



UNIVERSITÀ  
DI PAVIA

# *Deep Learning*

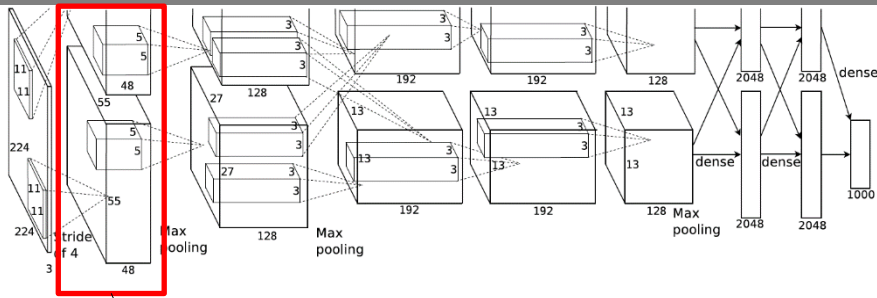
## *07-Deep Convolutional Neural Networks and Beyond*

Marco Piastra

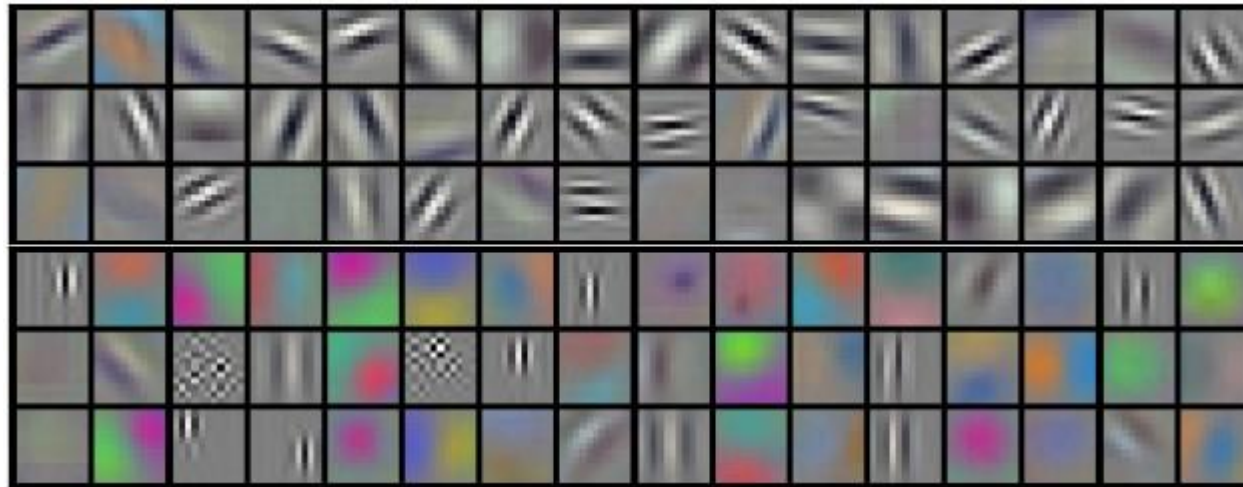
*This presentation can be downloaded at:*  
<http://vision.unipv.it/DL>

# *Inside AlexNet (after training)*

# AlexNet Filters (after training)



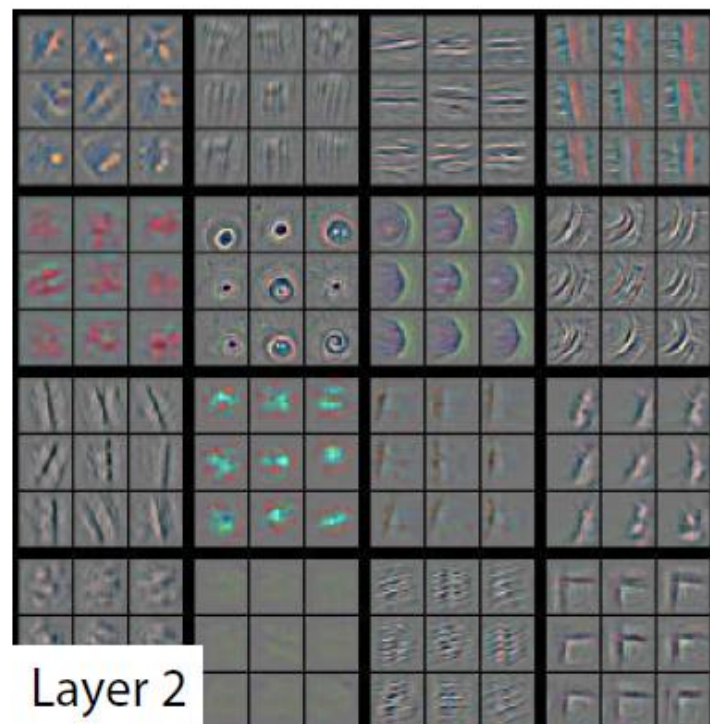
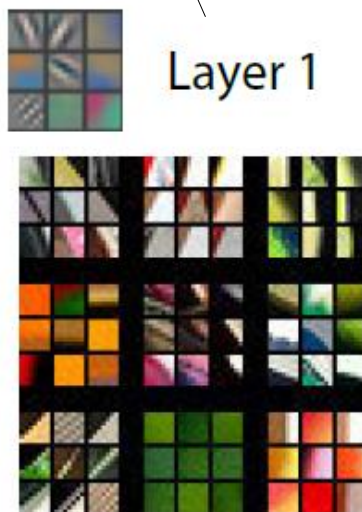
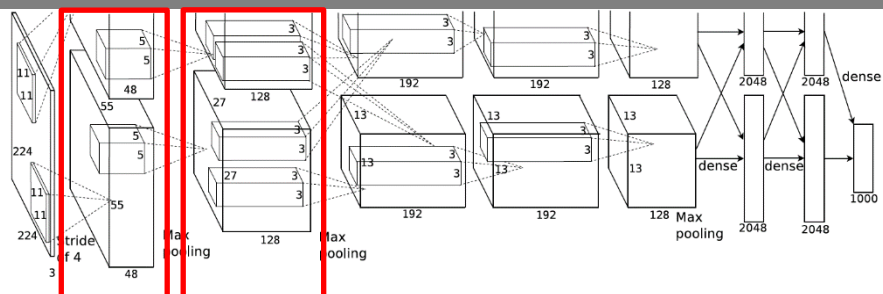
Layer 1



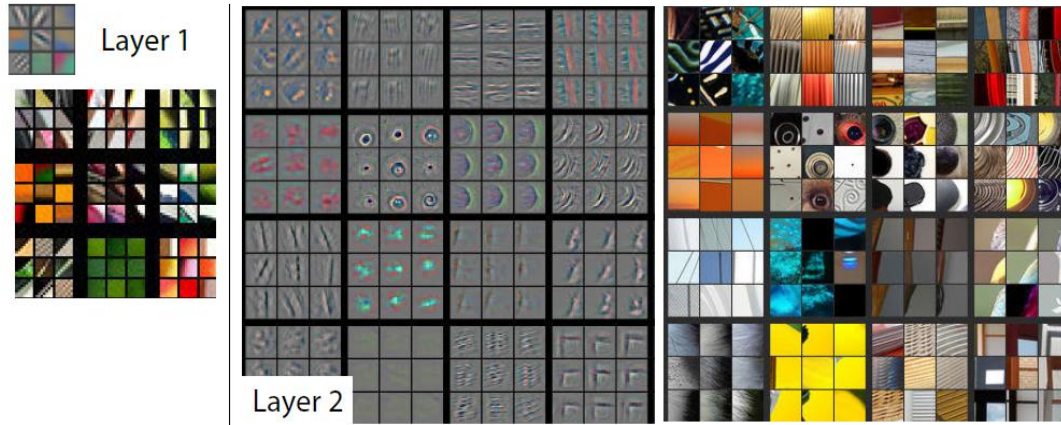
These are 96 real examples of convolutive filters for RGB images

[image from <http://cs231n.github.io/convolutional-networks/>]

# AlexNet Filters- DeconvNet

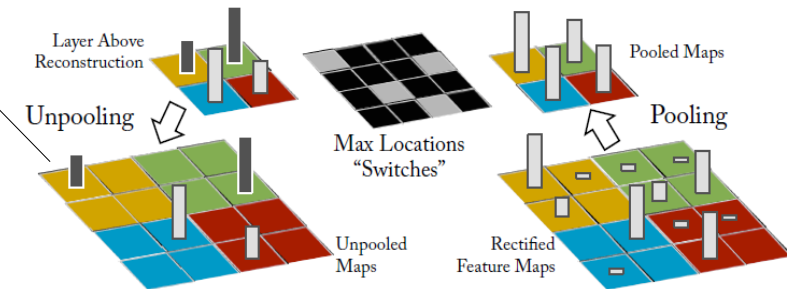
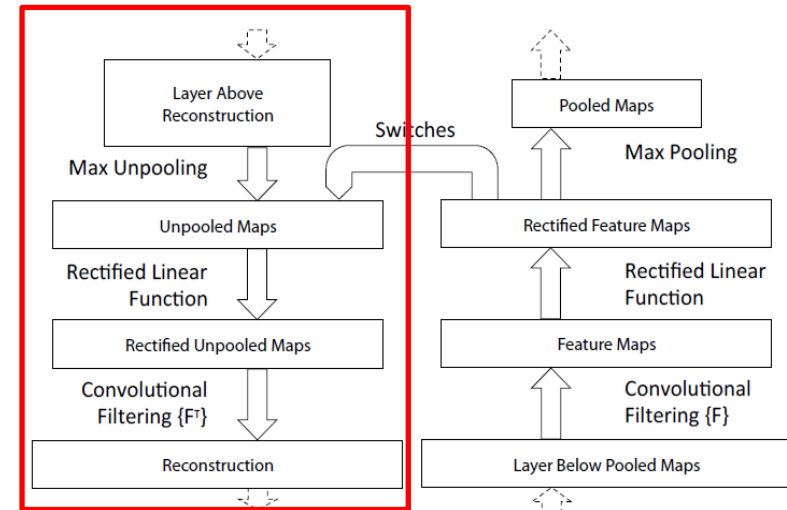


# AlexNet Filters- DeconvNet

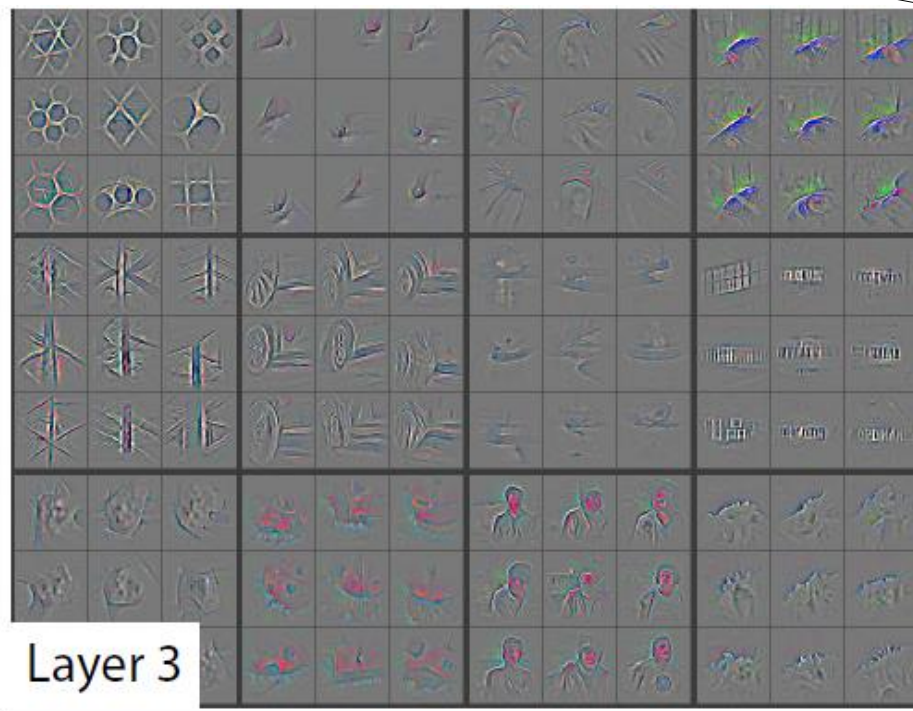
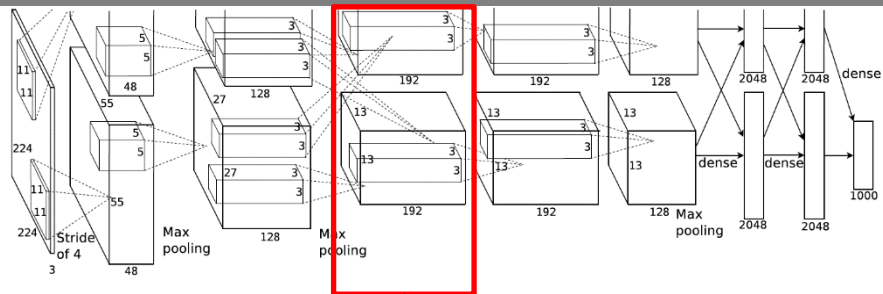


*DeconvNet:*  
using a DCNN in reverse

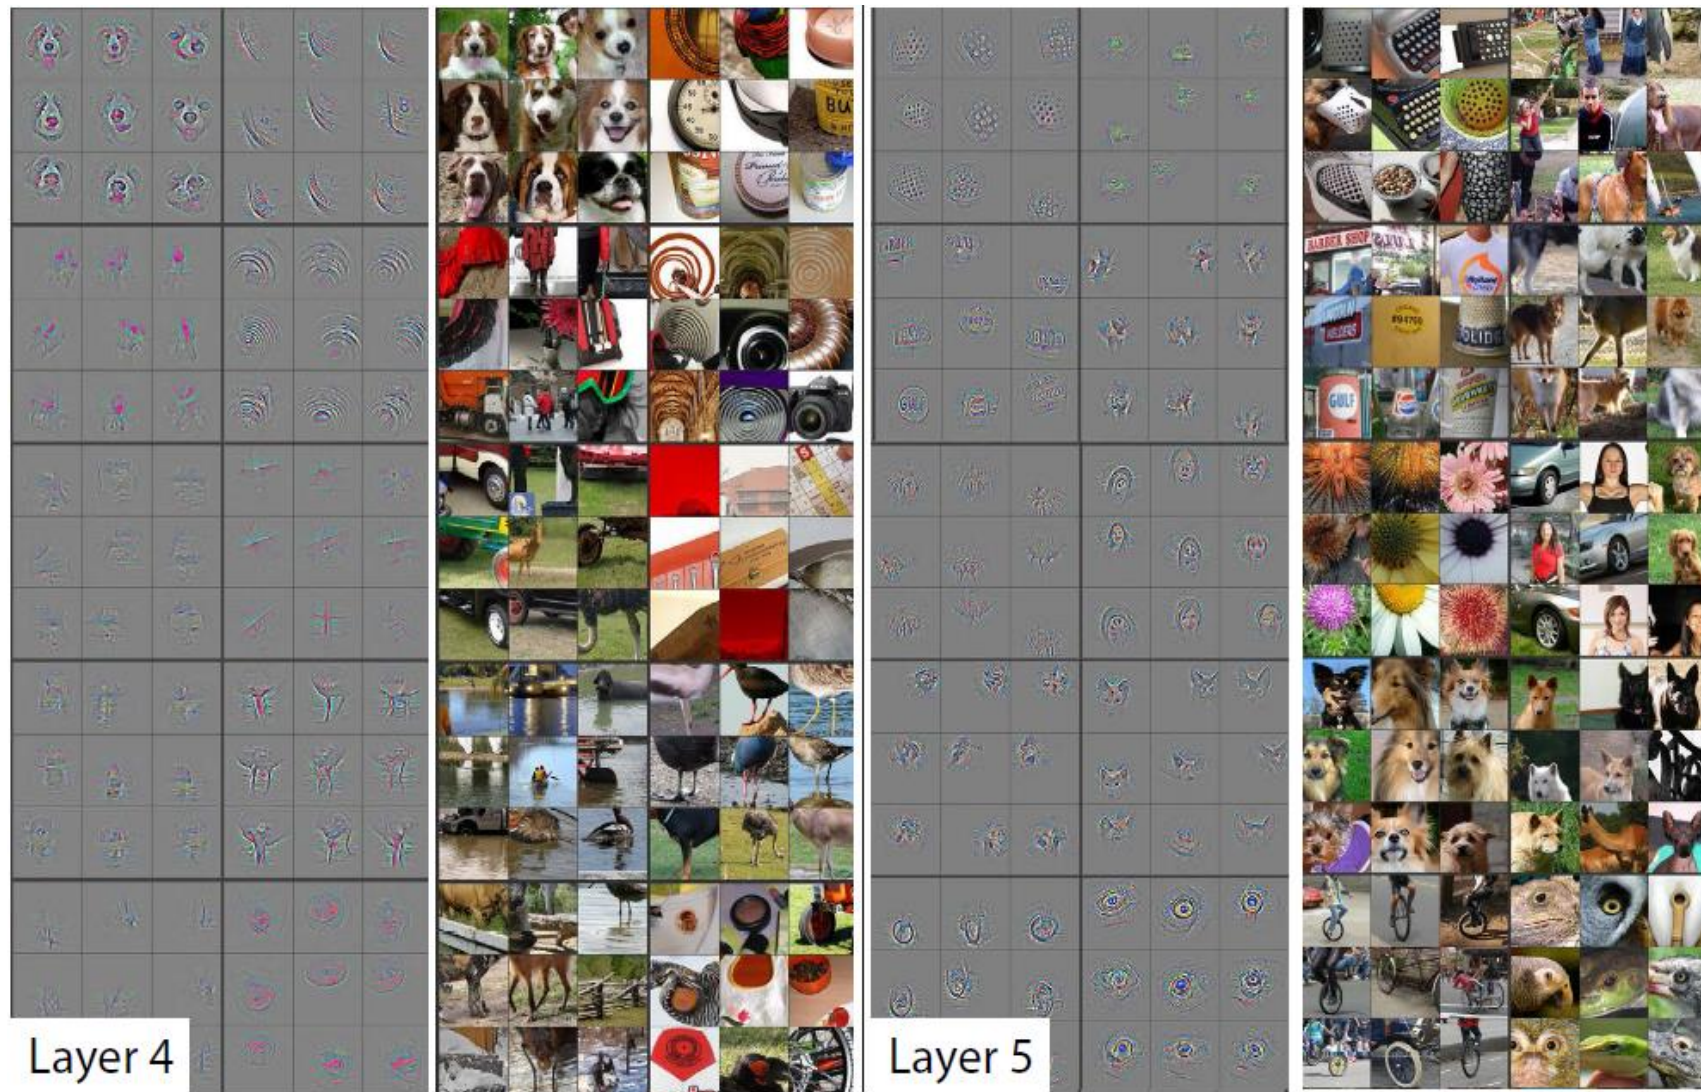
## DeconvNet



# AlexNet Filters- DeconvNet



# AlexNet Filters- DeconvNet

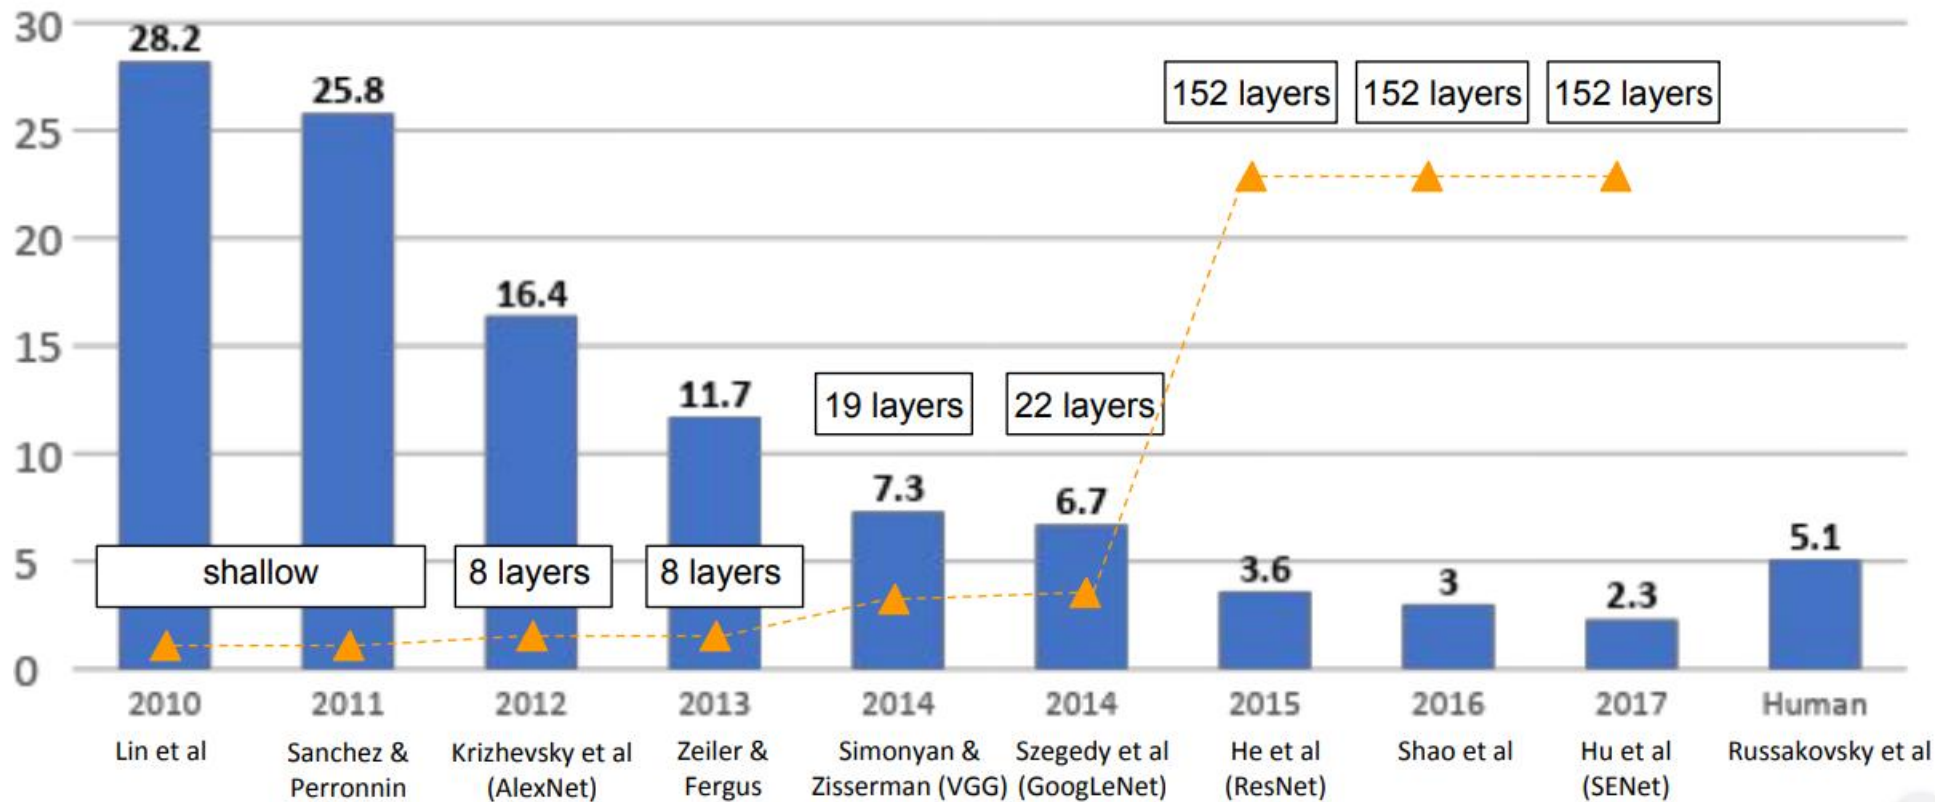


# *Beyond AlexNet: The DCNN storm*



# ImageNet: the full story

## ImageNet Large Scale Visual Recognition Challenge (ILSVRC) winners



***The challenge is now over***

Image from  
[[http://cs231n.stanford.edu/slides/2019/cs231n\\_2019\\_lecture09.pdf](http://cs231n.stanford.edu/slides/2019/cs231n_2019_lecture09.pdf)]

