Deep Learning in Biomedical Imaging: Analysis and Reconstruction

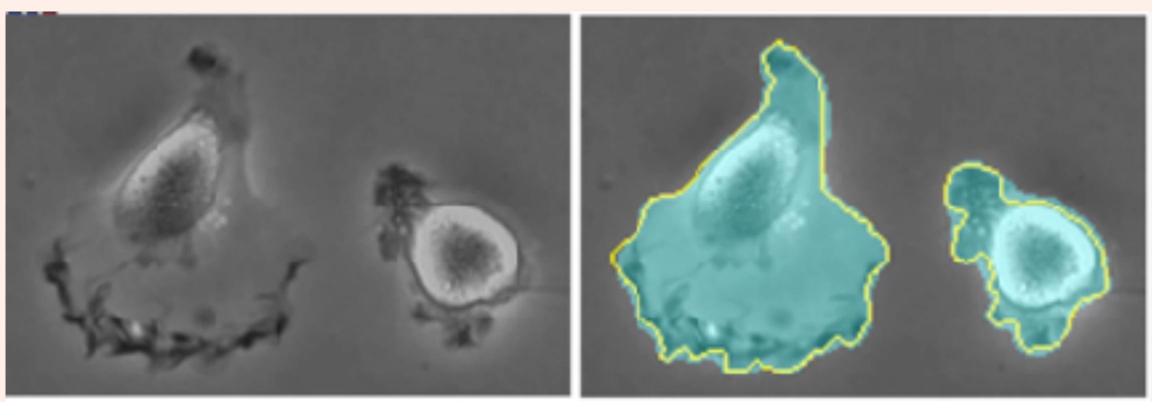

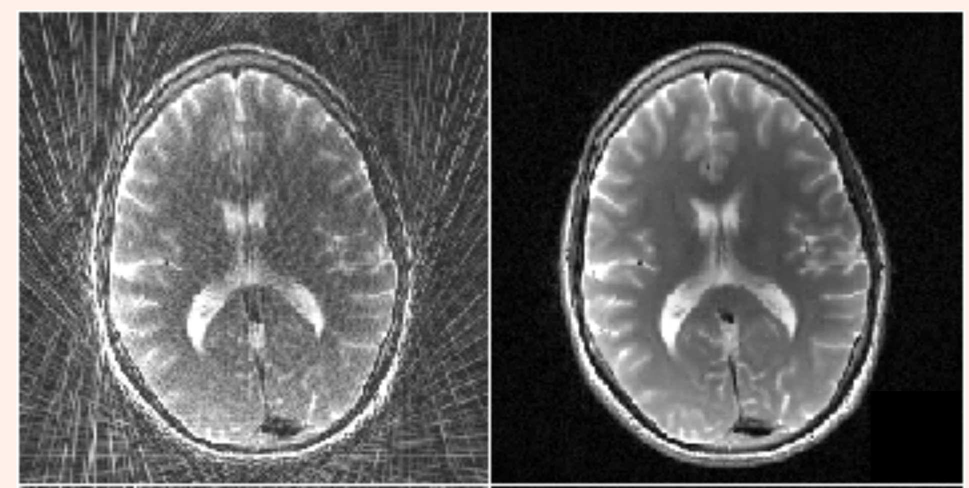

Greg Ongie
Postdoc, Department of Statistics
University of Chicago

Guest Lecture
Machine Learning for Biomedical Informatics
August 22, 2019

Biomedical imaging pipeline

Aquire Data → Reconstruct Images → Process Images → Analyze Diagnose Quantify Interpret Intervene

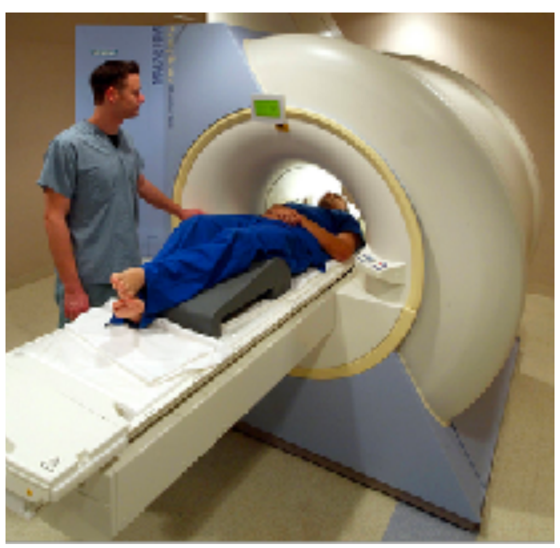

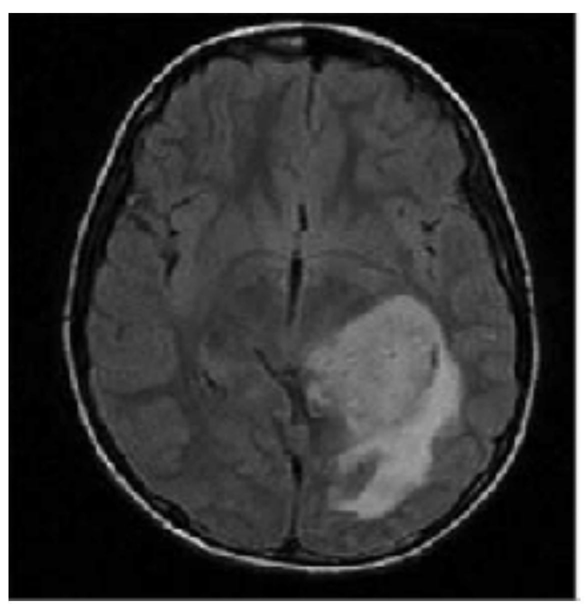

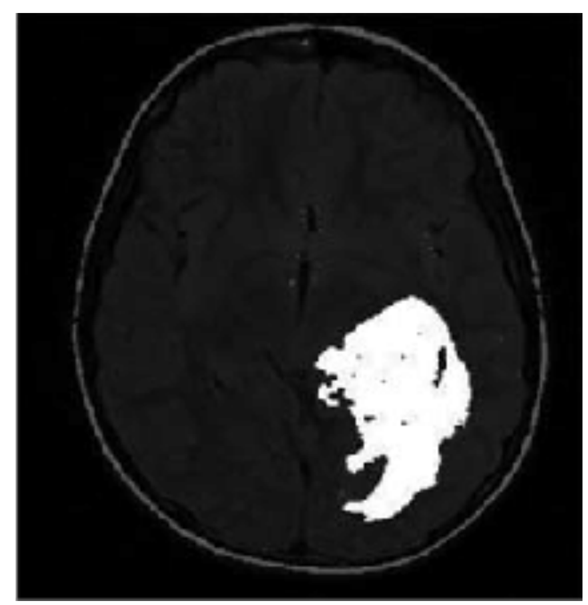

Machine learning in biomedical imaging

Most obvious place for machine learning in post-processing

reconstructed images

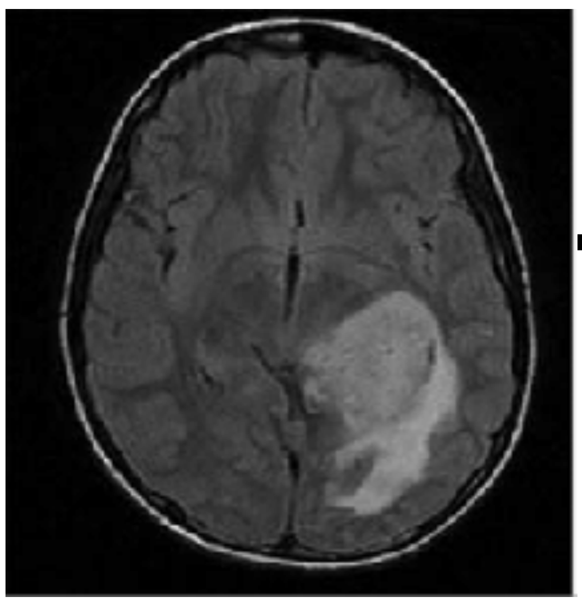

ML-based image processing (classification or regression)

interpretation

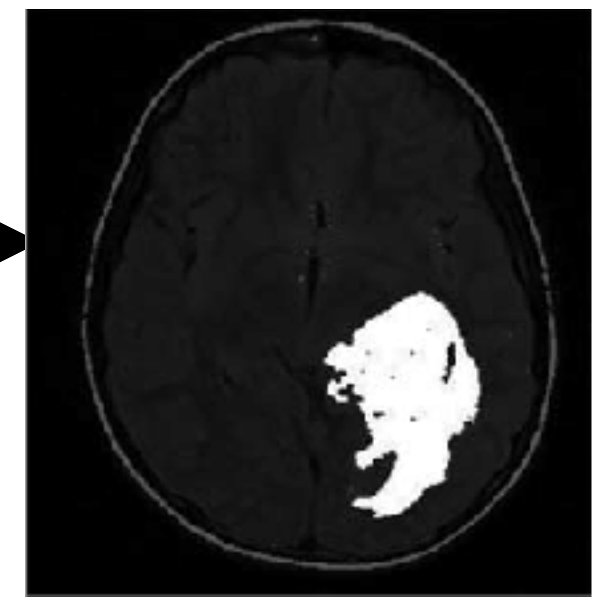

Tumor: Yes or No?

Deep learning in medical image analysis

Deep learning methods achieve state-of-the-art results on a wide variety of **image analysis** tasks:

- mammography mass classification
- segmentation of lesions in the brain
- leak detection in airway tree segmentation
- diabetic retinopathy classification
- prostate classification
- lung nodule classification

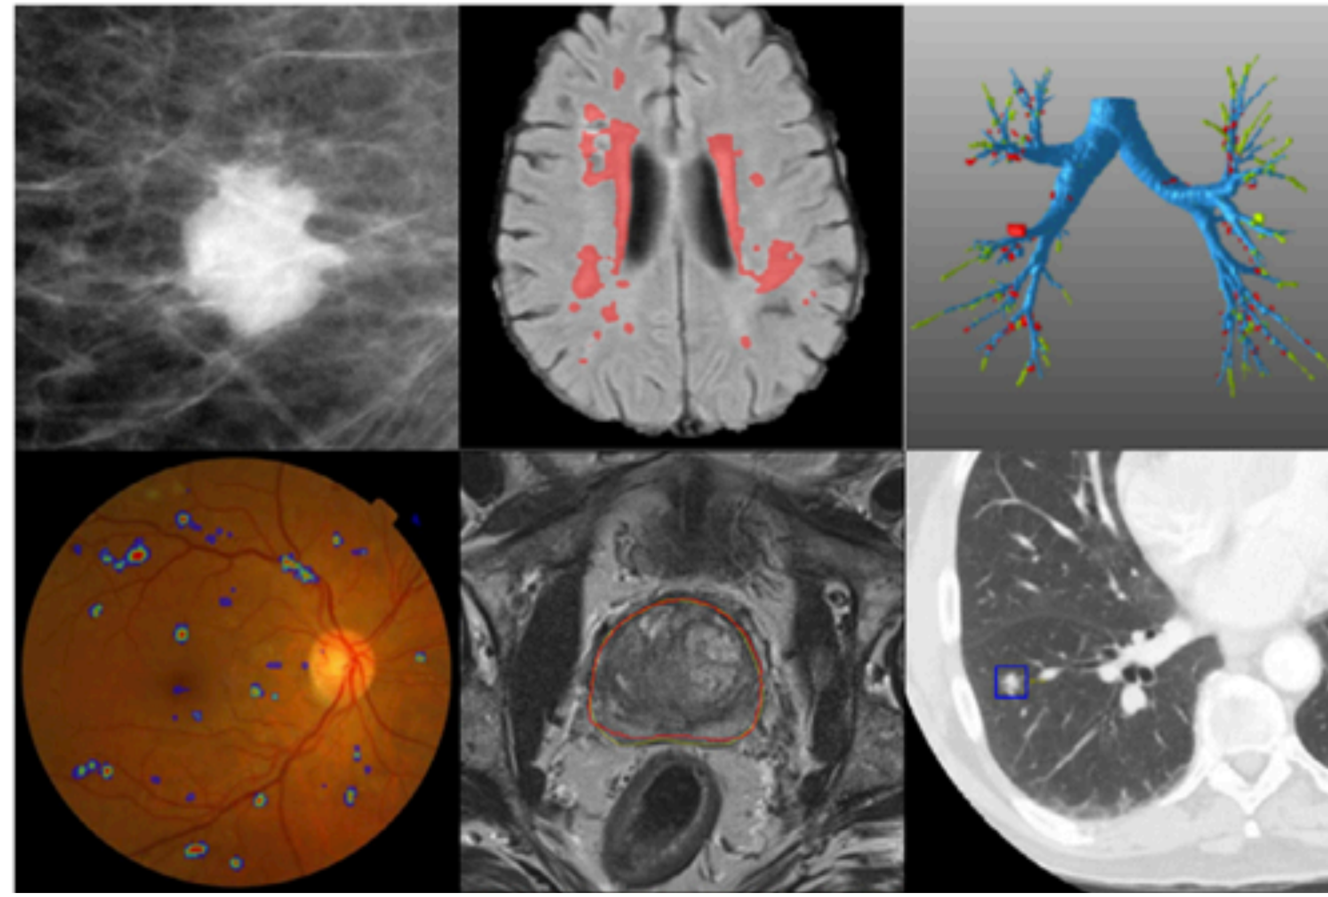

ELSEVIER

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Survey Paper

A survey on deep learning in medical image analysis

Geert Litjens*, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A.W.M. van der Laak, Bram van Ginneken, Clara I. Sánchez

Diagnostic Image Analysis Group, Radboud University Medical Center, Nijmegen, The Netherlands

Machine learning for image recon?

Another (initially less obvious?) place for machine learning: **image recon**

raw data

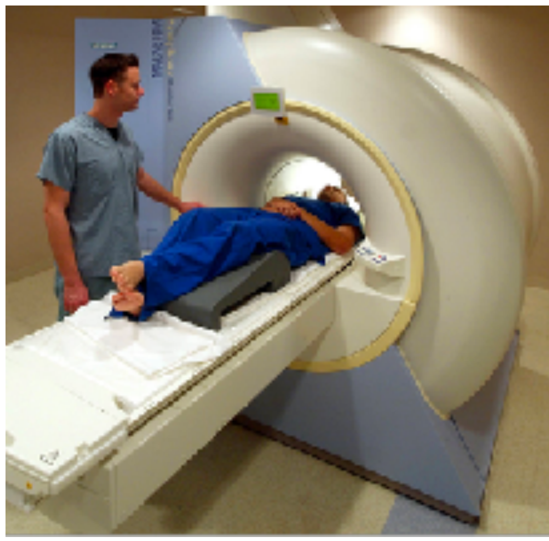

ML-based image reconstruction

images

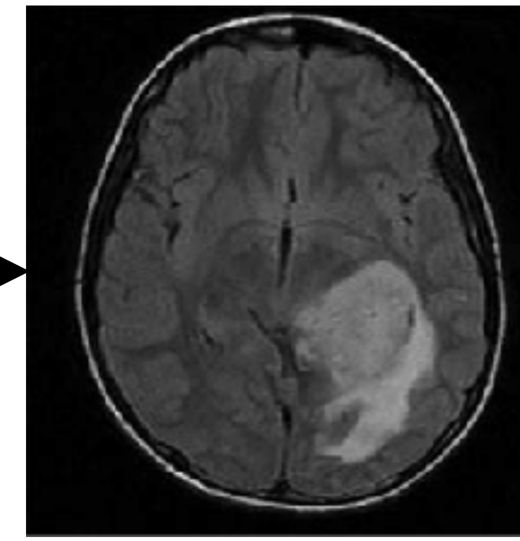

Possibly easier (than diagnosis) due to lower bar:

- current reconstruction methods based on simplistic image models
- human eyes are better at detection (tumor vs. no tumor) than they are at converting raw data to images

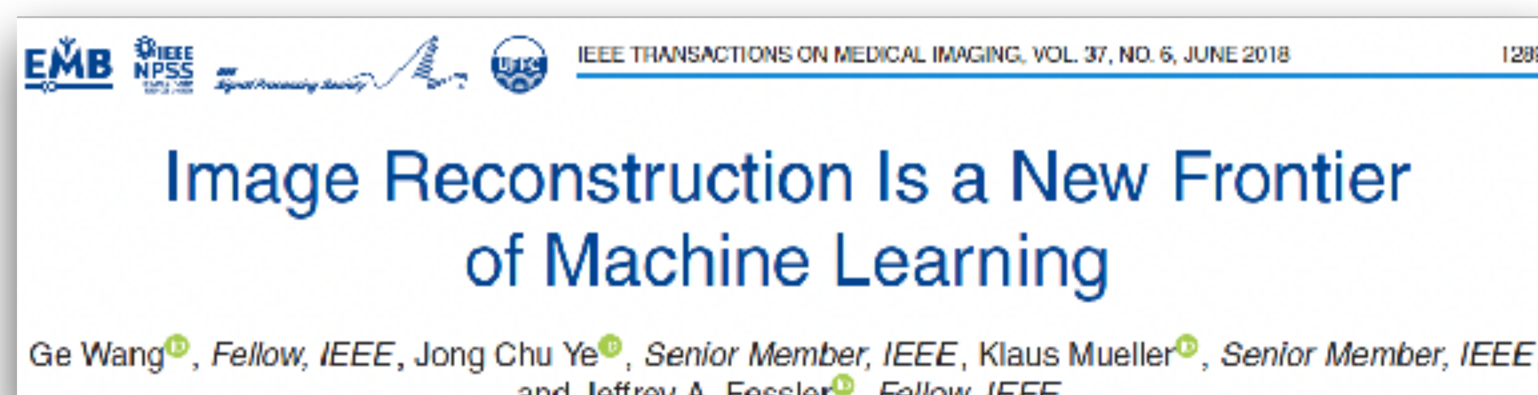

IEEE TRANSACTIONS ON MEDICAL IMAGING, VOL. 37, NO. 6, JUNE 2018 1289

Image Reconstruction Is a New Frontier of Machine Learning

Ge Wang, *Fellow, IEEE*, Jong Chu Ye, *Senior Member, IEEE*, Klaus Mueller, *Senior Member, IEEE*, and Jeffrey A. Fessler, *Fellow, IEEE*

Outline:

I. Deep learning for biomedical image analysis (60 min)
   1. The CNN zoo
   2. Image classification/detection tasks
   3. Image segmentation with the U-net

II. Deep learning for medical image reconstruction (60 min)
   1. Medical image reconstruction basics
   2. Learning to “enhance”
   3. Training generative models
   4. Unrolling of optimization algorithms

