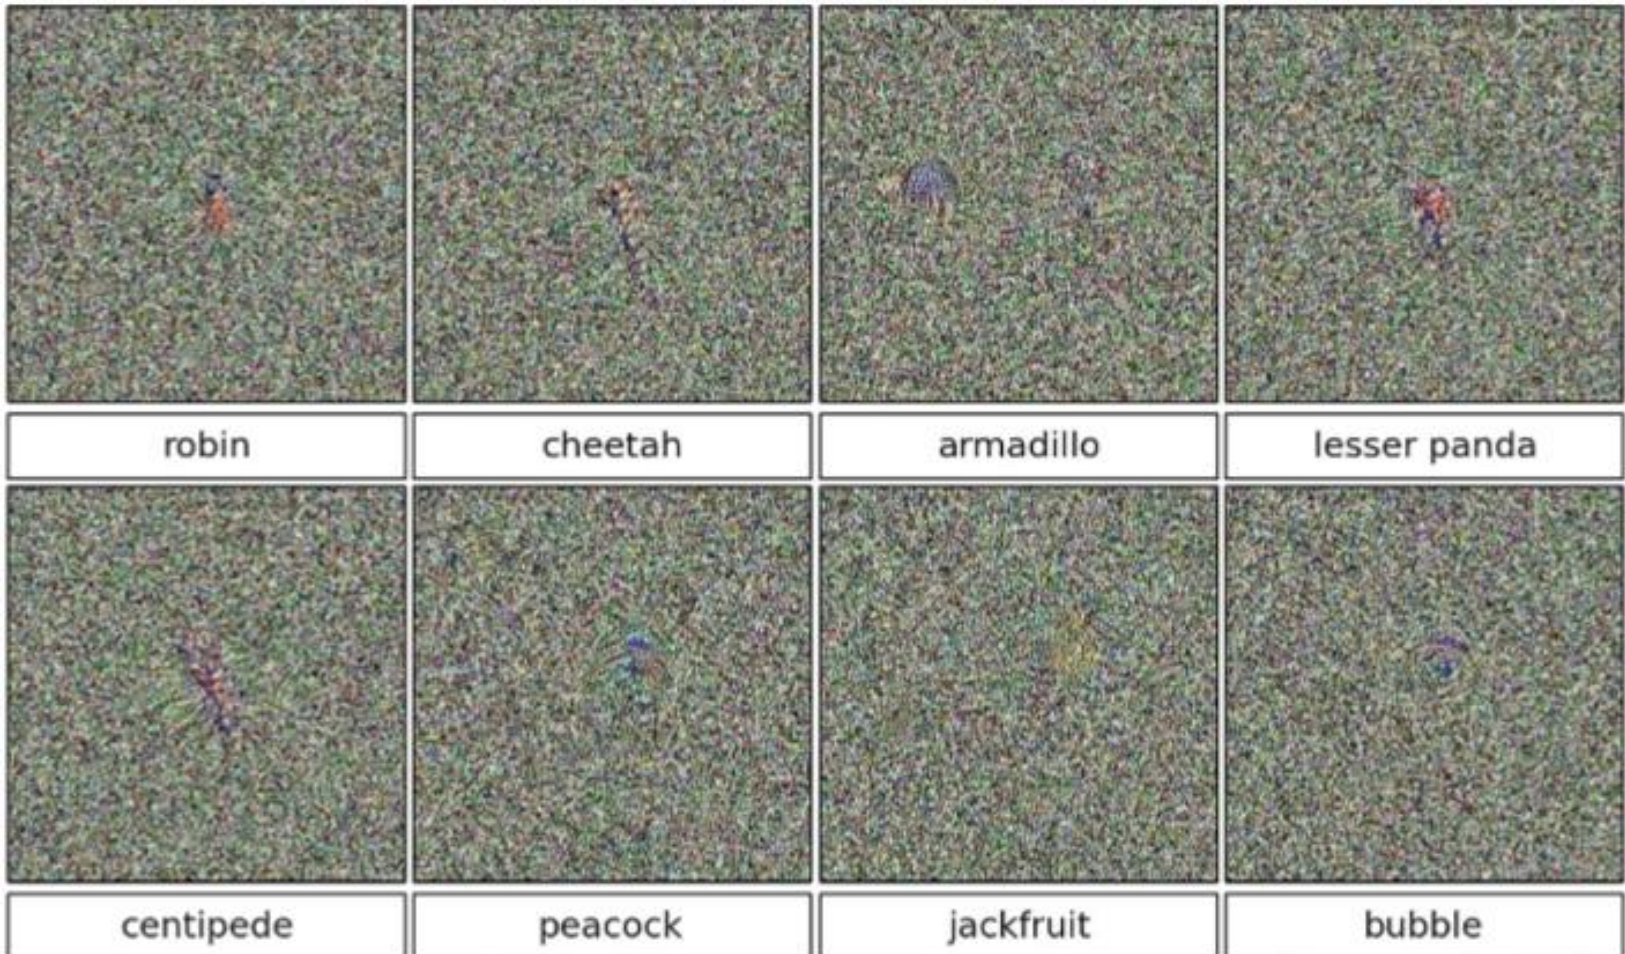
# Images that both CNN and Human can recognize

|  |  |  |  |
|--|--|--|--|
|                                   |   |                                      |                                       |
| parot  | coffee pot   | race car   | sea  |
| <br>toucan<br>ocarina<br>pinwheel | <br>cup<br>coffee mug<br>water jug              | <br>racer<br>sports car<br>car wheel | <br>seashore<br>pier<br>sandbar       |
|                                  |    |                                     |                                      |
| black panther  | fly  | red crayon   | latte  |
| <br>mask<br>mouse<br>loupe      | <br>ground beetle<br>fly<br>rhinoceros beetle | <br>syringe<br>lipstick<br>maraca  | <br>cup<br>CD player<br>stethoscope |

Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images [Nguyen et al. CVPR 2015]