# VGG Architecture

Several variants

Only 3x3 convolutional filters used (each with ReLU)

LRN used in only one variant

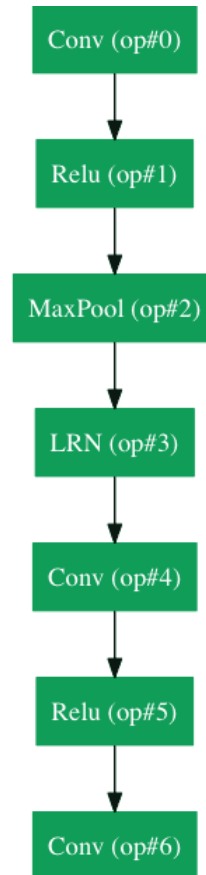
Image from [https://arxiv.org/pdf/1409.1556.pdf]

ConvNet Configuration					
A	A-LRN	B	C	D	E
11 weight layers	11 weight layers	13 weight layers	16 weight layers	16 weight layers	19 weight layers
input (224 × 224 RGB image)					
conv3-64	conv3-64 <b>LRN</b>	conv3-64 <b>conv3-64</b>	conv3-64 conv3-64	conv3-64 conv3-64	conv3-64 conv3-64
maxpool					
conv3-128	conv3-128	conv3-128 <b>conv3-128</b>	conv3-128 conv3-128	conv3-128 conv3-128	conv3-128 conv3-128
maxpool					
conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256 <b>conv1-256</b>	conv3-256 conv3-256 <b>conv3-256</b>	conv3-256 conv3-256 conv3-256 <b>conv3-256</b>
maxpool					
conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512 <b>conv1-512</b>	conv3-512 conv3-512 <b>conv3-512</b>	conv3-512 conv3-512 conv3-512 <b>conv3-512</b>
maxpool					
conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512 <b>conv1-512</b>	conv3-512 conv3-512 <b>conv3-512</b>	conv3-512 conv3-512 conv3-512 <b>conv3-512</b>
maxpool					
FC-4096					
FC-4096					
FC-1000					
soft-max					

# Inception Architecture

- The ImageNet Large Scale Visual Recognition Challenge

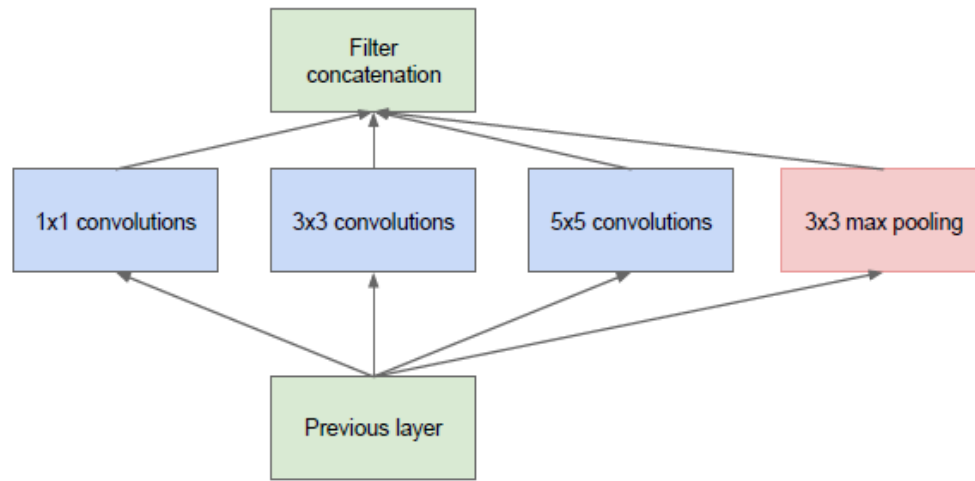
*How deep is a deep neural network, for a task like this?*



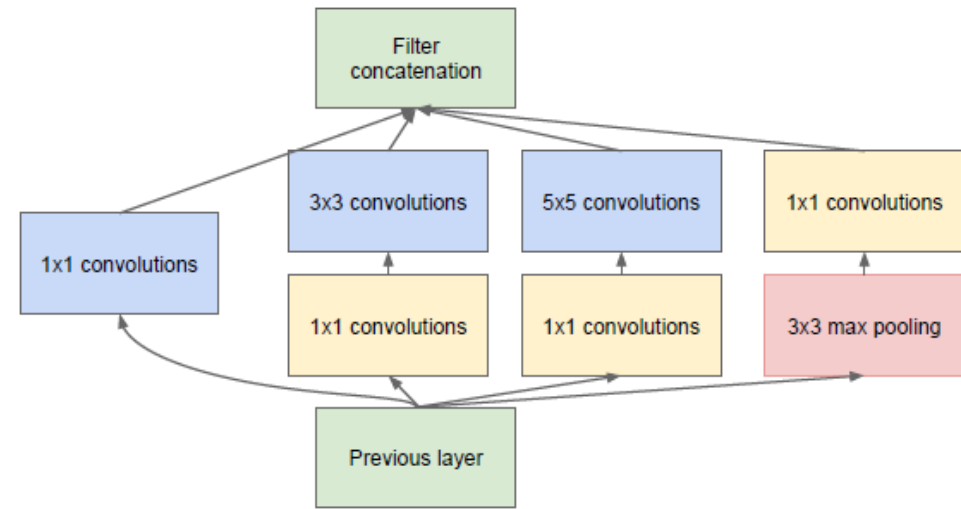
*GoogLeNet (Inception v4) winner of two out of three categories in 2014: 154 network layers*

# Inception Architecture

- Inception modules



(a) Inception module, naïve version

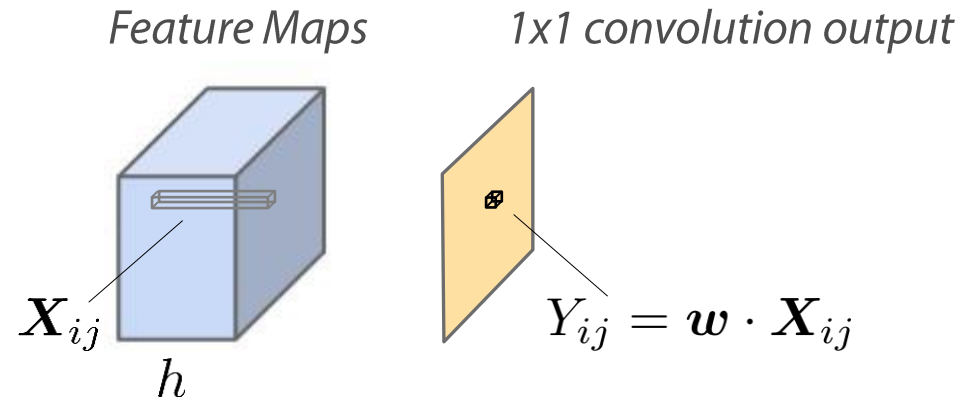


(b) Inception module with dimension reductions

Image from [<https://arxiv.org/pdf/1409.4842.pdf>]

# Inception Architecture

- 1x1 convolution?



*(It is a kind of misnomer)*

Each filter has dimension  $1 \times 1 \times h$   
where  $h$  is the depth of the set of filter maps

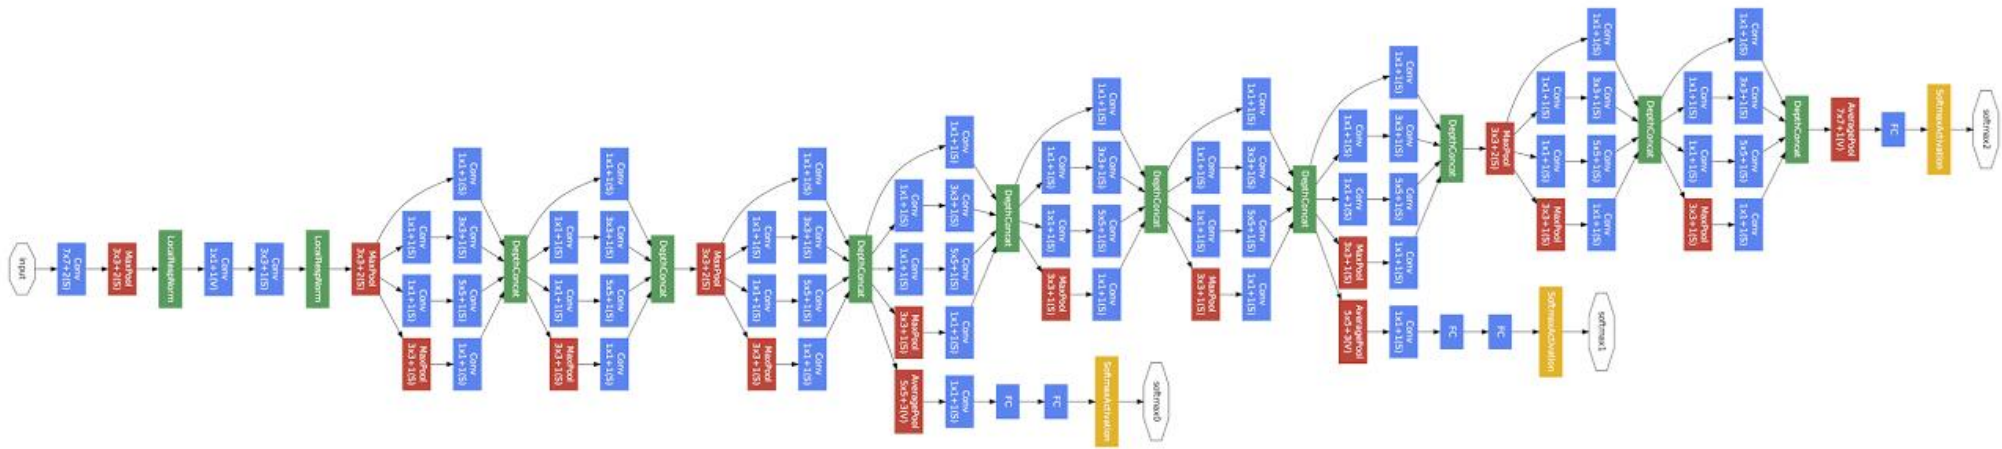
Using  $d$  1x1 convolution filters allows changing depth  $h$  into  $d$

Clearly the assumption is  $d < h$

*It mimics a fully connected layer (across channels)*

# Inception Architecture

- GoogLeNet architecture



Convulsive  
Max Pool  
Softmax  
Filter Concat

Image from [https://arxiv.org/pdf/1409.4842.pdf]

# Inception Architecture

## GoogLeNet architecture

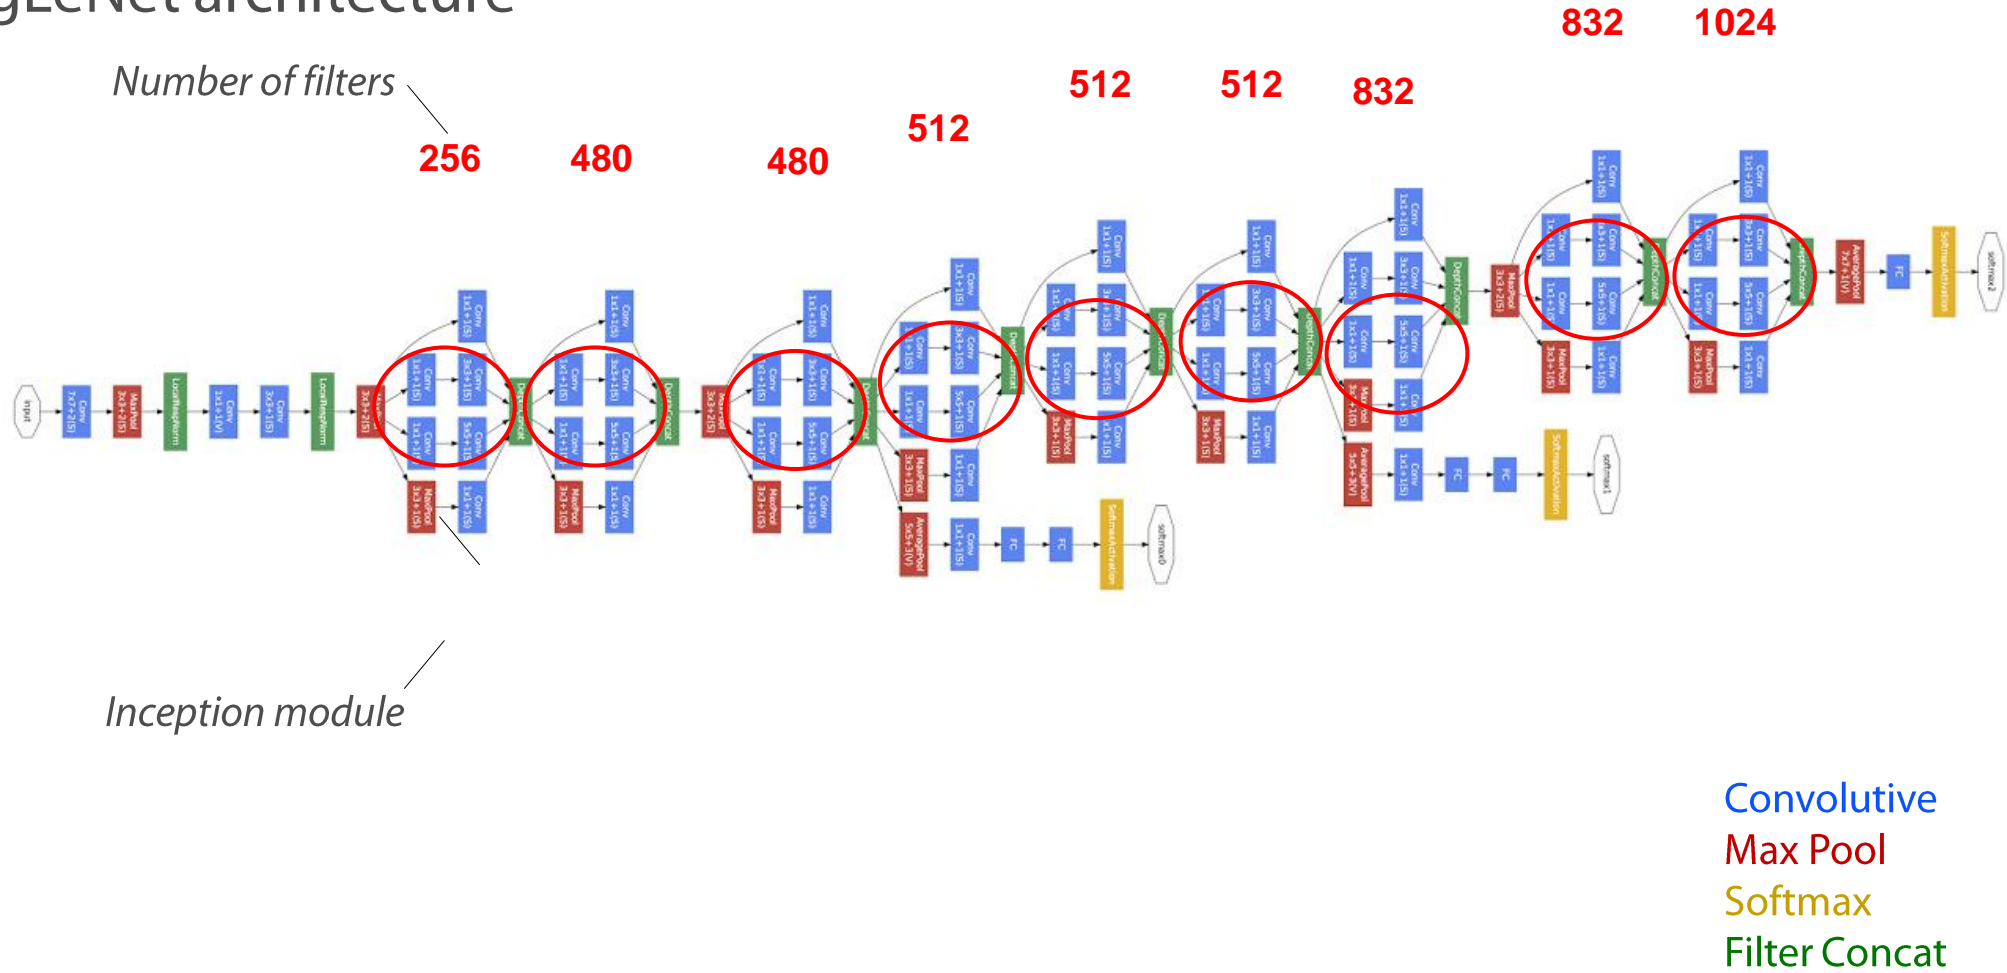
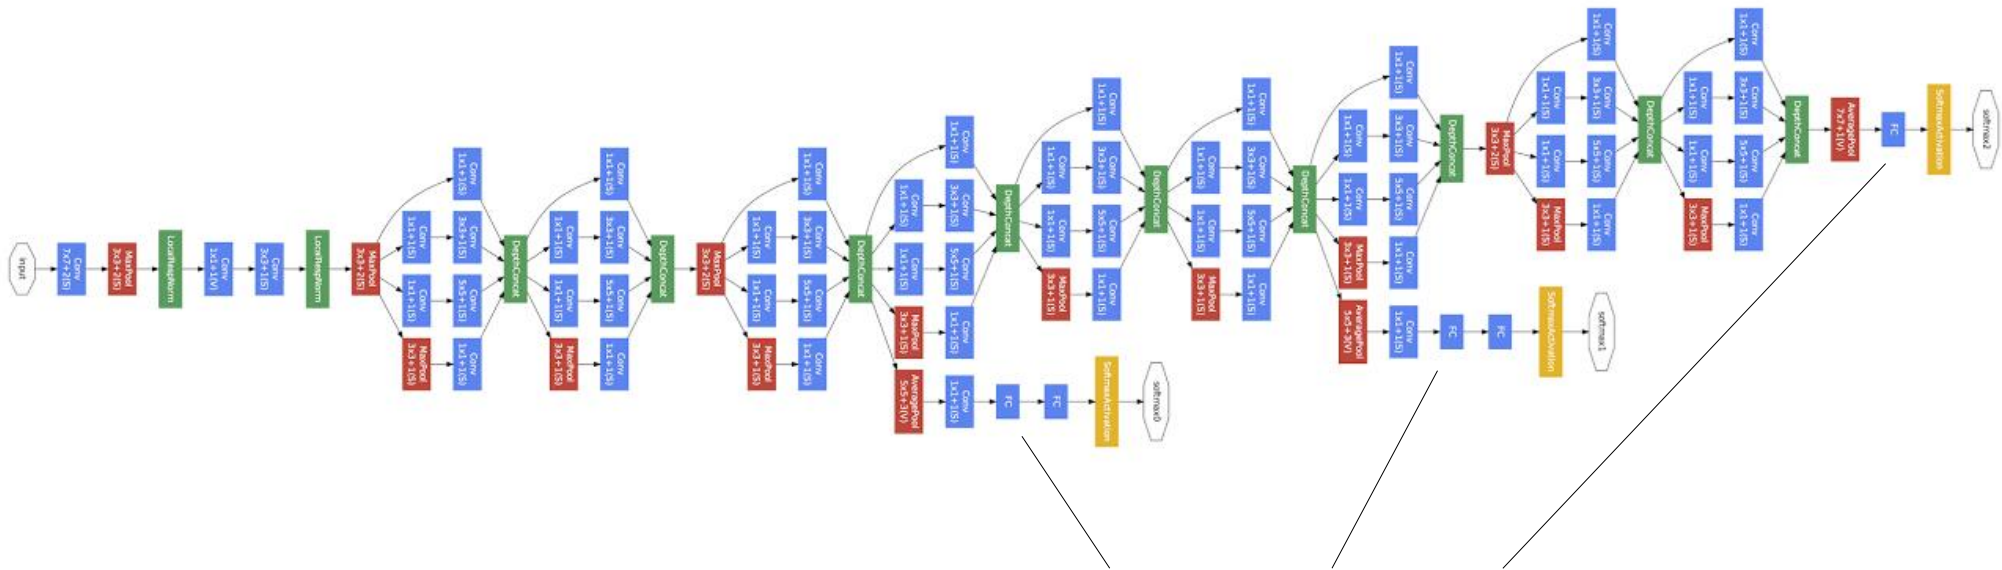


Image from [https://arxiv.org/pdf/1409.4842.pdf]