Part I:
Deep learning for biomedical image analysis

The CNN Zoo


ZOO
Inception-v3
ResNet
Inception-v1
FractalNet
VGGNet
AlexNet
DenseNet

TensorFlow demo app

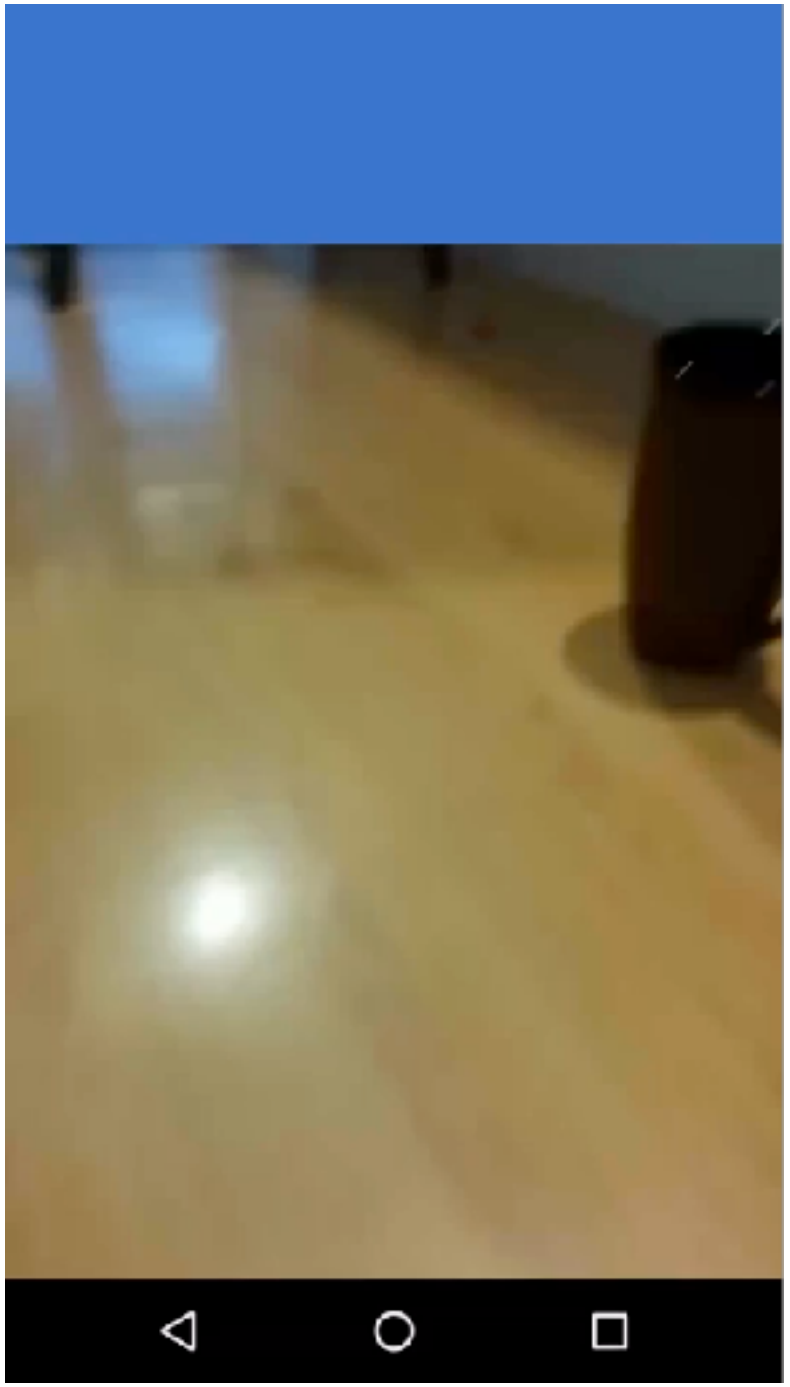

https://www.youtube.com/watch?v=4oU4N6bAjR4

Convolutional Neural Networks (CNN)

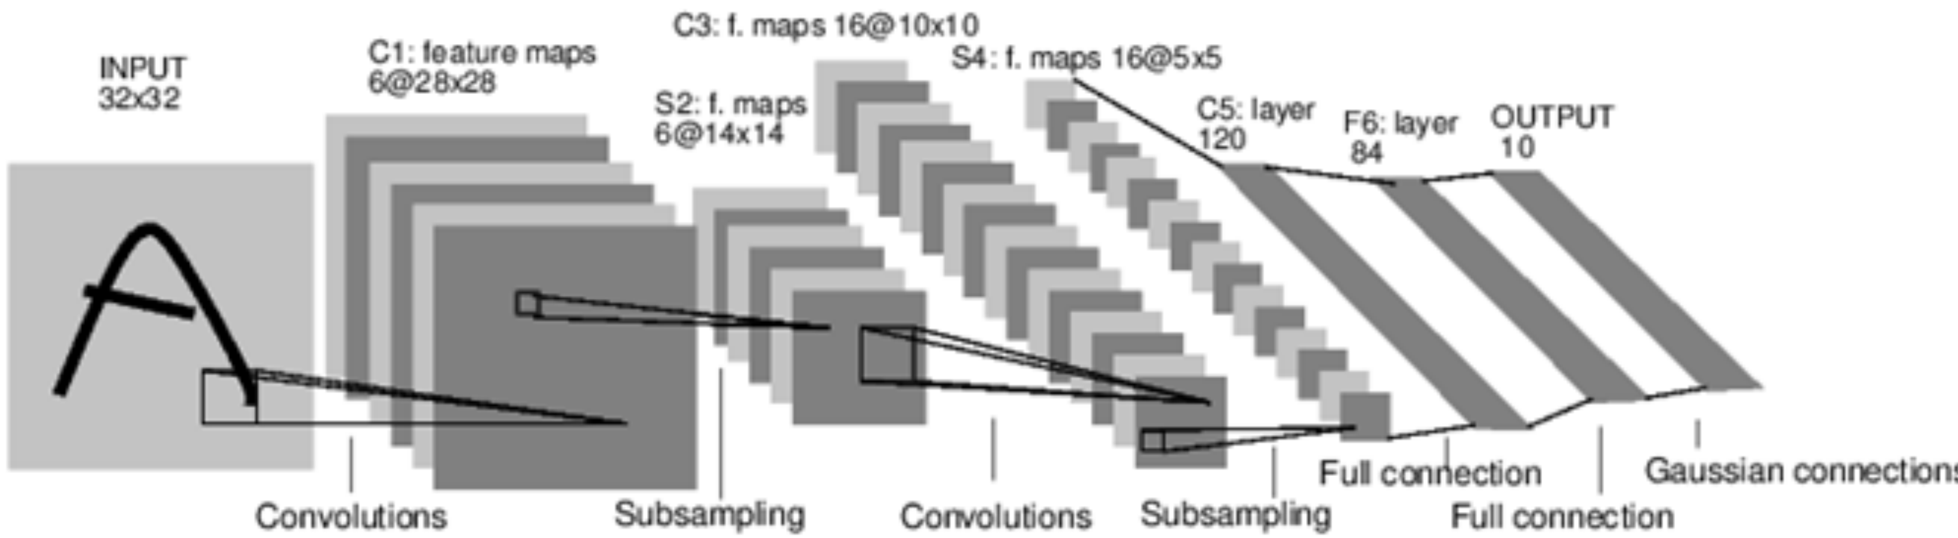

A Full Convolutional Neural Network (LeNet)

convolution:

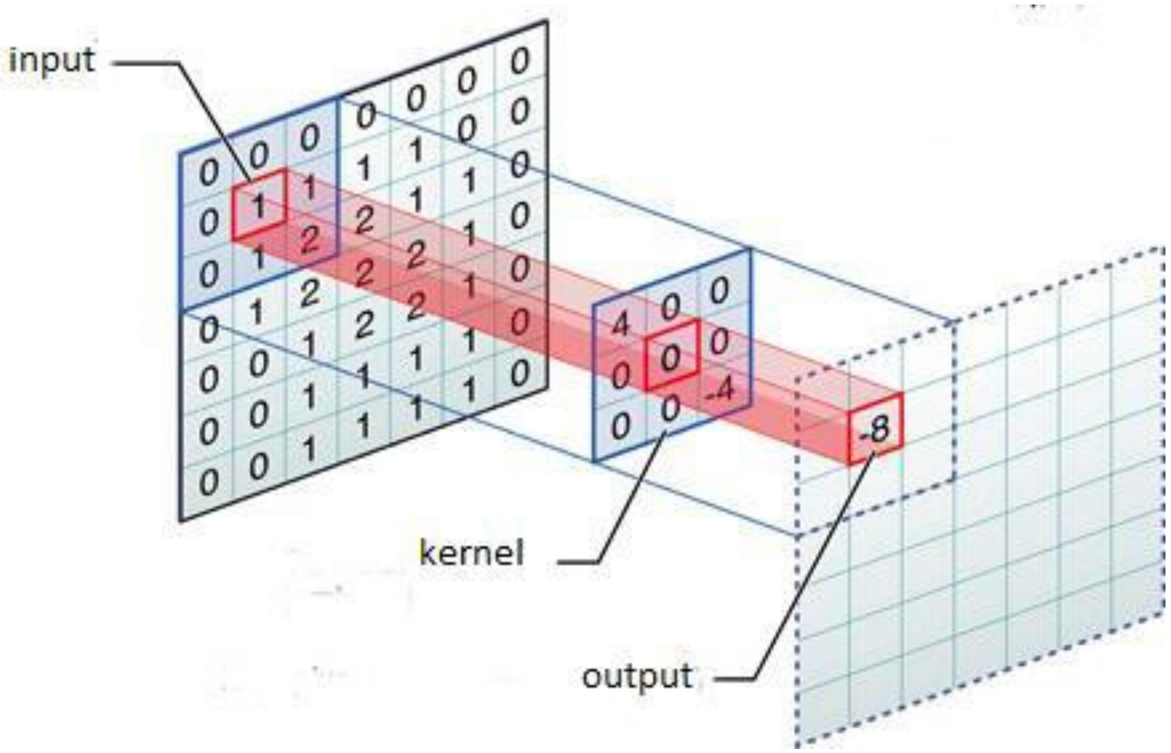

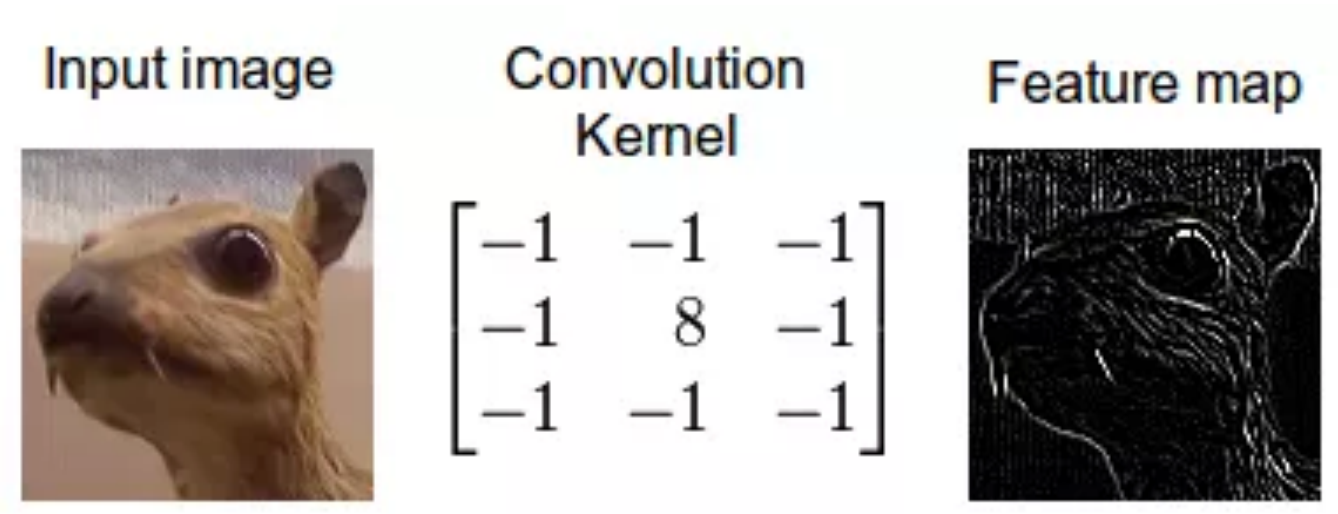

filters learned from data

CNN Example

8 Layer Architecture!

5 convolutional layers

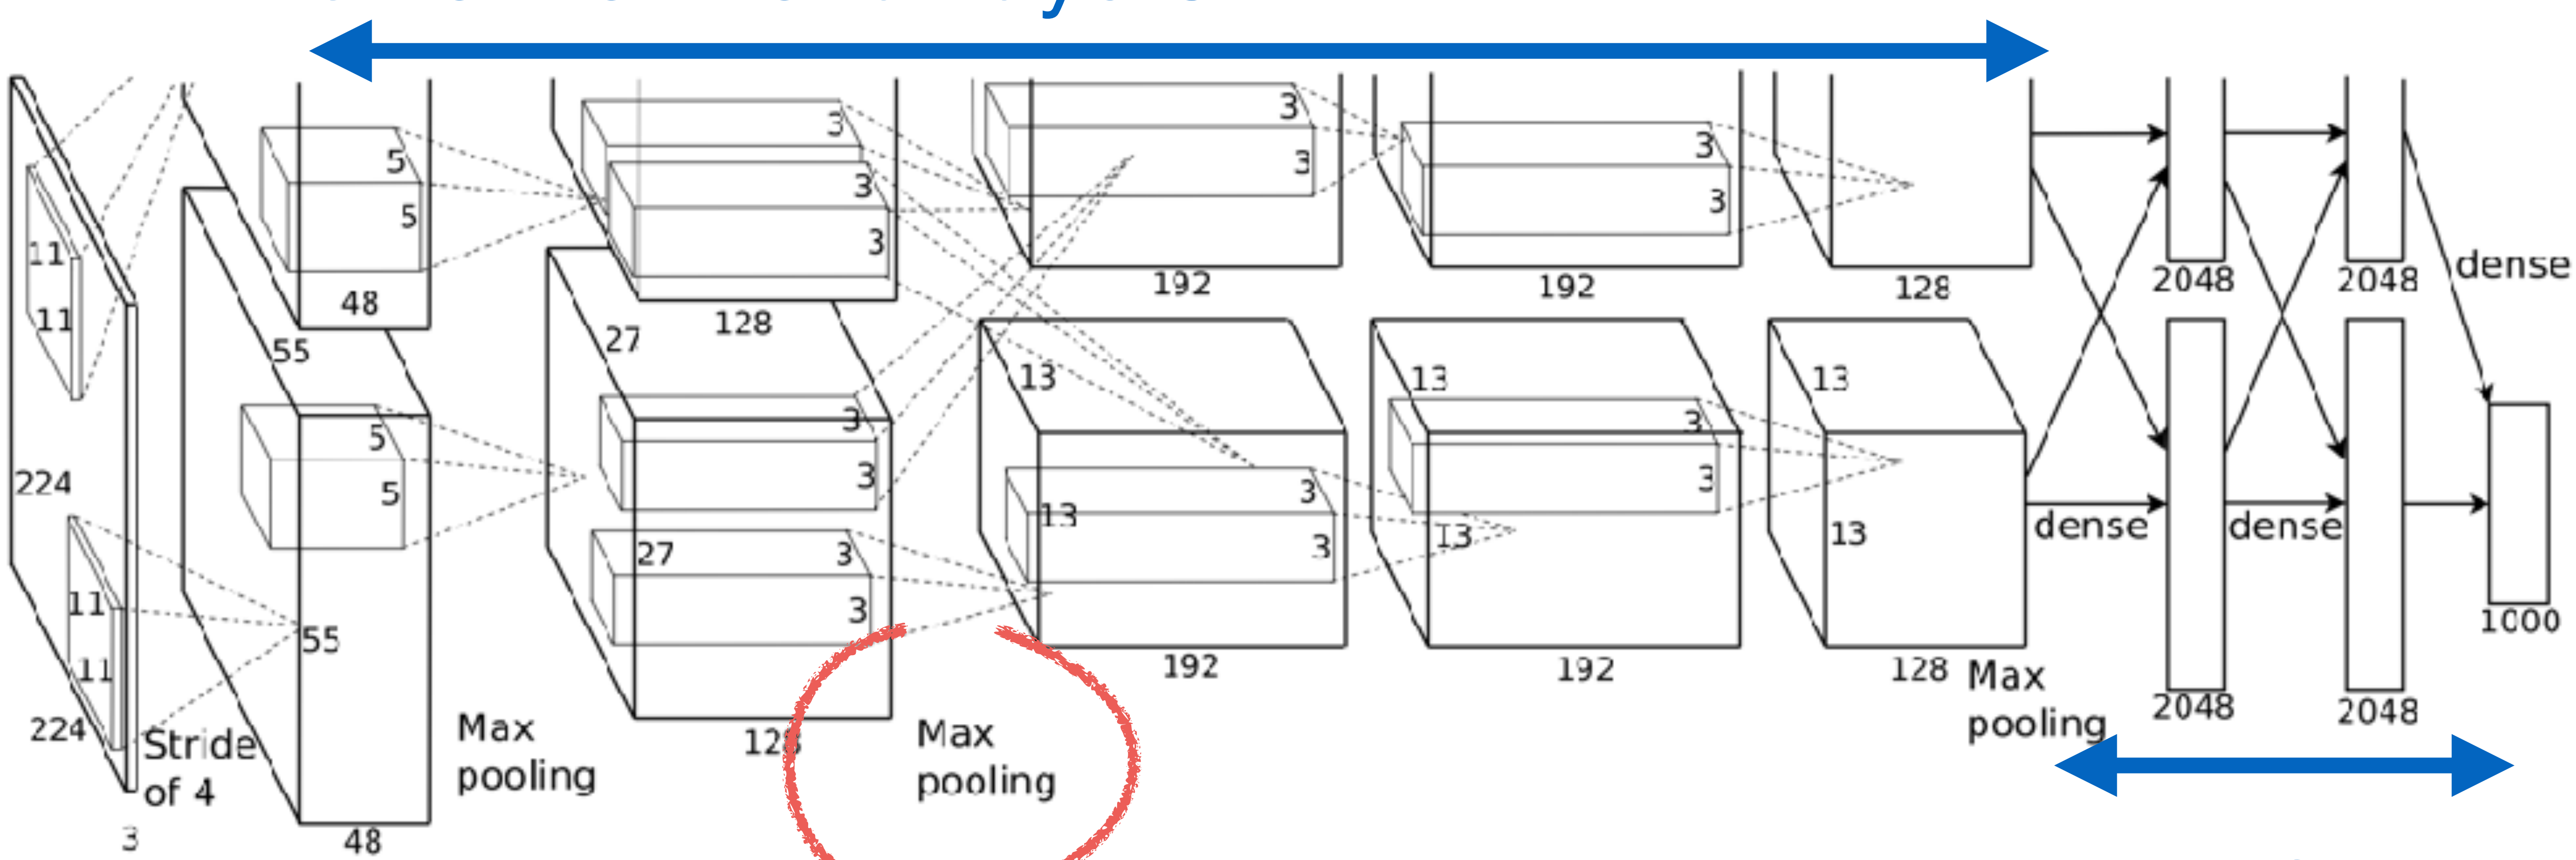

+ 3 fully connected layers

downsampling

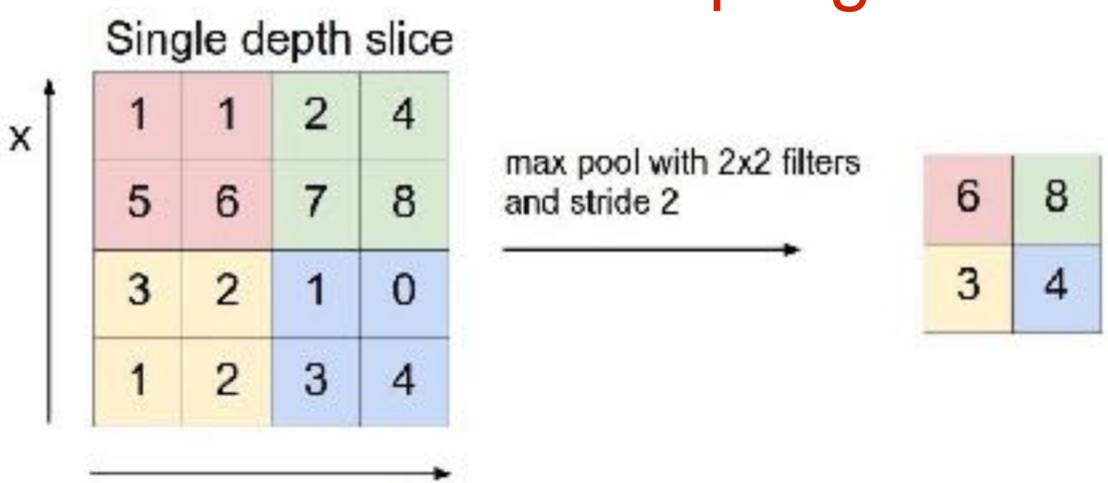

