


Deep Learning in Image Processing


Topics:

- Image Filtering 101
- CNNs 101
- Image Processing Pipelines

Frank Dellaert
CS 4476 Computer Vision



2. Image Formation



3. Image Processing




4. Features




5. Segmentation




6-7. Structure from Motion




8. Motion



9. Stitching




10. Computational Photography



11. Stereo



12. 3D Shape



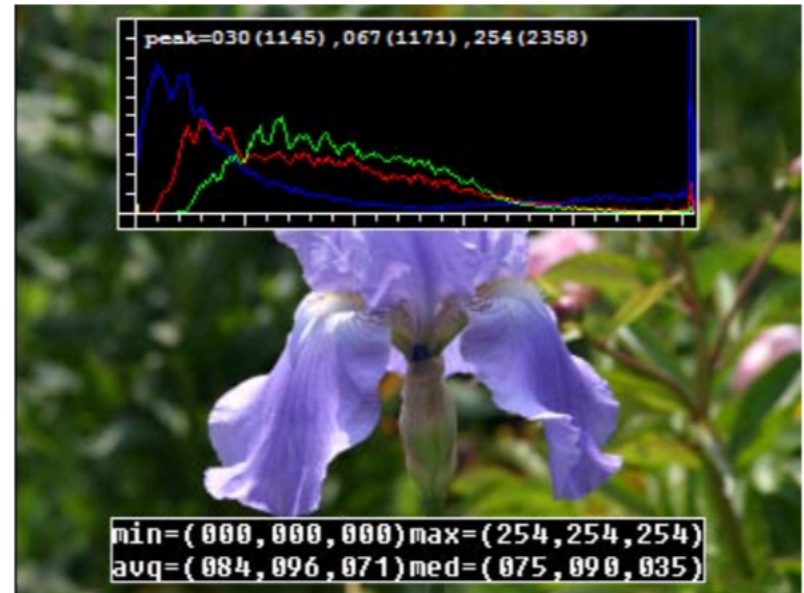

13. Image-based Rendering



14. Recognition

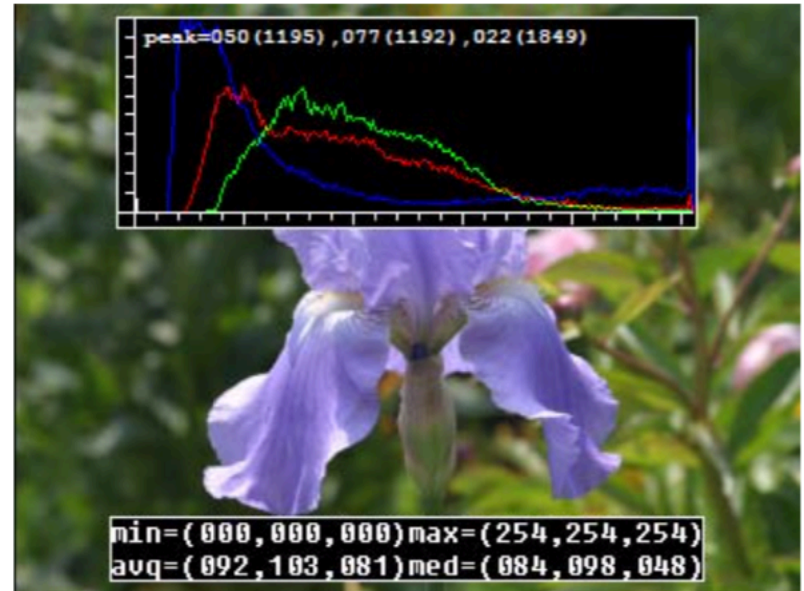
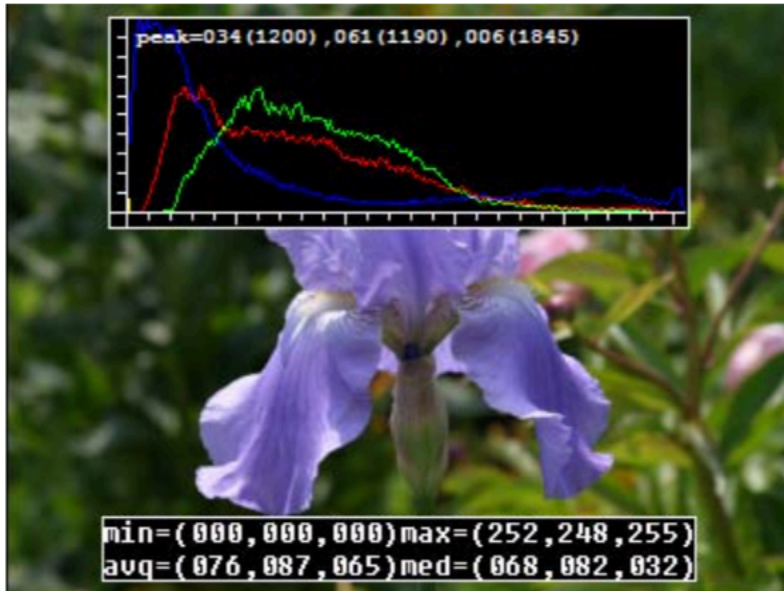
3.1	Point operators	101
3.1.1	Pixel transforms	103
3.1.2	Color transforms	104
3.1.3	Compositing and matting	105
3.1.4	Histogram equalization	107
3.1.5	<i>Application: Tonal adjustment</i>	111
3.2	Linear filtering	111
3.2.1	Separable filtering	115
3.2.2	Examples of linear filtering	117
3.2.3	Band-pass and steerable filters	118
3.3	More neighborhood operators	122
3.3.1	Non-linear filtering	122
3.3.2	Morphology	127
3.3.3	Distance transforms	129
3.3.4	Connected components	131
3.4	Fourier transforms	132
3.4.1	Fourier transform pairs	136
3.4.2	Two-dimensional Fourier transforms	140
3.4.3	Wiener filtering	140
3.4.4	<i>Application: Sharpening, blur, and noise removal</i>	144
3.5	Pyramids and wavelets	144
3.5.1	Interpolation	145
3.5.2	Decimation	148
3.5.3	Multi-resolution representations	150
3.5.4	Wavelets	154
3.5.5	<i>Application: Image blending</i>	160

Contrast



- $g(x) = a f(x)$, $a=1.1$

Brightness



- $g(x) = f(x) + b, b=16$

3.1	Point operators	101
3.1.1	Pixel transforms	103
3.1.2	Color transforms	104
3.1.3	Compositing and matting	105
3.1.4	Histogram equalization	107
3.1.5	<i>Application: Tonal adjustment</i>	111
3.2	Linear filtering	111
3.2.1	Separable filtering	115
3.2.2	Examples of linear filtering	117
3.2.3	Band-pass and steerable filters	118
3.3	More neighborhood operators	122
3.3.1	Non-linear filtering	122
3.3.2	Morphology	127
3.3.3	Distance transforms	129
3.3.4	Connected components	131
3.4	Fourier transforms	132
3.4.1	Fourier transform pairs	136
3.4.2	Two-dimensional Fourier transforms	140
3.4.3	Wiener filtering	140
3.4.4	<i>Application: Sharpening, blur, and noise removal</i>	144
3.5	Pyramids and wavelets	144
3.5.1	Interpolation	145
3.5.2	Decimation	148
3.5.3	Multi-resolution representations	150
3.5.4	Wavelets	154
3.5.5	<i>Application: Image blending</i>	160


Image filtering

- Compute a function of the local neighborhood at each pixel in the image
 - Function specified by a “filter” or mask saying how to combine values from neighbors.
- Uses of filtering:
 - Enhance an image (denoise, resize, etc)
 - Extract information (texture, edges, etc)
 - Detect patterns (template matching)

Image filtering

- Compute a function of the local neighborhood at each pixel in the image
 - Function specified by a “filter” or mask saying how to combine values from neighbors.
- Uses of filtering:
 - Enhance an image (denoise, resize, etc)
 - Extract information (texture, edges, etc)
 - Detect patterns (template matching)

Motivation: noise reduction



- Even multiple images of the **same static scene** will not be identical.




Common types of noise

- **Salt and pepper noise:** random occurrences of black and white pixels
- **Impulse noise:** random occurrences of white pixels
- **Gaussian noise:** variations in intensity drawn from a Gaussian normal distribution



Original

Gaussian noise



$$f(x, y) = \underbrace{\hat{f}(x, y)}_{\text{Ideal Image}} + \underbrace{\eta(x, y)}_{\text{Noise process}}$$


Gaussian i.i.d. ("white") noise:
 $\eta(x, y) \sim \mathcal{N}(\mu, \sigma)$

```
>> noise = randn(size(im)).*sigma;
```

```
>> output = im + noise;
```

What is impact of the sigma?