# Inception Architecture

## ■ GoggleNet architecture



Much smaller FC layers

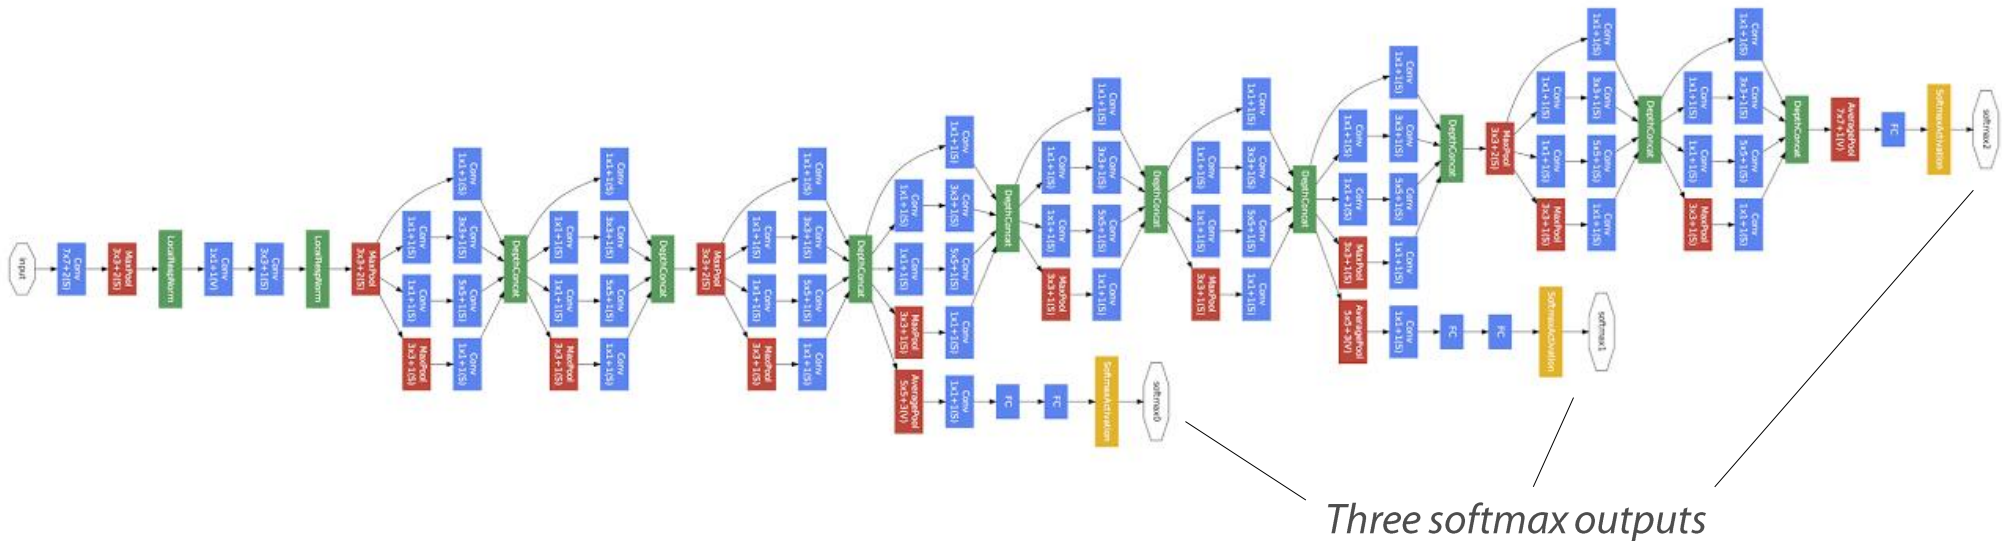
Convulsive  
Max Pool  
Softmax  
Filter Concat

Image from [<https://arxiv.org/pdf/1409.4842.pdf>]



# Inception Architecture

- GoogLeNet architecture



They are trained to produce the same output, simultaneously

Convolutional  
Max Pool  
Softmax  
Filter Concat

Image from [<https://arxiv.org/pdf/1409.4842.pdf>]

# ResNet Architecture

- ResNet block

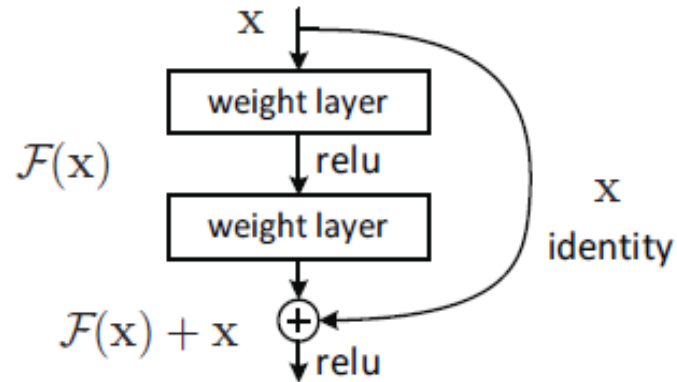


Figure 2. Residual learning: a building block.

Image from [<https://arxiv.org/pdf/1512.03385.pdf>]

# ResNet Architecture

- ResNet architecture

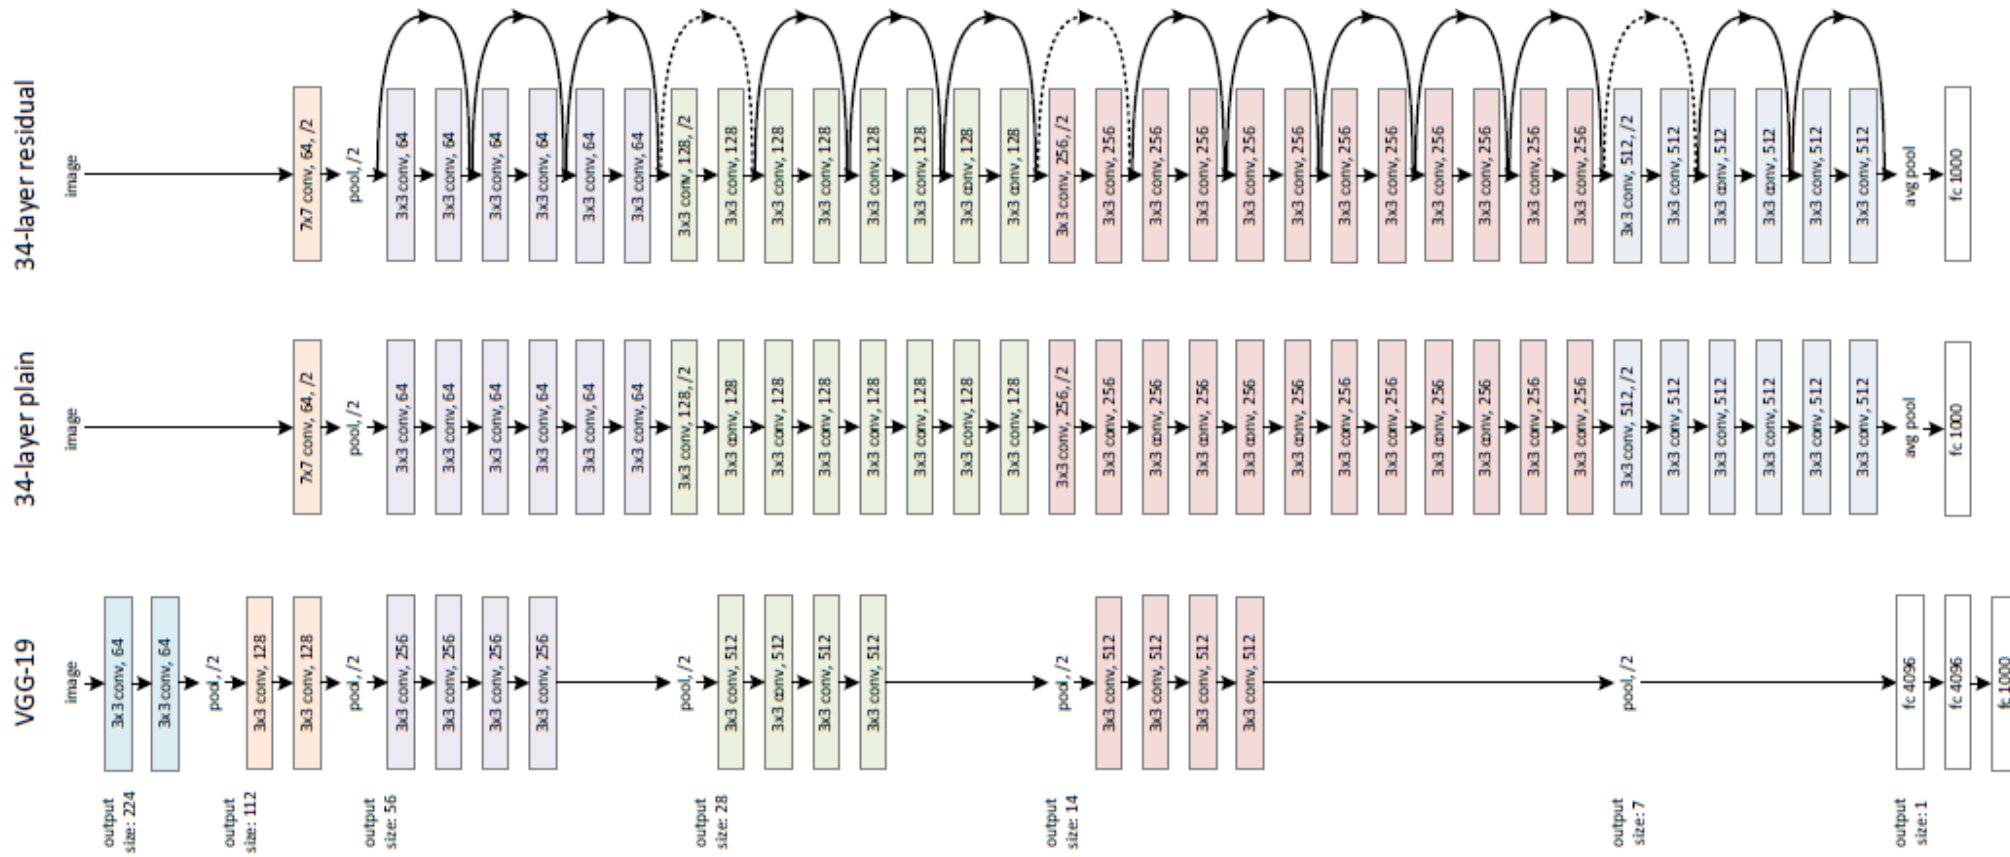


Image from [<https://arxiv.org/pdf/1512.03385.pdf>]

# Comparing Different DCNNs

## ■ Comparative charts at Top-1 accuracy

*i.e. how often the DCNN is right with ImageNet with its top prediction*

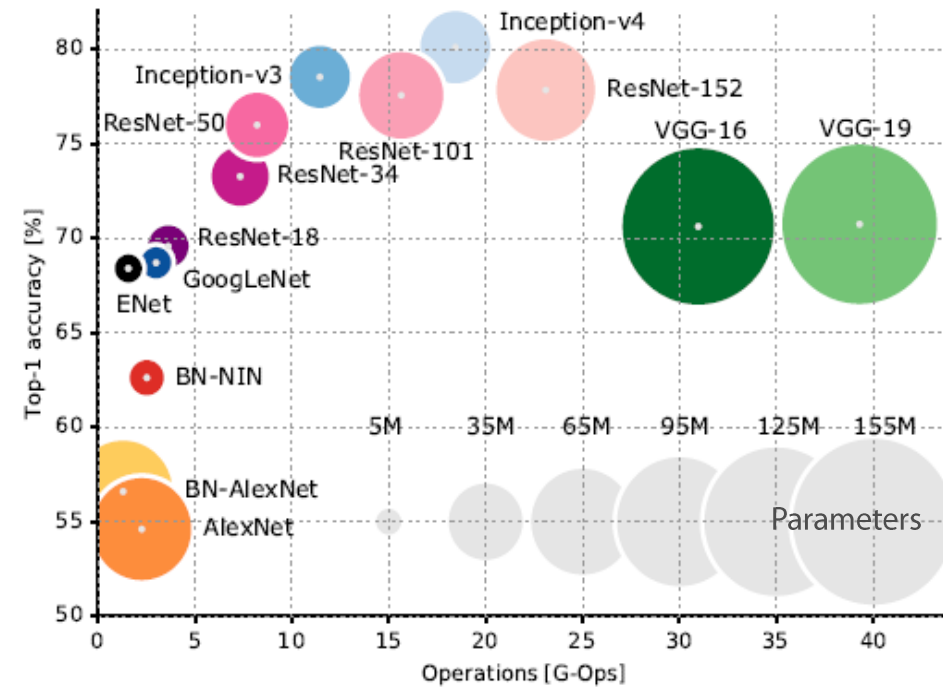
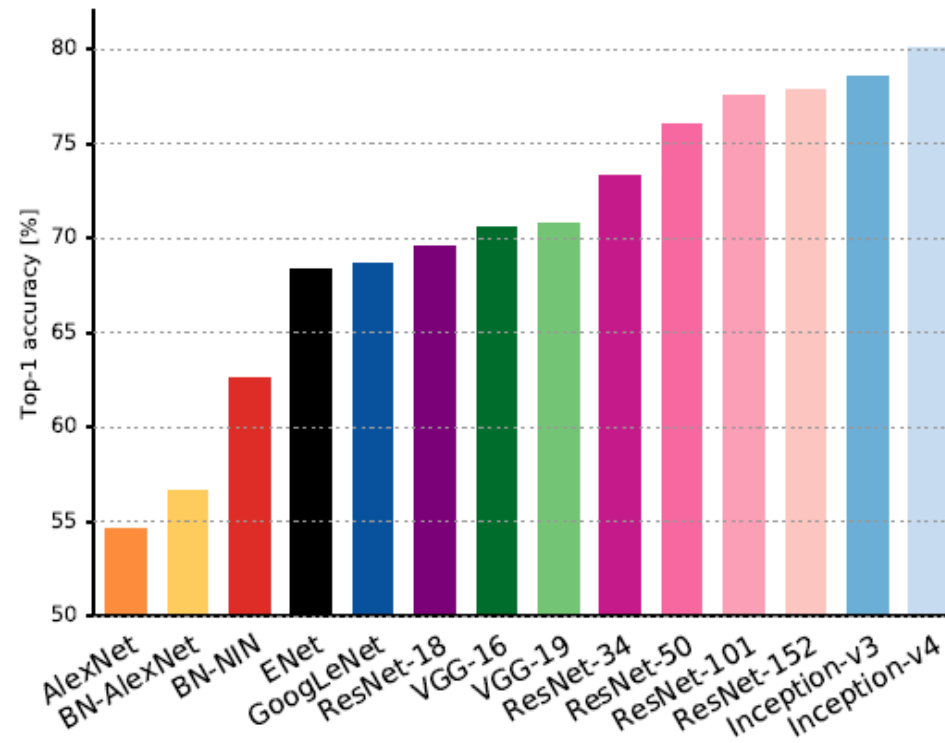


Image from [<https://arxiv.org/abs/1605.07678>, 2017]

# Comparing Different DCNNs

(Same chart, a more recent version)

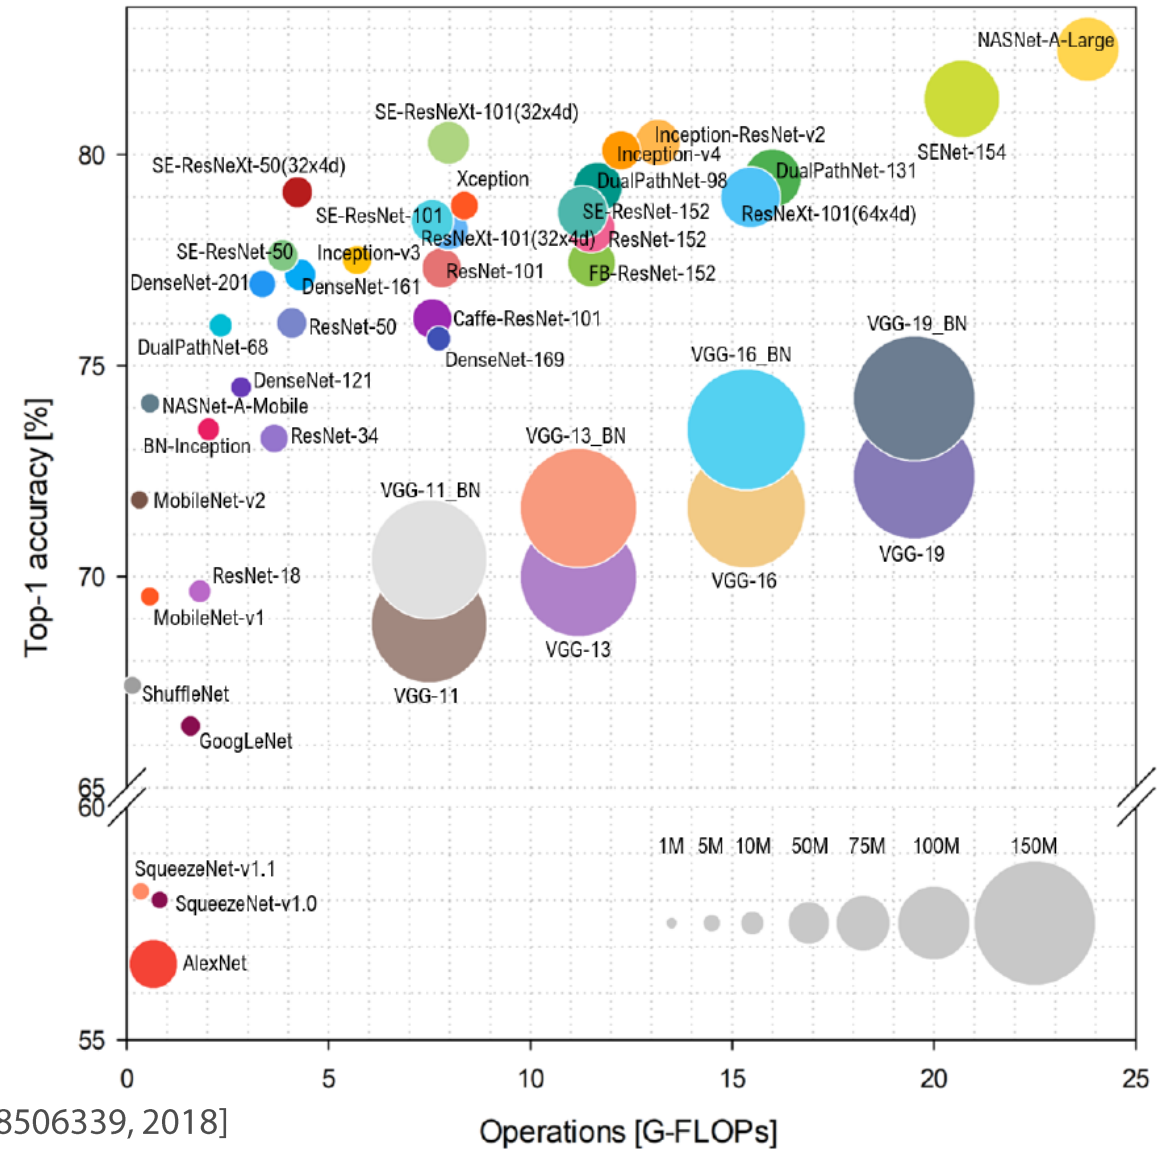
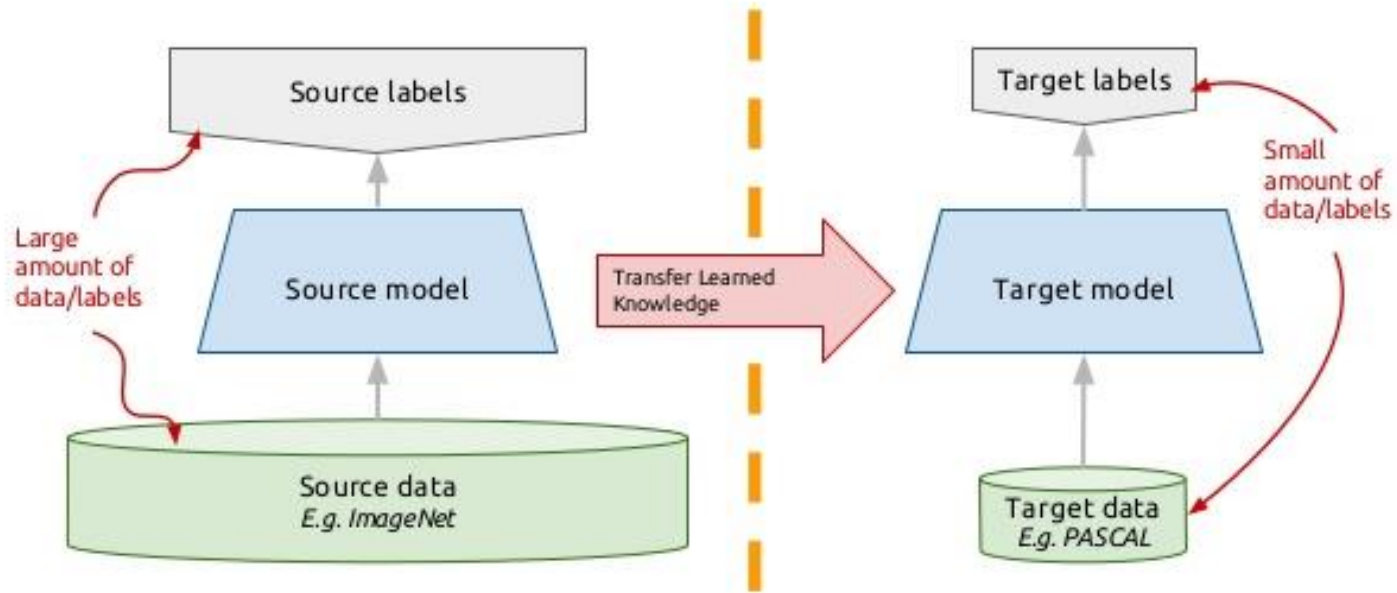


Image from [<https://ieeexplore.ieee.org/document/8506339>, 2018]

# *Transfer Learning*

# Transfer Learning

## Transfer learning: idea



*Do DCNNs Dream  
of Electric Sheep?*



# Can DCNNs 'dream'?

sign in | become a supporter | subscribe | search

jobs | dating | more | International edition

# theguardian

UK | world | sport | football | opinion | culture | business | lifestyle | fashion | environment | tech | travel

home > tech

## Artificial intelligence (AI) | Yes, androids do dream of electric sheep

Google sets up feedback loop in its image recognition neural network - which looks for patterns in pictures - creating hallucinatory images of animals, buildings and landscapes which veer from beautiful to terrifying

This article is 1 year old

109,591 | 445

Alex Hern | @alexhern

Thursday 18 June 2015 12.57 BST



A hallucinatory filter over a red tree. Spot the animals. Photograph: Google

Advertisement

Office 365 Business

Passa a Outlook 2016 e accedi alla posta ovunque, online e offline.