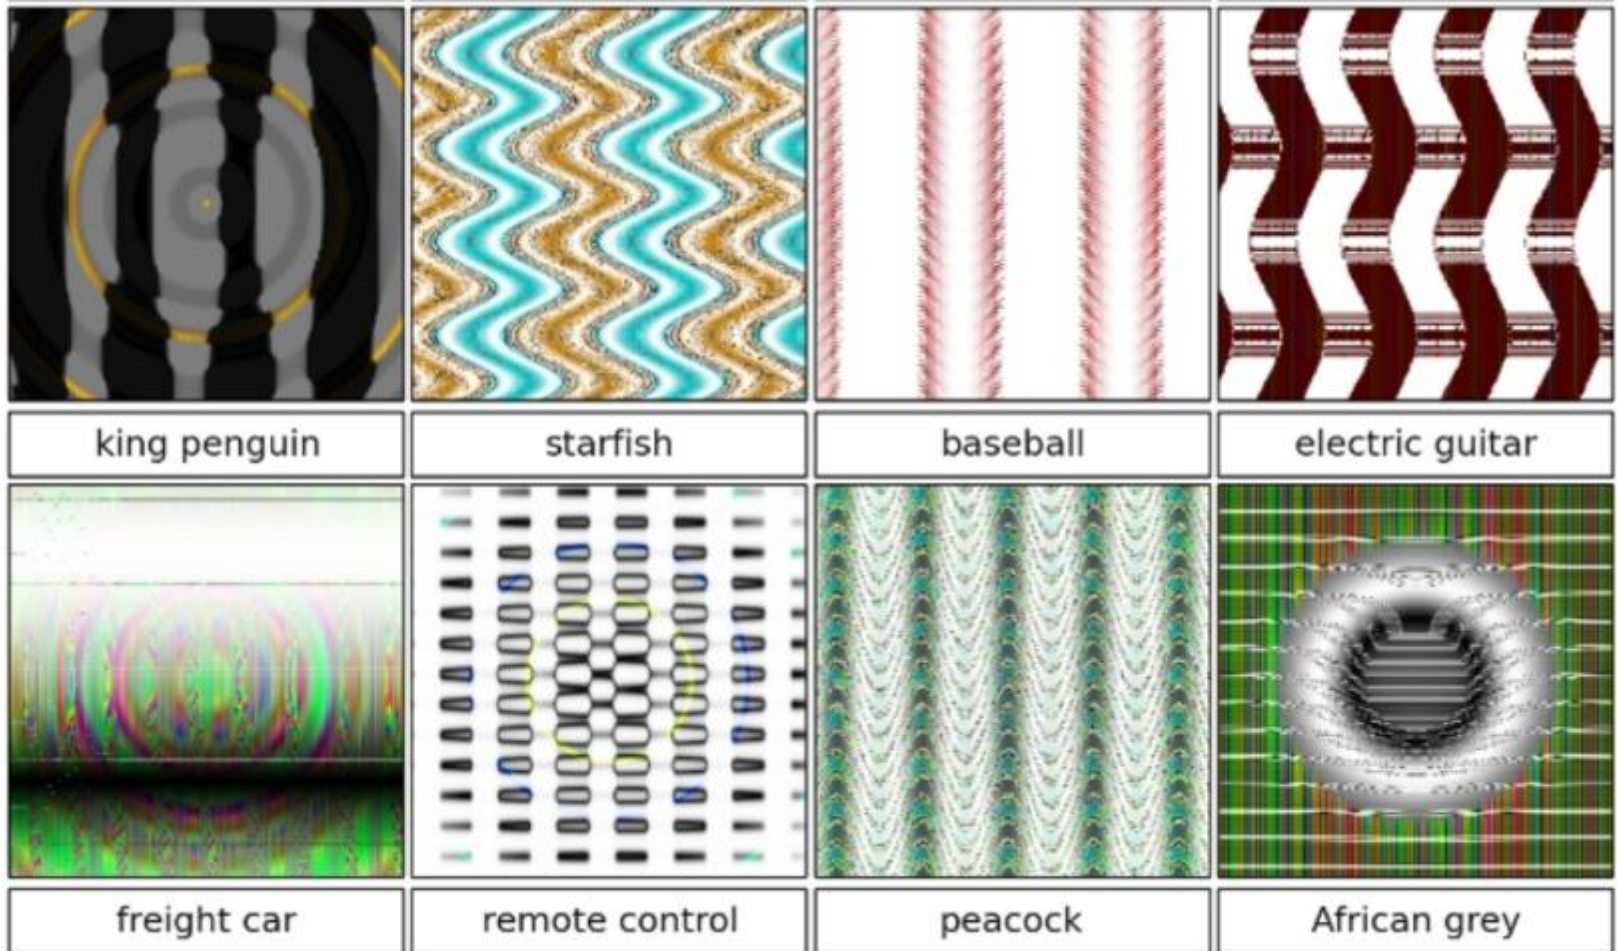


# Direct Encoding



Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images [Nguyen et al. CVPR 2015]

# Indirect Encoding



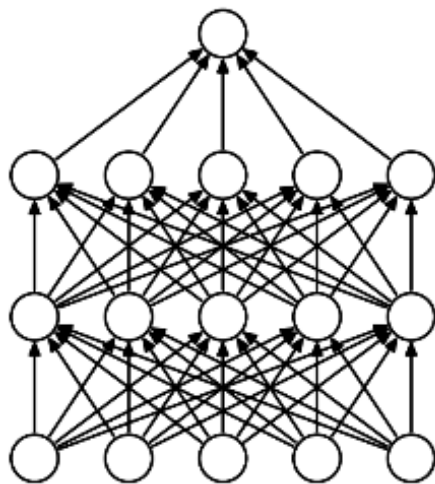
Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images [Nguyen et al. CVPR 2015]

# Training Convolutional Neural Networks

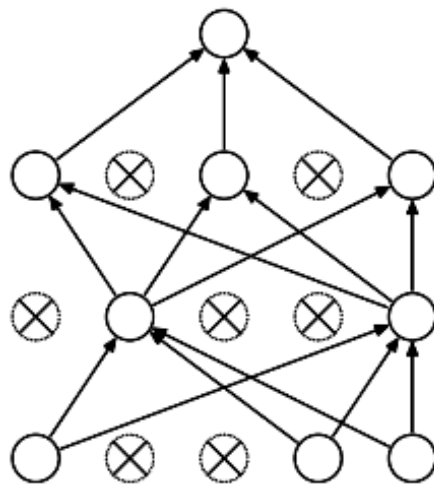
- Backpropagation + stochastic gradient descent with momentum
  - [Neural Networks: Tricks of the Trade](#)
- Dropout
- Data augmentation
- Batch normalization
- Initialization
  - Transfer learning



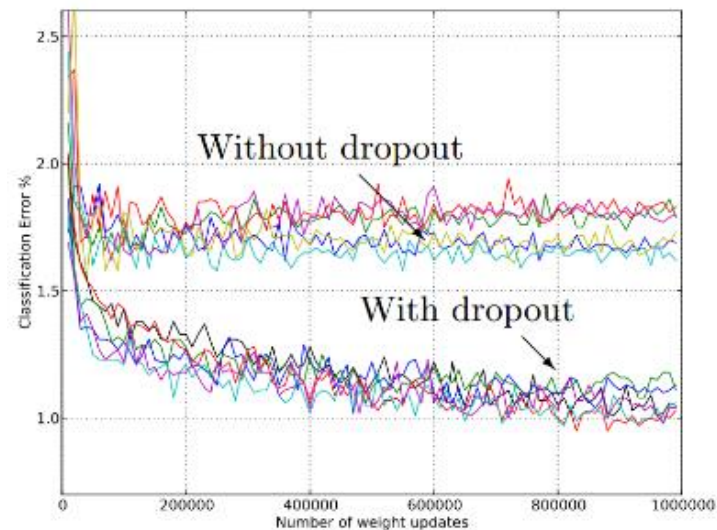
# Dropout



(a) Standard Neural Net



(b) After applying dropout.