Motivation: noise reduction




- Even multiple images of the same static scene will not be identical.
- How could we reduce the noise, i.e., give an estimate of the true intensities?
- **What if there's only one image?**

First attempt at a solution

- Let's replace each pixel with an average of all the values in its neighborhood
- Assumptions:
 - Expect pixels to be like their neighbors
 - Expect noise processes to be independent from pixel to pixel


First attempt at a solution

- Let's replace each pixel with an average of all the values in its neighborhood
- Moving average in 1D:



Weighted Moving Average

- Can add weights to our moving average
- *Weights* [1, 1, 1, 1, 1] / 5



Weighted Moving Average

- Non-uniform weights [1, 4, 6, 4, 1] / 16




Image filtering

- Image filtering: compute function of local neighborhood at each position
- Really important!
 - Enhance images
 - Denoise, resize, increase contrast, etc.
 - Extract information from images
 - Texture, edges, distinctive points, etc.
 - Detect patterns
 - Template matching
 - Deep Convolutional Networks

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	0	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0

$G[x, y]$

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	0	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0									

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0								

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10							

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10							

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10	20						

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10	20						

Moving Average In 2D

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10	20	30					

Moving Average In 2D

$F[x, y]$

$G[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0	10	20	30					

Moving Average In 2D

$F[x, y]$

$G[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0	10	20	30	30				

Moving Average In 2D


$F[x, y]$

$G[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0



	0	10	20	30	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	60	90	90	90	60	30	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	10	20	30	30	30	30	20	10	
	10	10	10	0	0	0	0	0	

Smoothing with box filter



Important filter: Gaussian

- Weight contributions of neighboring pixels by nearness




0.003	0.013	0.022	0.013	0.003
0.013	0.059	0.097	0.059	0.013
0.022	0.097	0.159	0.097	0.022
0.013	0.059	0.097	0.059	0.013
0.003	0.013	0.022	0.013	0.003


5x5, $\sigma = 1$

$$G_{\sigma} = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$


Smoothing with Gaussian filter



Smoothing with box filter



Other filters




1	0	-1
2	0	-2
1	0	-1

Sobel




Vertical Edge
(absolute value)

Other filters



1	2	1
0	0	0
-1	-2	-1

Sobel



Horizontal Edge
(absolute value)

Gaussian filters

- Remove “high-frequency” components from the image (low-pass filter)
 - Images become more smooth
- Convolution with self is another Gaussian
 - So can smooth with small-width kernel, repeat, and get same result as larger-width kernel would have
 - Convoluting two times with Gaussian kernel of width σ is same as convoluting once with kernel of width $\sigma\sqrt{2}$
- *Separable* kernel
 - Factors into product of two 1D Gaussians

Separability of the Gaussian filter

$$\begin{aligned} G_{\sigma}(x, y) &= \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right)\right) \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{y^2}{2\sigma^2}\right)\right) \end{aligned}$$

The 2D Gaussian can be expressed as the product of two functions, one a function of x and the other a function of y

In this case, the two functions are the (identical) 1D Gaussian

Separability example

2D convolution
(center location only)

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 2 & 3 & 3 \\ \hline 3 & 5 & 5 \\ \hline 4 & 4 & 6 \\ \hline \end{array}$$


The filter factors
into a product of 1D
filters:

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array} = \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline 1 \\ \hline \end{array} \times \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array}$$


Perform convolution
along rows:

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 2 & 3 & 3 \\ \hline 3 & 5 & 5 \\ \hline 4 & 4 & 6 \\ \hline \end{array} = \begin{array}{|c|c|c|} \hline & 11 & \\ \hline & 18 & \\ \hline & 18 & \\ \hline \end{array}$$

Followed by convolution
along the remaining column:



2. Image Formation



3. Image Processing




4. Features




5. Segmentation




6-7. Structure from Motion




8. Motion



9. Stitching




10. Computational Photography



11. Stereo



12. 3D Shape


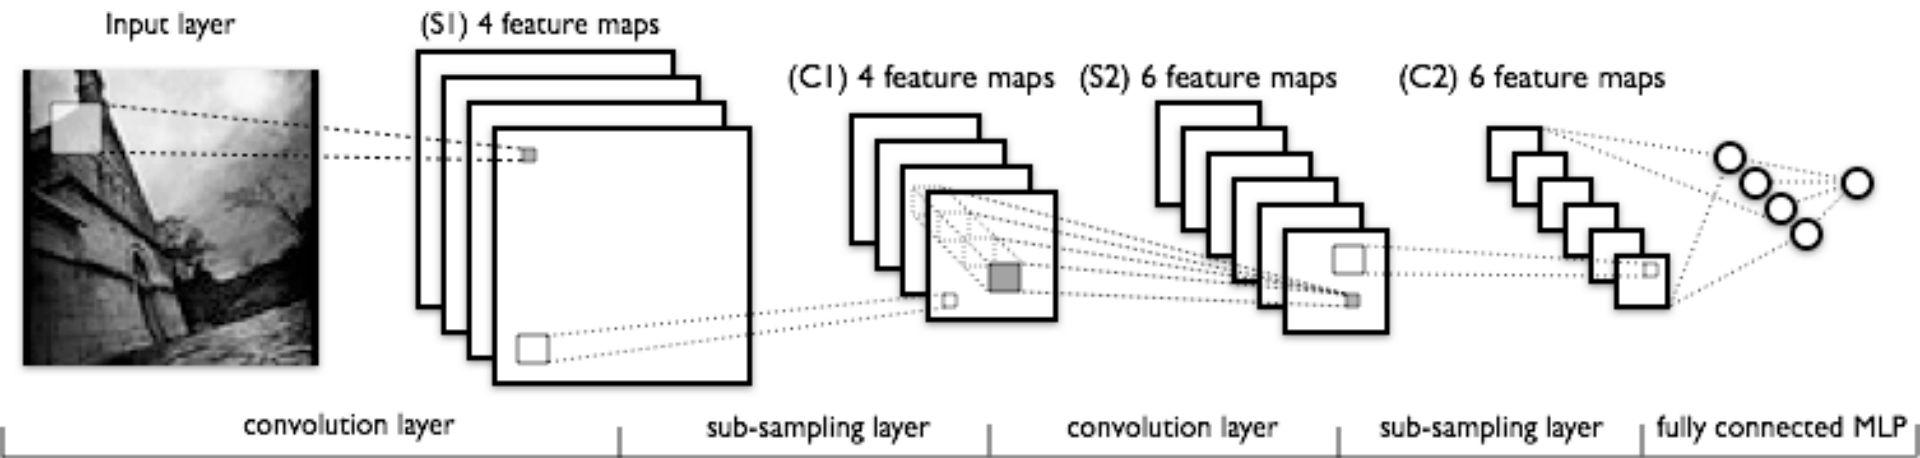


13. Image-based Rendering

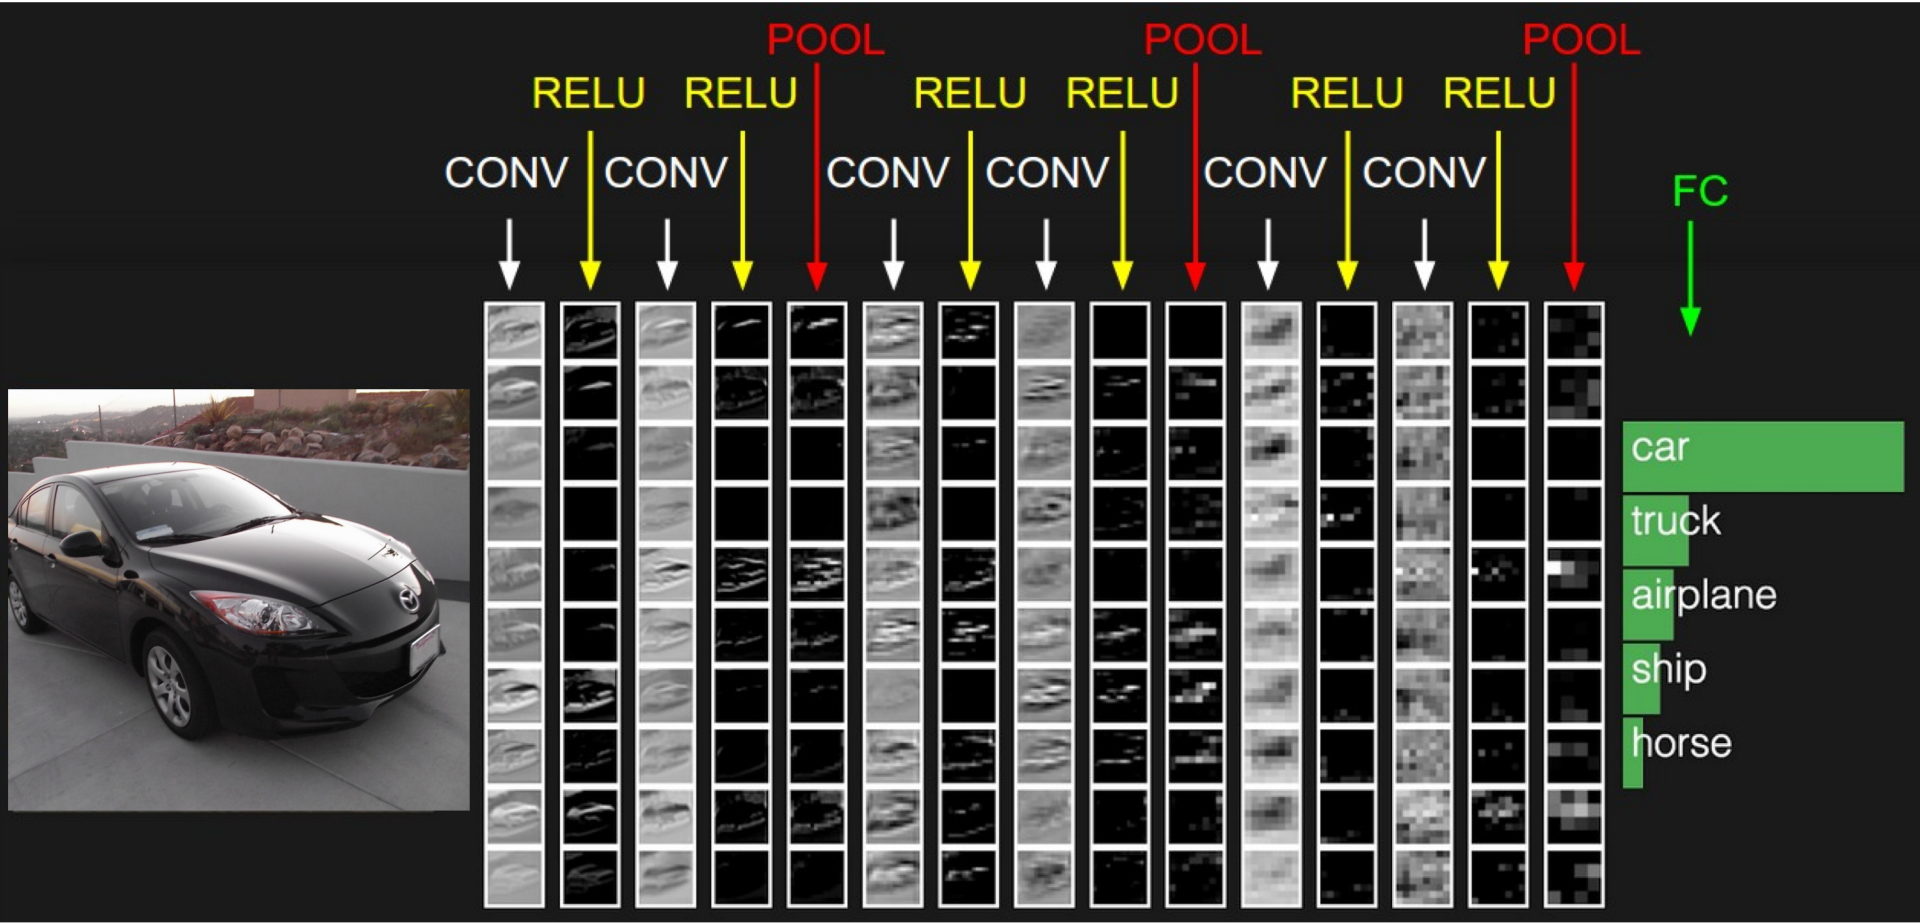


14. Recognition

Convolutional Neural Networks




preview:



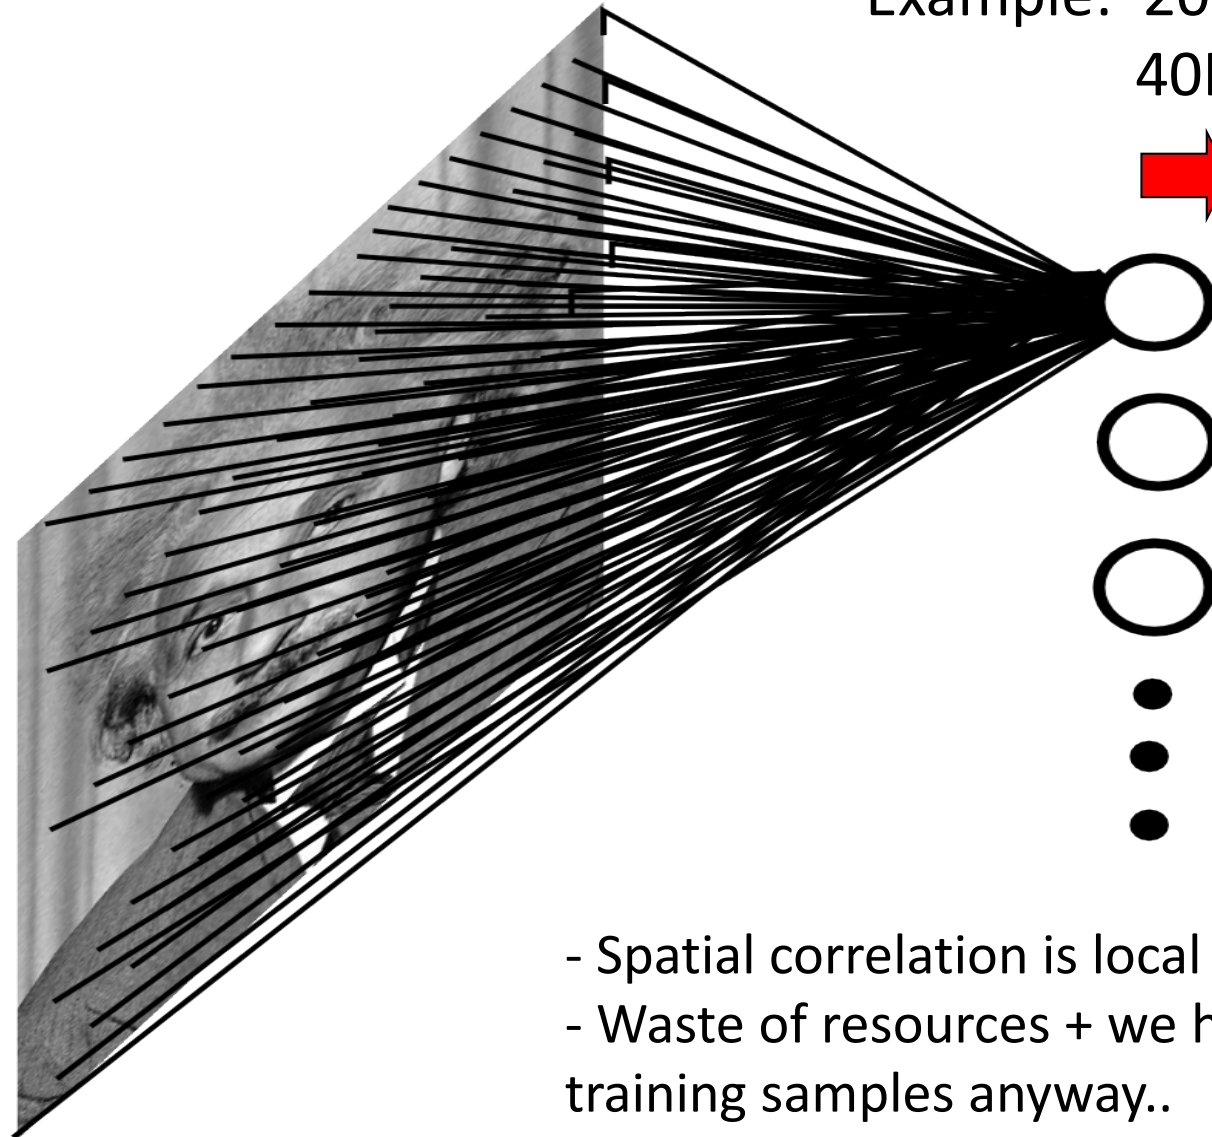
Fully Connected Layer

Example: 200x200 image

40K hidden units




~2B parameters!!!