data -> **conv1** -> pool1 -> conv2 -> pool2 -> conv3 -> conv4 -> conv5 -> pool3

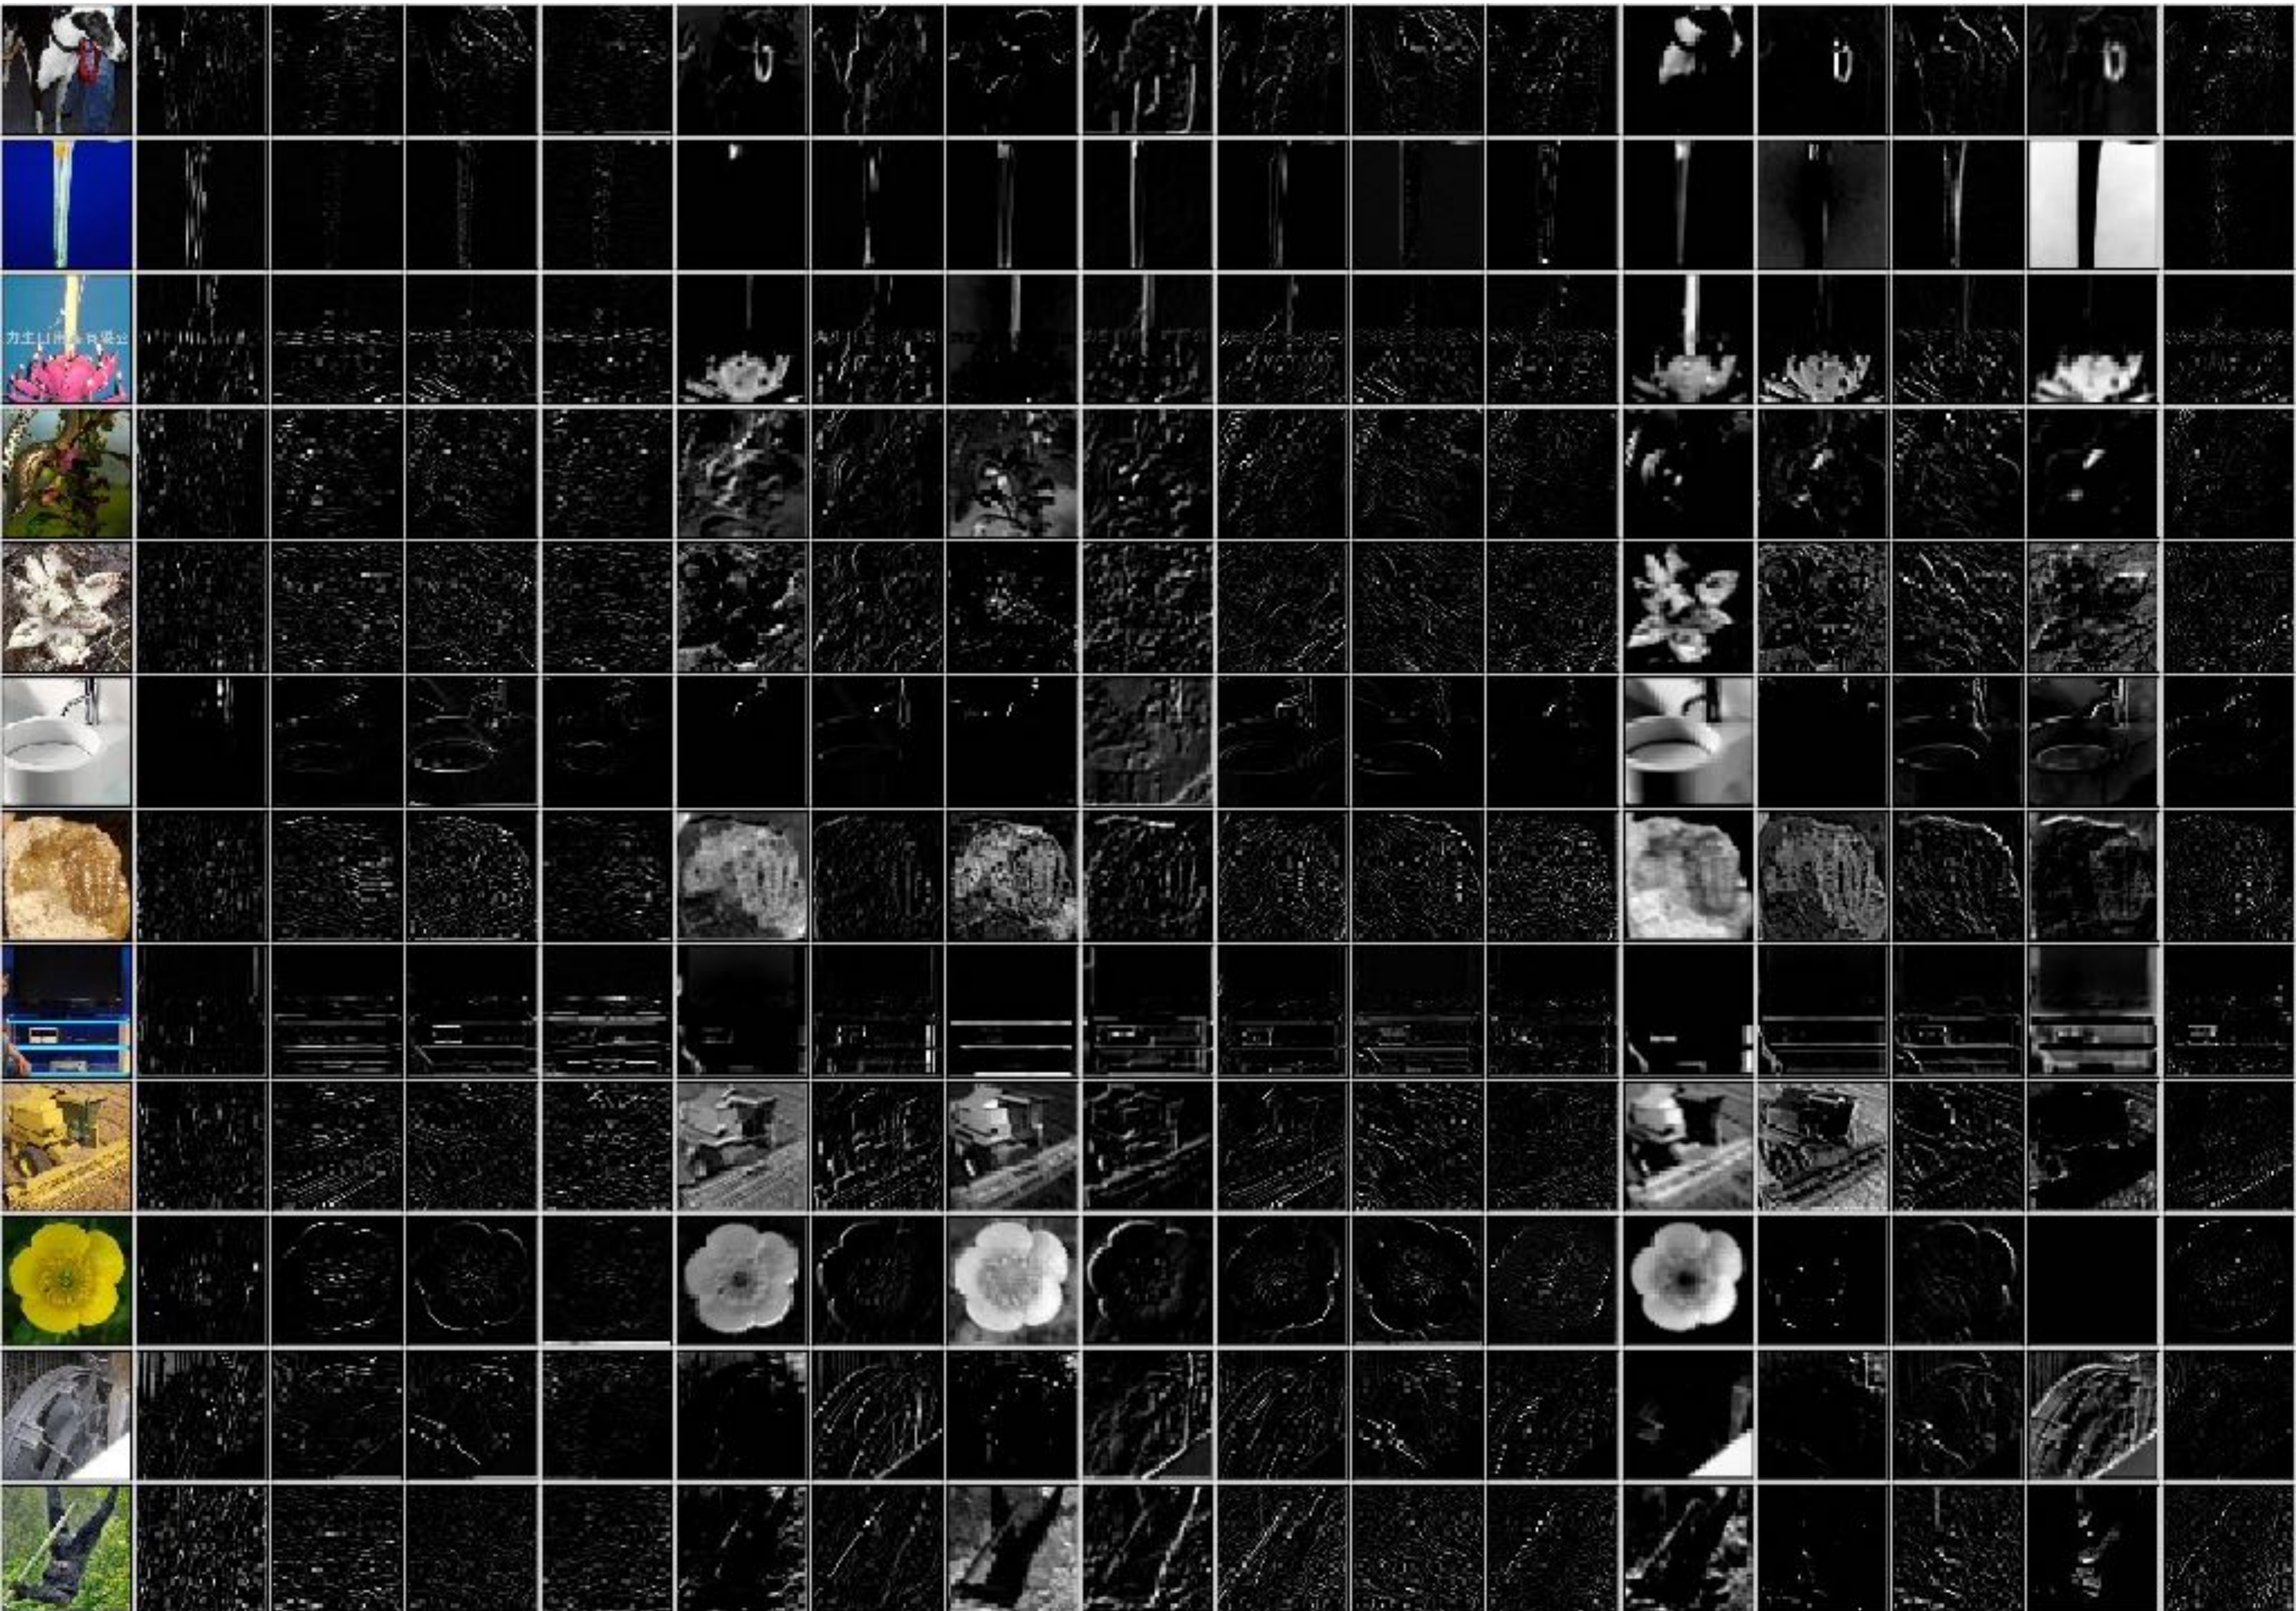

data -> conv1 -> **pool1** -> conv2 -> pool2 -> conv3 -> conv4 -> conv5 -> pool3

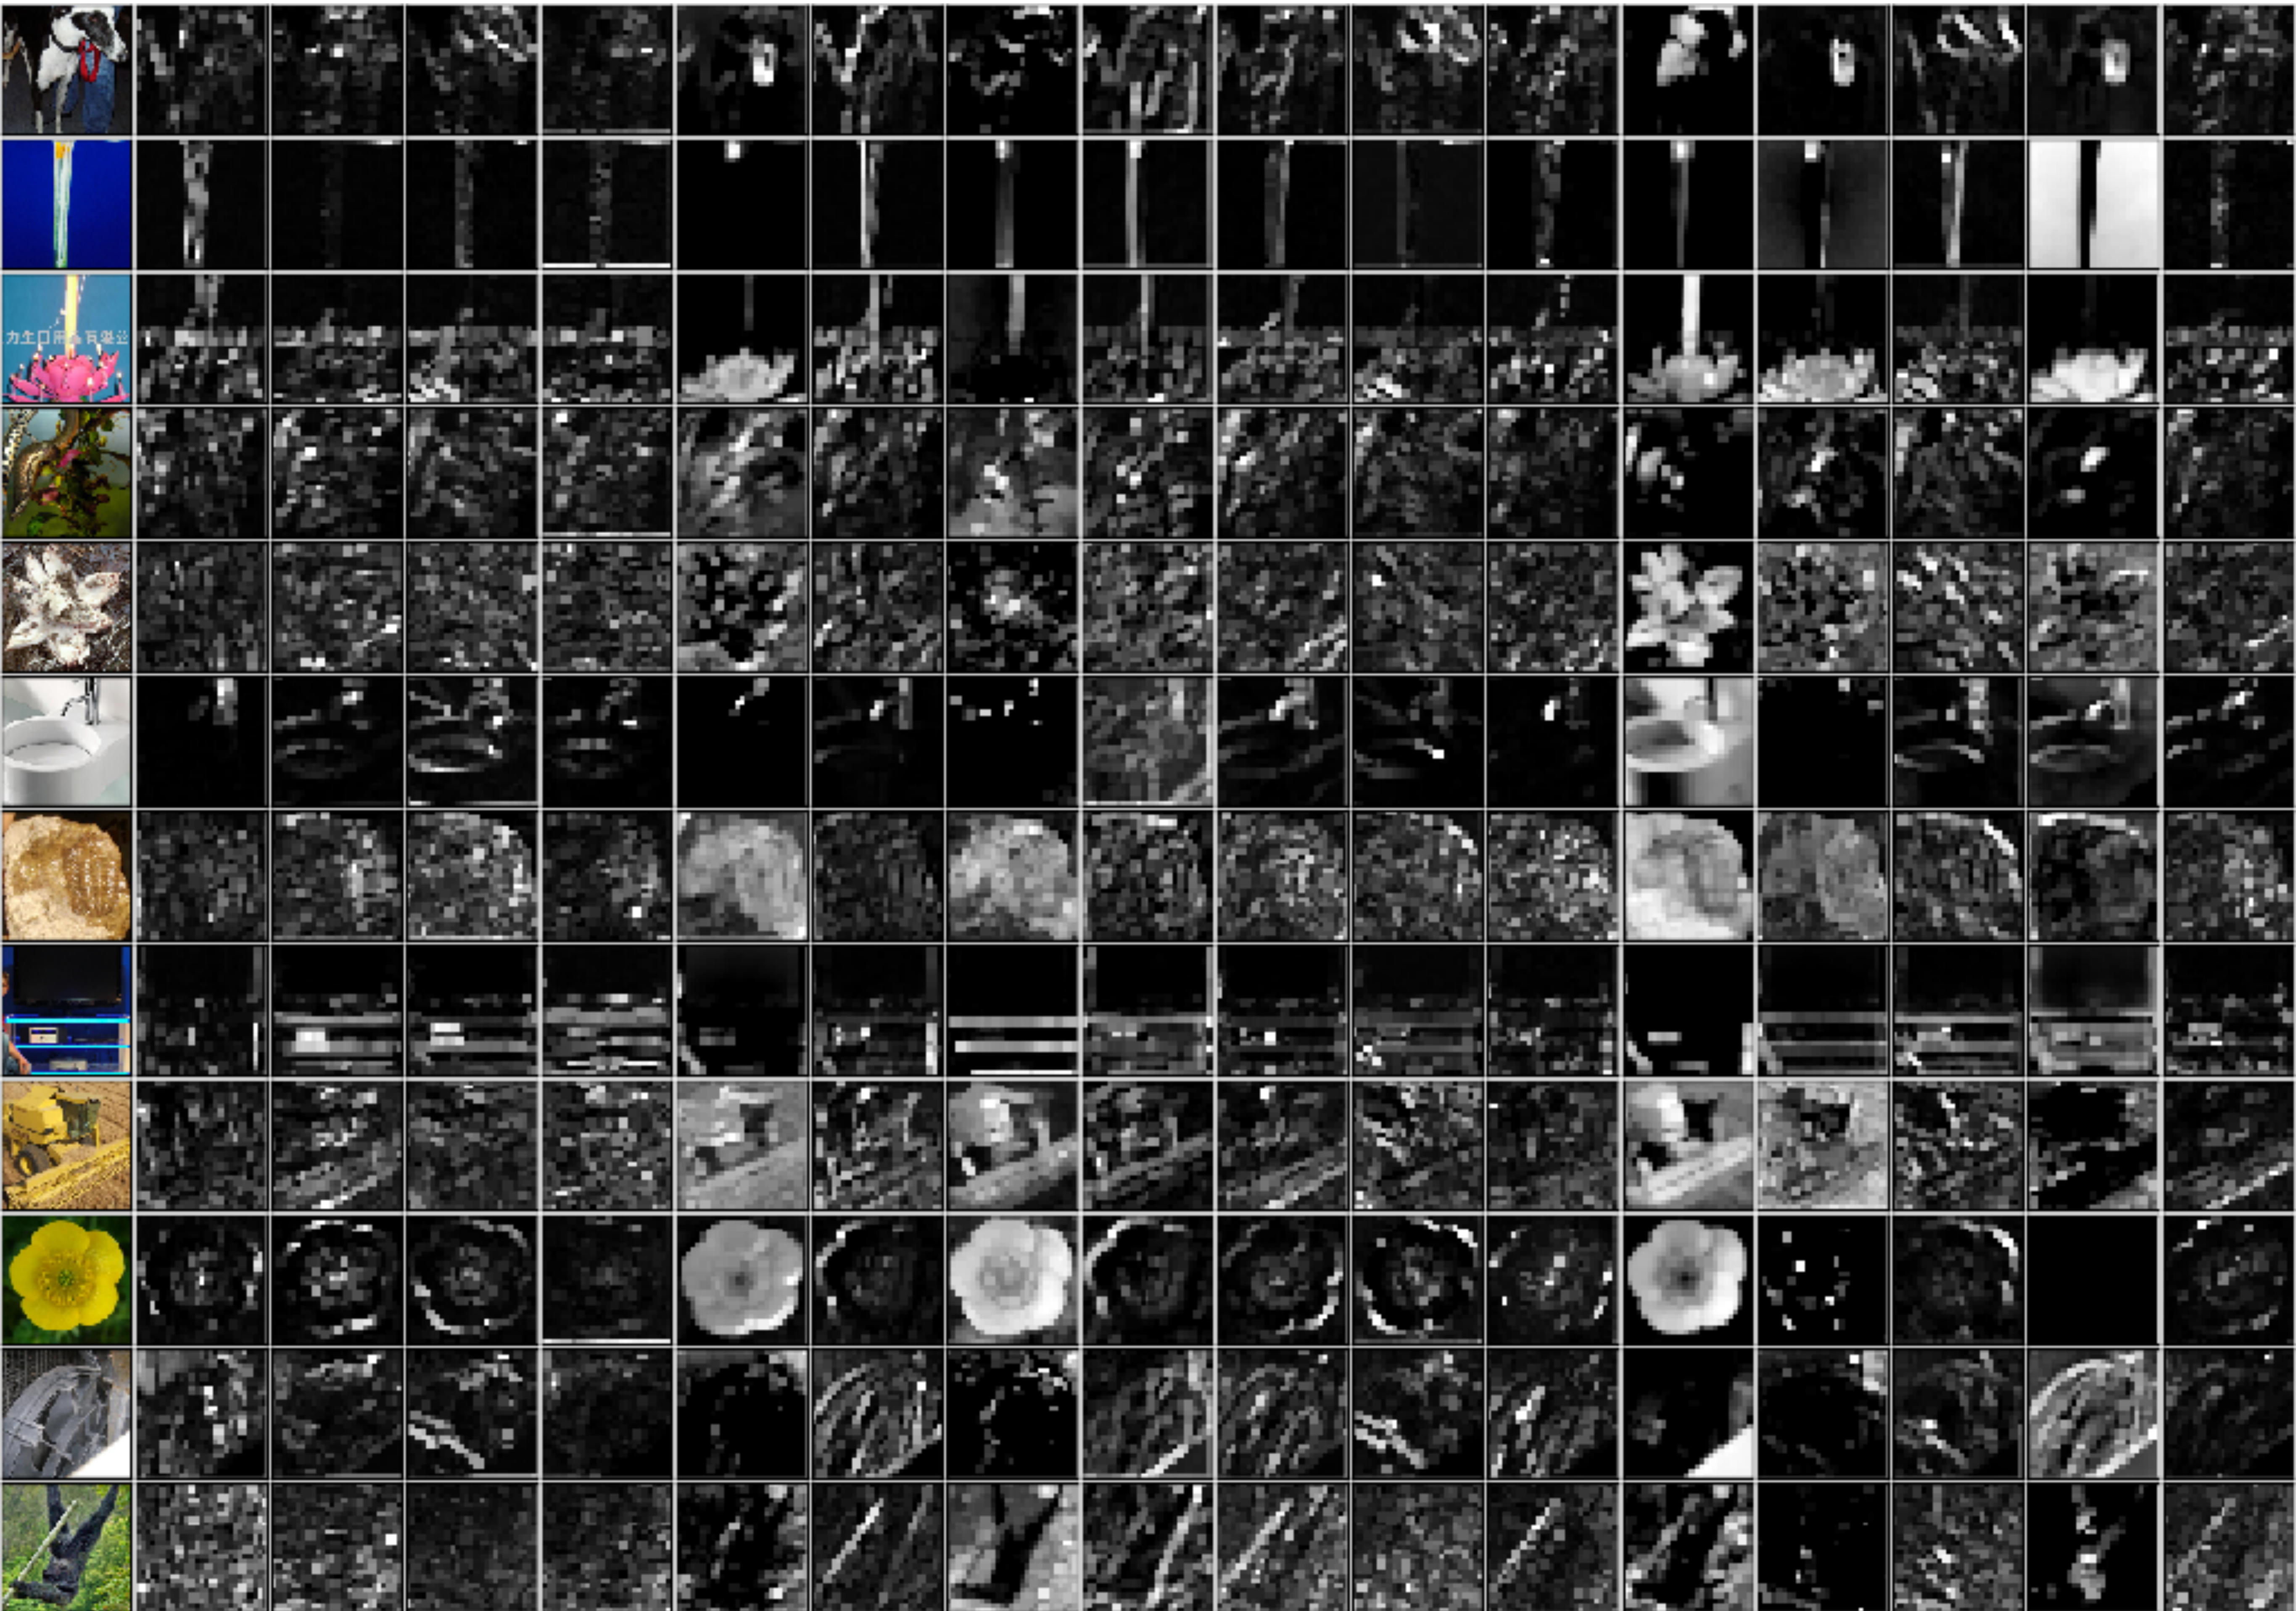

data -> conv1 -> pool1 -> **conv2** -> pool2 -> conv3 -> conv4 -> conv5 -> pool3