**Main Idea:** approximately combining exponentially many different neural network architectures efficiently

| Model  | Top-1 (val) | Top-5 (val) | Top-5 (test) |
|--|-------------|-------------|--------------|
| SVM on Fisher Vectors of Dense SIFT and Color Statistics | -           | -           | 27.3         |
| Avg of classifiers over FVs of SIFT, LBP, GIST and CSIFT | -           | -           | 26.2         |
| Conv Net + dropout (Krizhevsky et al., 2012)             | 40.7        | 18.2        | -            |
| Avg of 5 Conv Nets + dropout (Krizhevsky et al., 2012)   | 38.1        | 16.4        | 16.4         |

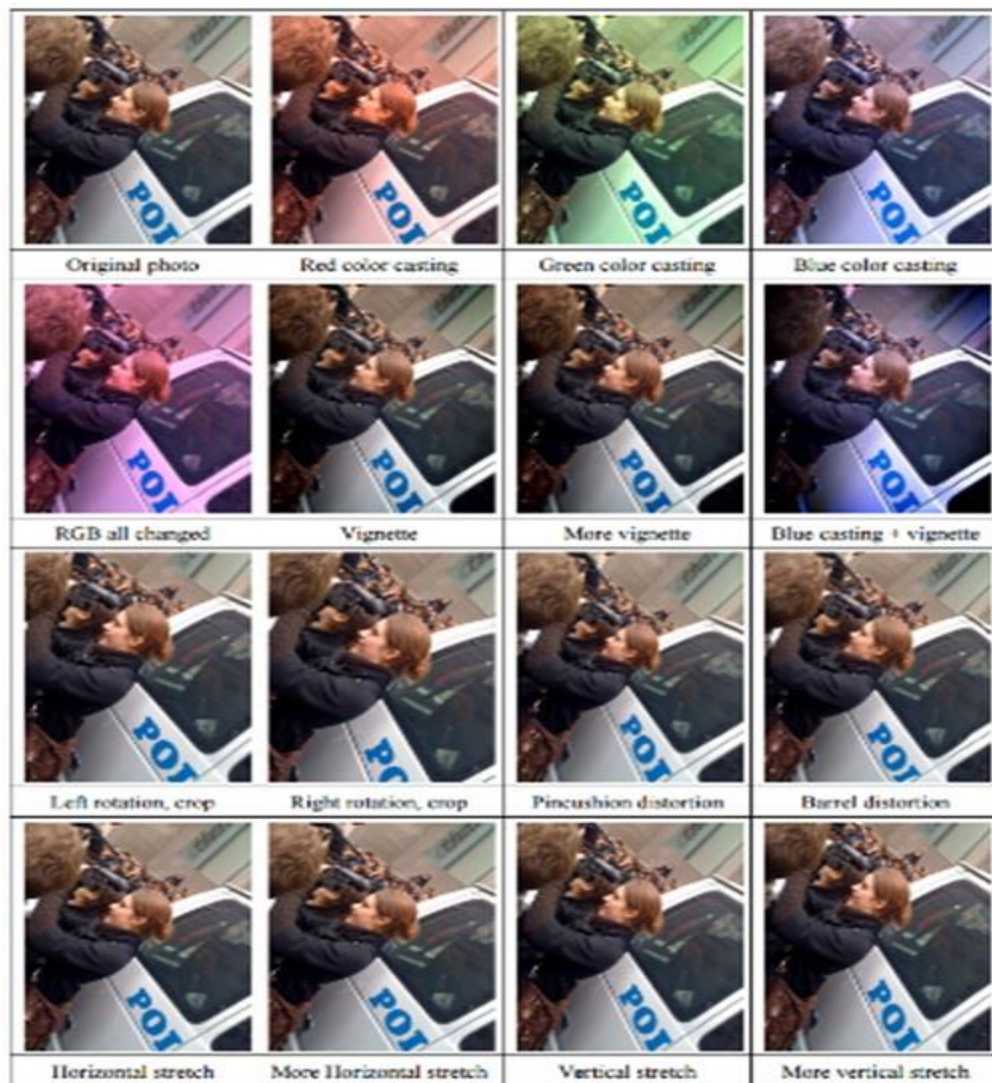
Table 6: Results on the ILSVRC-2012 validation/test set.

Dropout: A simple way to prevent neural networks from overfitting [Srivastava JMLR 2014]

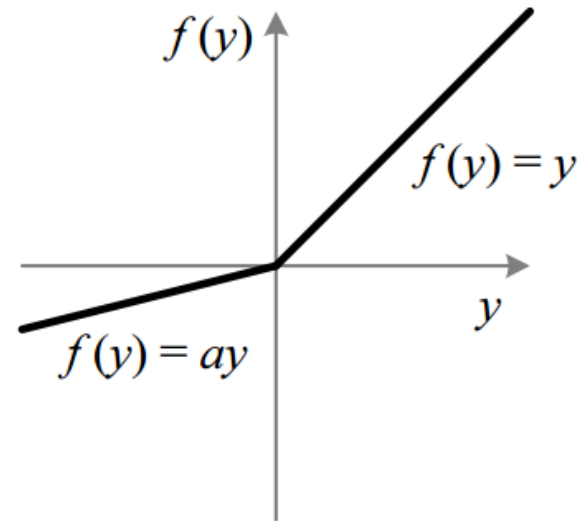
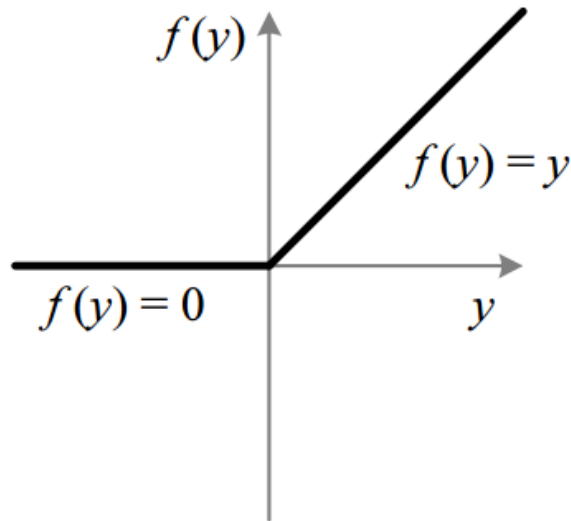


# Data Augmentation (Jittering)

- Create *virtual* training samples
  - Horizontal flip
  - Random crop
  - Color casting
  - Geometric distortion



# Parametric Rectified Linear Unit



|                             | team                    | top-5 (test) |
|-----------------------------|-------------------------|--------------|
| in competition<br>ILSVRC 14 | MSRA, SPP-nets [11]     | 8.06         |
|                             | VGG [25]                | 7.32         |
|                             | GoogLeNet [29]          | 6.66         |
| post-competition            | VGG [25] (arXiv v5)     | 6.8          |
|                             | Baidu [32]              | 5.98         |
|                             | <b>MSRA, PReLU-nets</b> | <b>4.94</b>  |

Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification [He et al. 2015]