- Spatial correlation is local
- Waste of resources + we have not enough training samples anyway..

Locally Connected Layer




Example: 200x200 image
40K hidden units
Filter size: 10x10
4M parameters


Note: This parameterization is good when input image is registered (e.g., face recognition).

Locally Connected Layer

STATIONARITY? Statistics is similar at different locations




Convolutional Layer




Share the same parameters across different locations (assuming input is stationary):

Convolutions with learned kernels

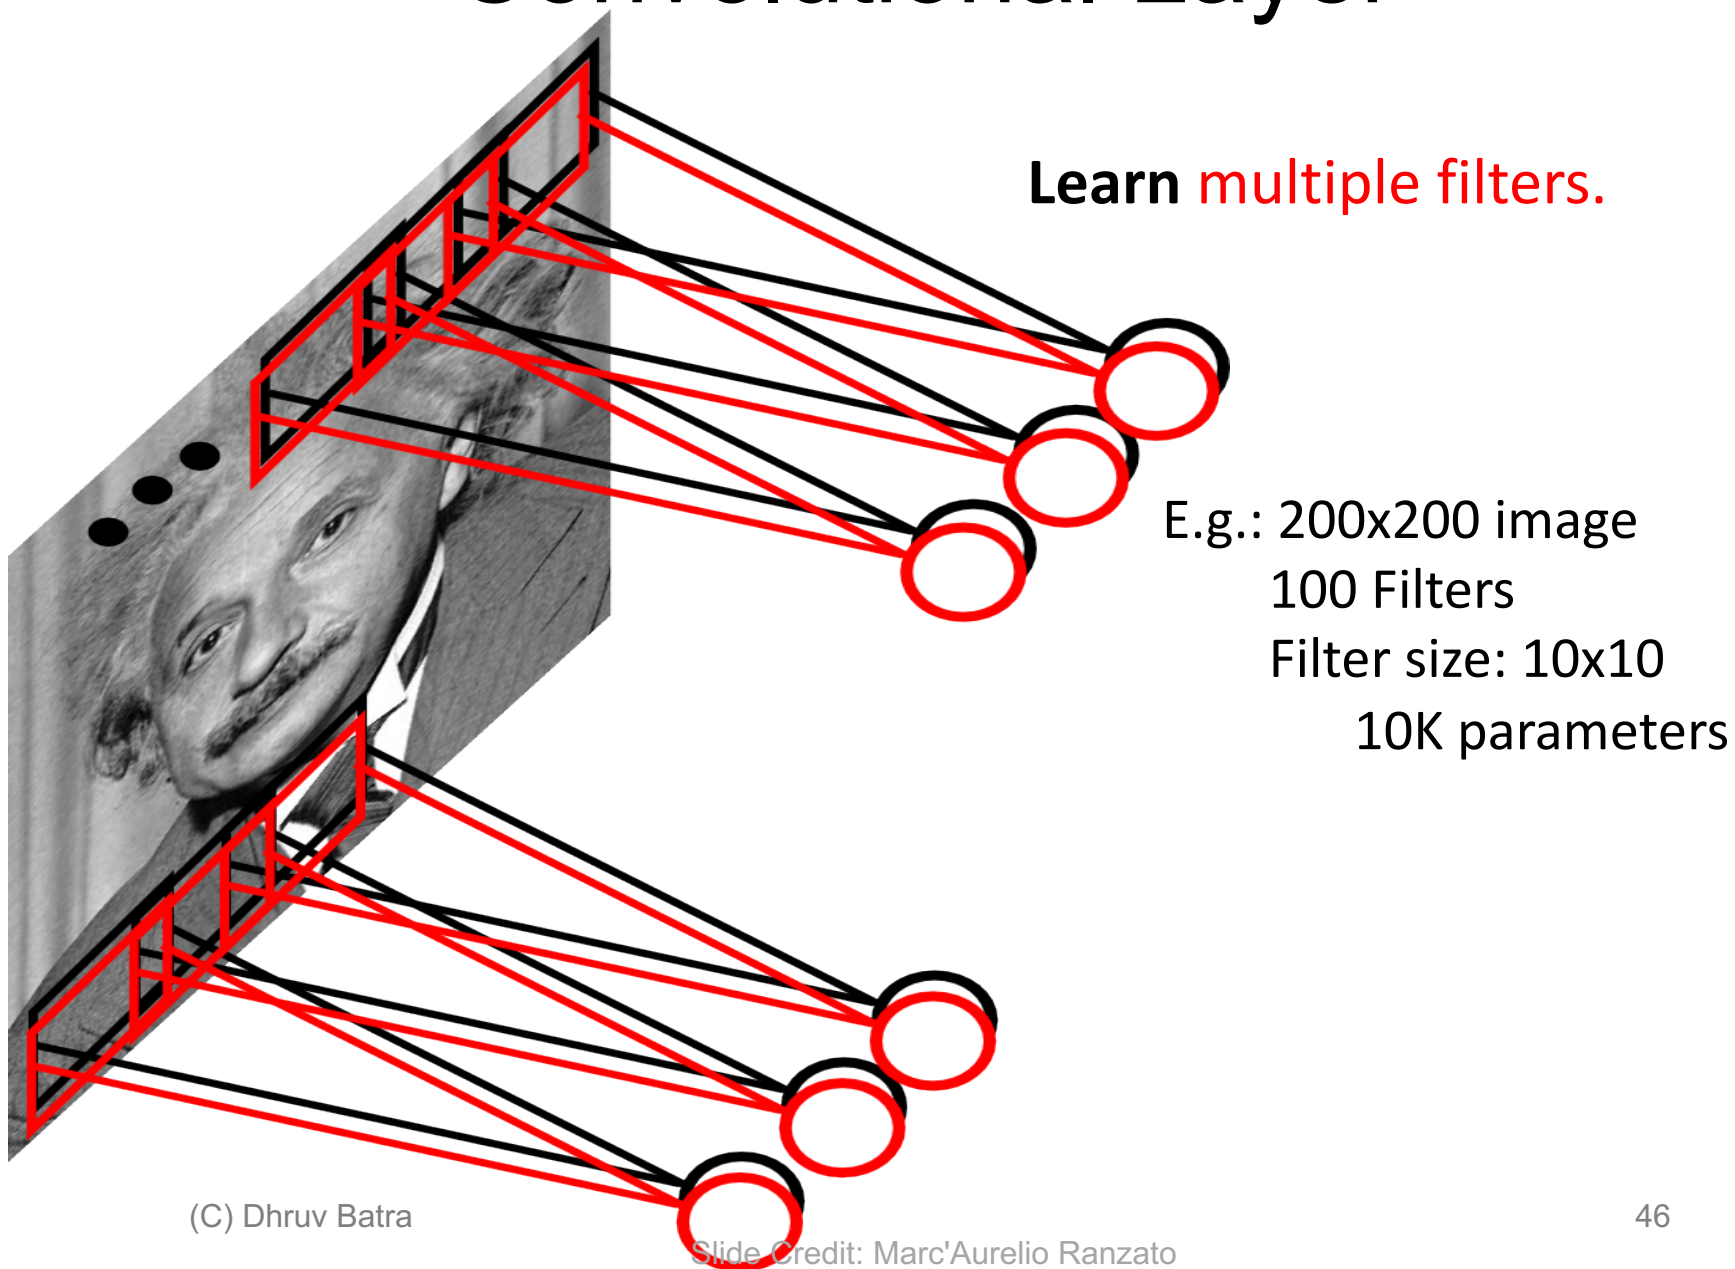
Convolutional Layer



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




Convolutional Layer




Convolution Layer


32x32x3 image -> preserve spatial structure



Convolution Layer




Multiple filters: if we have 6 5x5 filters, we'll get 6 separate activation maps:




We stack these up to get a “new image” of size 28x28x6!

Preview: ConvNet is a sequence of Convolution Layers, interspersed with activation functions




Preview: ConvNet is a sequence of Convolutional Layers, interspersed with activation functions



Preview

[Zeiler and Fergus 2013]

Visualization of VGG-16 by Lane McIntosh. VGG-16 architecture from [Simonyan and Zisserman 2014].