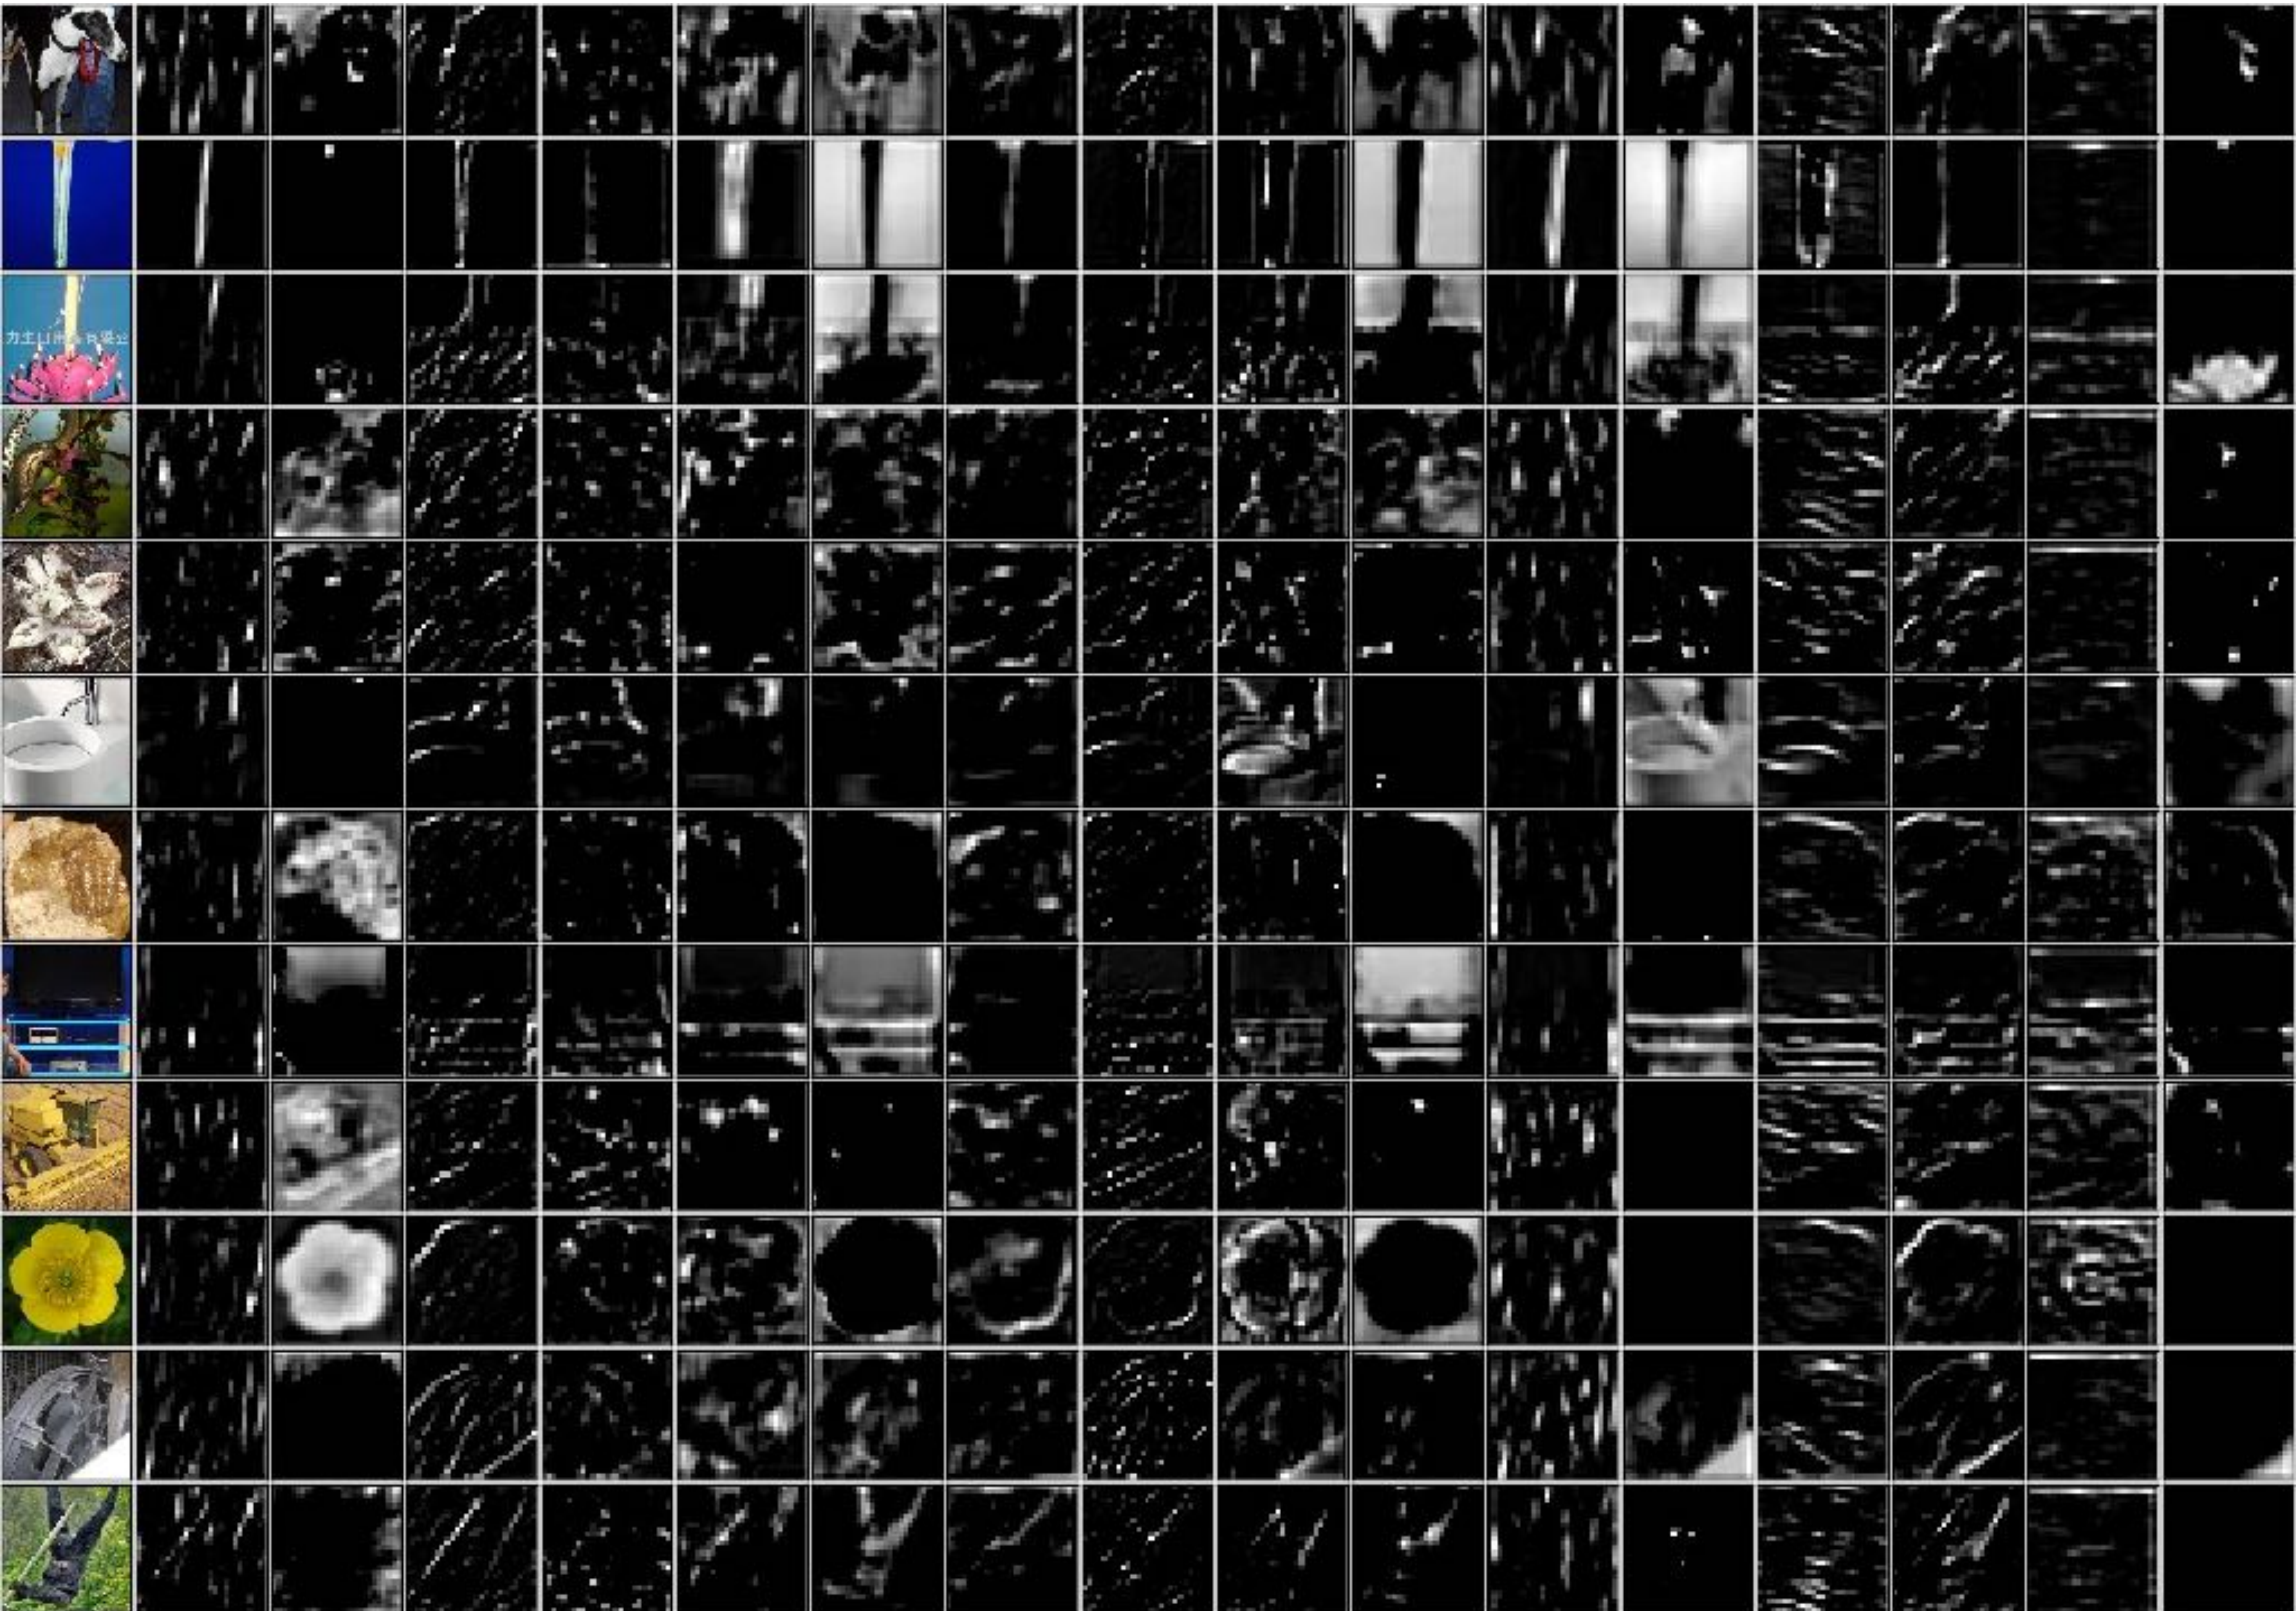

data -> conv1 -> pool1 -> conv2 -> **pool2** -> conv3 -> conv4 -> conv5 -> pool3

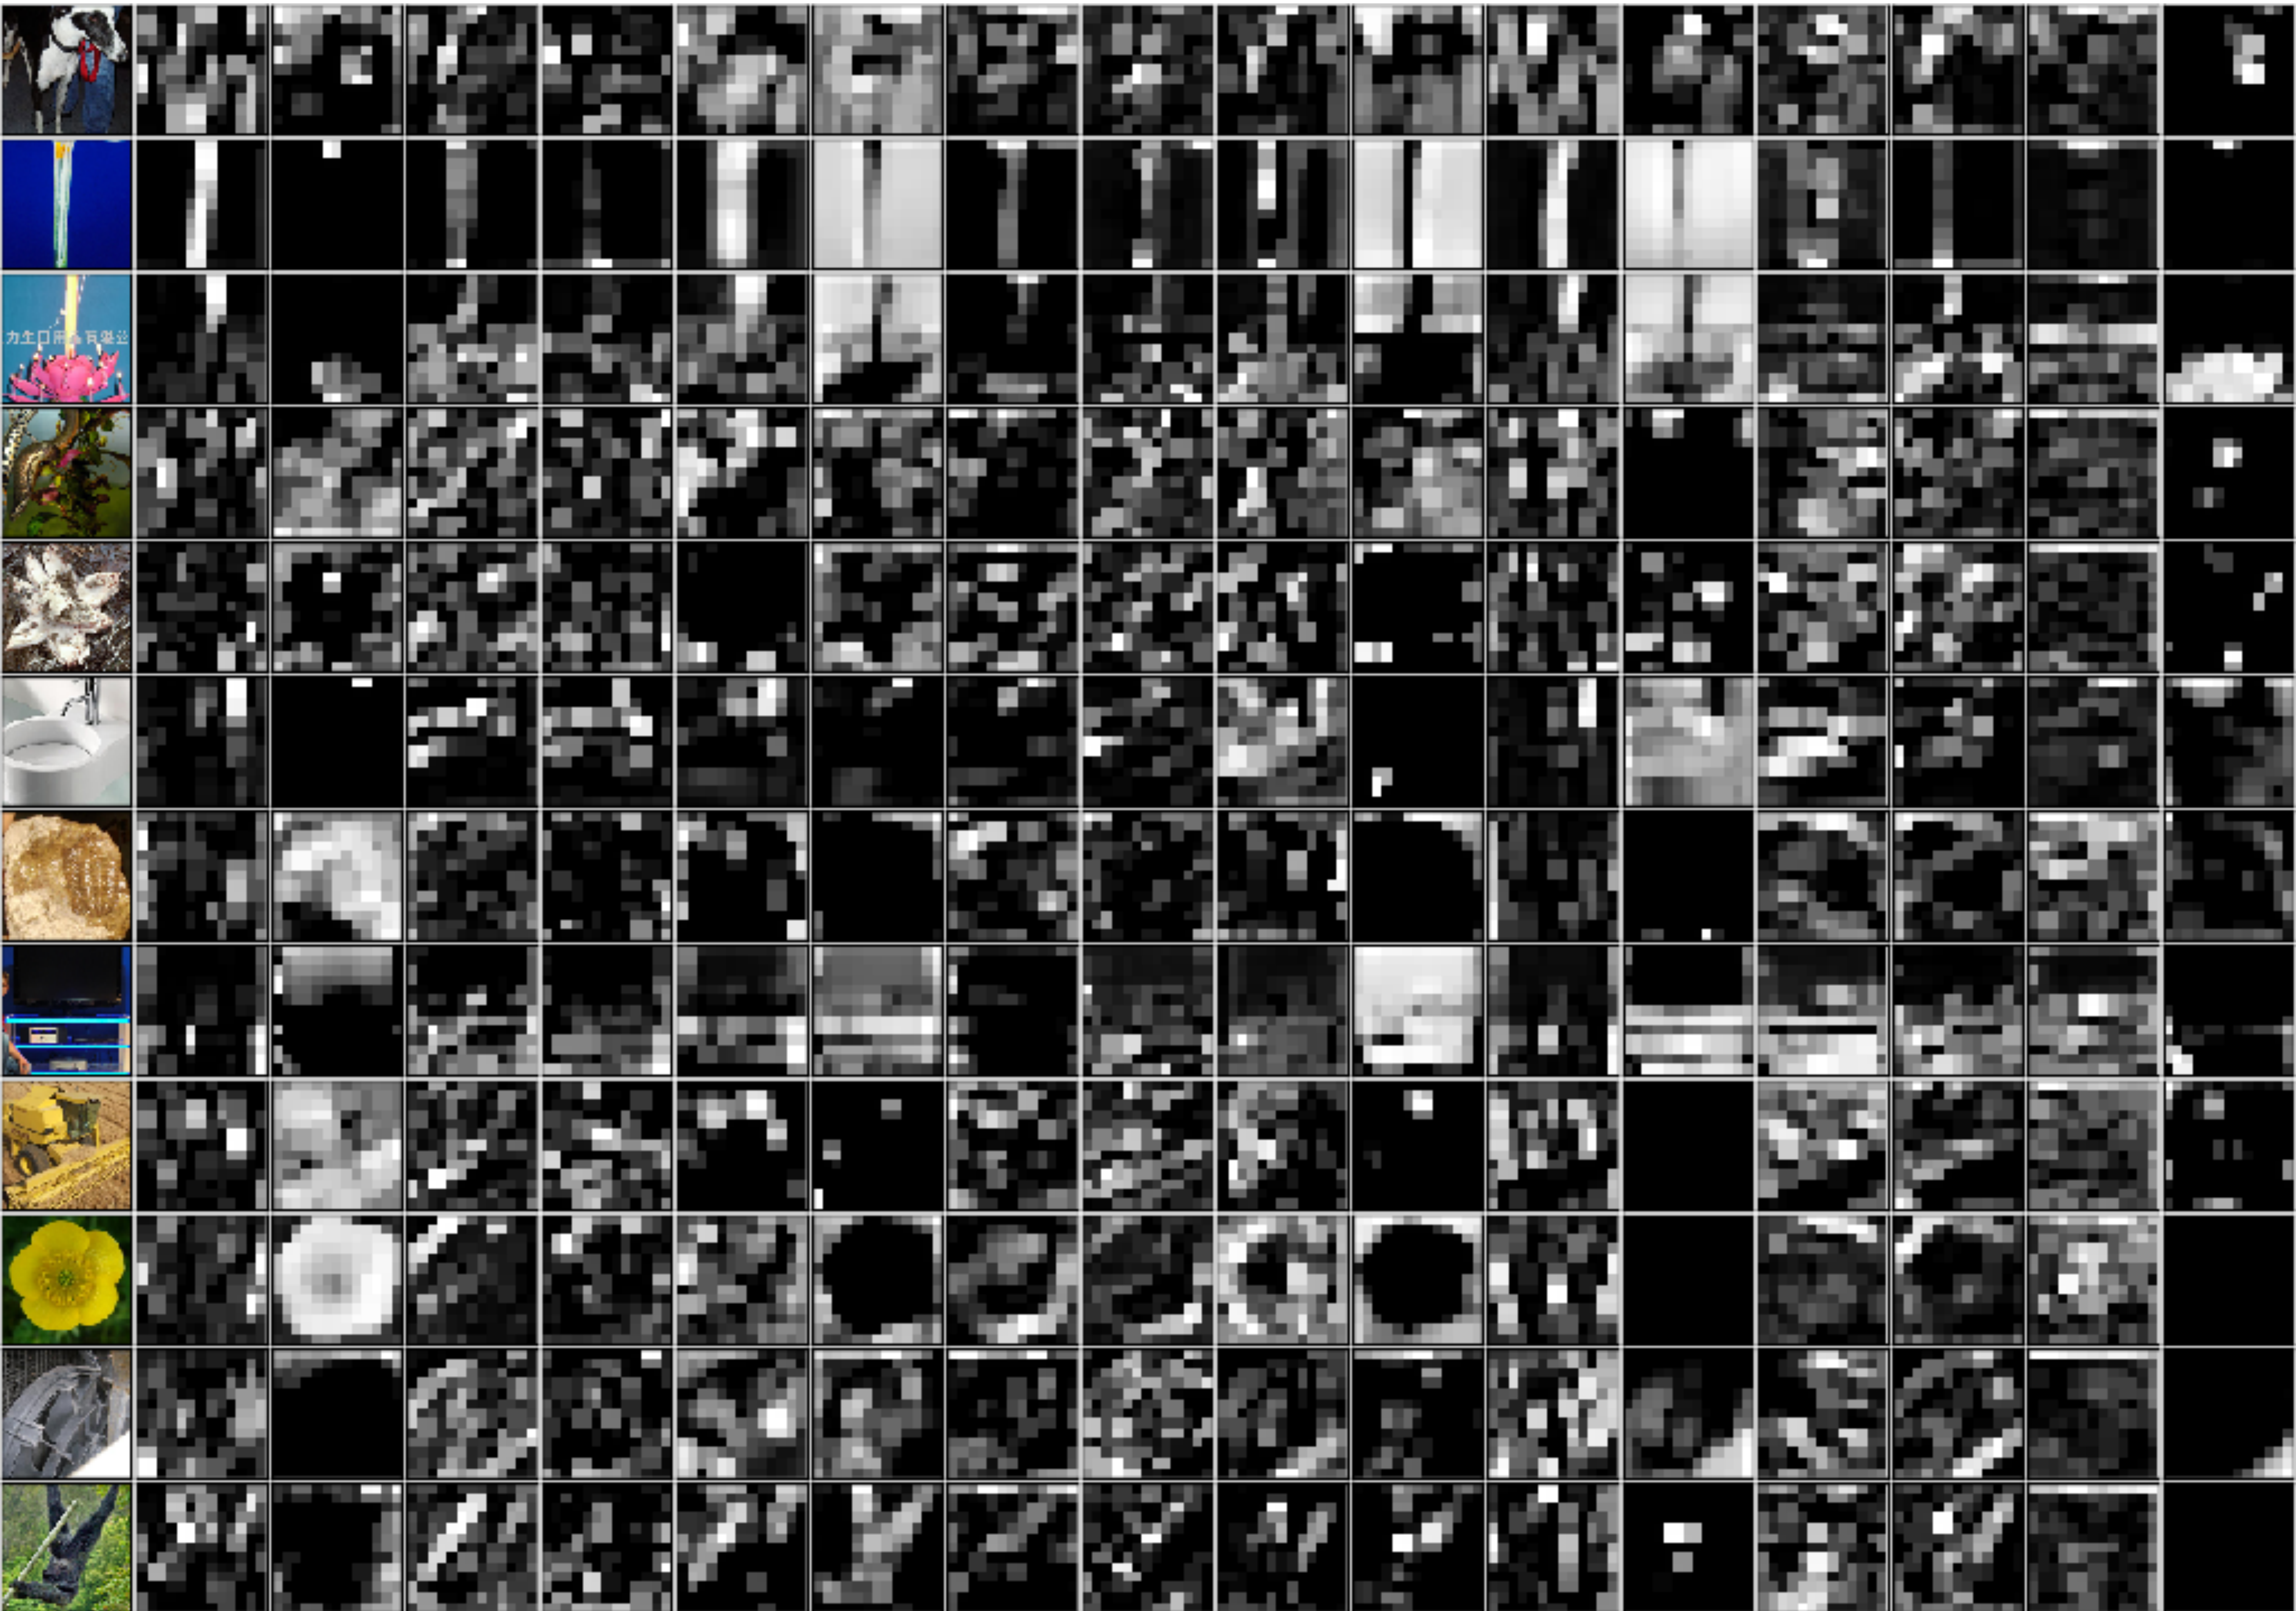

data -> conv1 -> pool1 -> conv2 -> pool2 -> **conv3** -> conv4 -> conv5 -> pool3

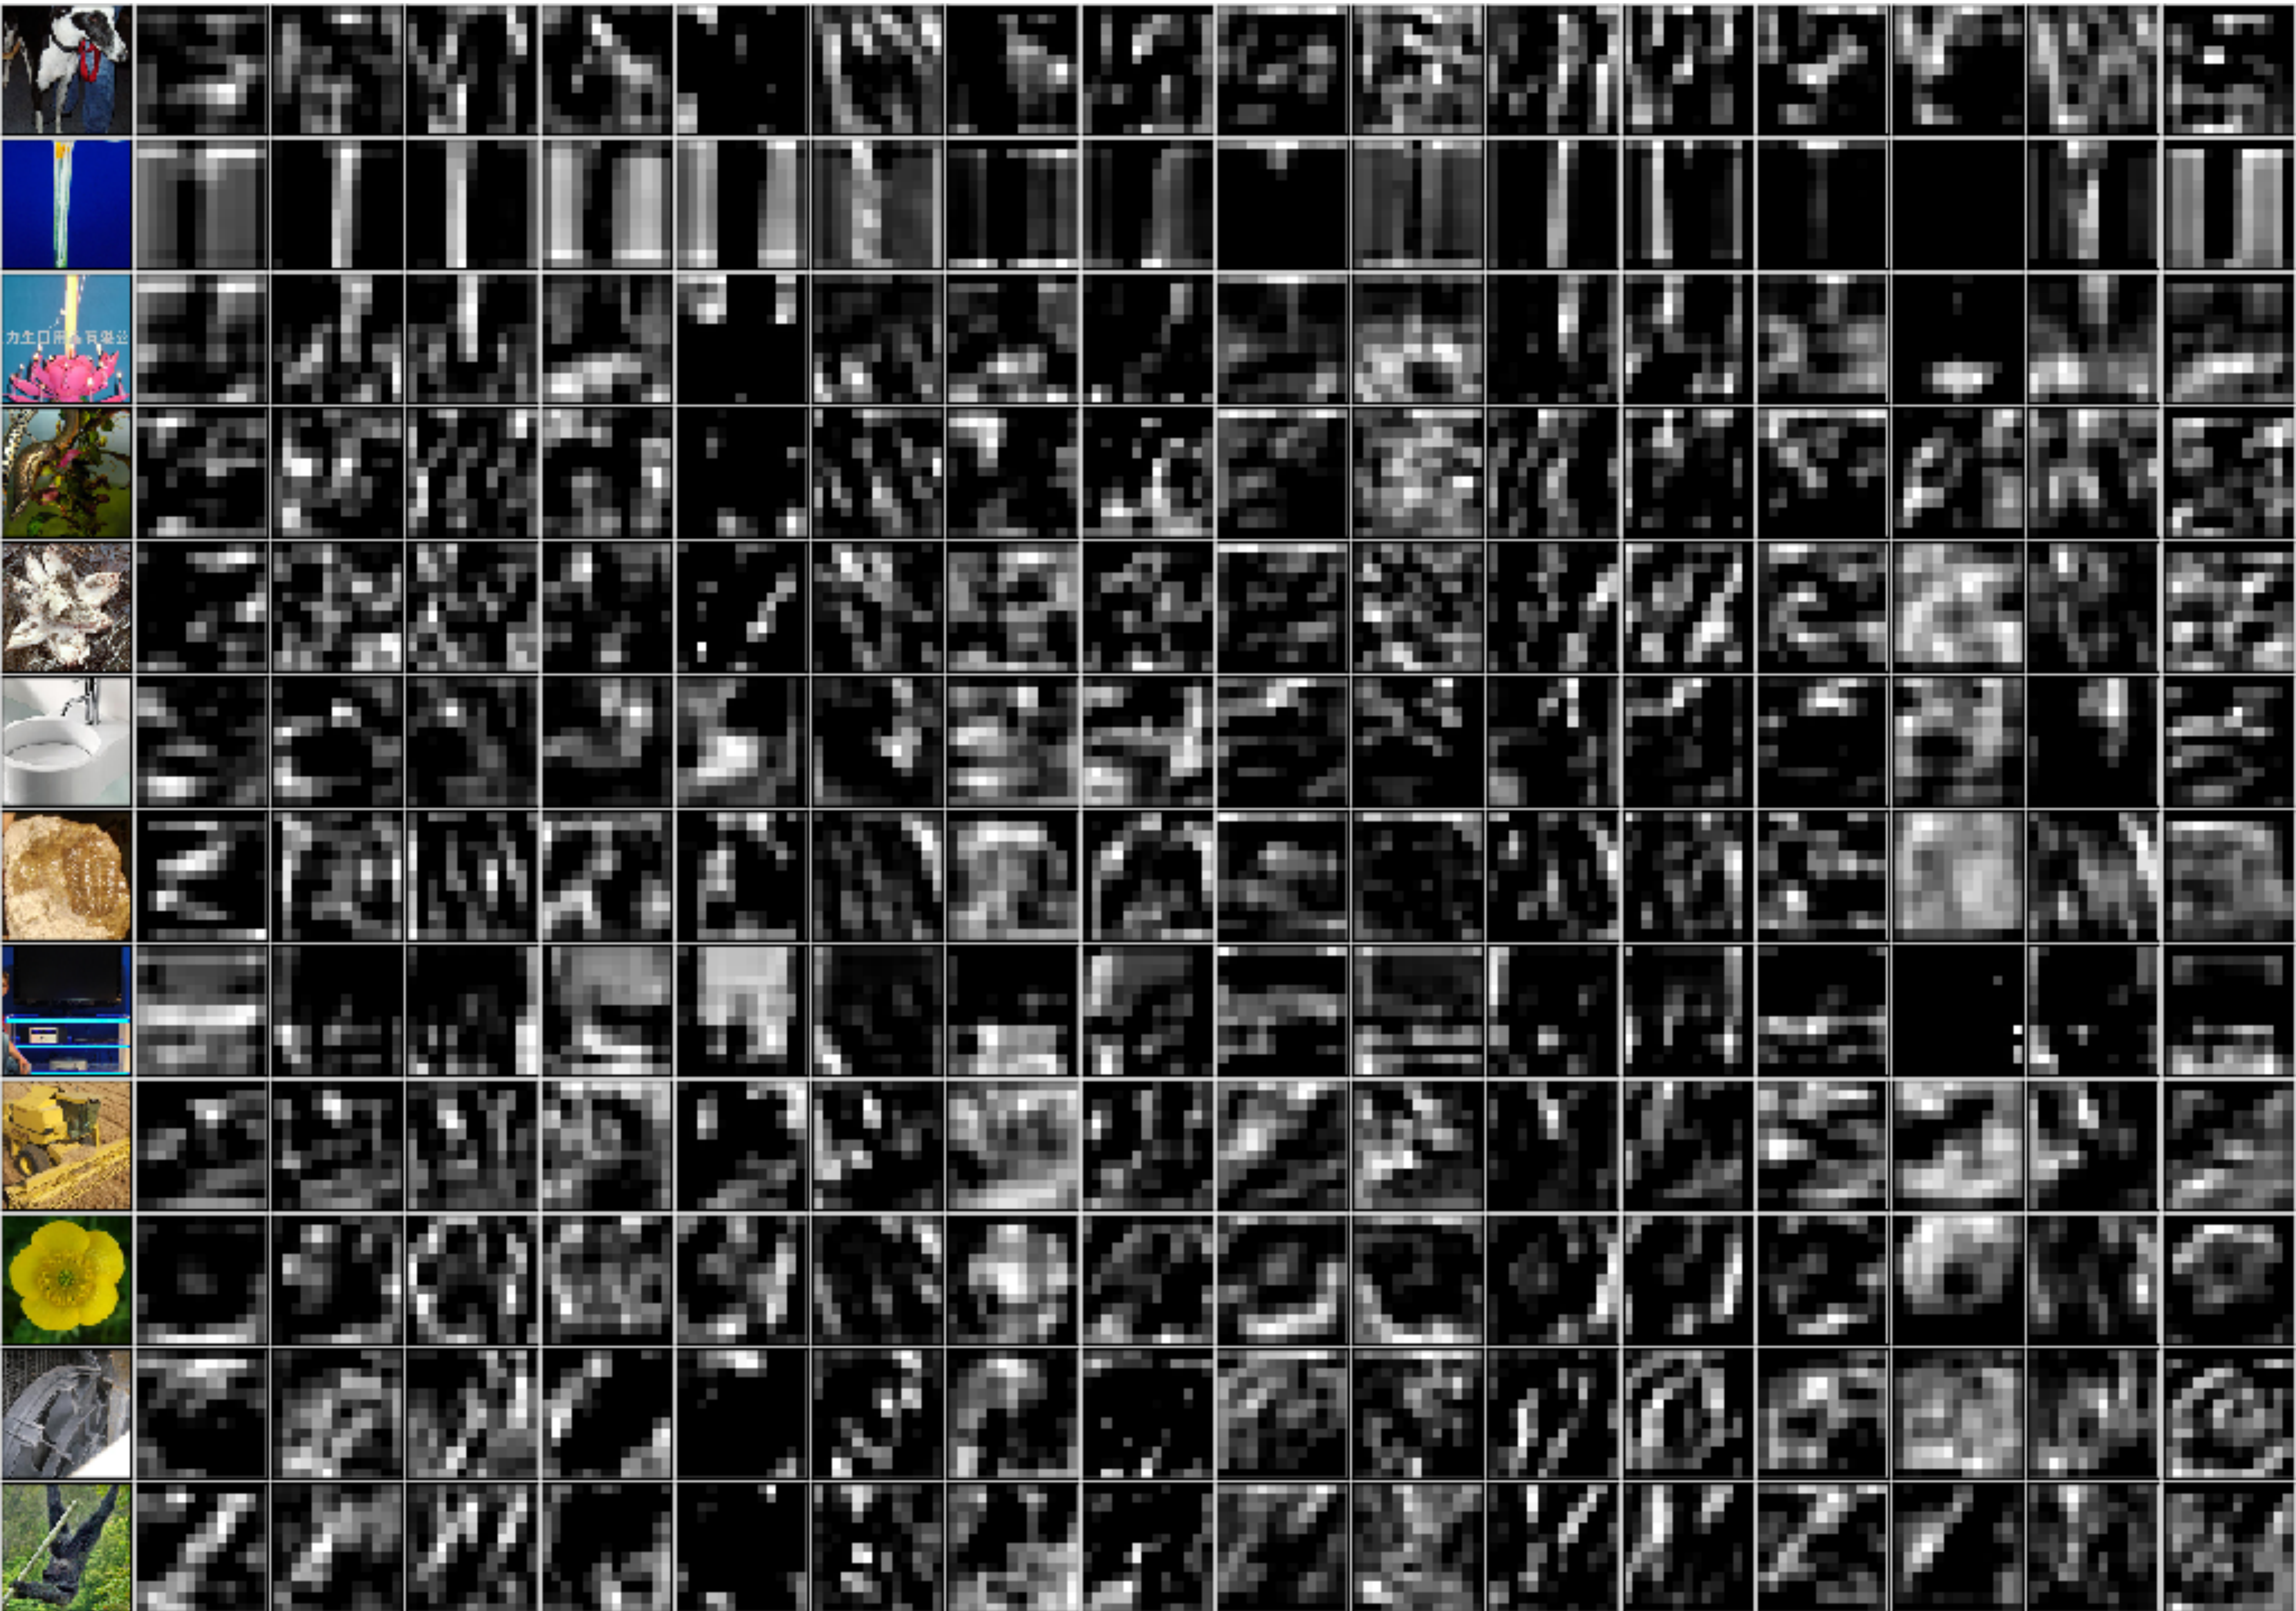