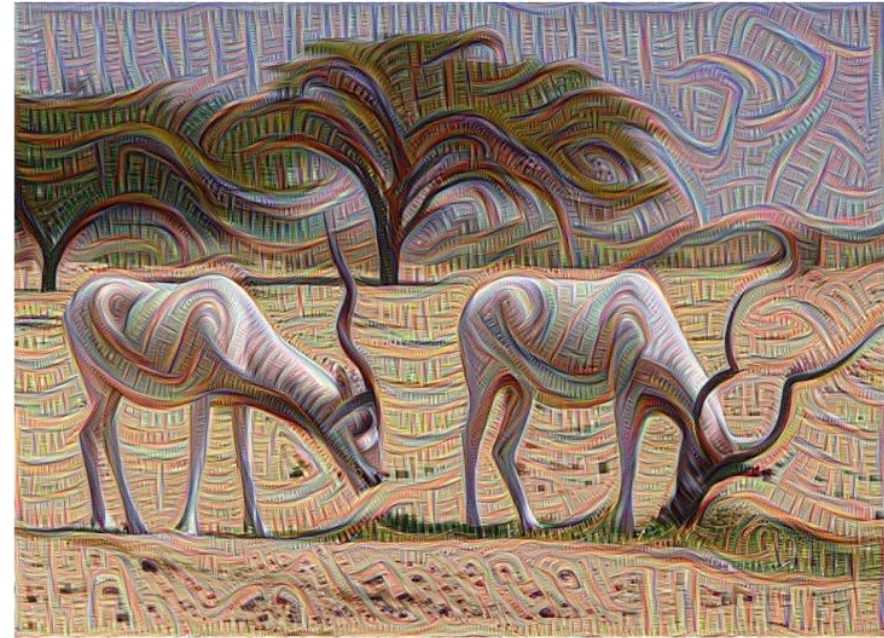
A soli €10,70 al mese\*

\*Il prezzo non include l'IVA

Microsoft | Scopri di più

# Can DCNNs 'dream'?

Enhancing lower layers



[images from <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>]

# Feature Enhancement

## ■ Image Space Gradient Descent

Define

$$\Phi_{k,l}(\mathbf{I})$$

as the response of a DCNN at a layer  $k$ , filter  $l$  to an image  $\mathbf{I}$

Given a specific image  $\hat{\mathbf{I}}$ , we define the loss function

$$L(\hat{\mathbf{I}}, \mathbf{I}) := \|\gamma \Phi_{k,l}(\hat{\mathbf{I}}) - \Phi_{k,l}(\mathbf{I})\|^2$$

The optimization problem *Amplification factor*

$$\mathbf{I}^* := \operatorname{argmin}_{\mathbf{I}} \left( L(\hat{\mathbf{I}}, \mathbf{I}) + \lambda \|\mathbf{I}\|^2 \right)$$

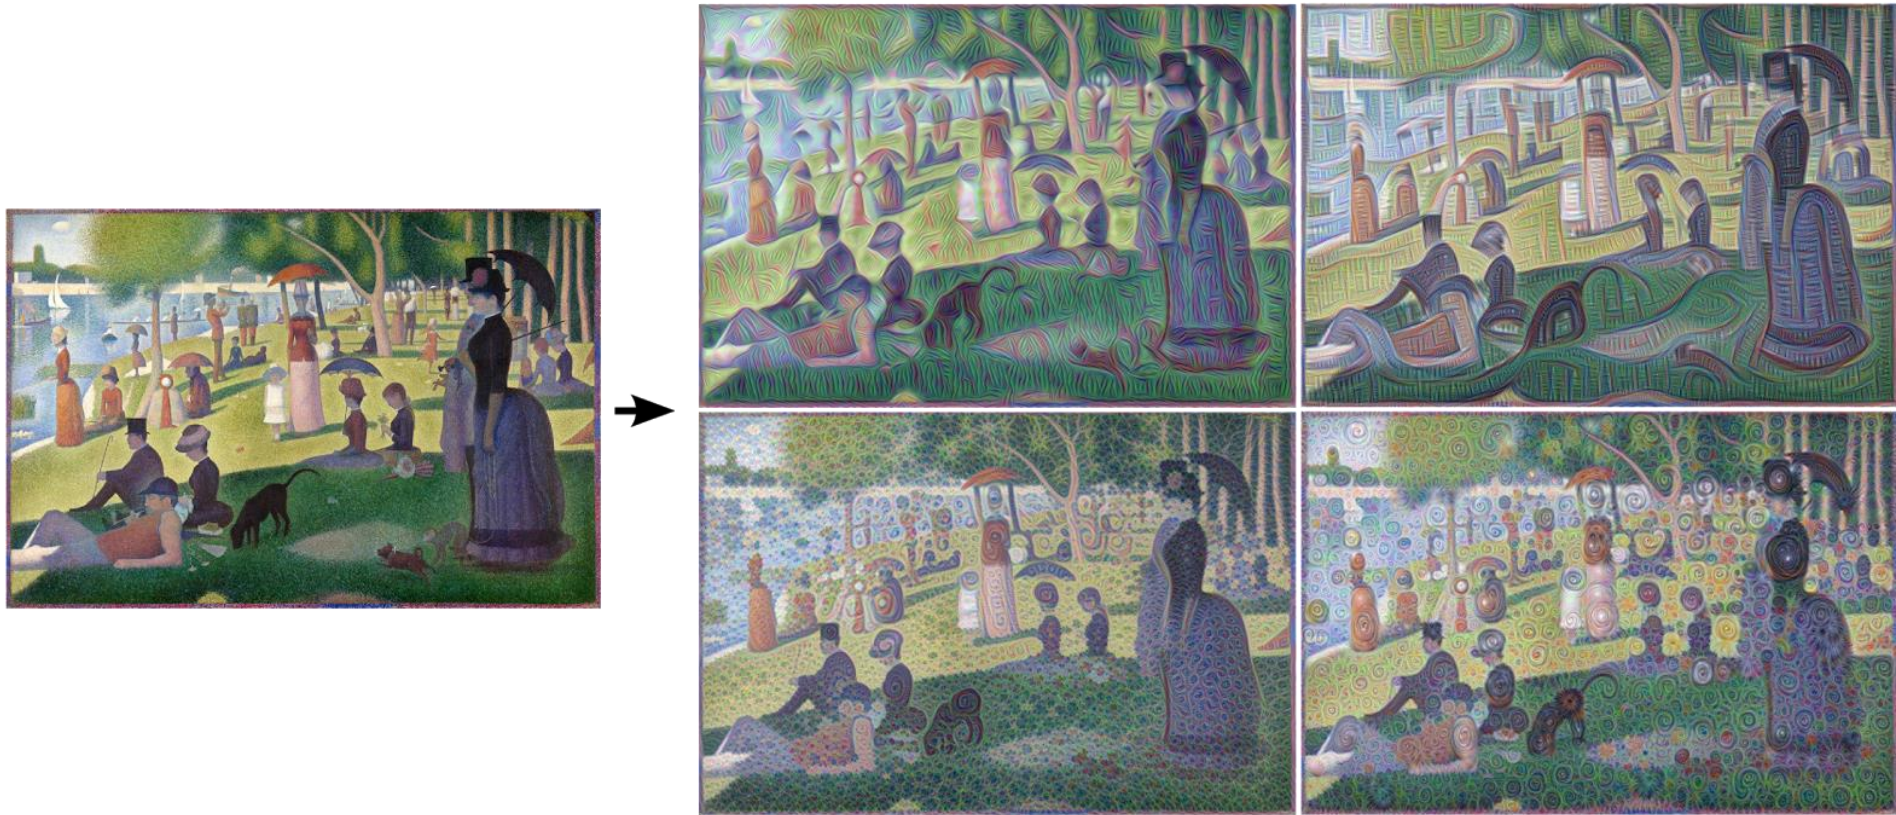
is solved via gradient descent by computing

$$\frac{\partial}{\partial \mathbf{I}} \left( L(\hat{\mathbf{I}}, \mathbf{I}) + \lambda \|\mathbf{I}\|^2 \right)$$

and starting from  $\mathbf{I}^{(0)} = \hat{\mathbf{I}}$

# Can DCNNs 'dream'?

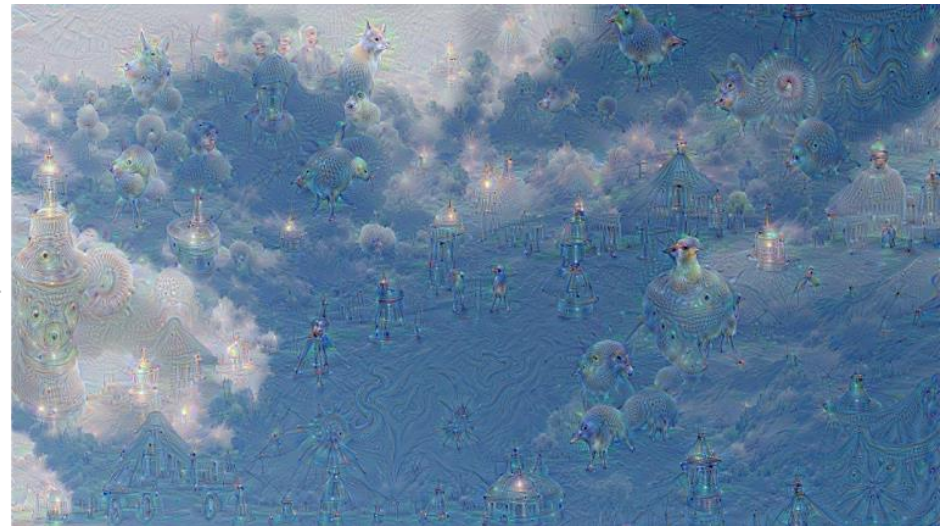
Enhancing lower layers



[images from <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>]

# Can DCNNs 'dream'?

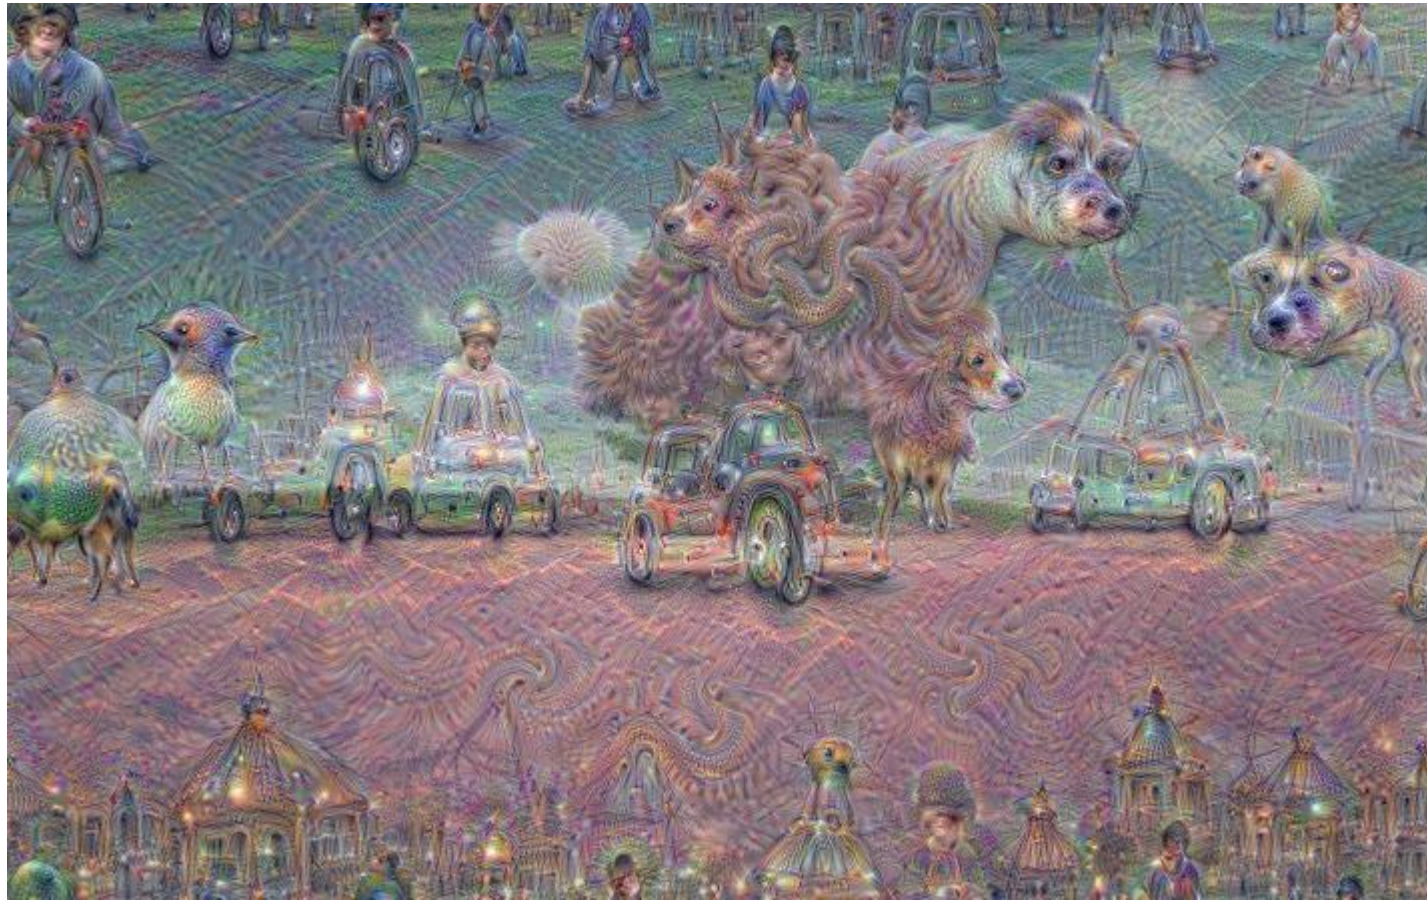
Enhancing upper layers



[images from <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>]

# Can DCNNs 'dream'?

Letting the DCNN go on its own



[images from <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>]

# *Can DCNNs 'dream'?*

*Letting the DCNN go on its own*



[images from <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>]

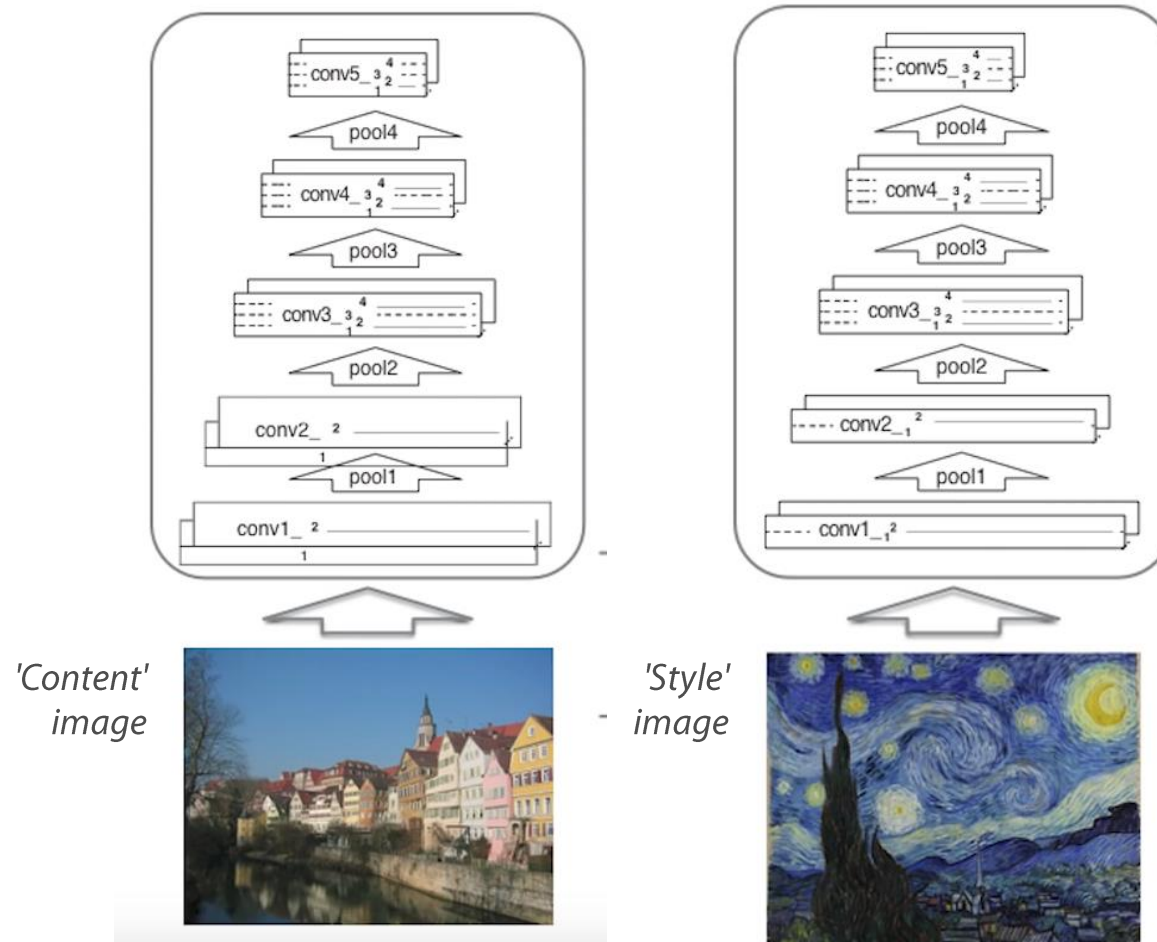
# *The Power of Abstraction (in layers)*



# The Power of Abstraction

## ■ Different Layers of a Deep Convolutional Neural Network

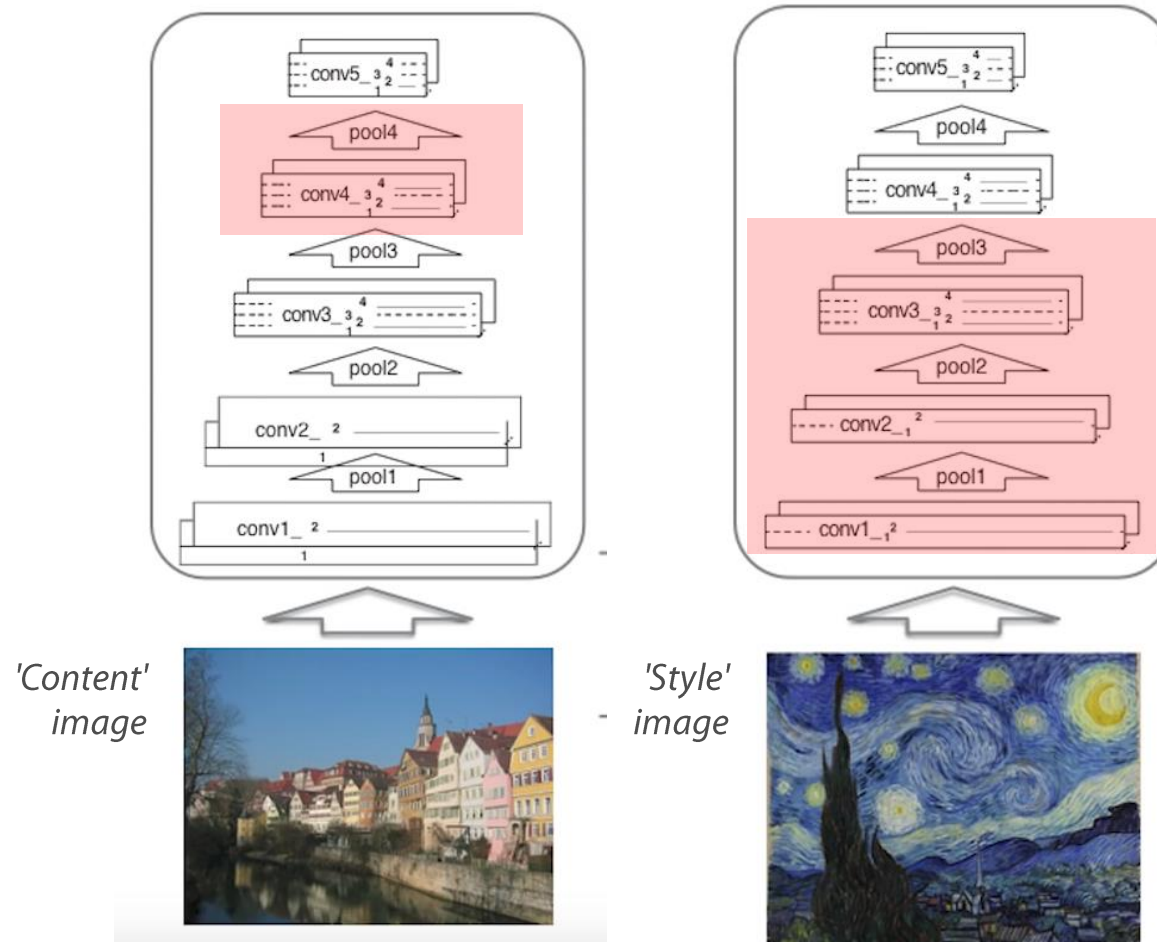
What kind of information does each layer 'store'?



# The Power of Abstraction

## ■ Different Layers of a Deep Convolutional Neural Network

What kind of information does each layer 'store'?