# Batch Normalization

**Input:** Values of  $x$  over a mini-batch:  $\mathcal{B} = \{x_1 \dots x_m\}$ ;

Parameters to be learned:  $\gamma, \beta$

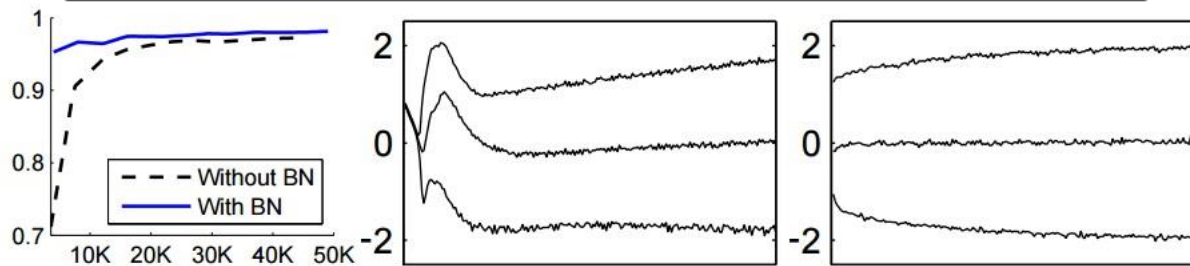
**Output:**  $\{y_i = \text{BN}_{\gamma, \beta}(x_i)\}$

$$\mu_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad // \text{ mini-batch mean}$$

$$\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2 \quad // \text{ mini-batch variance}$$

$$\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} \quad // \text{ normalize}$$

$$y_i \leftarrow \gamma \hat{x}_i + \beta \equiv \text{BN}_{\gamma, \beta}(x_i) \quad // \text{ scale and shift}$$



(a)

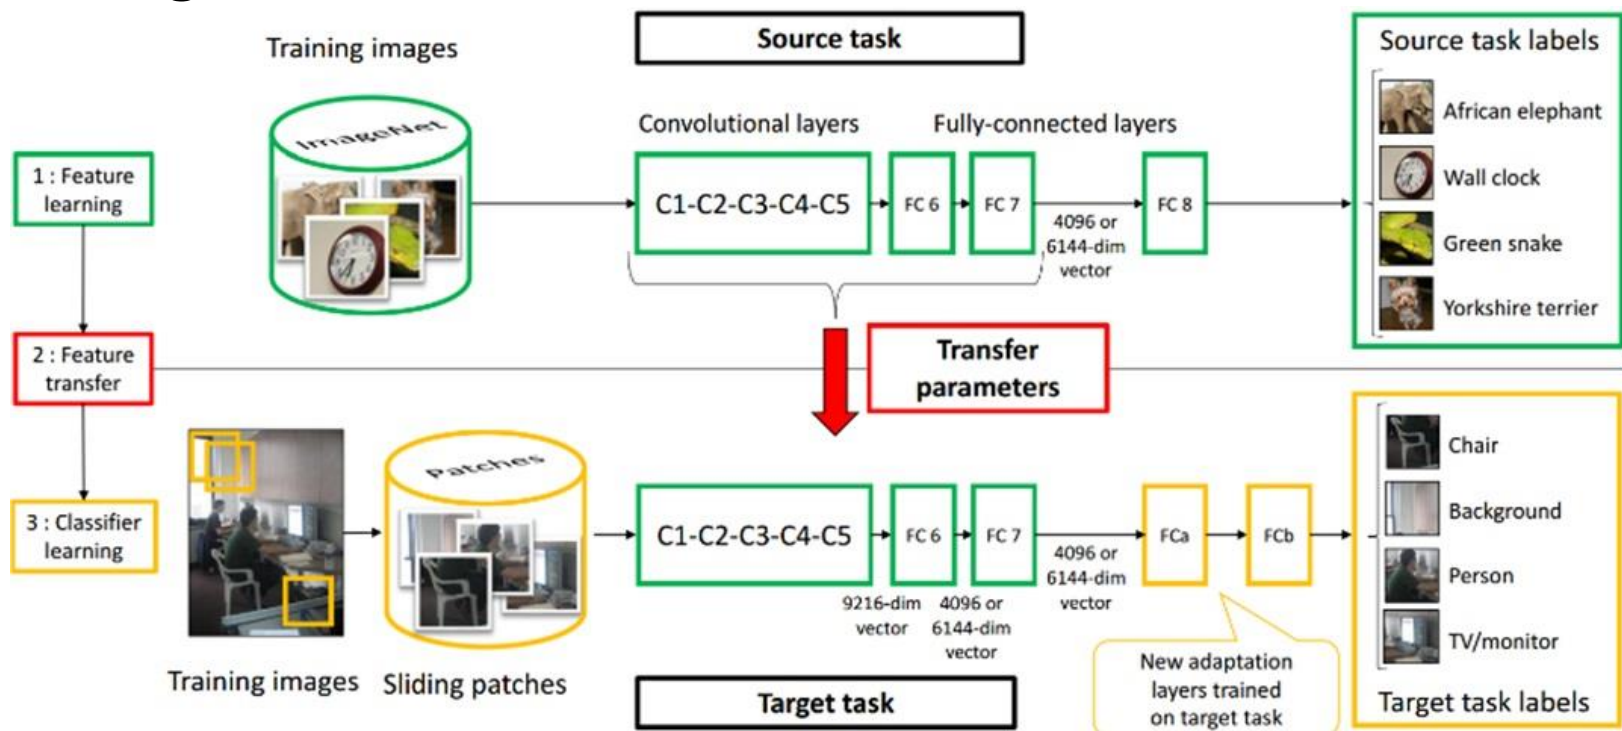
(b) Without BN

(c) With BN

Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift [[Ioffe and Szegedy 2015](#)]

# Transfer Learning

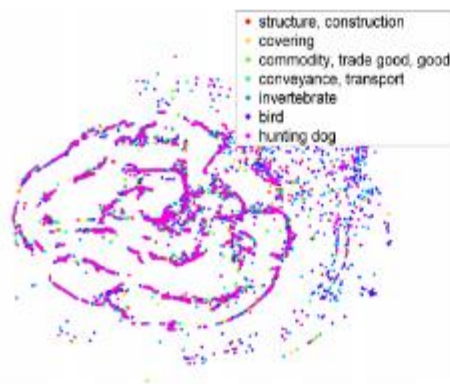
- Improvement of learning in a **new** task through the *transfer of knowledge* from a **related** task that has already been learned.
- Weight initialization for CNN



Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks [Oquab et al. CVPR 2014]