data -> conv1 -> pool1 -> conv2 -> pool2 -> conv3 -> **conv4** -> conv5 -> pool3

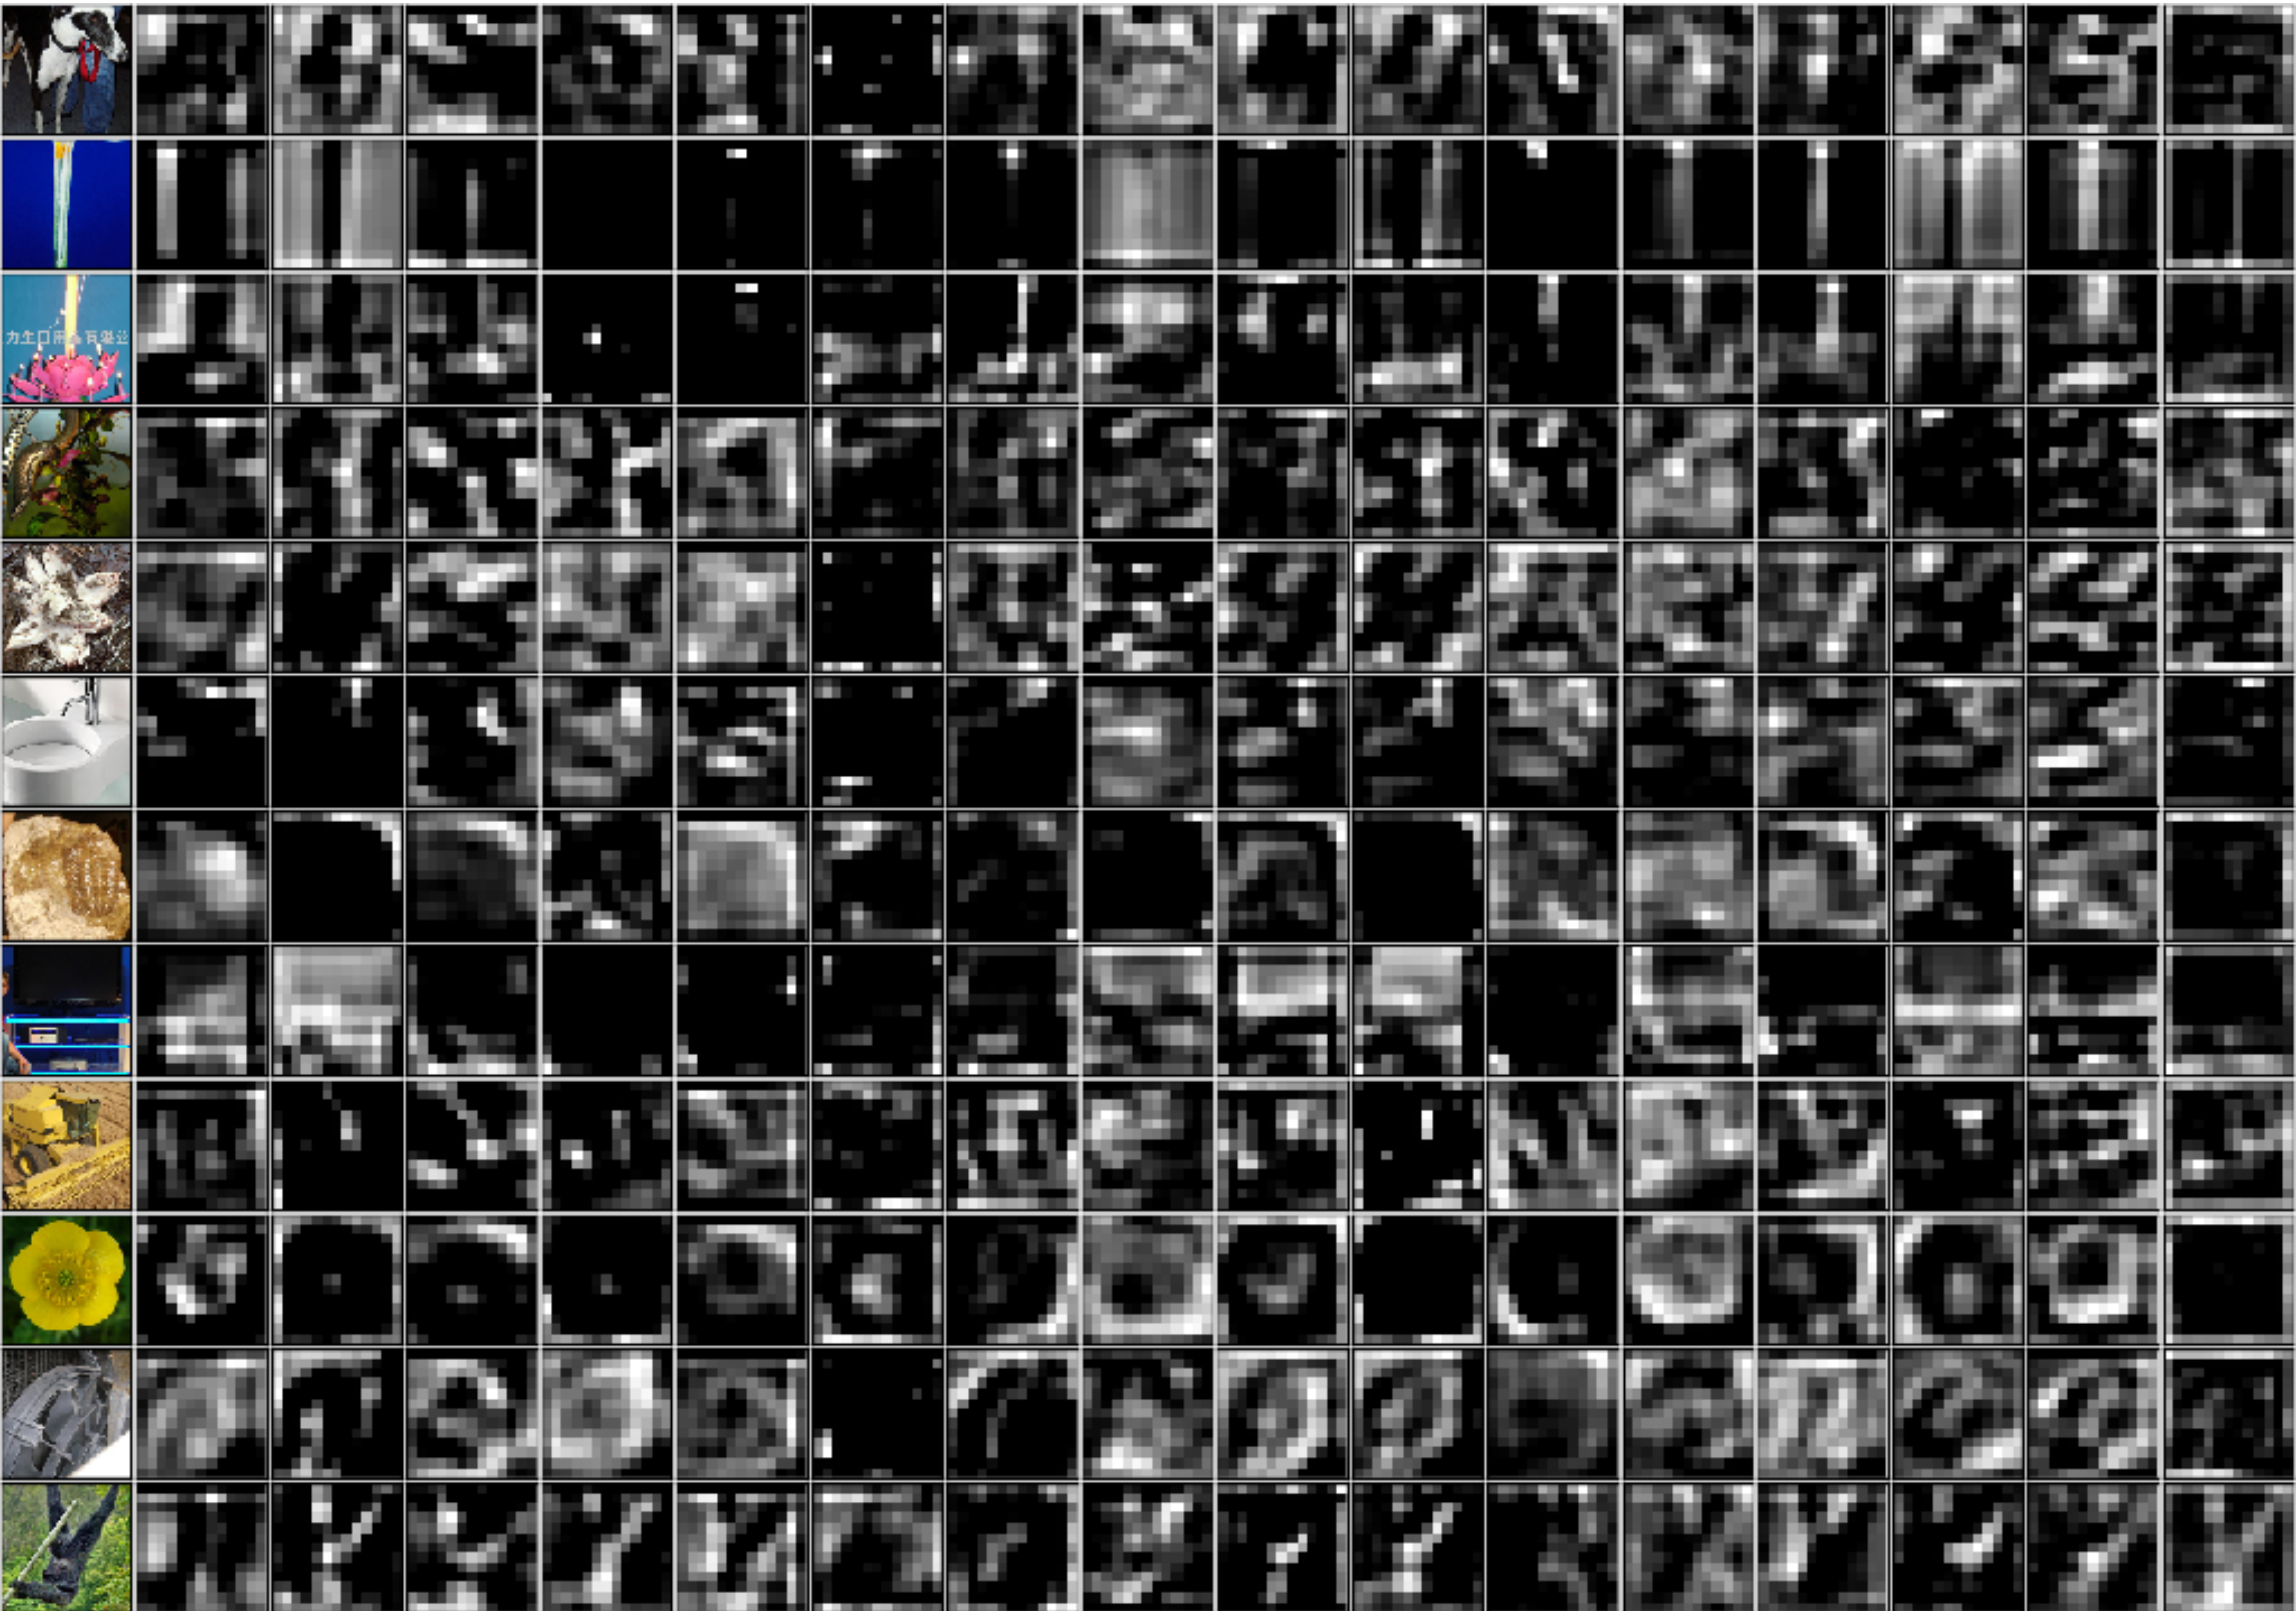

data -> conv1 -> pool1 -> conv2 -> pool2 -> conv3 -> conv4 -> **conv5** -> pool3

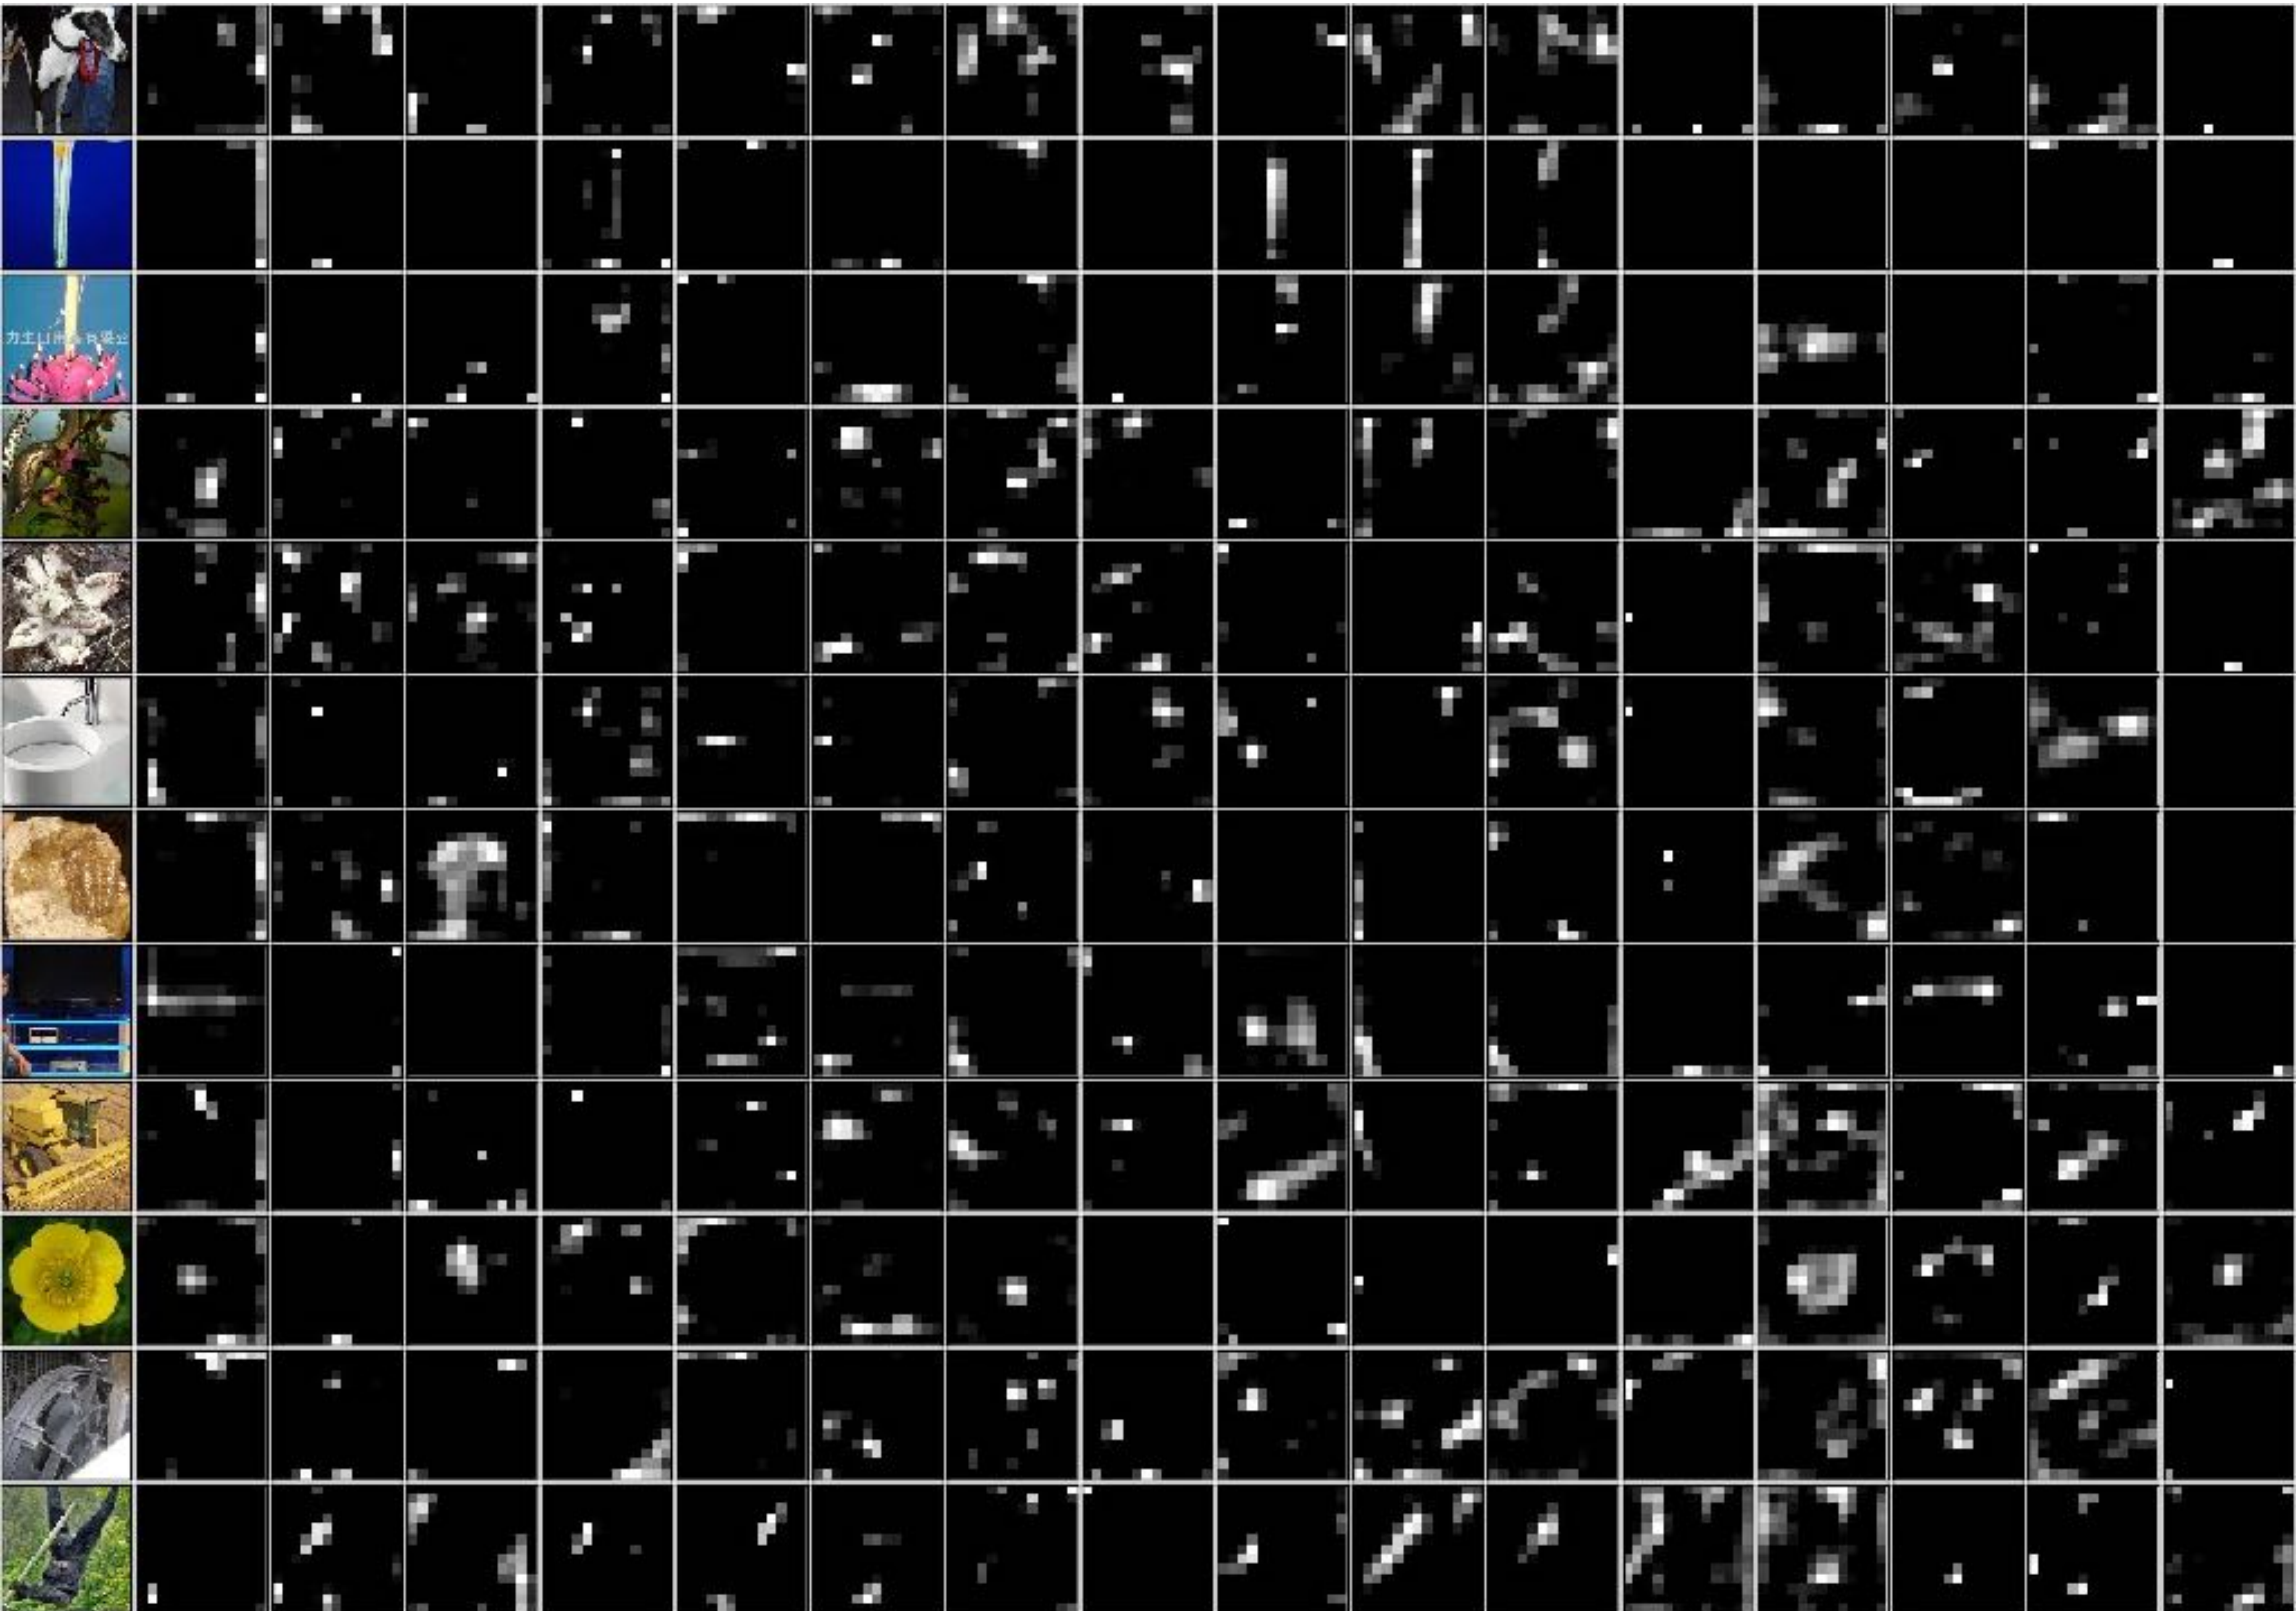

data -> conv1 -> pool1 -> conv2 -> pool2 -> conv3 -> conv4 -> conv5 -> **pool3**