Create a new image  
by combining more  
of the 'Content' top layer  
and more of 'Style' low layers

# Mixing Two Images

## ■ Image Space Gradient Descent

Define

$$\Phi_{k,l}(\mathbf{I})$$

as the response of a DCNN at a layer  $k$ , filter  $l$  to an image  $\mathbf{I}$

Given a specific image  $\hat{\mathbf{I}}_1$  and  $\hat{\mathbf{I}}_2$ , we define the loss function

$$L(\hat{\mathbf{I}}, \mathbf{I}) := \sum_{k,l} \left\| \underbrace{M_{k,l}(\Phi_{k,l}(\hat{\mathbf{I}}_2), \Phi_{k,l}(\hat{\mathbf{I}}_1))}_{\text{Weighted Merge Function}} - \Phi_{k,l}(\mathbf{I}) \right\|^2$$

The optimization problem

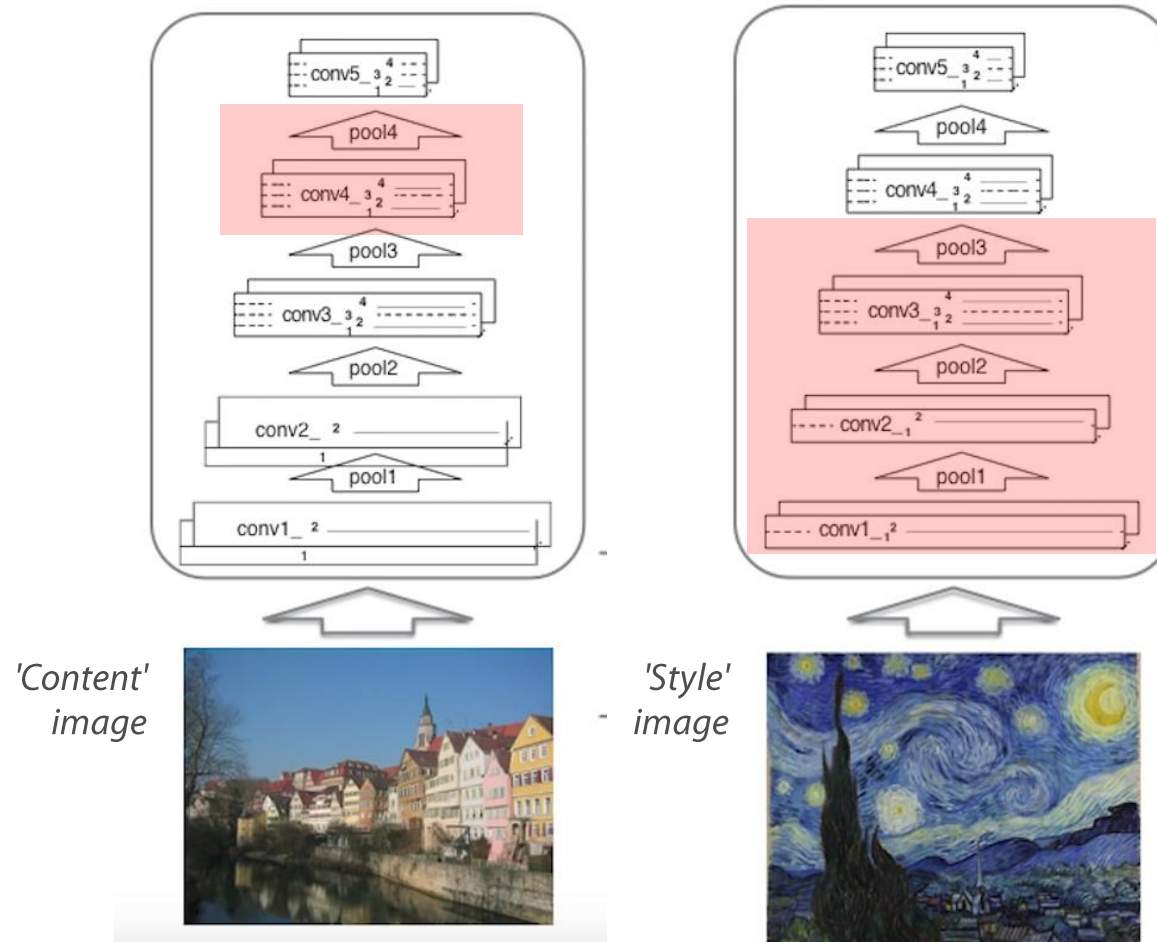
$$\mathbf{I}^* := \operatorname{argmin}_{\mathbf{I}} \left( L(\hat{\mathbf{I}}, \mathbf{I}) + \lambda \|\mathbf{I}\|^2 \right)$$

is solved via gradient descent starting from  $\mathbf{I}^{(0)} = \hat{\mathbf{I}}_1$

# The Power of Abstraction

## ■ Different Layers of a Deep Convolutional Neural Network

What kind of information does each layer 'store'?



Create a new image by combining more of the 'Content' top layer and more of 'Style' low layers

This is the result



# The Power of Abstraction

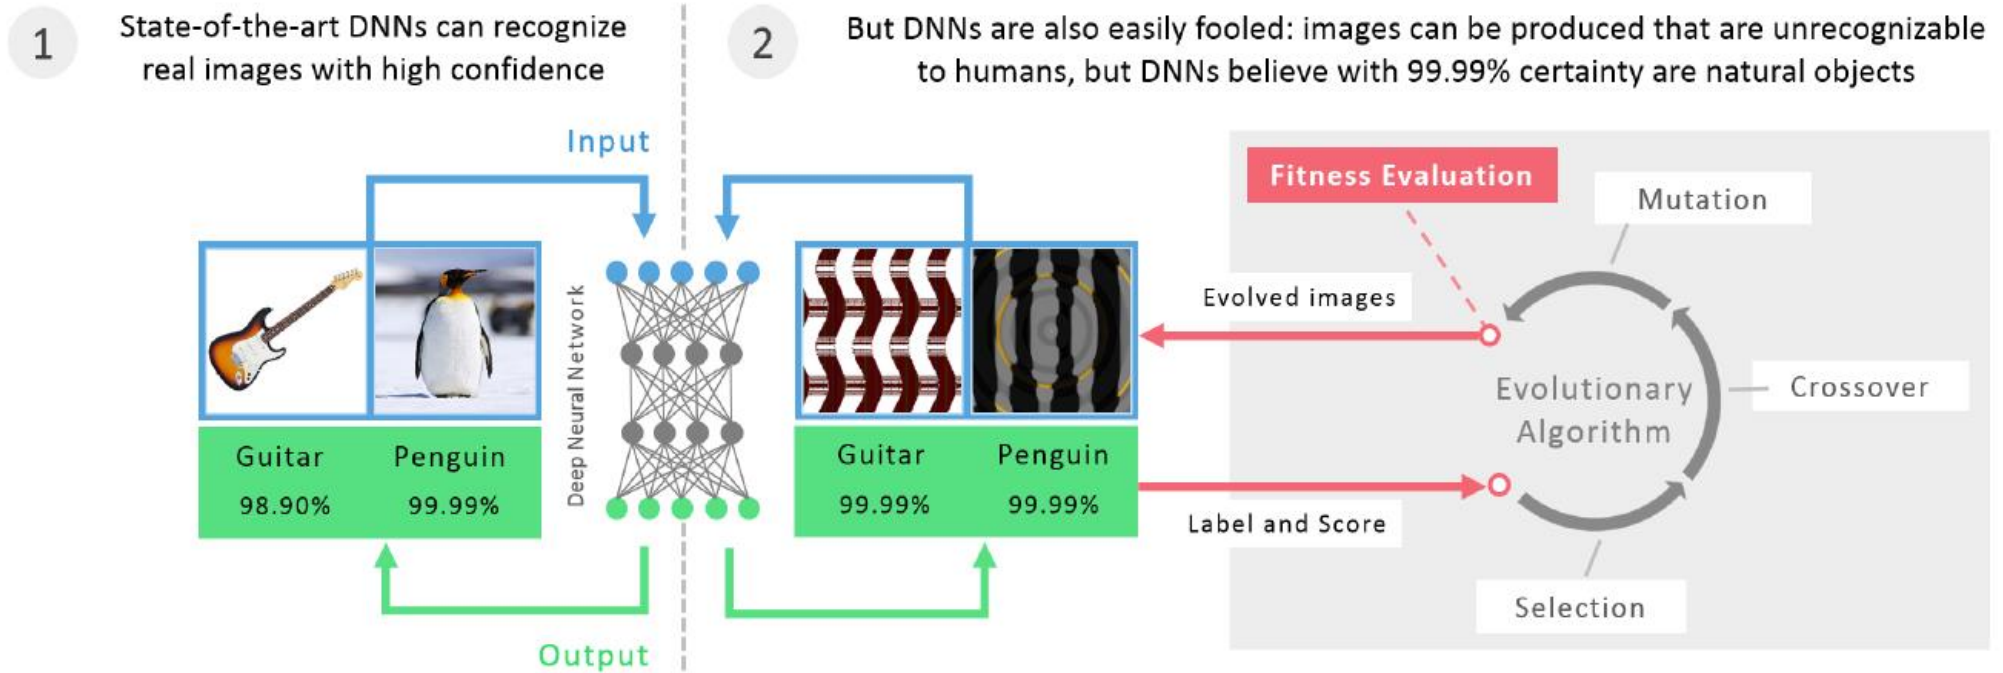
## ■ Different Layers of a Deep Convolutional Neural Network

Further examples:

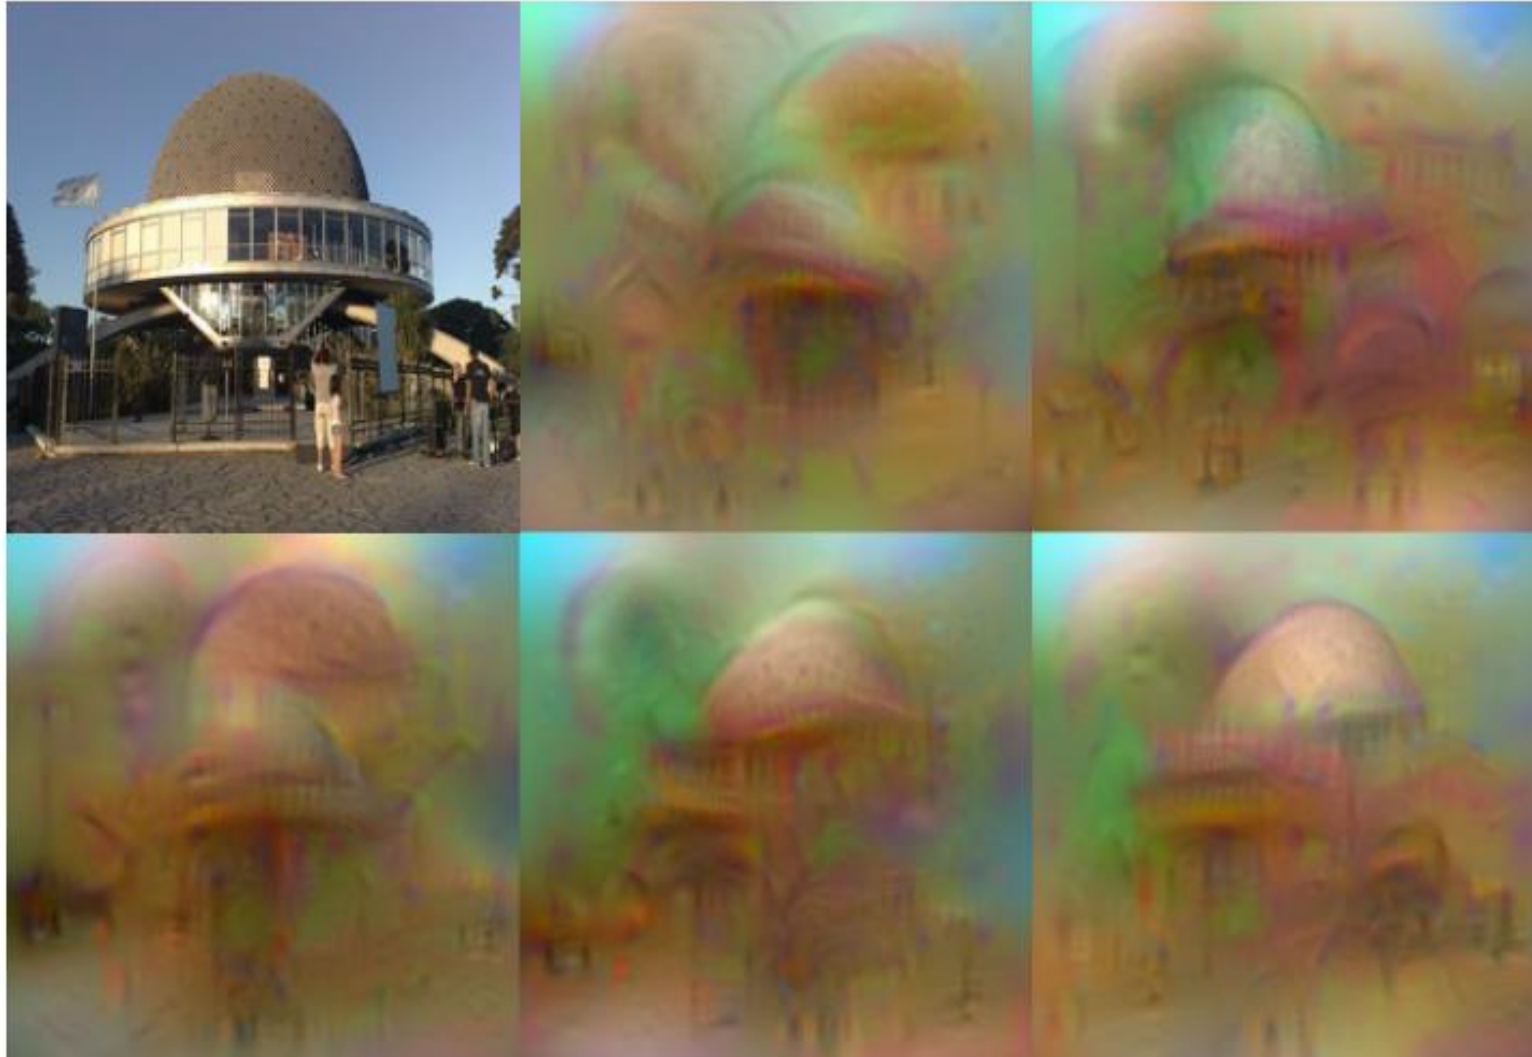


*Human-like Vision?*  
*No way*

# A DCNN can be fooled...



# Reconstructing Images from Feature Maps





# Reconstructing Images from Feature Maps

## ■ Image Space Gradient Descent

Define

$$\Phi_{k,l}(\mathbf{I})$$

as the response of a DCNN at a layer  $k$ , filter  $l$  to an image  $\mathbf{I}$

Given a specific image  $\hat{\mathbf{I}}$ , we define the loss function

$$L(\hat{\mathbf{I}}, \mathbf{I}) := \|\Phi_{k,l}(\hat{\mathbf{I}}) - \Phi_{k,l}(\mathbf{I})\|^2$$

and the optimization problem

$$\mathbf{I}^* := \operatorname{argmin}_{\mathbf{I}} \left( L(\hat{\mathbf{I}}, \mathbf{I}) + \rho P(\mathbf{I}) + \lambda \|\mathbf{I}\|^2 \right)$$

*L2 Regularization*

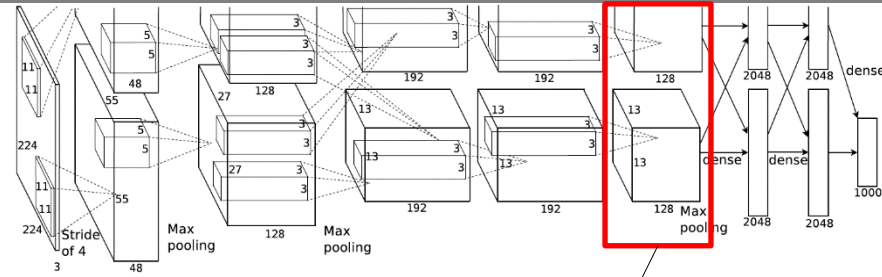
To solve this, we can compute

$$\frac{\partial}{\partial \mathbf{I}} \left( L(\hat{\mathbf{I}}, \mathbf{I}) + \rho P(\mathbf{I}) + \lambda \|\mathbf{I}\|^2 \right)$$

*'Statistical Realism'*

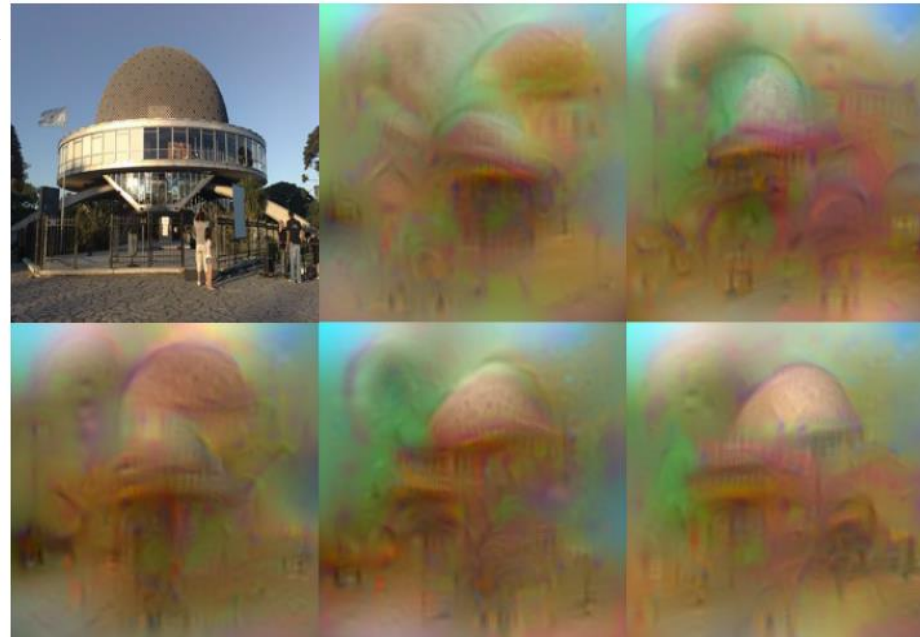
and apply a gradient descent procedure, starting from a random image  $\mathbf{I}^{(0)}$

# Reconstructing Images from Feature Maps



$\Phi_{k,l}(\hat{\mathbf{I}})$  is taken here

This is  $\hat{\mathbf{I}}$



The remaining five images were generated using image space gradient descent with different initial images  $\mathbf{I}^{(0)}$

# Just add some little noise ...

nature

Subscribe

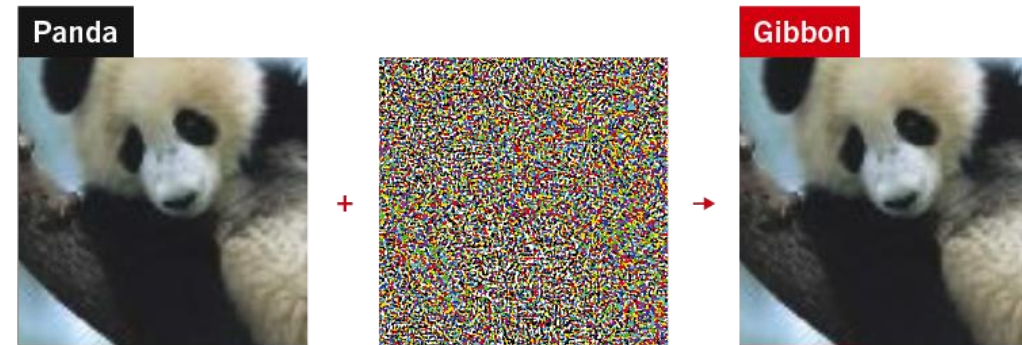
NEWS FEATURE · 09 OCTOBER 2019

## Why deep-learning AIs are so easy to fool

Artificial-intelligence researchers are trying to fix the flaws of neural networks.

### PERCEPTION PROBLEMS

Adding carefully crafted noise to a picture can create a new image that people would see as identical, but which a DNN sees as utterly different.



In this way, any starting image can be tweaked so a DNN misclassifies it as any target image a researcher chooses.



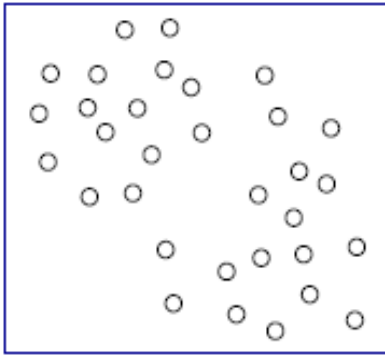
©nature

*No Free Lunch:  
creating an annotated dataset*

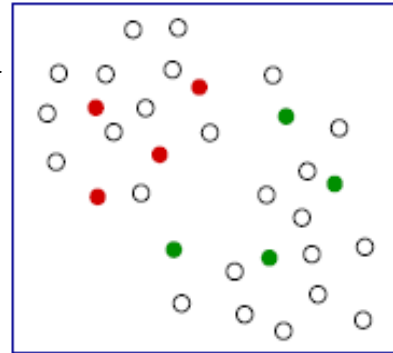
# Active Learning

When the network decides which annotations should be made

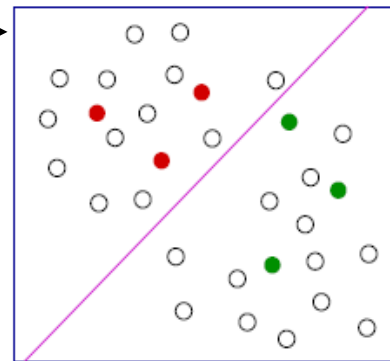
1) Consider a large non-annotated *dataset*