# Convolutional activation features



(a) LLC



(b) GIST

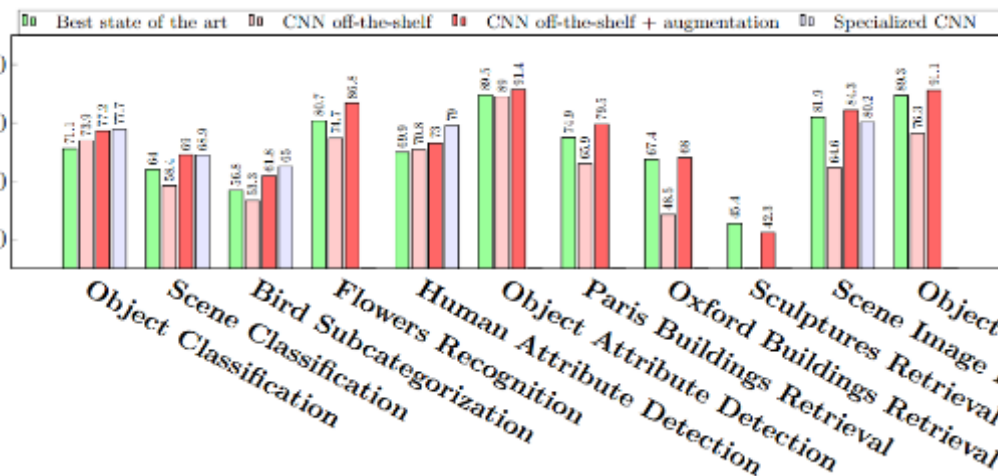
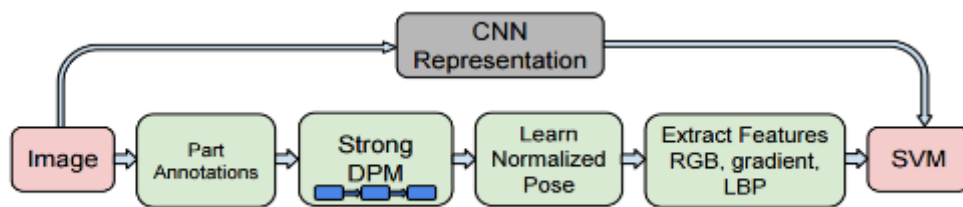


(c) DeCAF<sub>1</sub>



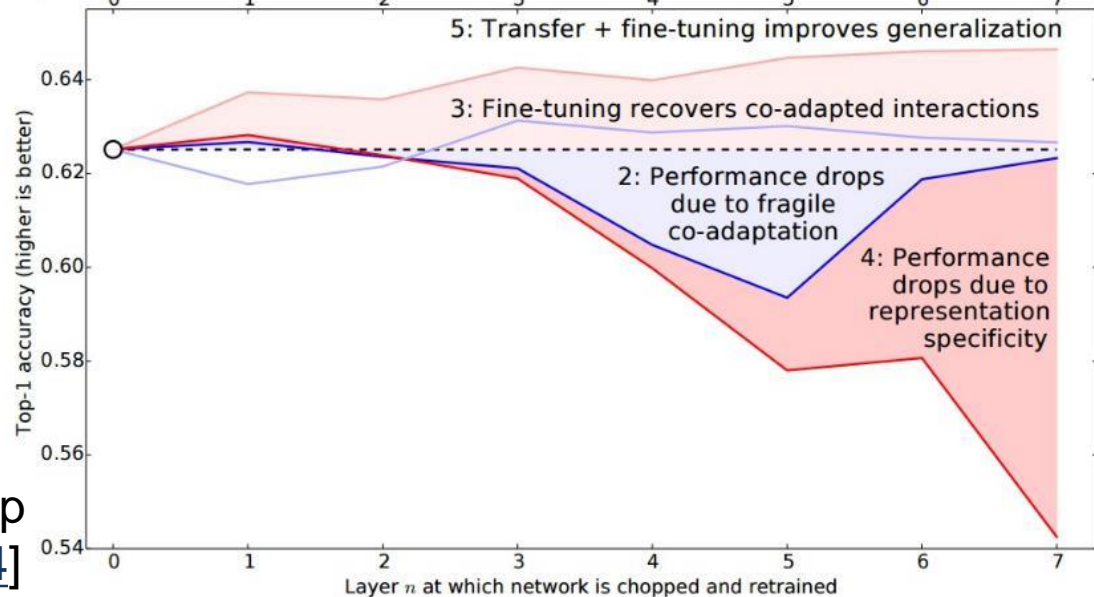
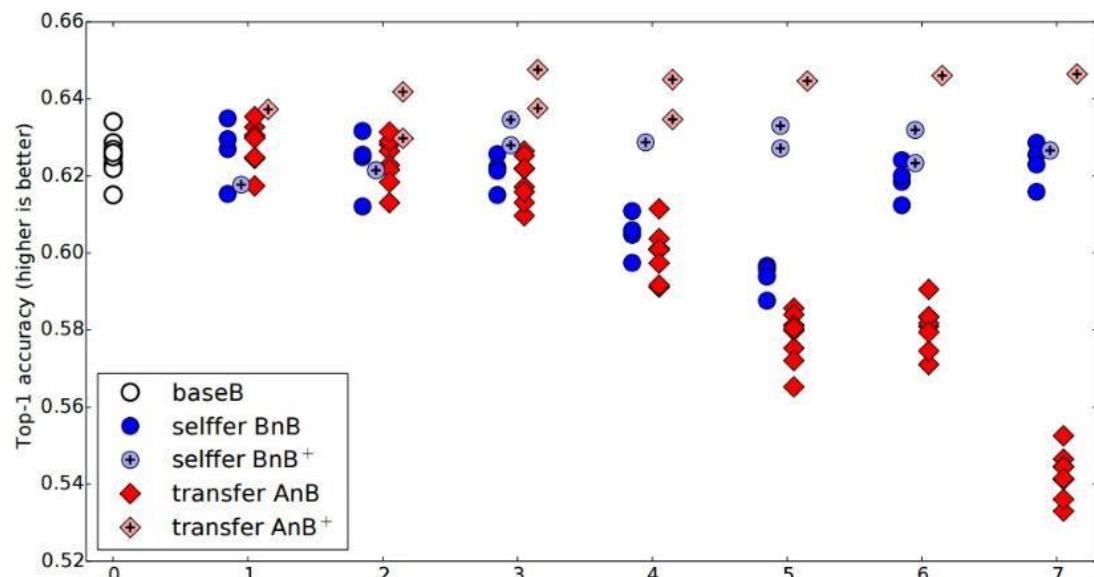
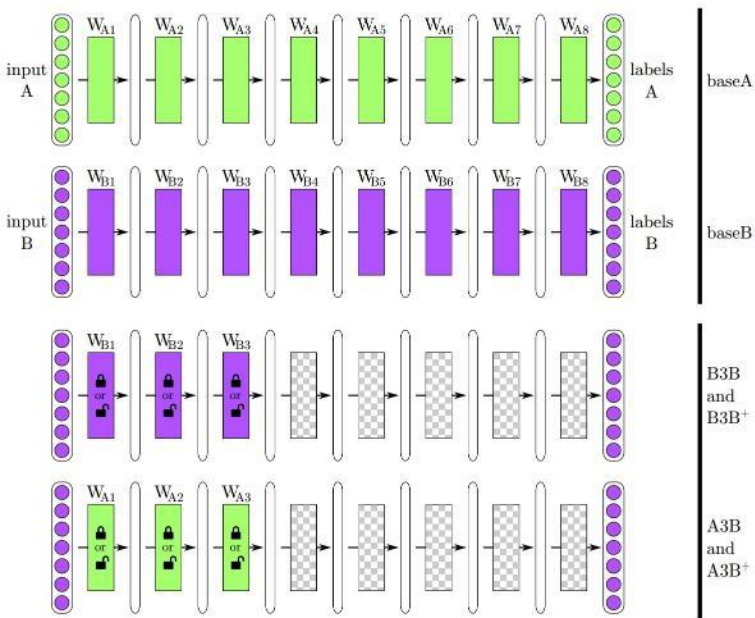
(d) DeCAF<sub>6</sub>