Low-level features


Mid-level features

High-level features


Linearly separable classifier



VGG-16 Conv1_1




VGG-16 Conv3_2

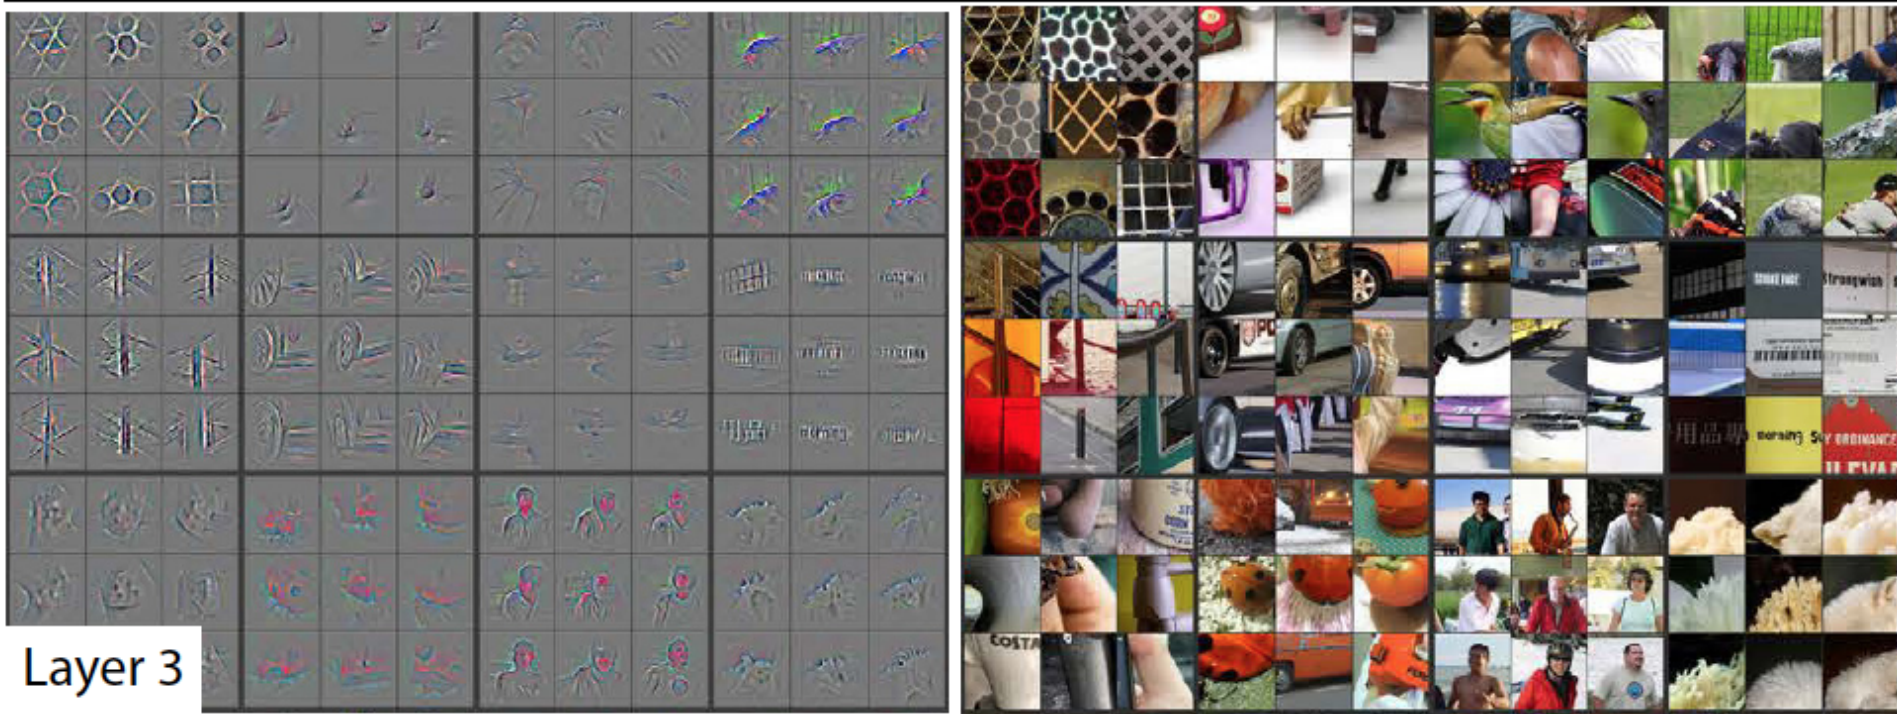


VGG-16 Conv5_3

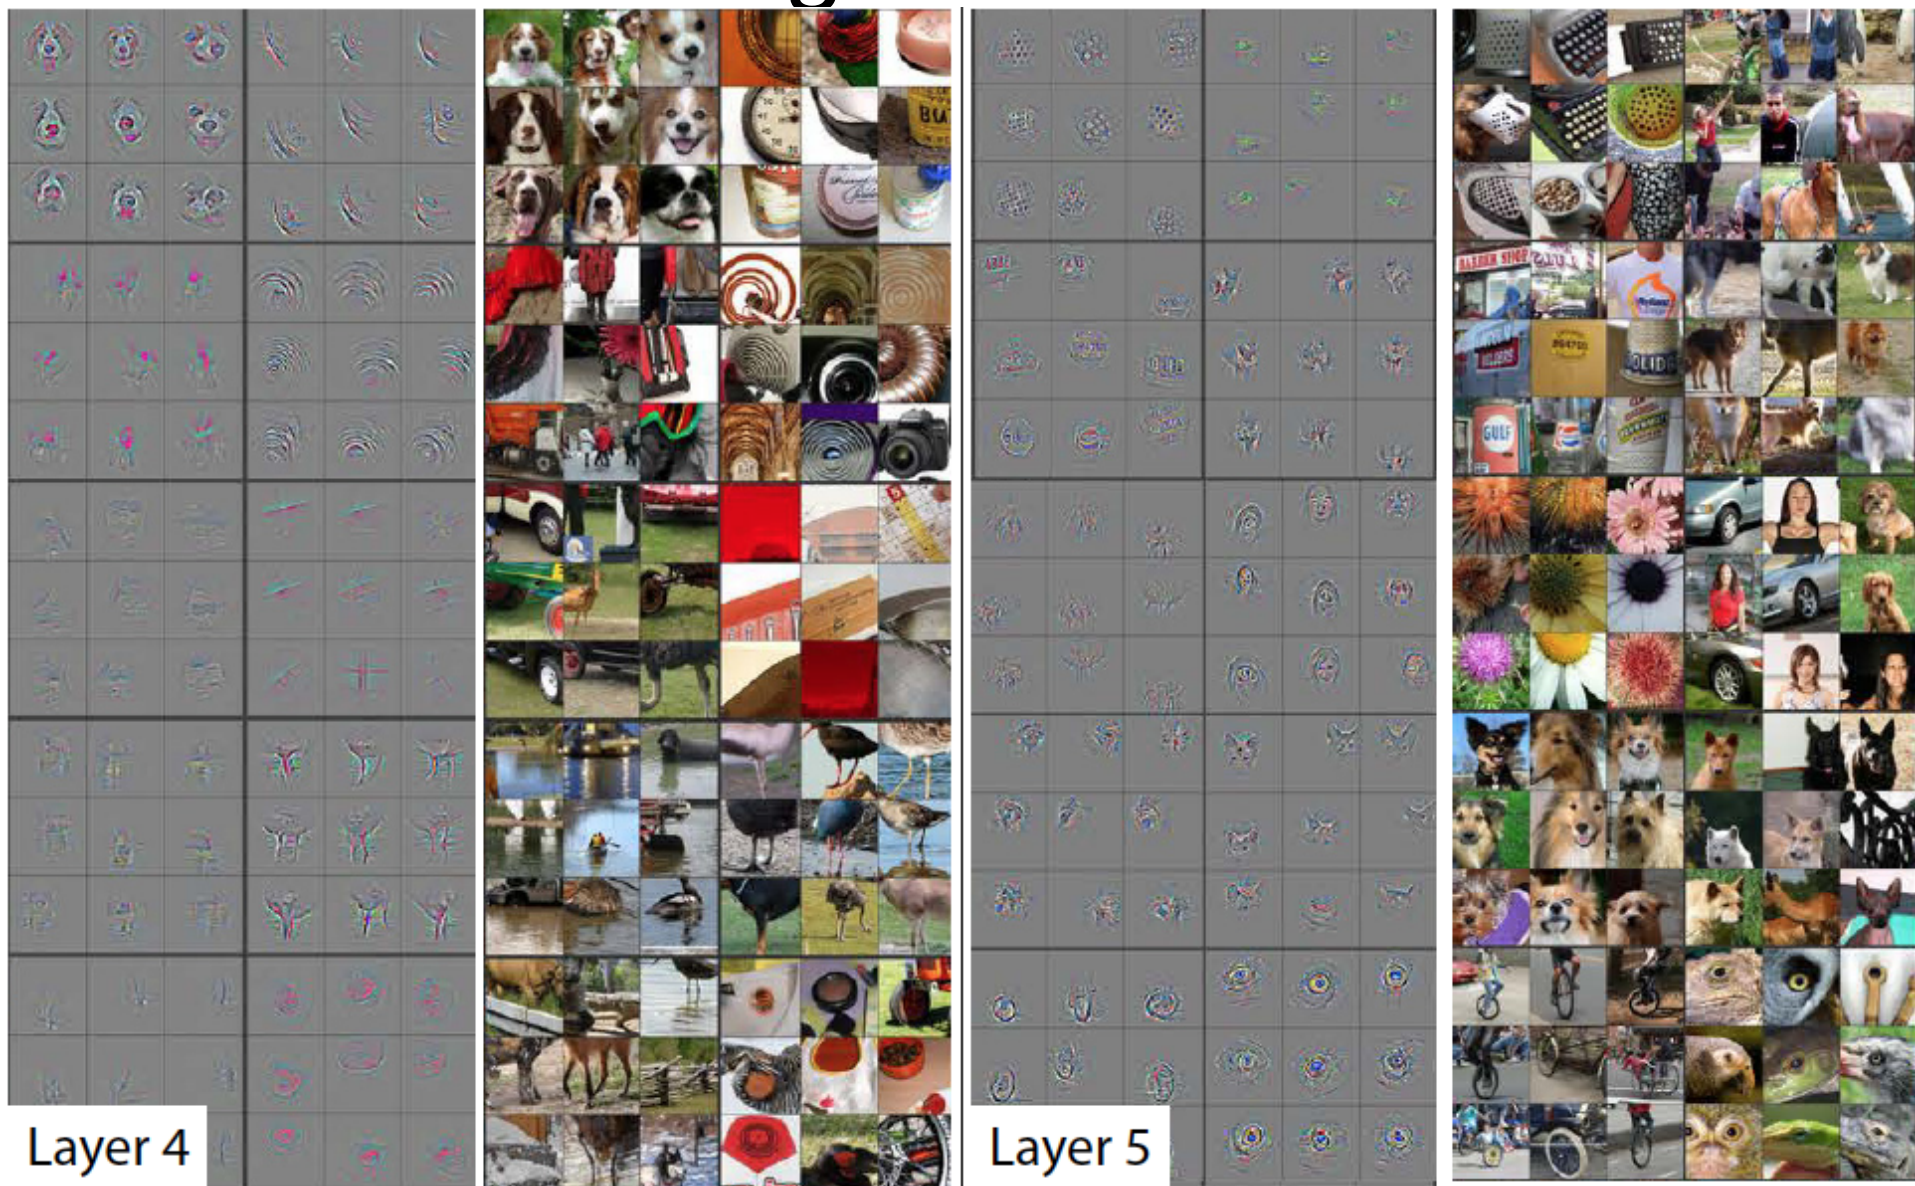
Visualizing Learned Filters



Visualizing Learned Filters