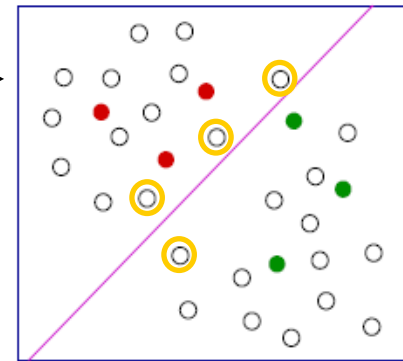
2) Annotate a few images at random



3) Train a classifier on annotated images only



4) Annotate borderline cases

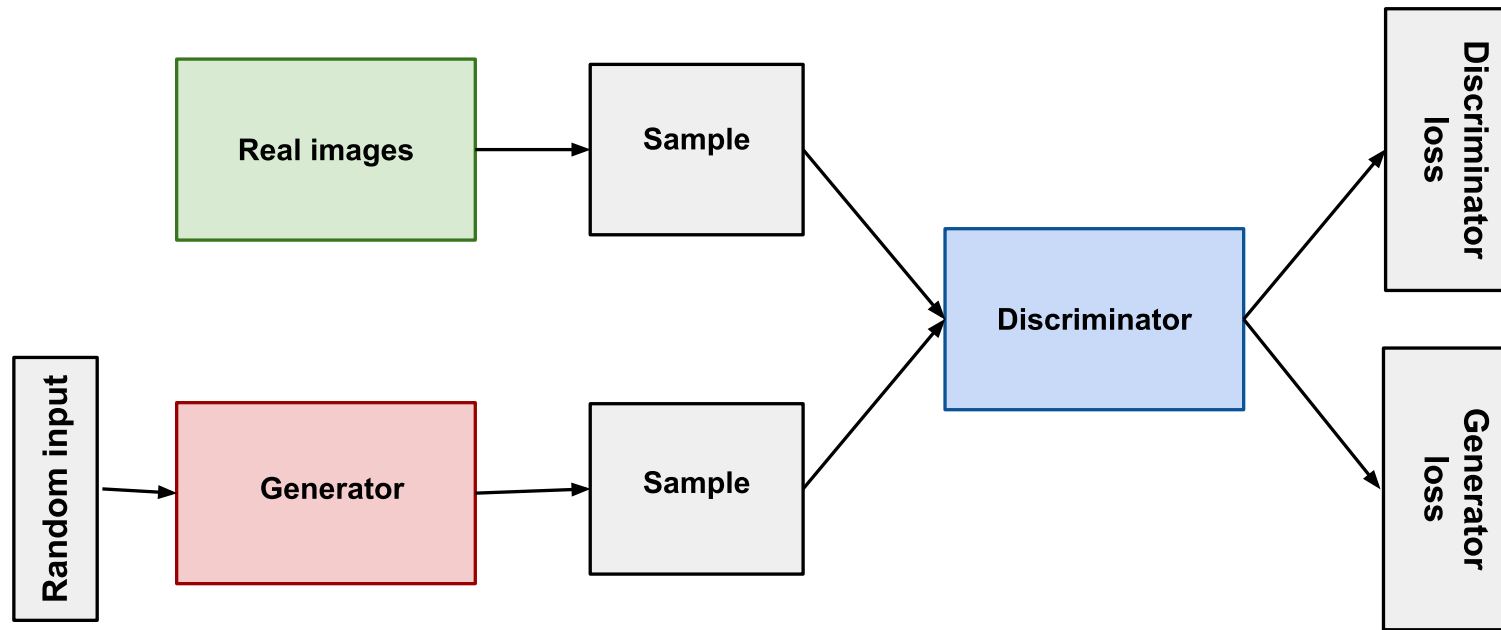


5) Repeat training

# Generative Adversarial Network

## ■ Two competing networks

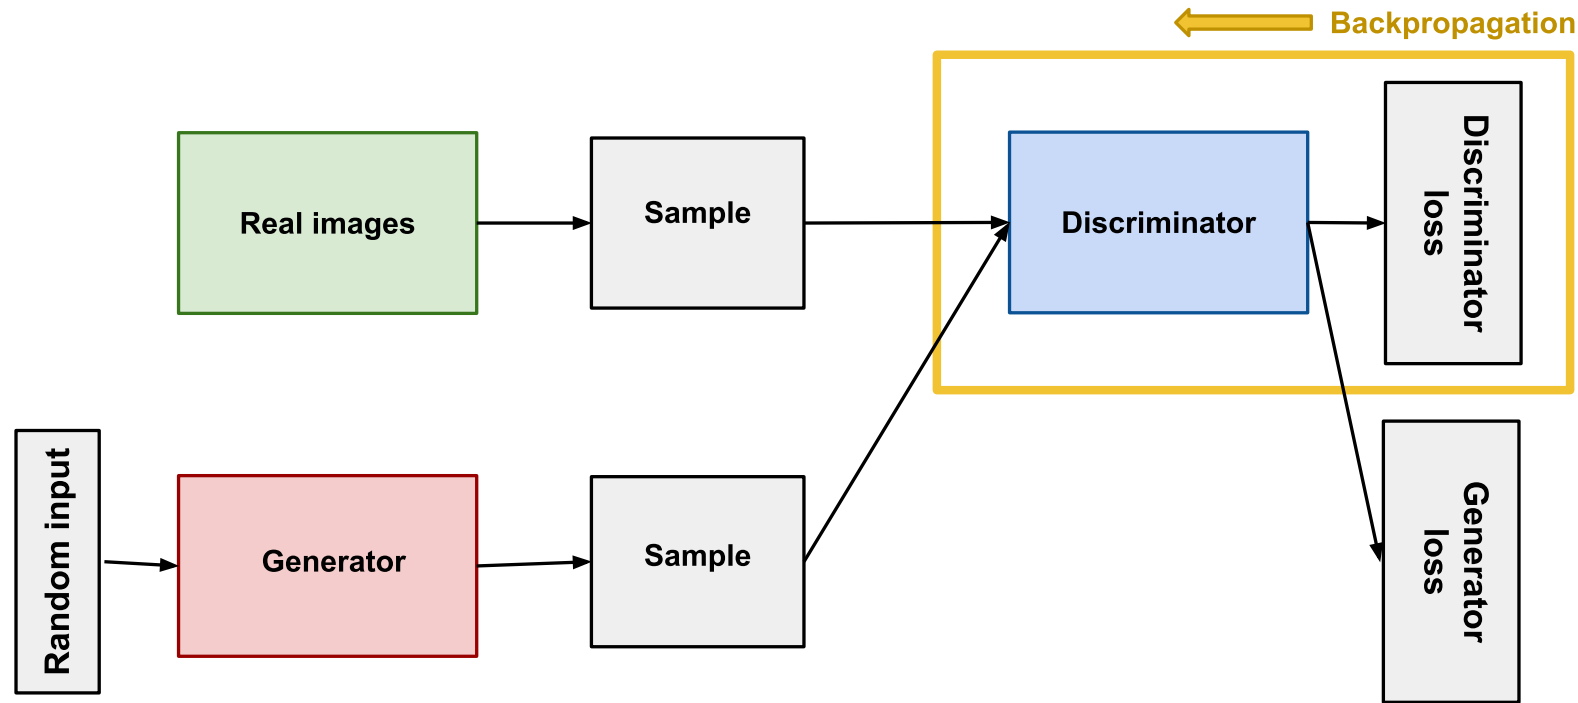
- a) A *discriminator* learns to classify images while detecting fake ones
- b) A *generator* learns how to fool the discriminator



# Generative Adversarial Network

## ■ Two competing networks

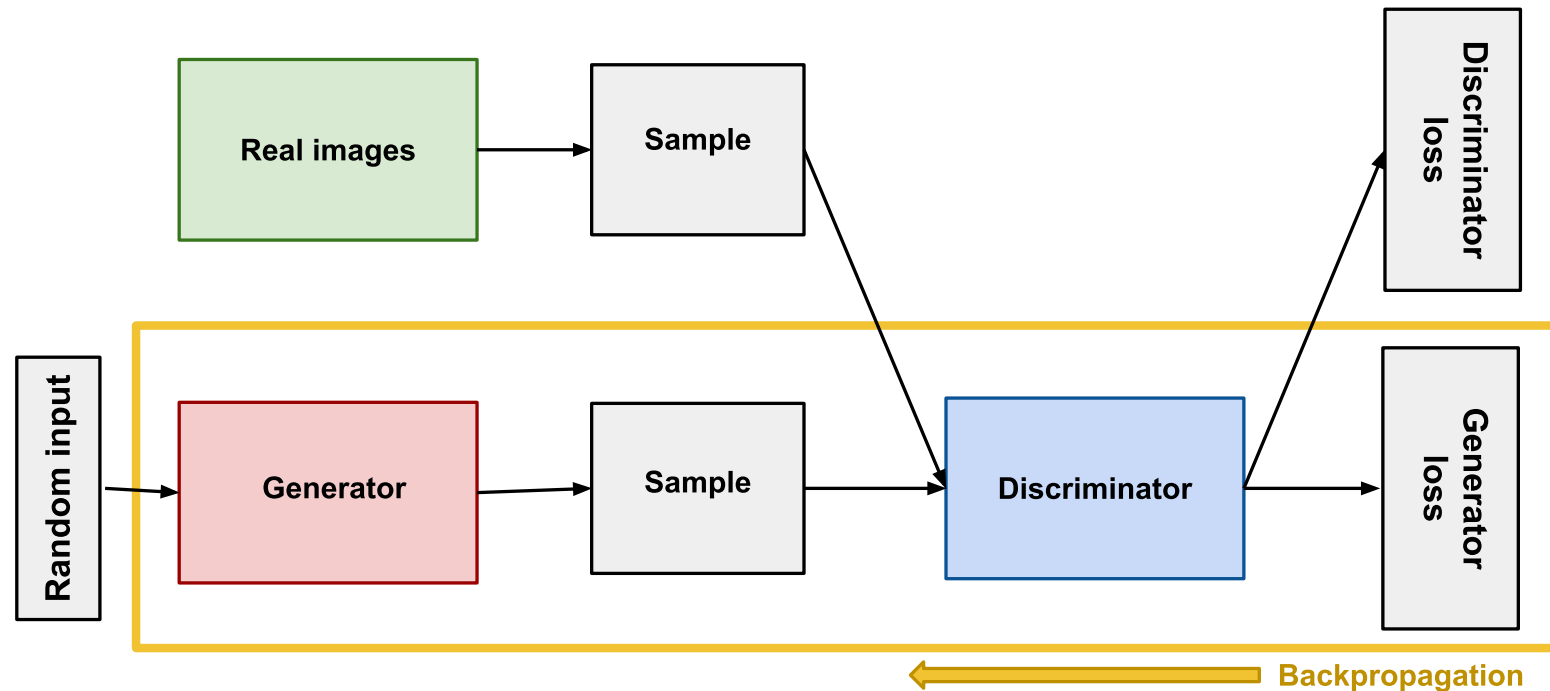
- A *discriminator* learns to classify images while detecting fake ones
- A *generator* learns how to fool the discriminator
- The two networks are trained in alternate turns



# Generative Adversarial Network

## ■ Two competing networks

- a) A *discriminator* learns to classify images while detecting fake ones
- b) A *generator* learns how to fool the discriminator
- c) The two networks are trained in alternate turns





# Generative Adversarial Network

## ■ Applications

- This method can be used to generate larger datasets from smaller ones
- Or to generate photo-realistic, yet synthetic images (even 'deep fakes')



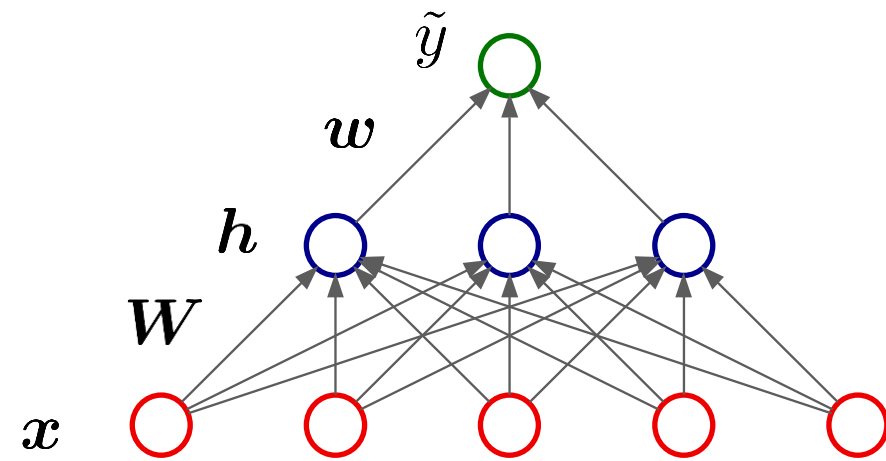
# *Unsupervised Learning: Auto-Encoders*

# Auto-Encoders

## ▪ Encoder

A feed-forward neural network with one hidden layer

$$\tilde{y} = w \cdot g(\mathbf{W}x + \mathbf{b}) + b$$



# Auto-Encoders

- **Encoder**

A feed-forward neural network with one hidden layer

$$\tilde{y} = w \cdot g(\mathbf{W}\mathbf{x} + \mathbf{b}) + b$$

- **Auto-encoder (basic idea): encoder + decoder**

$$\mathbf{x}^{[m]} = g(\mathbf{W}^{[m]} \cdot g(\mathbf{W}\mathbf{x} + \mathbf{b}) + \mathbf{b}^{[m]})$$

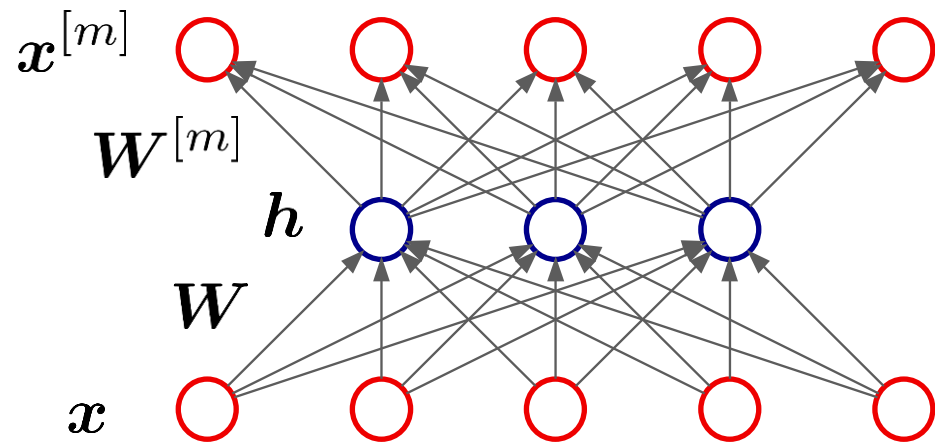
Loss function (MSE):

$$L(\mathbf{x}^{[m]}, \mathbf{x}) = (\mathbf{x}^{[m]} - \mathbf{x})^2$$

Initially:

$$\mathbf{W}^{[m]} = \mathbf{W}^T$$

then train the network with each data sample **onto itself**



# Auto-Encoders

- **Auto-encoder** (*More in general*)

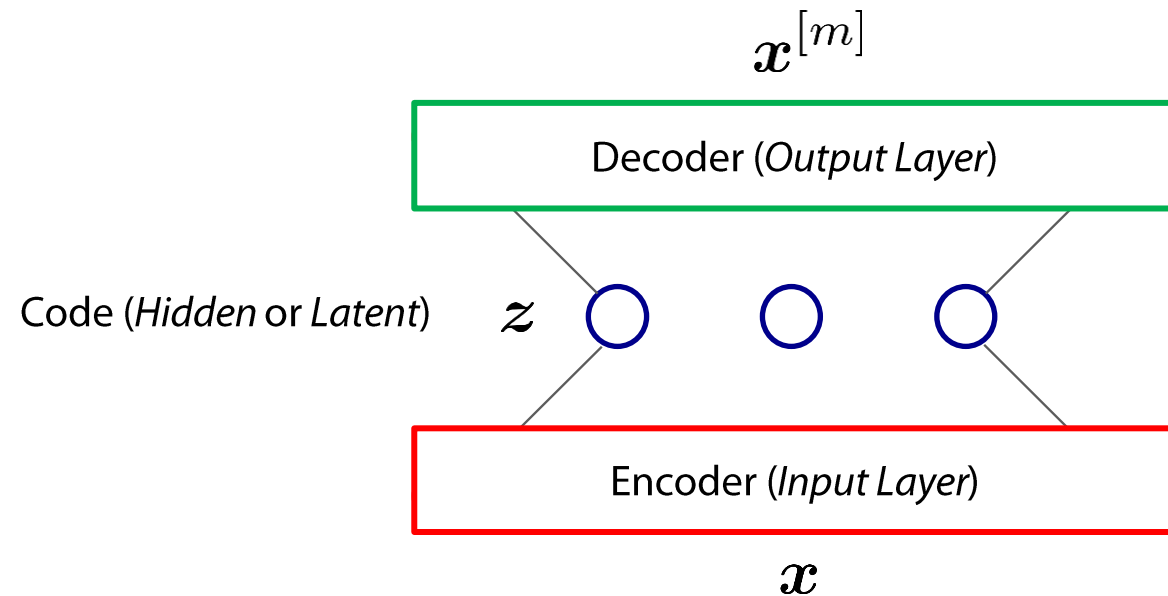
Two main (composite) layers: **encoder** and **decoder**

One **hidden** or **latent** layer  $z$

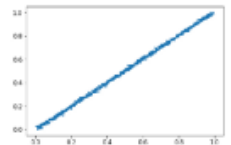
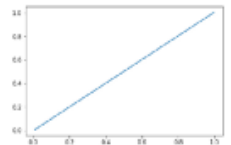
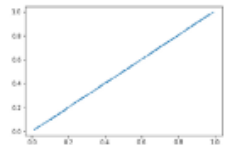
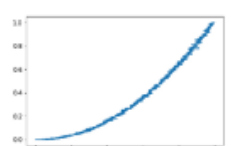
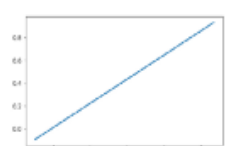
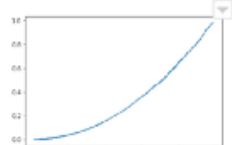



Each item in the dataset comprises the input only (*Unsupervised Learning*)

$$D := \{(\mathbf{x}^{(i)})\}_{i=1}^N,$$

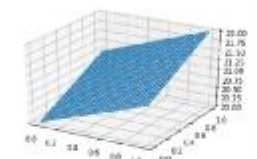
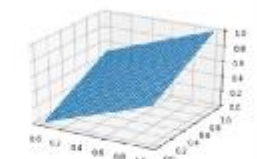
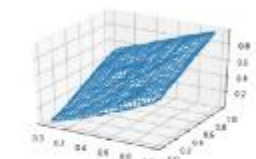
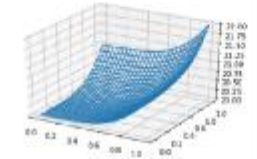
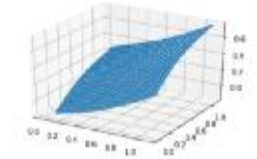
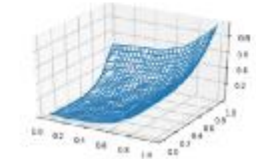
The result of the optimization is  $z$ :  
a compact (i.e. lower-dimensional)  
representation of the input  $\mathbf{x}$



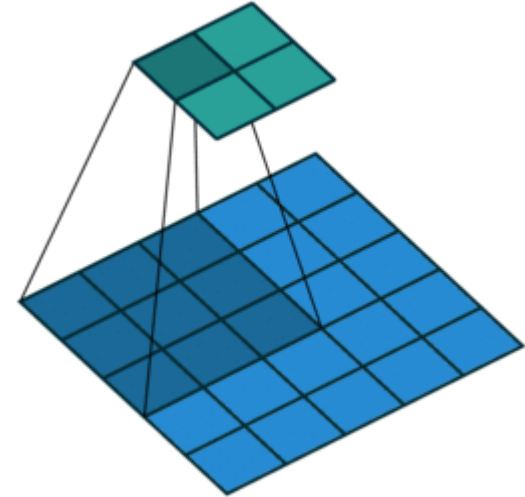
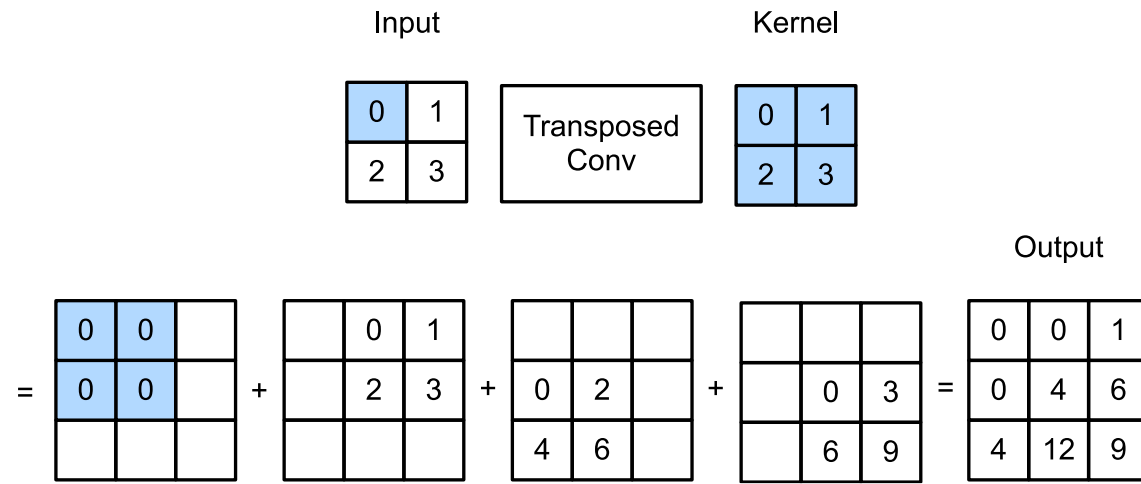
# Auto-Encoders vs PCA

Function	Feature Space	PCA Reconstruction	Auto Encoder Reconstruction
$y=mx+c$			
$y=mx^2+c$			
$y=mx^8+c$			

When non-linearity matters...