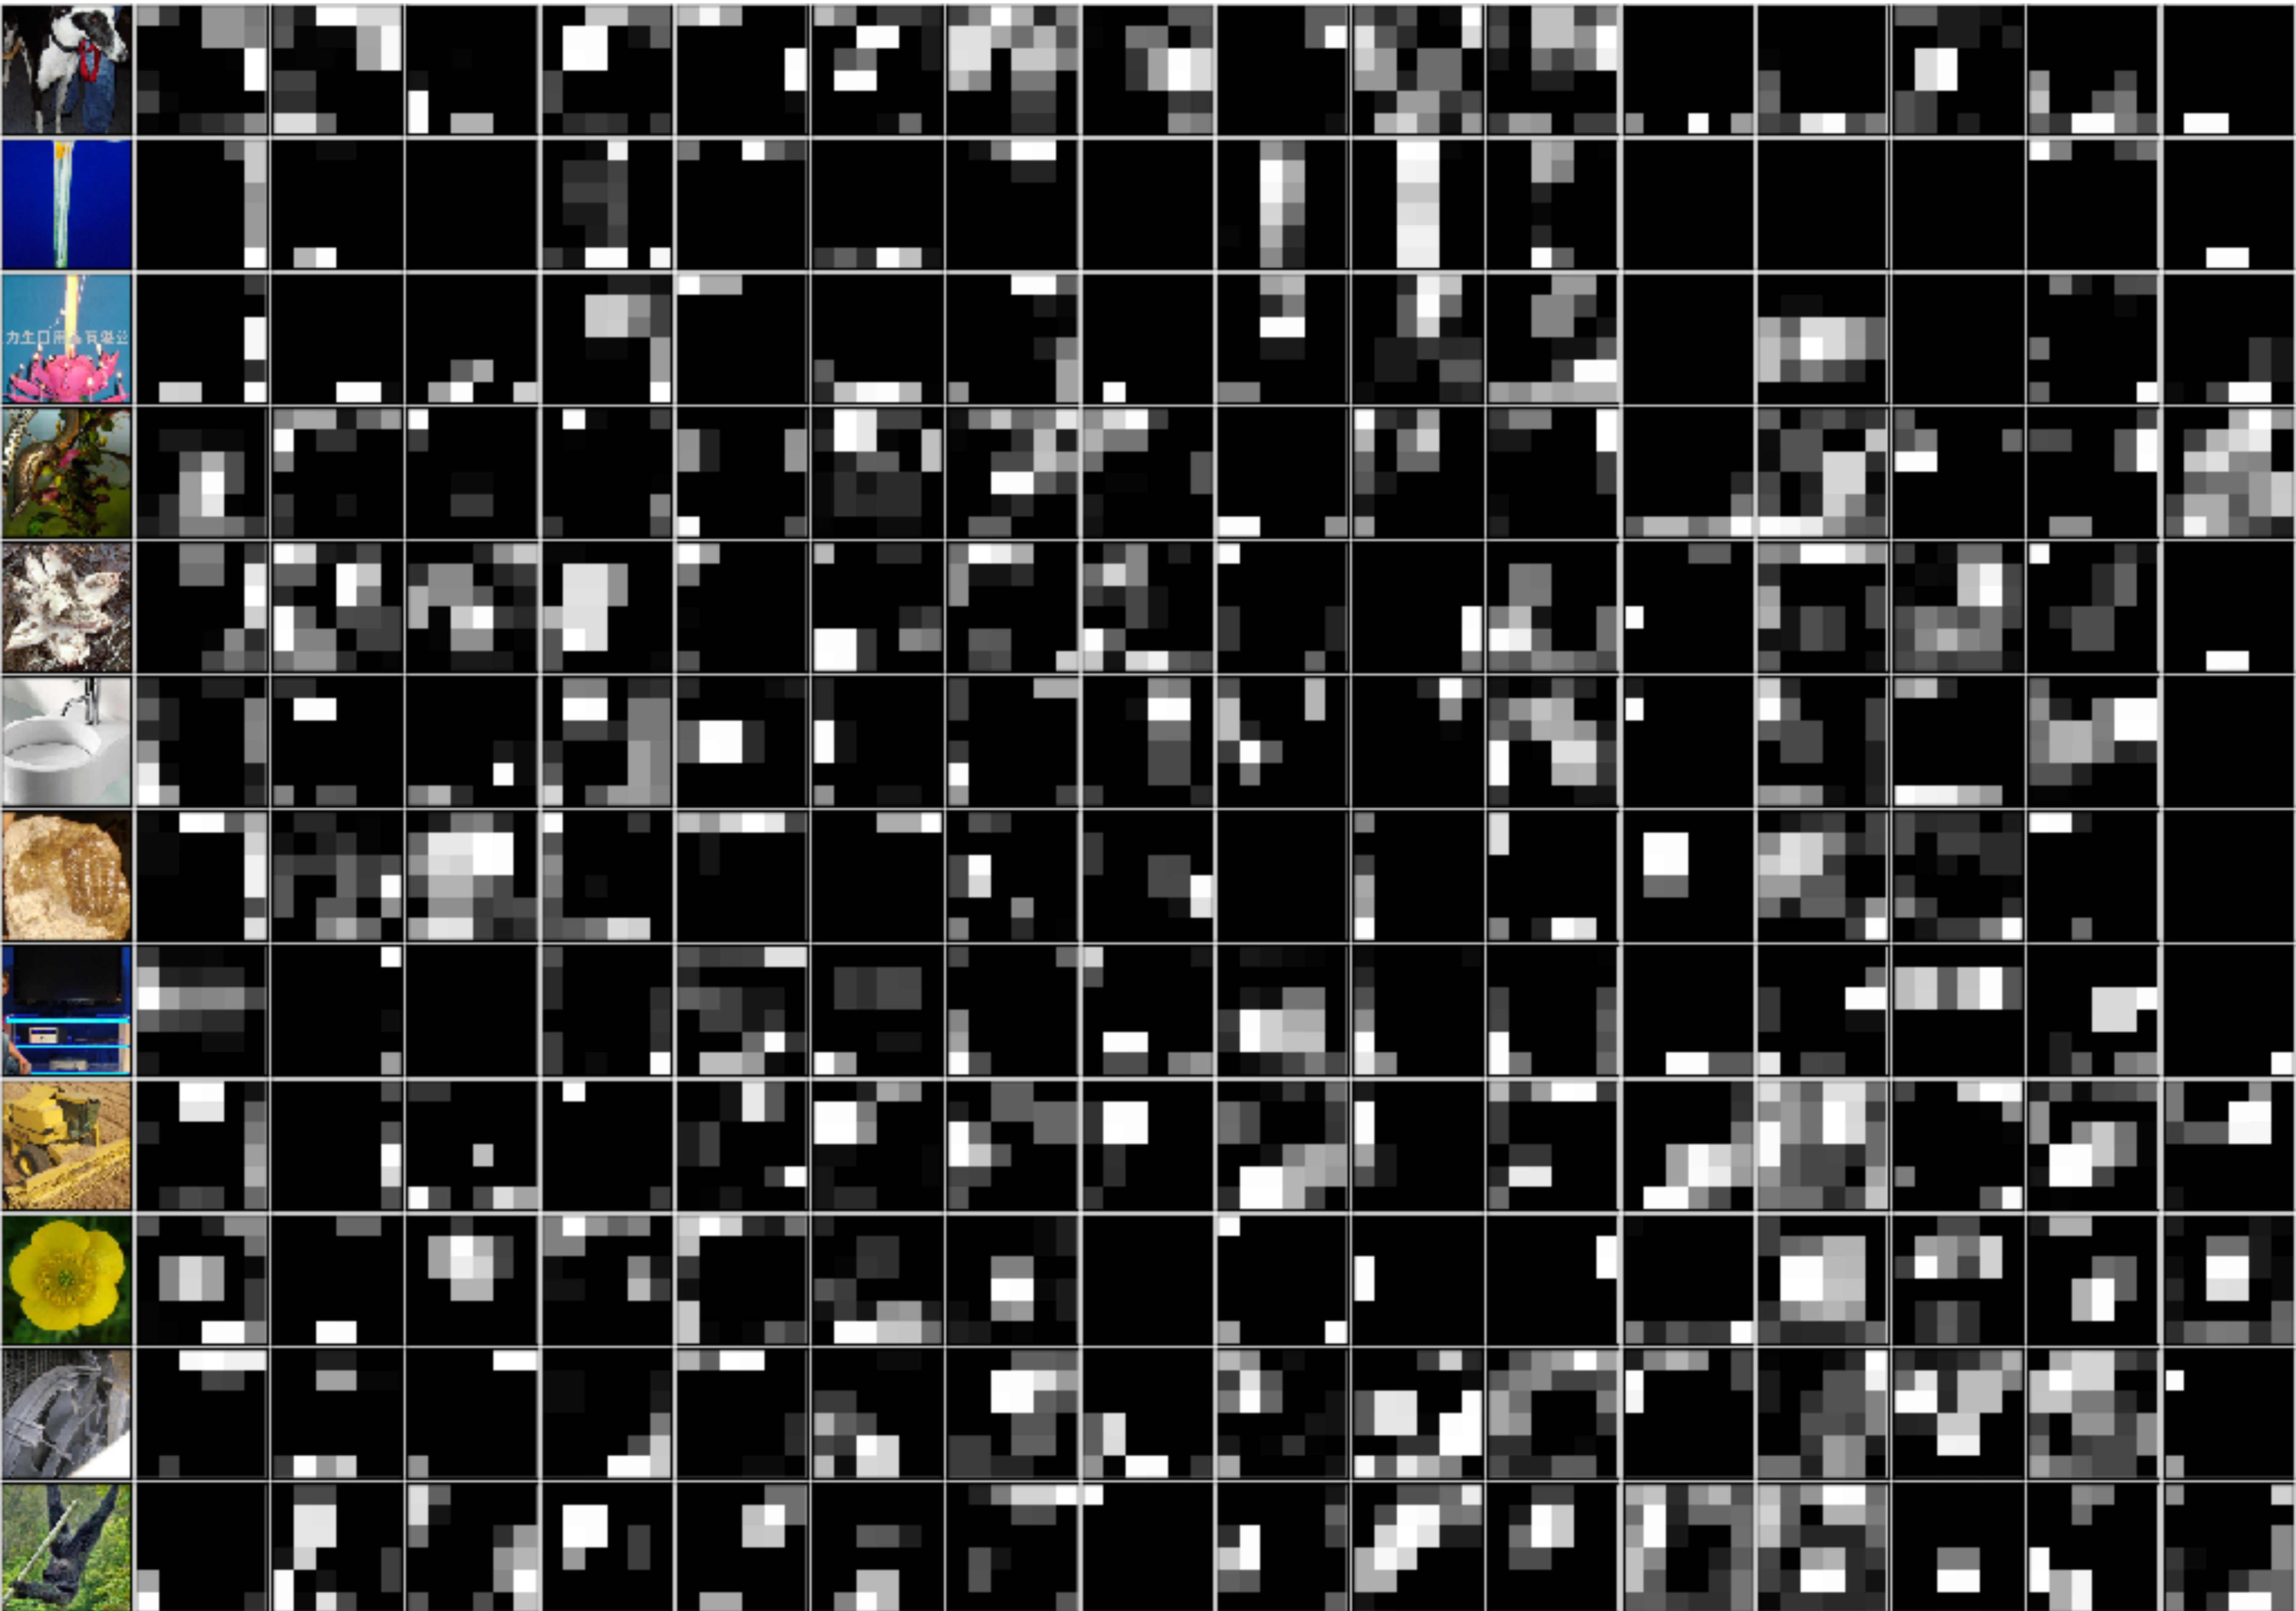

data -> conv1 -> pool1 -> conv2 -> pool2 -> conv3 -> conv4 -> conv5 -> **pool3**

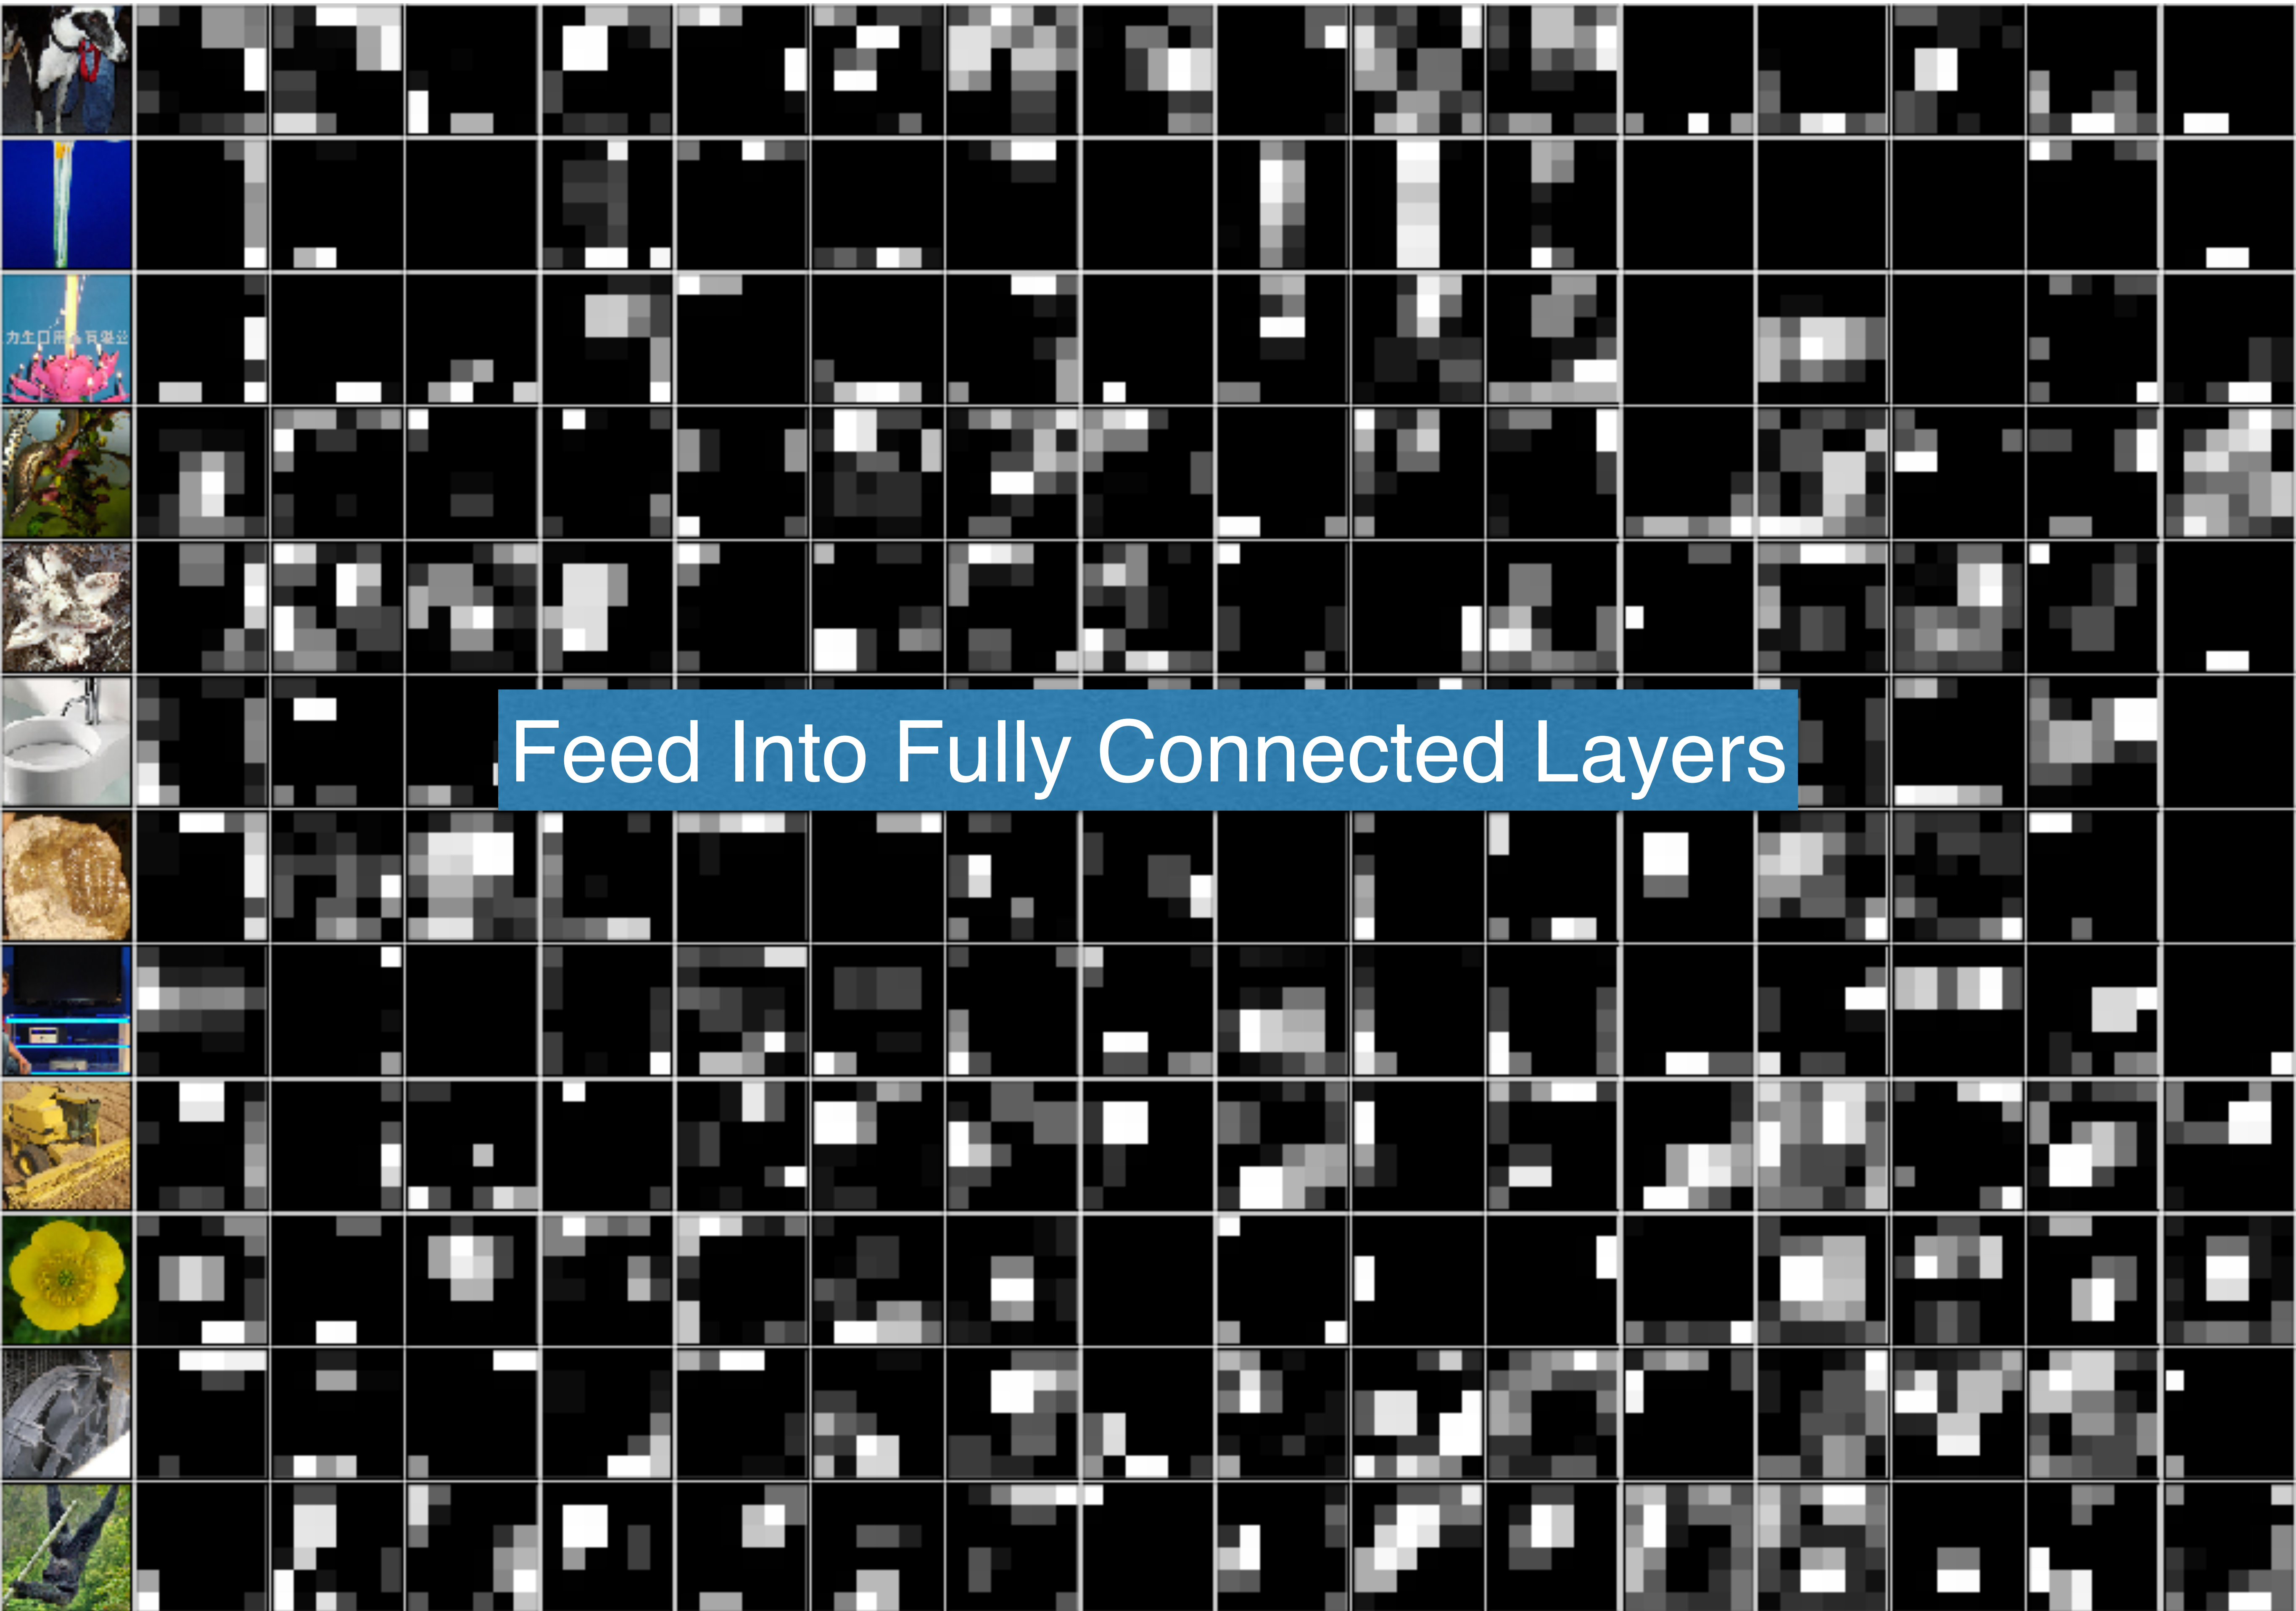

ImageNet Large Scale Visual Recognition Challenge (ILSVRC)

IMAGENET

1000 categories
1.2M train images,
150,000 test images