[Donahue et al. ICML 2013]



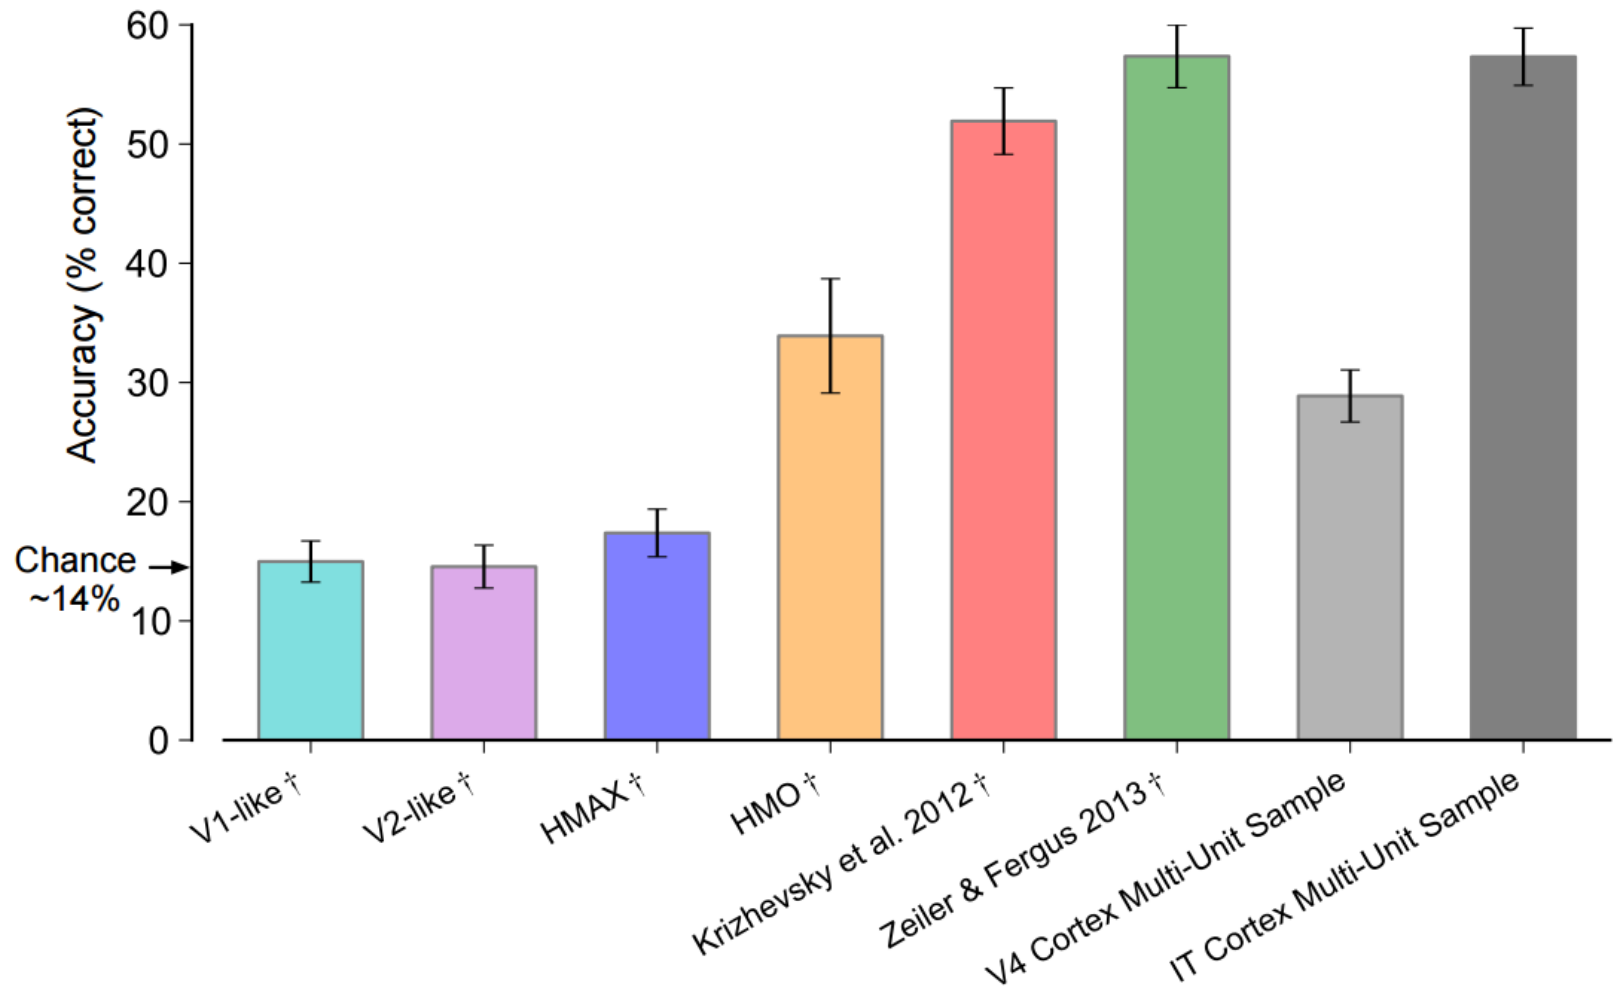
CNN Features off-the-shelf:  
an Astounding Baseline for Recognition  
[Razavian et al. 2014]

# How transferable are features in CNN?



How transferable are features in deep neural networks [Yosinski NIPS 2014]