Visualizing Learned Filters



(C) Dhruv Batra

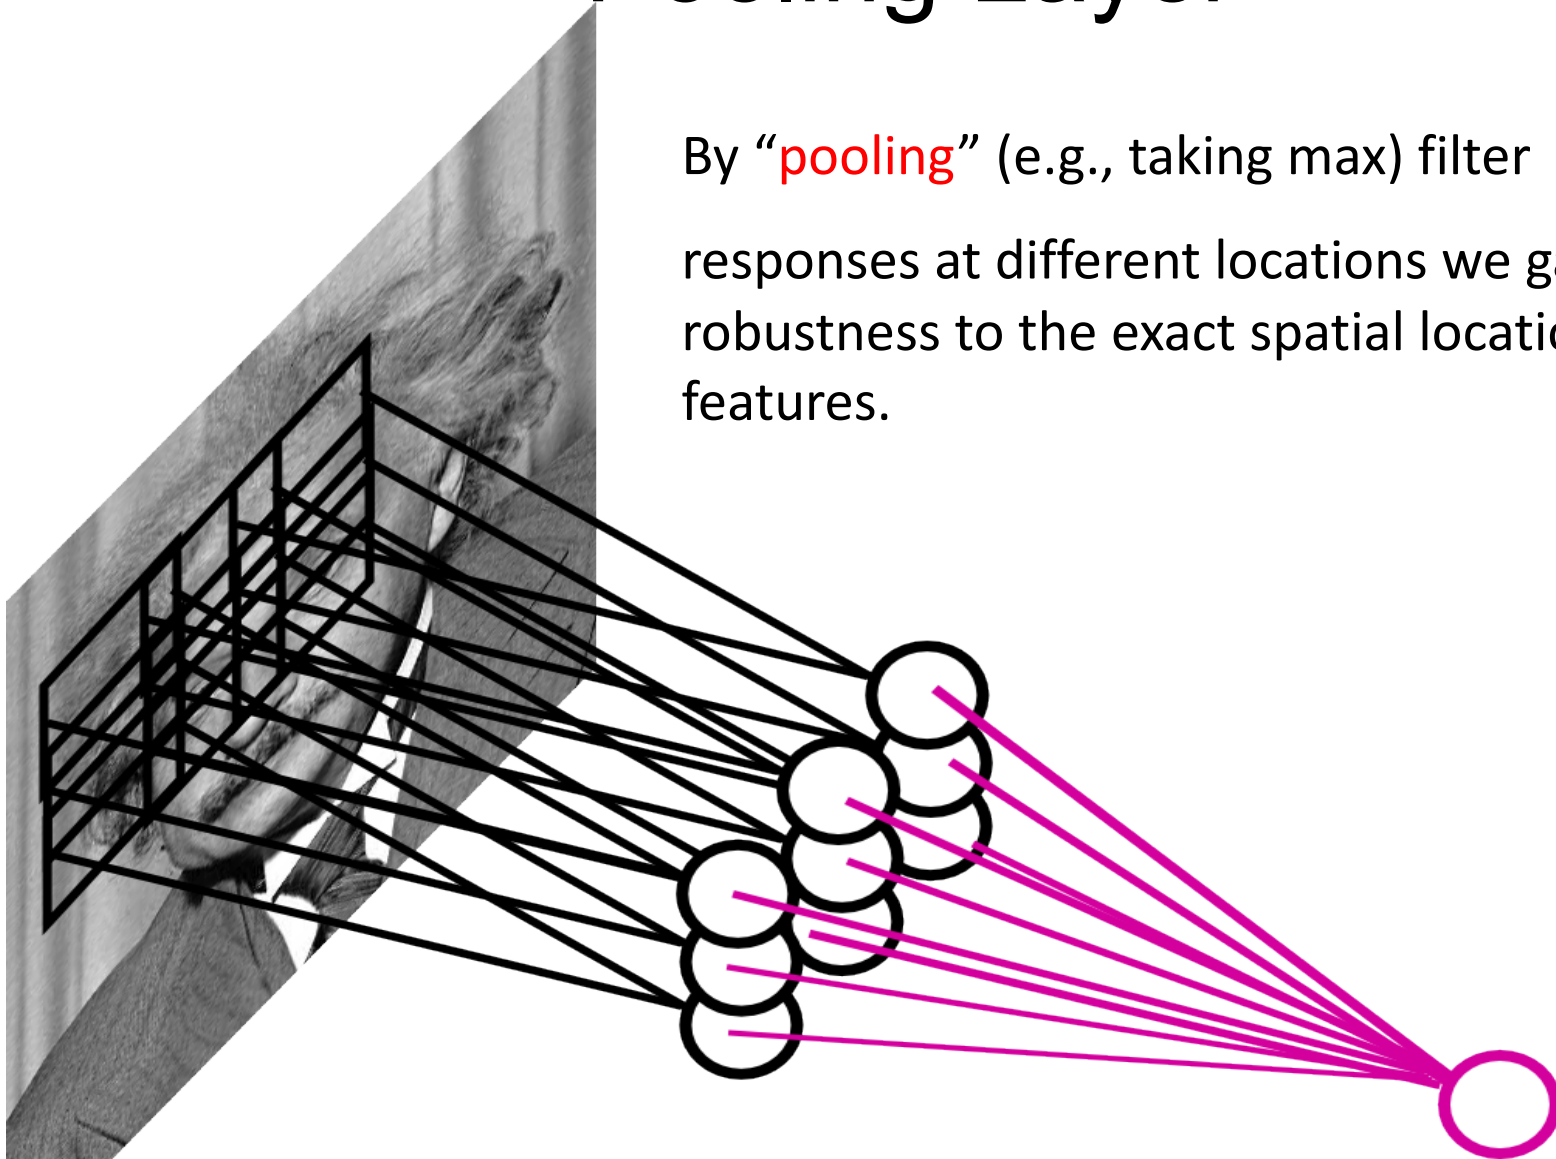
Figure Credit: [Zeiler & Fergus ECCV14]

two more layers to go: POOL/FC




Pooling Layer

By “pooling” (e.g., taking max) filter responses at different locations we gain robustness to the exact spatial location of features.