Function	Feature Space	PCA Reconstruction	Auto Encoder Reconstruction
Plane			
Curved Surface			

# Transposed Convolution (a.k.a. 'Deconvolution')



## ■ For autoencoders based on convolutional layers

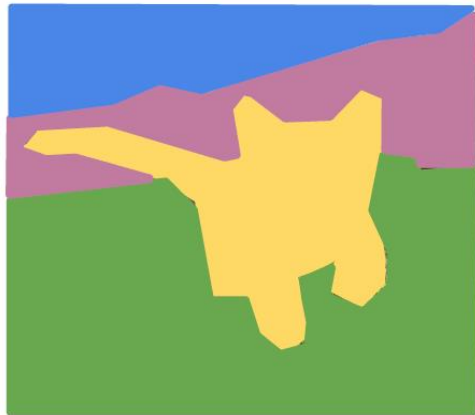
- Scalar input values are multiplied by the kernel tensor
- The output feature map is obtained by summing up all contributions

*Image Classification*  
*Object Detection*  
*Segmentation*



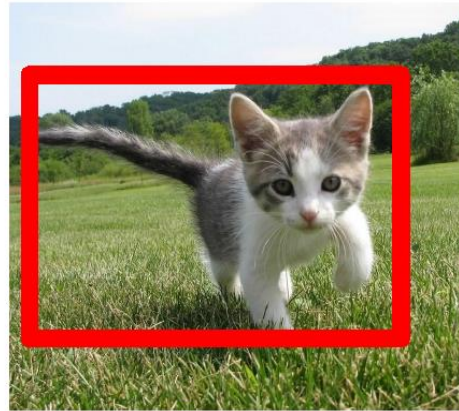
# Deep Learning for different imaging tasks

Beyond simple image classification



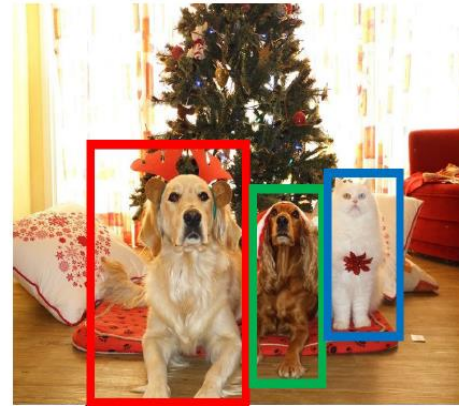
GRASS, CAT,  
TREE, SKY

No objects, just pixels



CAT

Single Object



DOG, DOG, CAT



DOG, DOG, CAT

Multiple Object

This image is CC0 public domain

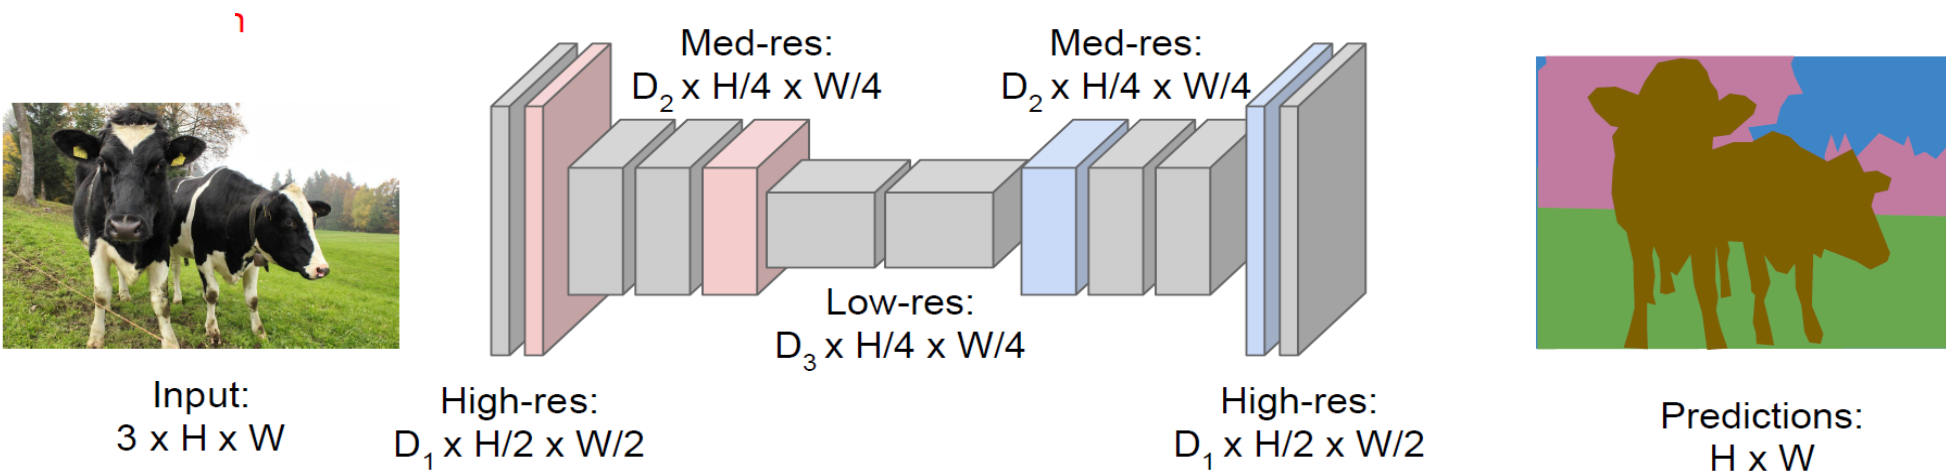
# Semantic segmentation

Beyond simple image classification

- **Similar network architecture, different arrangement**

Fully Convolutional Networks (FCN)

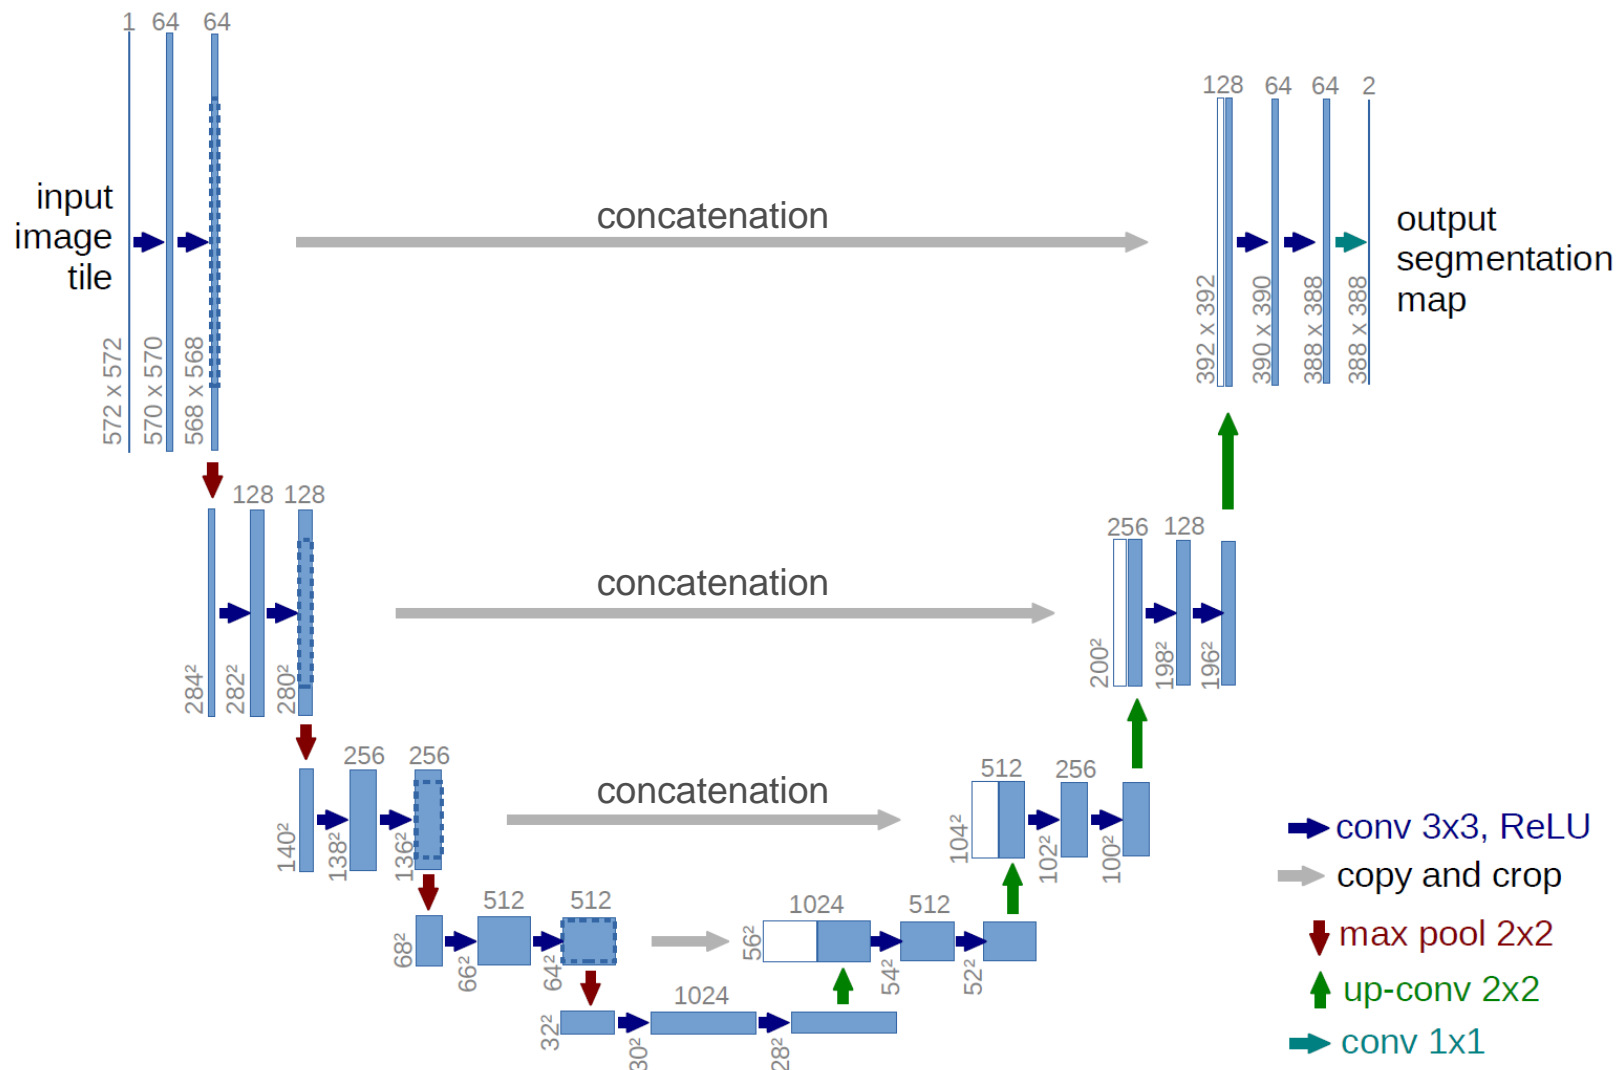
*Downsampling first, upsampling afterwards*



# Semantic segmentation

## U-Net [2015]

Great precision  
Fast to train



# Object detection and positioning

Generate boxes and classifications

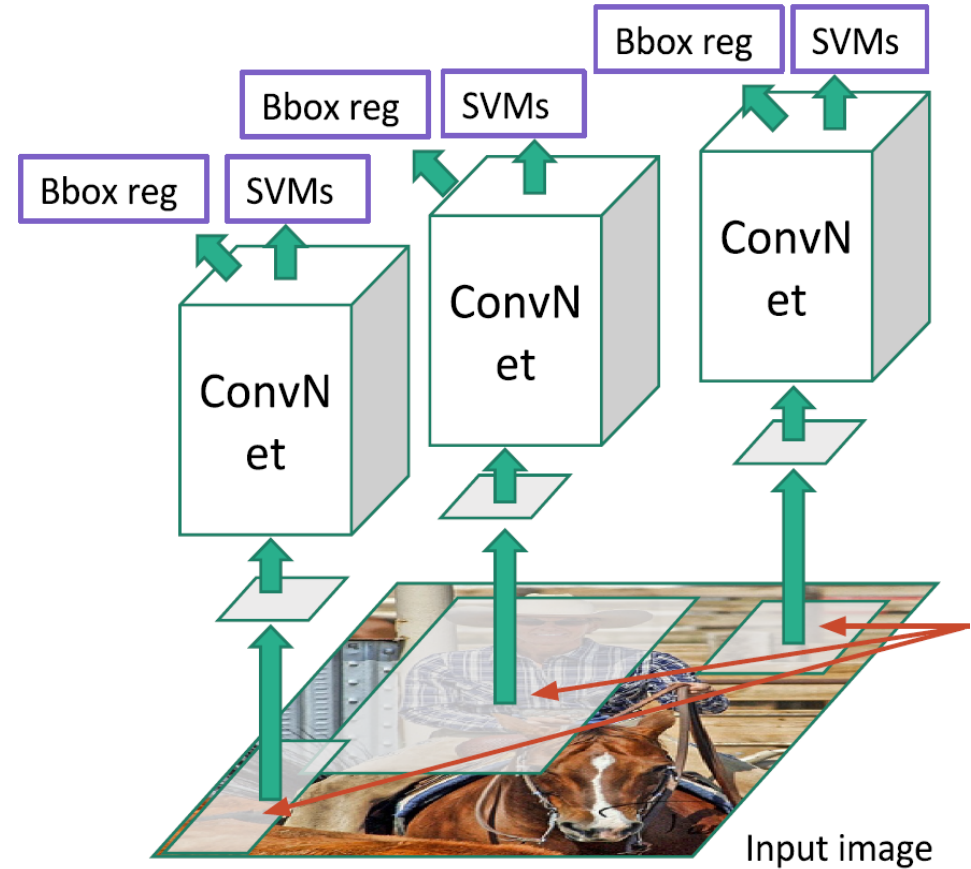
## ■ **Two-stage Process**

Generate bounding box candidates

Pass each candidate through a DCNN

Select those candidates that are classified with higher certainty

## R-CNN

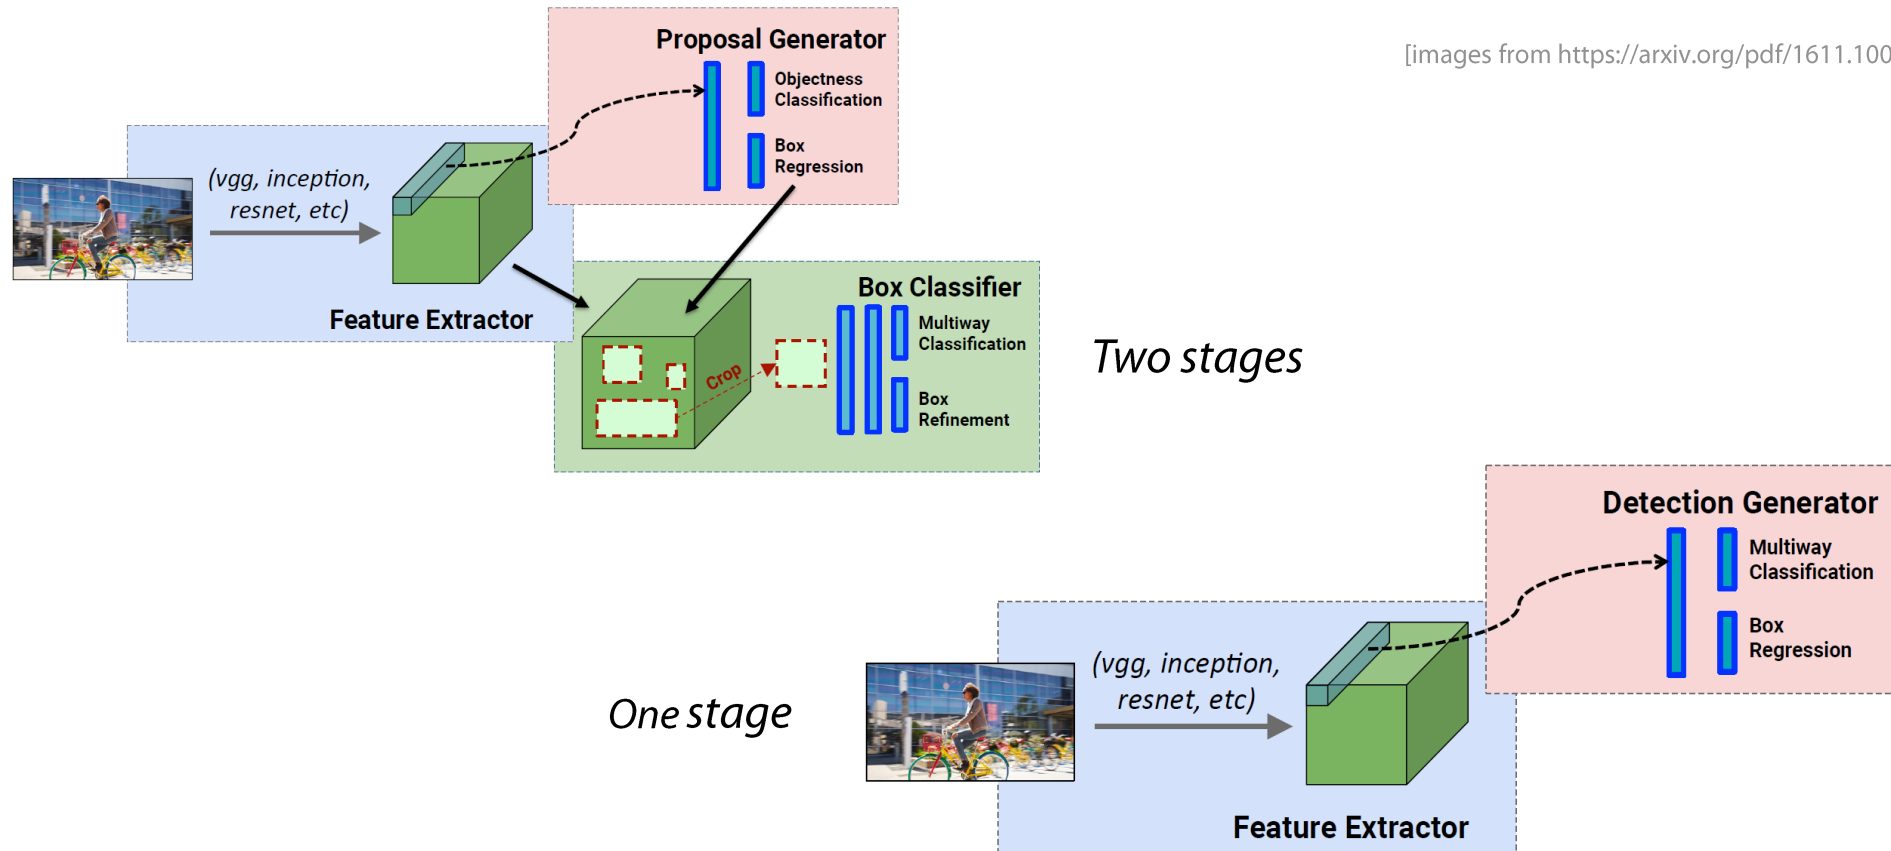


# Object detection and positioning

Generate boxes and classifications

## ■ Two-stage to One-stage process

Generate bounding box candidates and classifications in one go

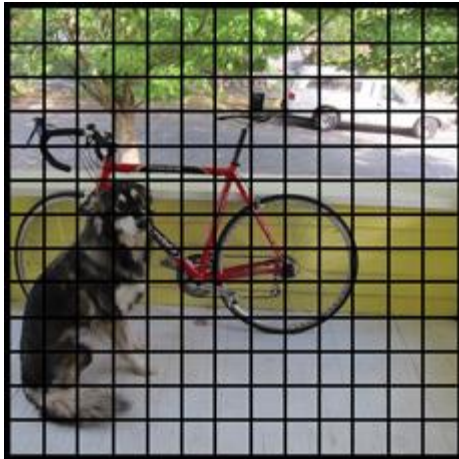


# Object detection and positioning

## ■ **YOLO and SSD: one-pass convolutional network for object detection**

Generate boxes and classifications at once

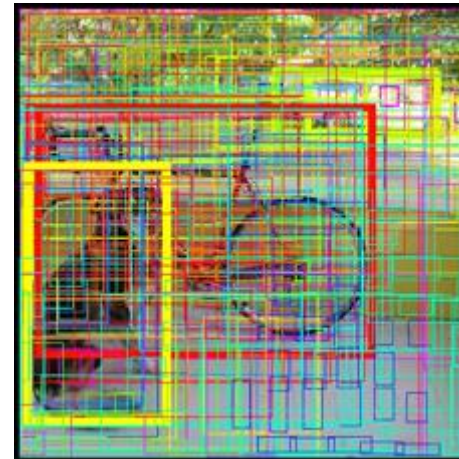
1) Impose a fixed grid over the input image



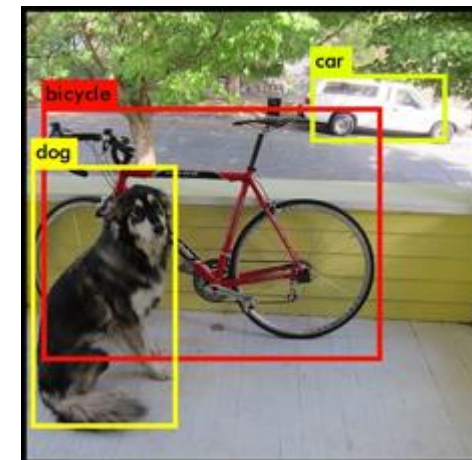
2) Generate possible bounding boxes



3) Classify each of them

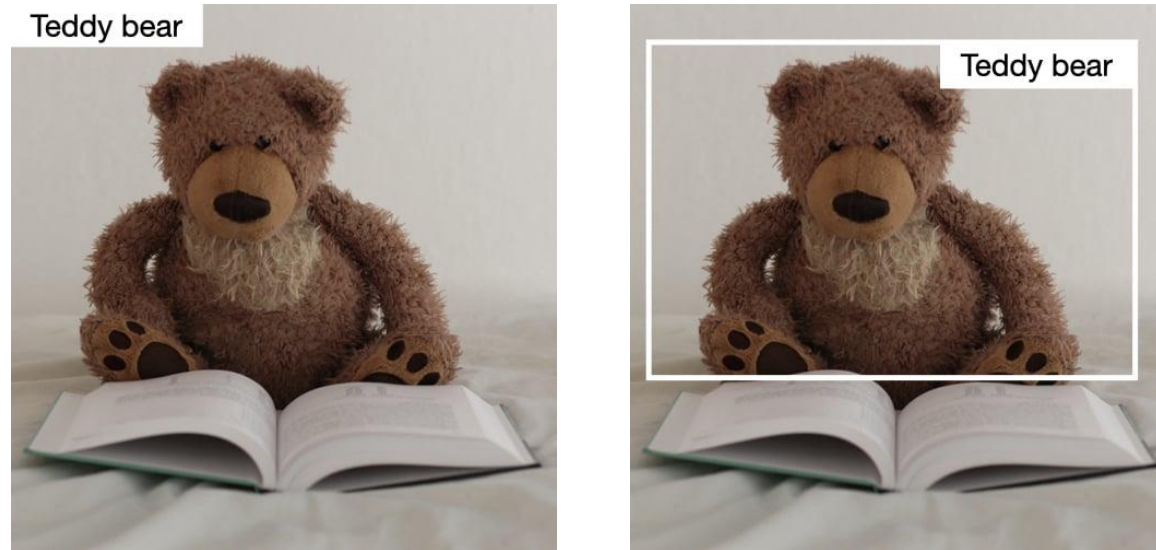


4) Keep the boxes at highest confidence

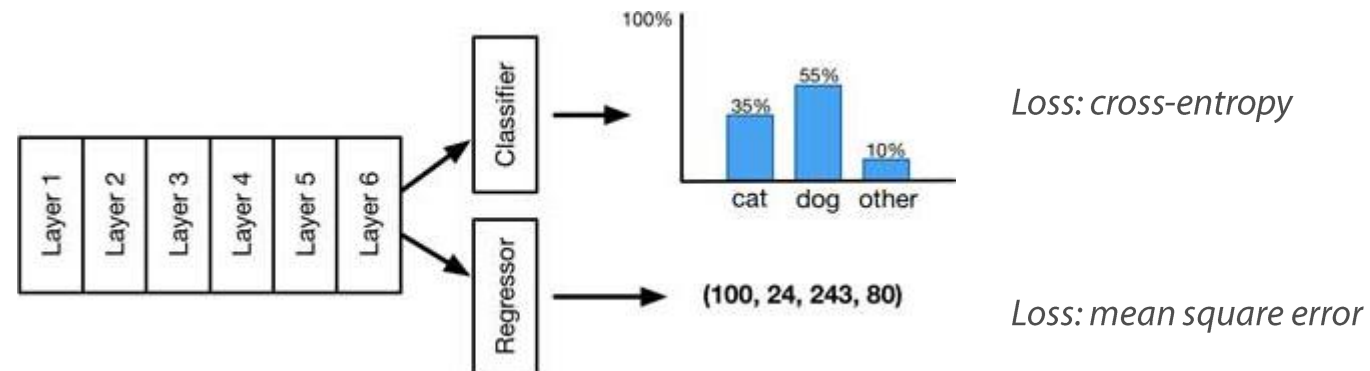


# Object detection and positioning

## ■ From classification to localization

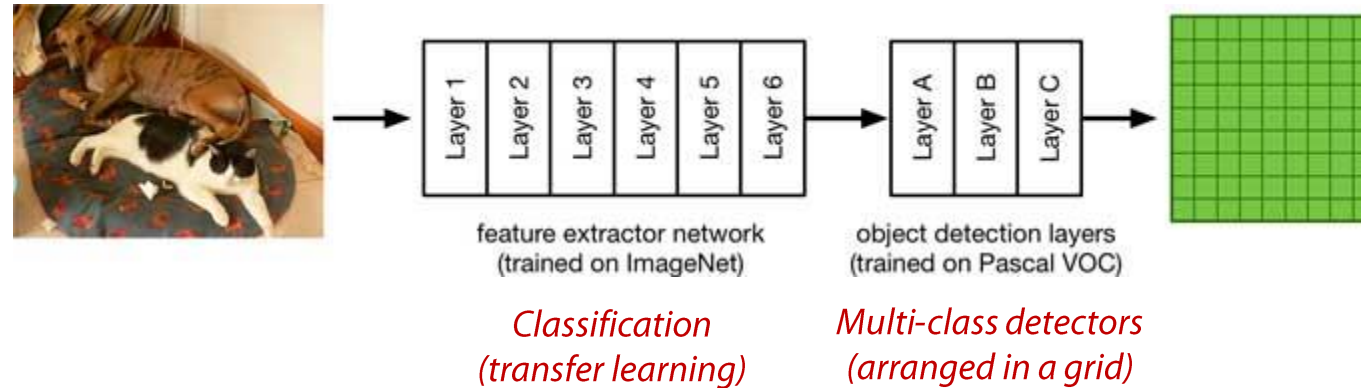


[images from <https://stanford.edu/~shervine/teaching/cs-230/cheatsheet-convolutional-neural-networks>]