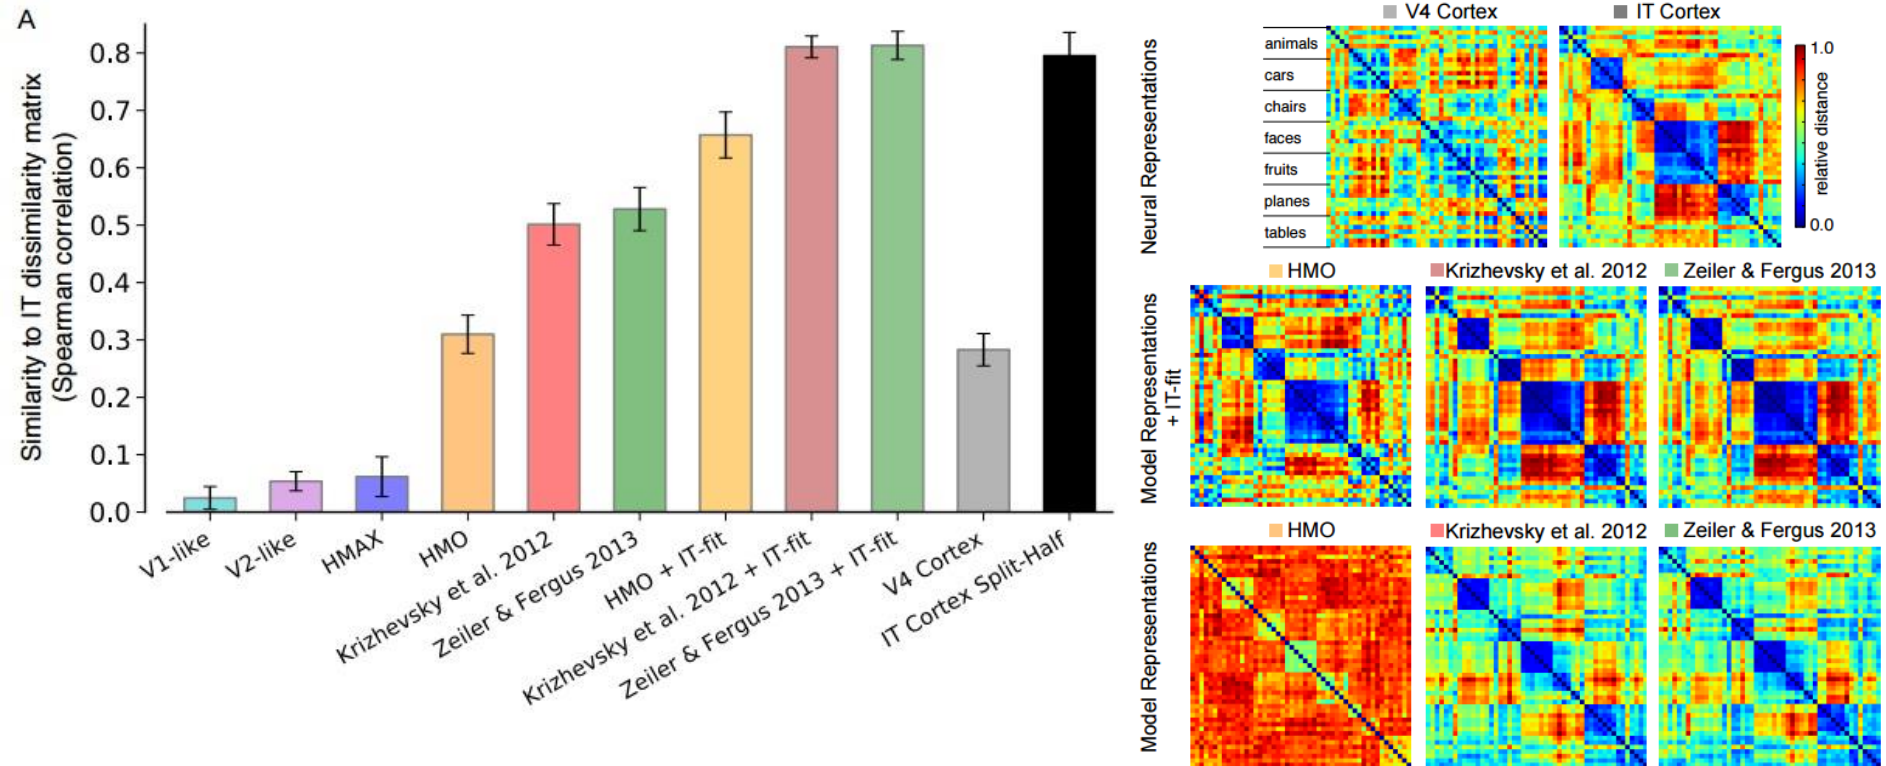
# Deep Neural Networks Rival the Representation of Primate Inferior Temporal Cortex



Deep Neural Networks Rival the Representation of Primate IT Cortex for Core Visual Object Recognition [Cadieu et al. PLOS 2014]



# Deep Neural Networks Rival the Representation of Primate Inferior Temporal Cortex



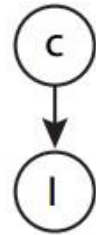
Deep Neural Networks Rival the Representation of Primate IT Cortex for Core Visual Object Recognition [Cadieu et al. PLOS 2014]