MAX POOLING

Single depth slice

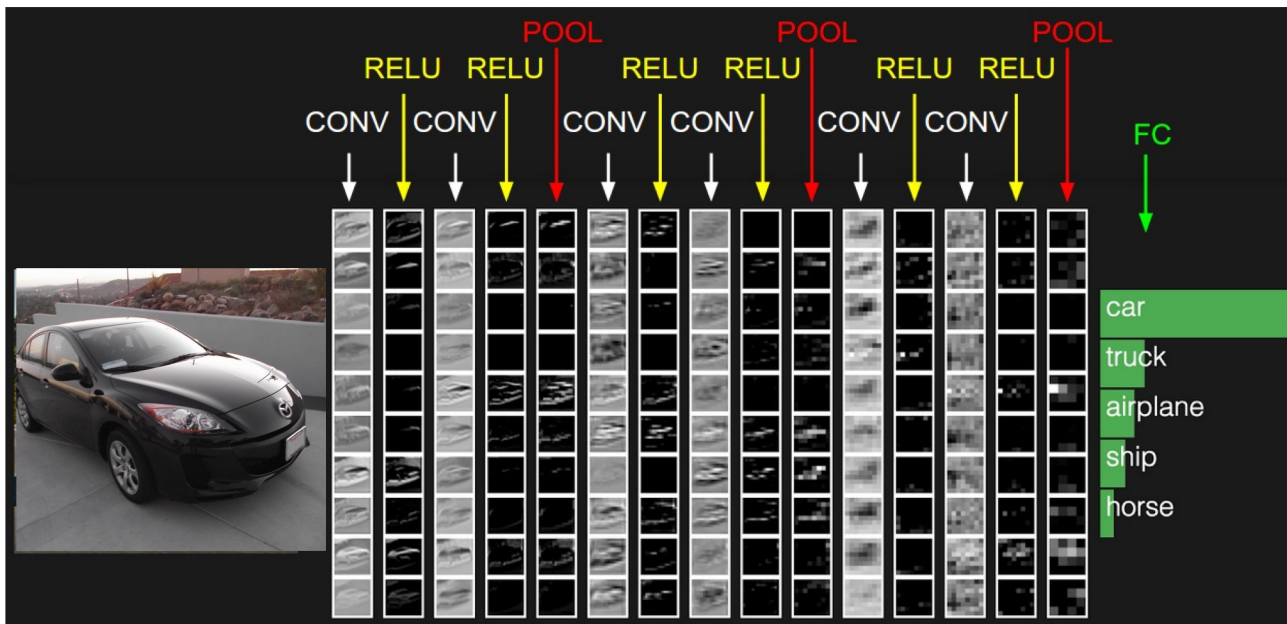


max pool with 2x2 filters
and stride 2



Fully Connected Layer (FC layer)

- Contains neurons that connect to the entire input volume, as in ordinary Neural Networks

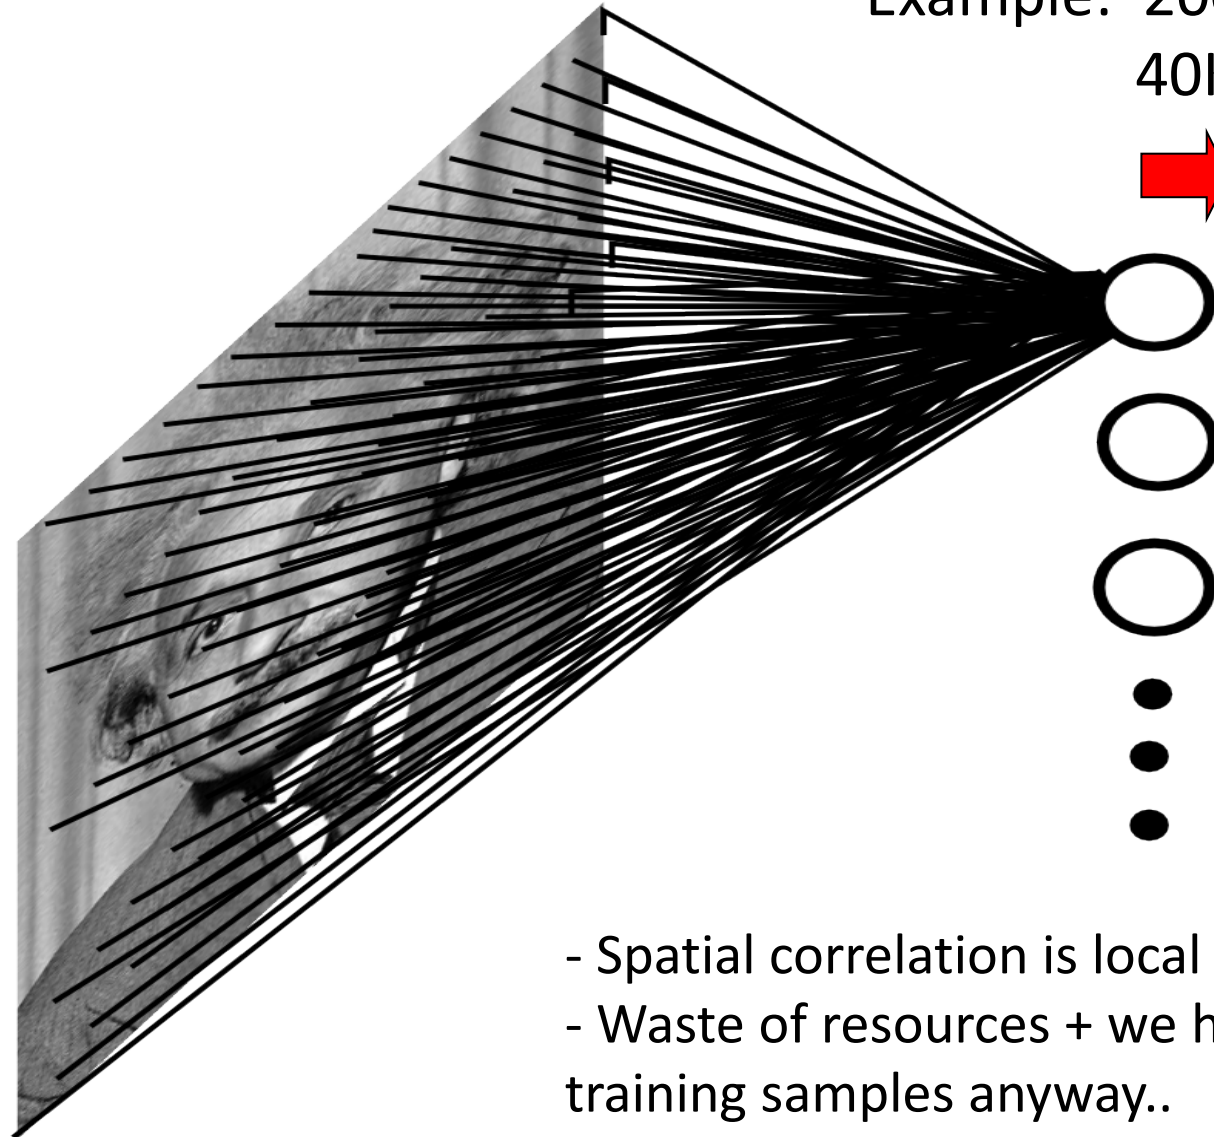


Fully Connected Layer

Example: 200x200 image

40K hidden units

→ ~2B parameters!!!