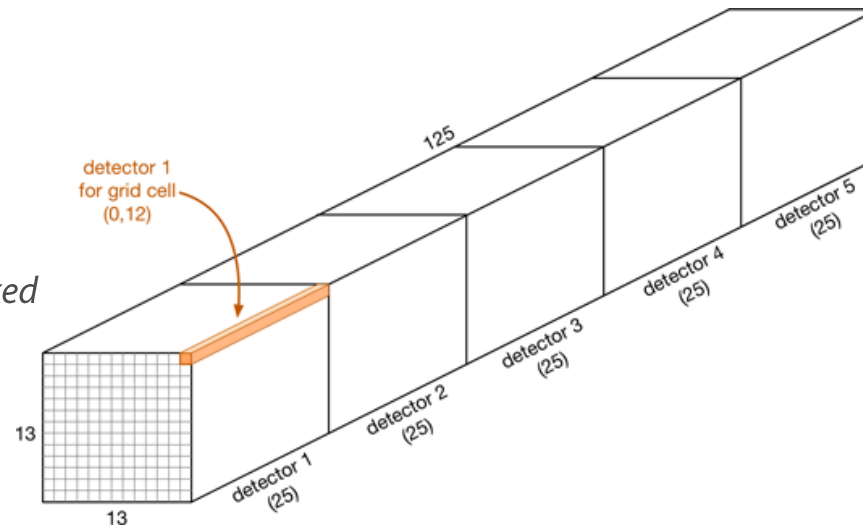


# Object detection and positioning

## ■ Grid detectors



(Detector blocks are shown as stacked but they work in parallel)



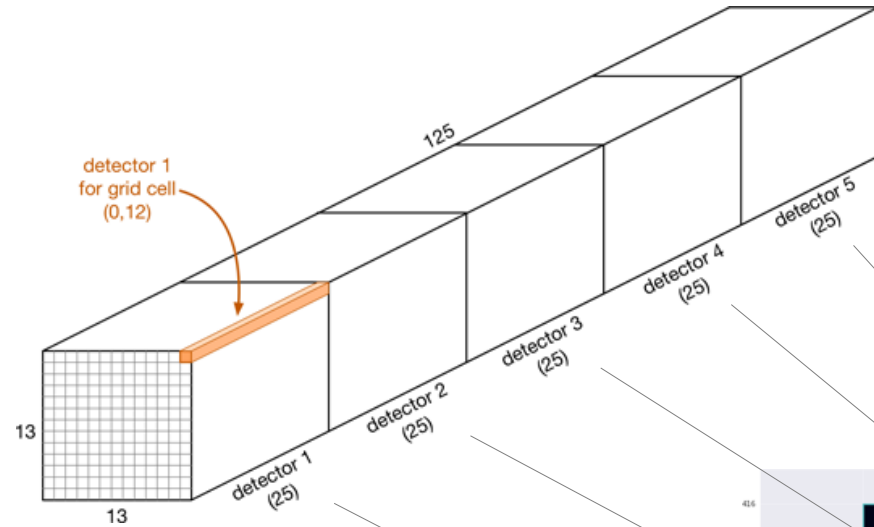
- 1) Impose a fixed 13 x 13 grid over the input image
- 2) Assign 5 multiclass detectors to each cell of the grid
- 3) Each multiclass detector works on a specific *anchor* shape (see next slide)

[images from <https://machinethink.net/blog/object-detection/>]



# Object detection and positioning

## ■ **Grid detectors: one per anchor**

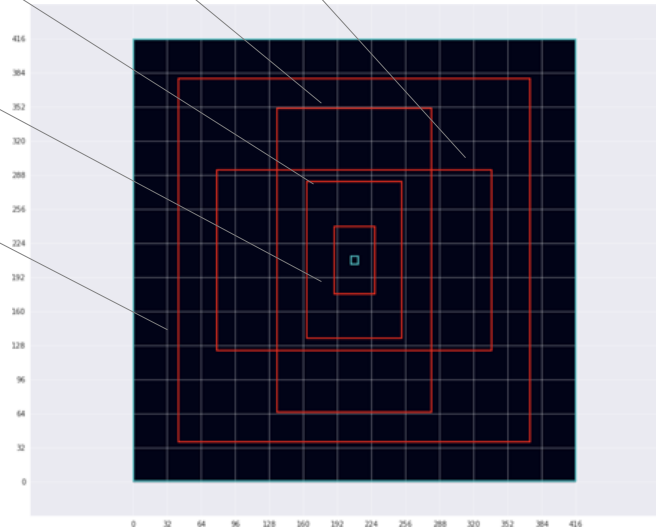


Anchor shapes and sizes are pre-computed from the training set (bounding box statistics)

Each detector produces:

- *class probabilities*
- *confidence score*

per its assigned anchor



[images from <https://machinethink.net/blog/object-detection/>]

# Object detection and positioning

- Given anchor, cell and class

$$\langle c_x, c_y, p_w, p_h \rangle$$

top-left cell  
coordinates

anchor  
sizes

- Each detector produces

$$\langle t_x, t_y, t_w, t_h, p_o, p_c \rangle$$

$$b_x = \sigma(t_x) + c_x$$

$$b_y = \sigma(t_y) + c_y$$

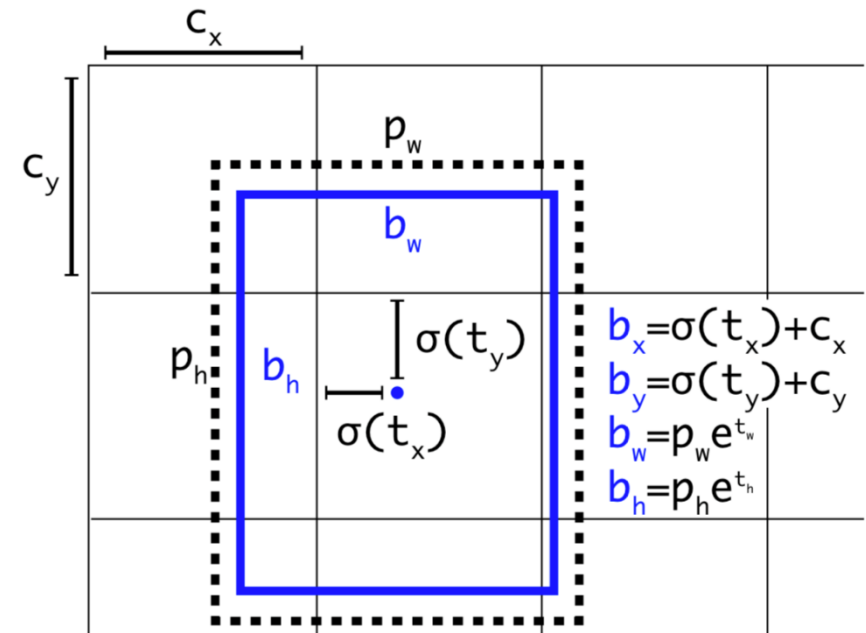
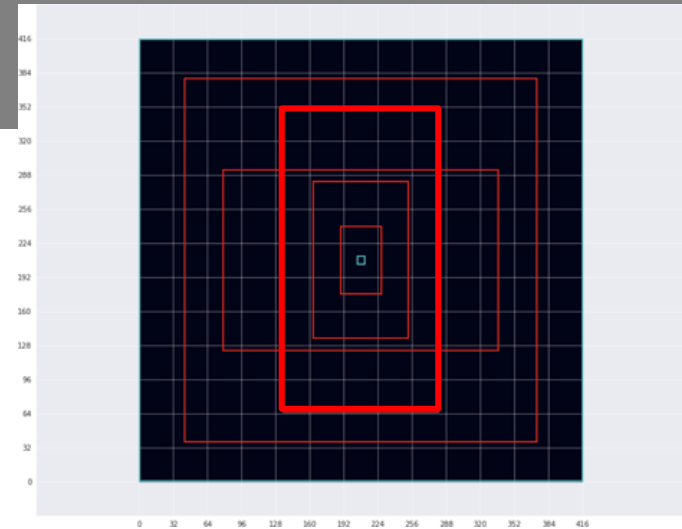
$$b_w = p_w e^{t_w}$$

$$b_h = p_h e^{t_h}$$

$p_o$  'objectness' probability

$p_c$  class probability

bounding box  
coordinates

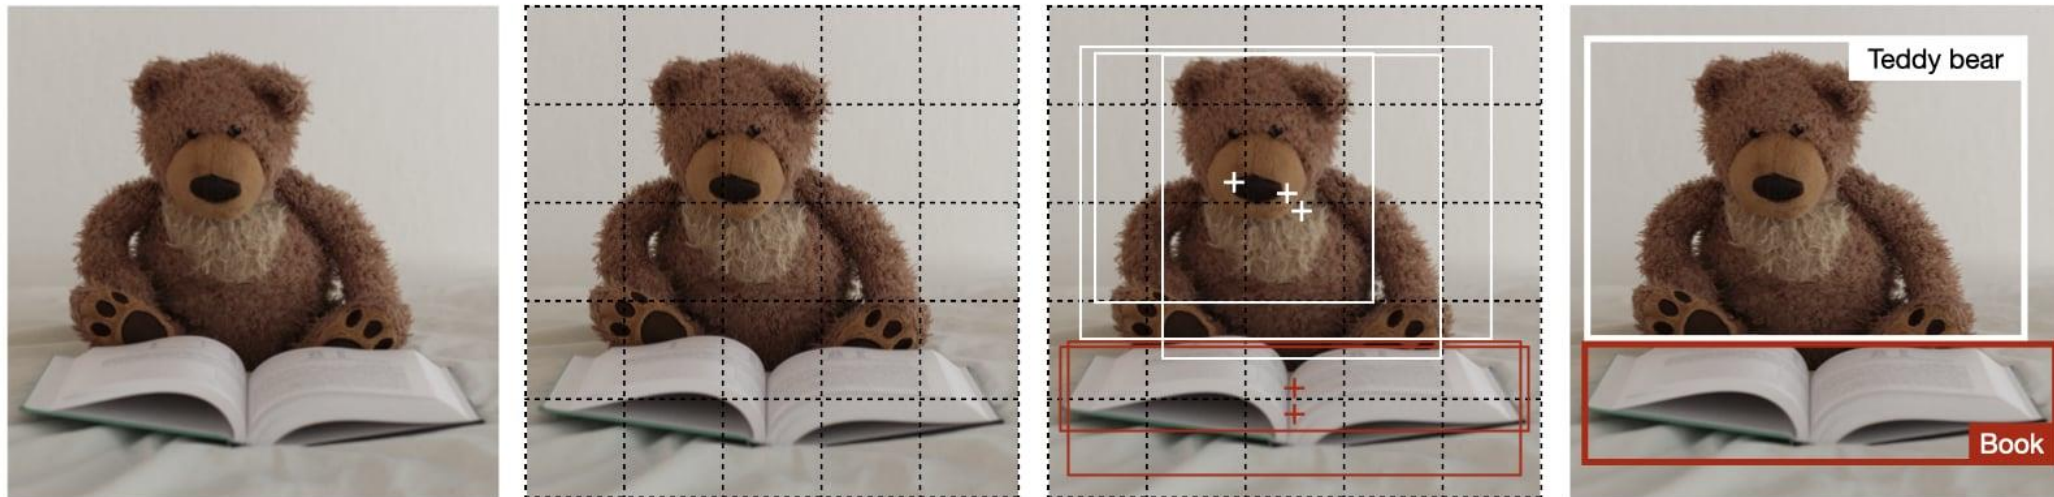


[images from <https://wikidocs.net/167697>]

# Object detection and positioning

## ■ From grid boxes to candidate boxes

Merging predictions

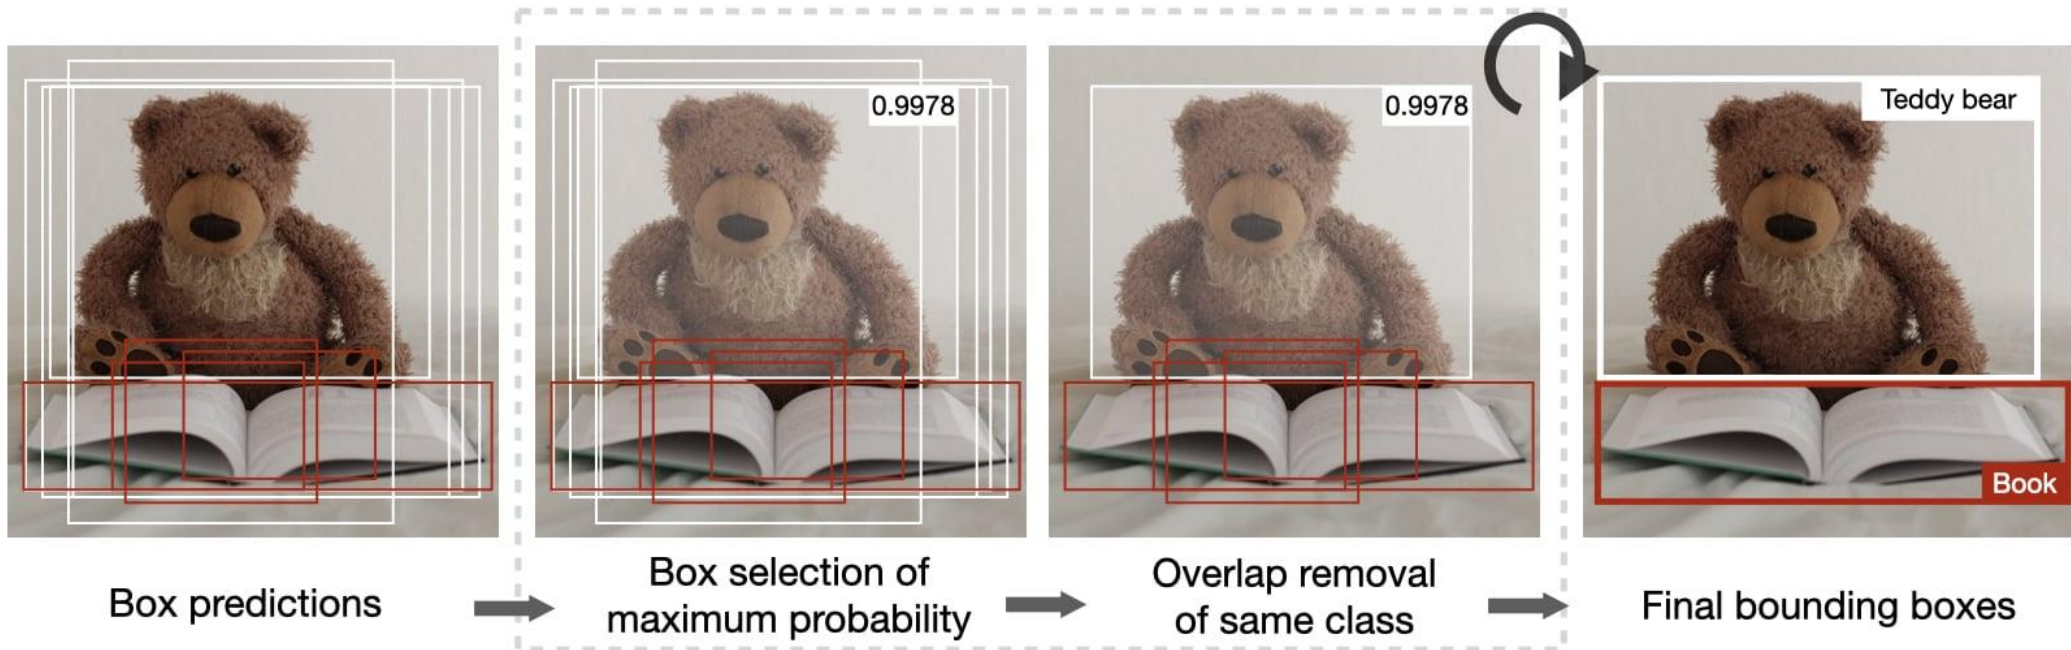


Original image → Division in  $G \times G$  grid → Bounding box prediction → Non-max suppression

[images from <https://stanford.edu/~shervine/teaching/cs-230/cheatsheet-convolutional-neural-networks>]

# Object detection and positioning

## ▪ Further processing



[images from <https://stanford.edu/~shervine/teaching/cs-230/cheatsheet-convolutional-neural-networks>]

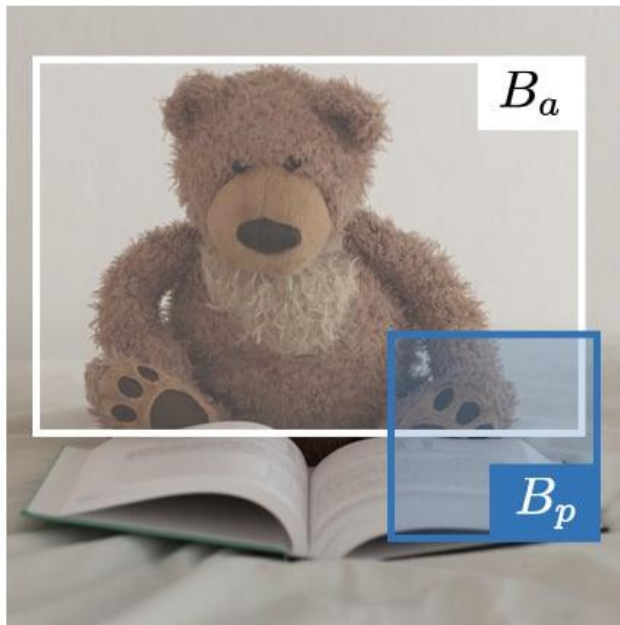
# Object detection and positioning

## ■ Measuring object detection accuracy

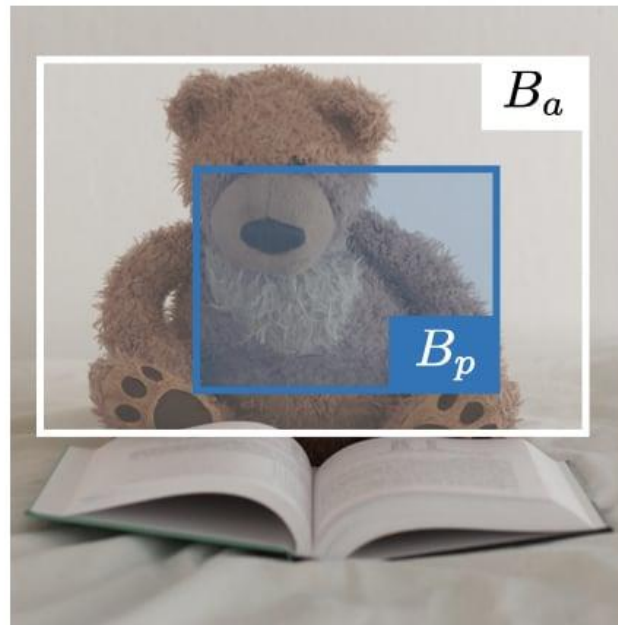
Intersection over Union (IoU)

$$\text{IoU}(B_p, B_a) := \frac{B_p \cap B_a}{B_p \cup B_a}$$

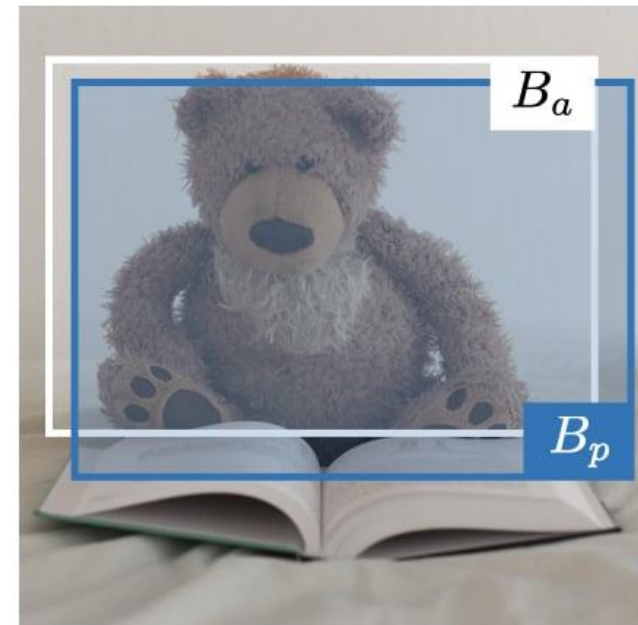
It's a post-localization accuracy measure  
(not a loss function)



$$\text{IoU}(B_p, B_a) = 0.1$$



$$\text{IoU}(B_p, B_a) = 0.5$$



$$\text{IoU}(B_p, B_a) = 0.9$$

[images from <https://stanford.edu/~shervine/teaching/cs-230/cheatsheet-convolutional-neural-networks>]