# Deep Rendering Model (DRM)

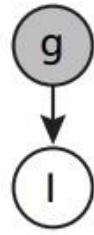
**A**

Naive Bayes Mixture

Classifier



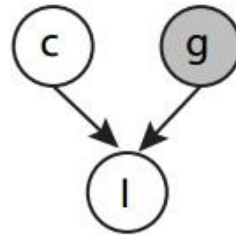
Model



**B**

Rendering

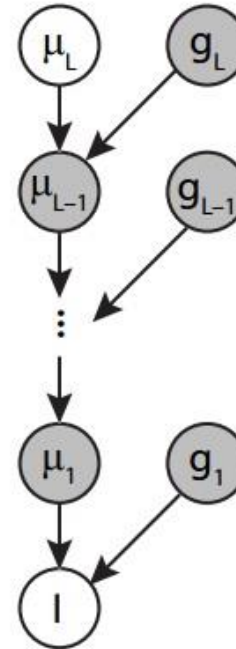
Model



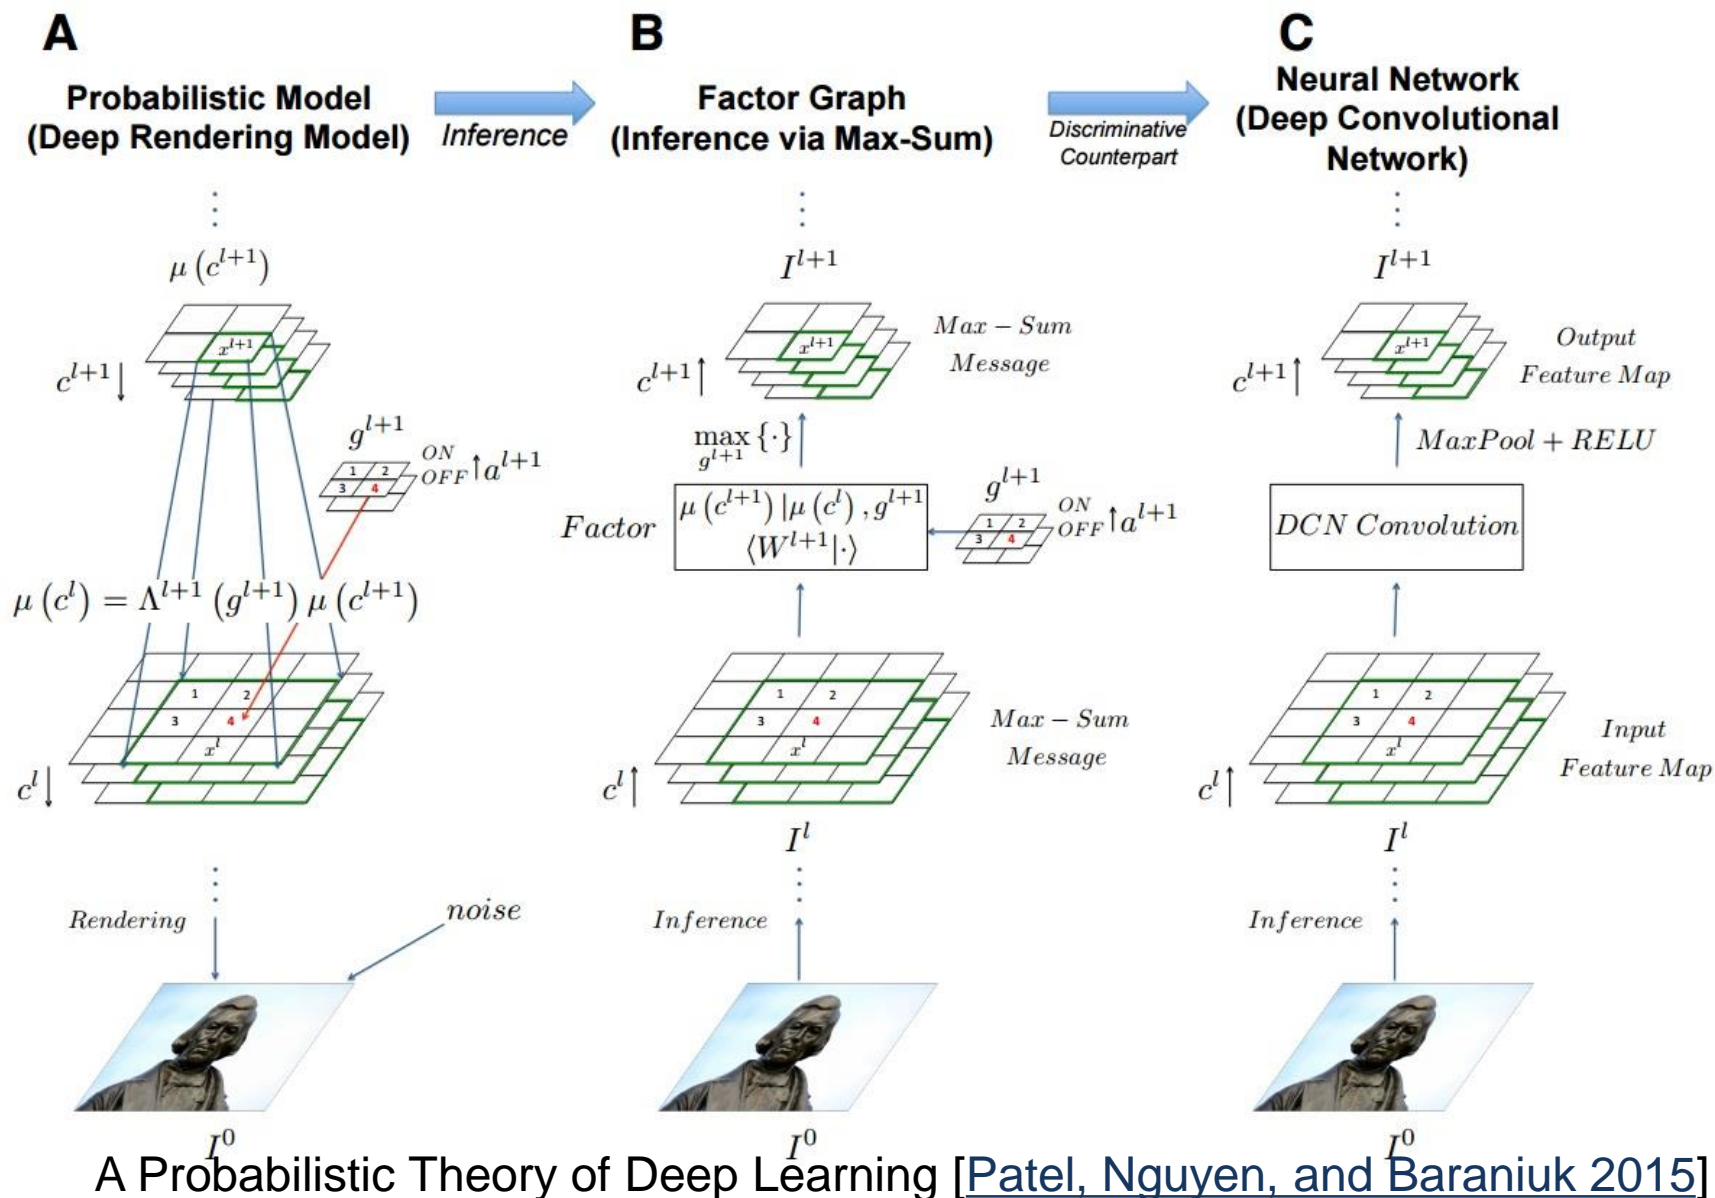
**C**

Deep Rendering

Model



# CNN as a Max-Sum Inference



# Model

# Inference

# Learning

| Aspect           | Neural Nets Perspective<br><i>Deep Convnets (DCNs)</i>   | Probabilistic Perspective<br><i>Deep Rendering Model (DRM)</i>  |
|------------------|--|---|
| <i>Model</i>     | Weights and biases of filters at a given layer   | Partial Rendering at a given abstraction level/scale  |
|                  | Number of Layers   | Number of Abstraction Levels  |
|                  | Number of Filters in a layer   | Number of Clusters/Classes at a given abstraction level   |
| <i>Inference</i> | Implicit in network weights; can be computed by product of weights over all layers or by activity maximization | Category prototypes are finely detailed versions of coarser-scale super-category prototypes. Fine details are modeled with affine nuisance transformations. |
|                  | Forward propagation thru DCN   | Exact bottom-up inference via Max-Sum Message Passing (with Max-Product for Nuisance Factorization).  |
|                  | Input and Output Feature Maps  | Probabilistic Max-Sum Messages (real-valued functions of variables nodes)   |
|                  | Template matching at a given layer (convolutional, locally or fully connected)                                 | Local computation at factor node (log-likelihood of measurements)   |
|                  | Max-Pooling over local pooling region  | Max-Marginalization over Latent Translational Nuisance transformations  |
| <i>Learning</i>  | Rectified Linear Unit (ReLU). Sparsifies output activations.   | Max-Marginalization over Latent Switching state of Renderer. Low prior probability of being ON.   |
|                  | Stochastic Gradient Descent  | Batch Discriminative EM Algorithm with Fine-to-Coarse E-step + Gradient M-step. <i>No coarse-to-fine pass in E-step.</i>                                    |
|                  | N/A  | Full EM Algorithm   |
|                  | Batch-Normalized SGD (Google state-of-the-art [BN])  | Discriminative Approximation to Full EM (assumes Diagonal Pixel Covariance)   |

# Tools

- [Caffe](#)
- [cuda-convnet2](#)
- [Torch](#)
- [MatConvNet](#)
- [Pylearn2](#)



# Resources

- <http://deeplearning.net/>
- <https://github.com/ChristosChristofidis/awesome-deep-learning>

# Things to remember

- Overview
  - Neuroscience, Perceptron, multi-layer neural networks
- Convolutional neural network (CNN)
  - Convolution, nonlinearity, max pooling
  - CNN for classification and beyond
- Understanding and visualizing CNN
  - Find images that maximize some class scores; visualize individual neuron activation, input pattern and images; breaking CNNs
- Training CNN
  - Dropout; data augmentation; batch normalization; transfer learning
- Probabilistic interpretation
  - Deep rendering model; CNN forward-propagation as max-sum inference; training as an EM algorithm