- Spatial correlation is local
- Waste of resources + we have not enough training samples anyway..

Fully Connected Layer

32x32x3 image -> stretch to 3072 x 1


Each neuron looks at the full input volume



1 number:
the result of taking a dot product between a row of W and the input (a 3072-dimensional dot product)


CNNs for Image Processing

Hybrid Images



- A. Oliva, A. Torralba, P.G. Schyns, [“Hybrid Images,”](#) SIGGRAPH 2006

Why do we get different, distance-dependent interpretations of hybrid images?







Colorization


- Given a grayscale image, colorize the image realistically
- Zhang et al. pose colorization as classification task and use class-rebalancing to improve results
- Demonstrate higher rates of fooling humans using “colorization Turing test”



Colorful Image Colorization. Richard Zhang, Phillip Isola, Alexei A. Efros. ECCV 2016.

Colorization

- Training data: decompose any RGB image into L^*a^*b color space
 - L : grayscale input (lightness channel)
 - ab : color channels
- Train CNN with **one million color images** and a new objective function to incorporate more diverse colors. Many possible correct colorizations!



Colorful Image Colorization. Richard Zhang, Phillip Isola, Alexei A. Efros. ECCV 2016.

How to convert the inferred distribution to an image?

- 313-way classification over discretized ab color bins
- Network will output a distribution \mathbf{z} over colors at each pixel. Need to convert to a single pixel value
 - Mode: vibrant but sometimes spatially inconsistent (e.g., the red splotches on the bus)
 - Mean: produces spatially consistent but desaturated results, exhibiting an unnatural sepia tone



$$\mathcal{H}(\mathbf{Z}_{h,w}) = \mathbb{E}[f_T(\mathbf{Z}_{h,w})], \quad f_T(\mathbf{z}) = \frac{\exp(\log(\mathbf{z})/T)}{\sum_q \exp(\log(\mathbf{z}_q)/T)}$$

DeOldify



Super-Resolution

Low resolution



High resolution



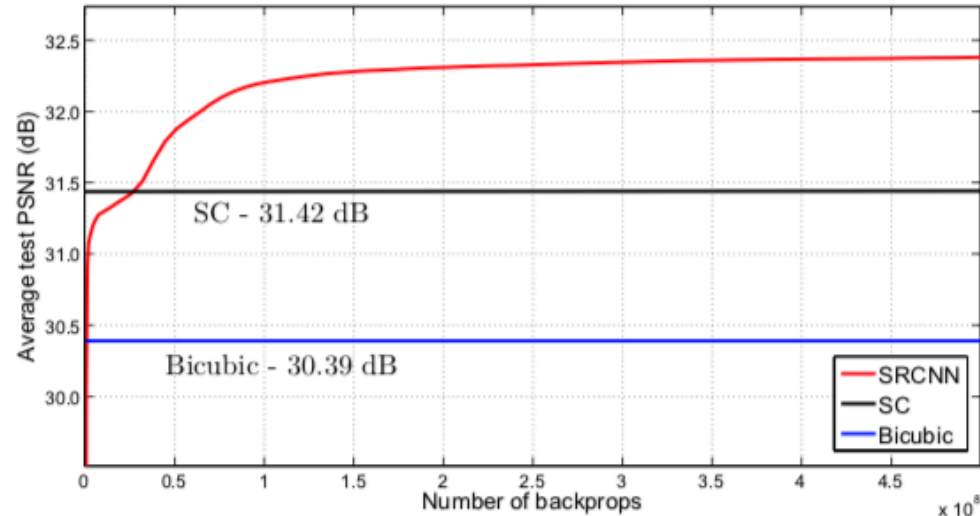
Super-Resolution as a task

- Quality-degrading factors / sources of noise:
 - Camera shake, shadows, motion blur, radial distortion from fisheye/GoPro type cameras, poor contrast, poor lighting, lossy compression, transmission defects, dust, haze, smoke, and mist, motion of the camera sensor platform, moving objects captured within the observed scene, e.g. people and vehicles.
- How to measure super-resolution?
 - Peak signal-to-noise ratio (PSNR), higher is better. Relies upon the Mean Square Error (MSE) error metric to evaluate image compression quality between two images:

$$MSE = \frac{1}{MN} \sum_M \sum_N [I_1(m, n) - I_2(m, n)]^2 = \|I_1 - I_2\|_F$$

$$PSNR = 10 \log_{10} \left(\frac{R^2}{MSE} \right)$$

An early CNN paper (2016)



An early CNN paper (2016)




Upscaling factor of 3 !


Dong, Chao, et al. "Learning a deep convolutional network for image super-resolution." *European conference on computer vision*. Springer, Cham, 2014.

Underexposed Photo Enhancement


- Goal: enhance extreme low-light imaging with severely limited illumination (e.g., moonlight) and short exposure (exposure time is set to 1/30 second)
- The less light there is, the more ISO you need
 - High ISO can be used to increase brightness, but amplifies noise
 - Leads to low signal-to-noise ratio (SNR) due to low photon counts




(a) Camera output with ISO 8,000




(b) Camera output with ISO 409,600




(c) Our result from the raw data of (a)



(a) JPEG image produced by camera



(b) Raw data via traditional pipeline



(c) Our result

Solution? Collect dataset and train a deep network

- See-in-the-Dark (SID) dataset contains 5094 raw short exposure images, each with a corresponding long-exposure reference image
- Corresponding reference (ground truth) images captured with 100-300x longer exposure (i.e. 10 to 30 seconds)
- Overcome low photon counts!
- Train deep neural networks to learn the image processing pipeline w/ L1 loss.




Figure 2. Example images in the SID dataset. Outdoor images in the top two rows, indoor images in the bottom rows. Long-exposure reference (ground truth) images are shown in front. Short-exposure input images (essentially black) are shown in the back. The illuminance at the camera is generally between 0.2 and 5 lux outdoors and between 0.03 and 0.3 lux indoors.

Underexposed Photo Enhancement

- Learn image-to-image mapping? Too hard!
- Instead estimate an image-to-illumination mapping (model varying-lighting conditions)
 - Illumination maps for natural images typically have relatively simple forms with known priors
- Then take illumination map to light up the underexposed photo.
- Minimize (reconstruction loss + smoothness loss + color loss)

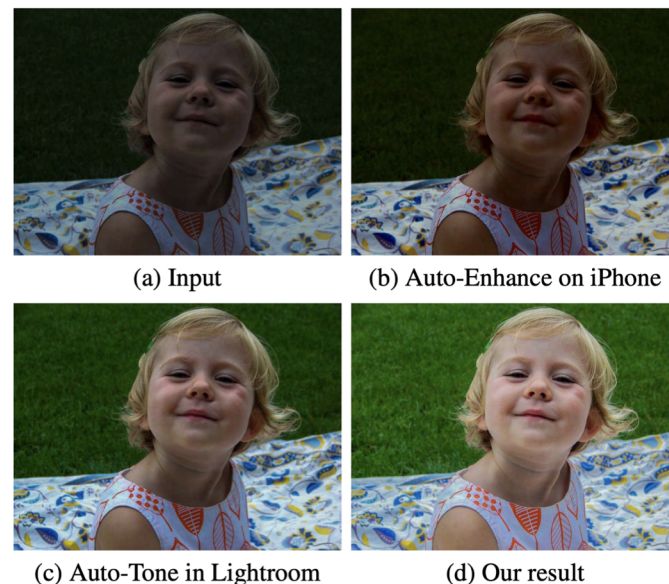


Figure 1: A challenging underexposed photo (a) enhanced by various tools (b)-(d). Our result contains more details, distinct contrast, and more natural color.

Image Inpainting

- Perceptual loss is added to ELBO, the typical objective function used in variational autoencoders, to increase the sharpness and overall quality of inpainted images
- Demonstrate results on attribute-guided image completion

$$\mathcal{L}_{recon} = \|x_{gen} - x_{gt}\|^2 + \sum_l \lambda_l \|\eta_l(x_{gen}) - \eta_l(x_{gt})\|^2$$

x_{gen} : generated image

x_{gt} : ground truth image

η_l : activation of the l^{th} layer of a pre-trained VGG




Image Inpainting

- Proposes partial convolutions, comprised of a masked & re-normalized convolution operator
- Updates mask automatically after partial convolutions, removing any masking where partial convolution was able to operate on unmasked value




Image Inpainting for Irregular Holes Using Partial Convolutions. Liu et al. ECCV 2018.