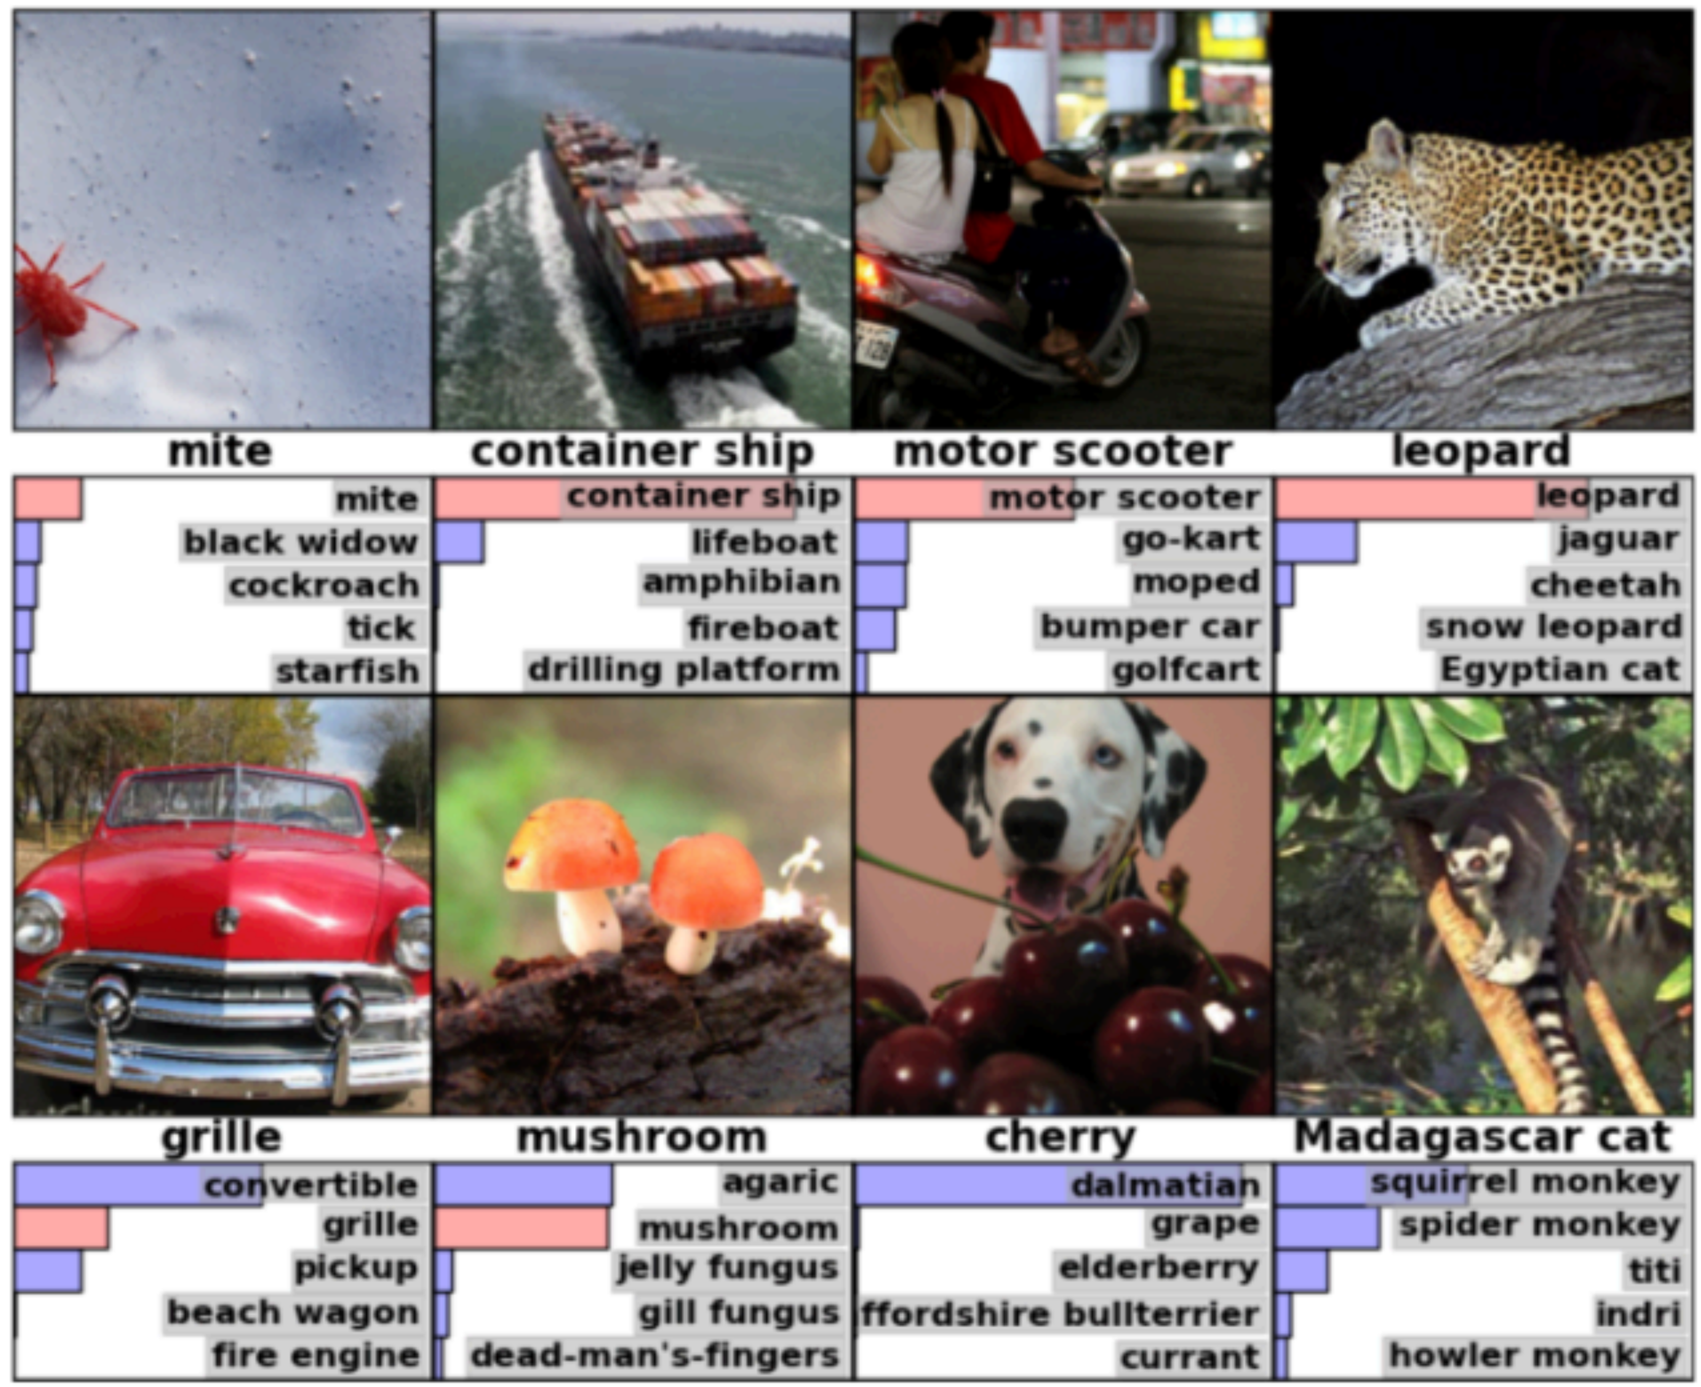

http://image-net.org/challenges/LSVRC/

CNN's Then and Now

1995:

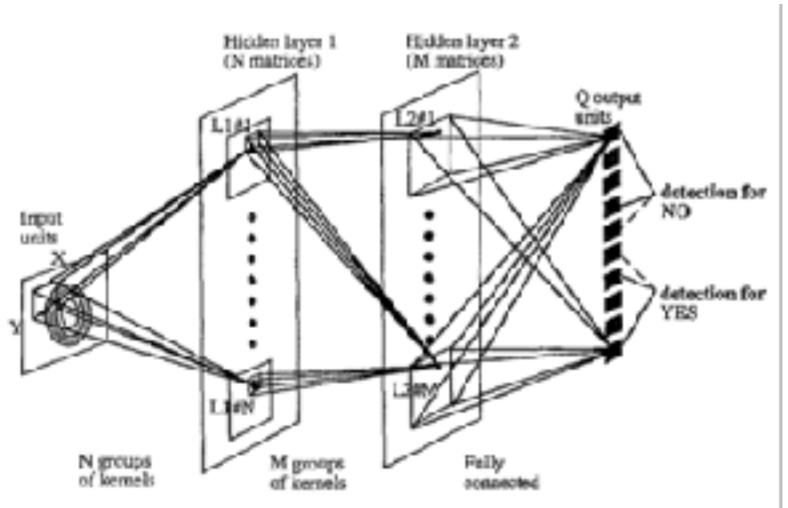

Fig. 3. Artificial convolution neural network for detection of lung nodule.

~2 layers

2019:

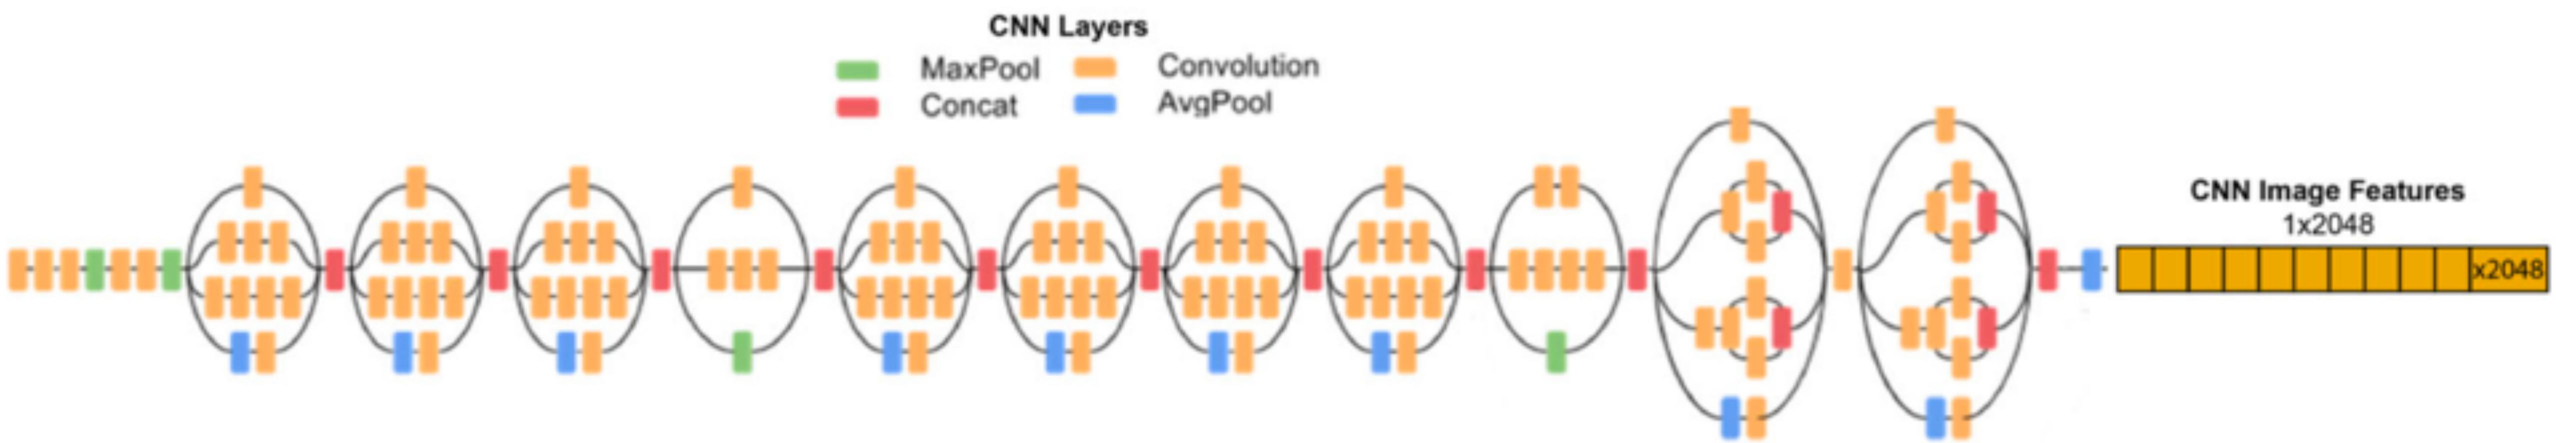

~100 layers

ImageNet Challenge Winners

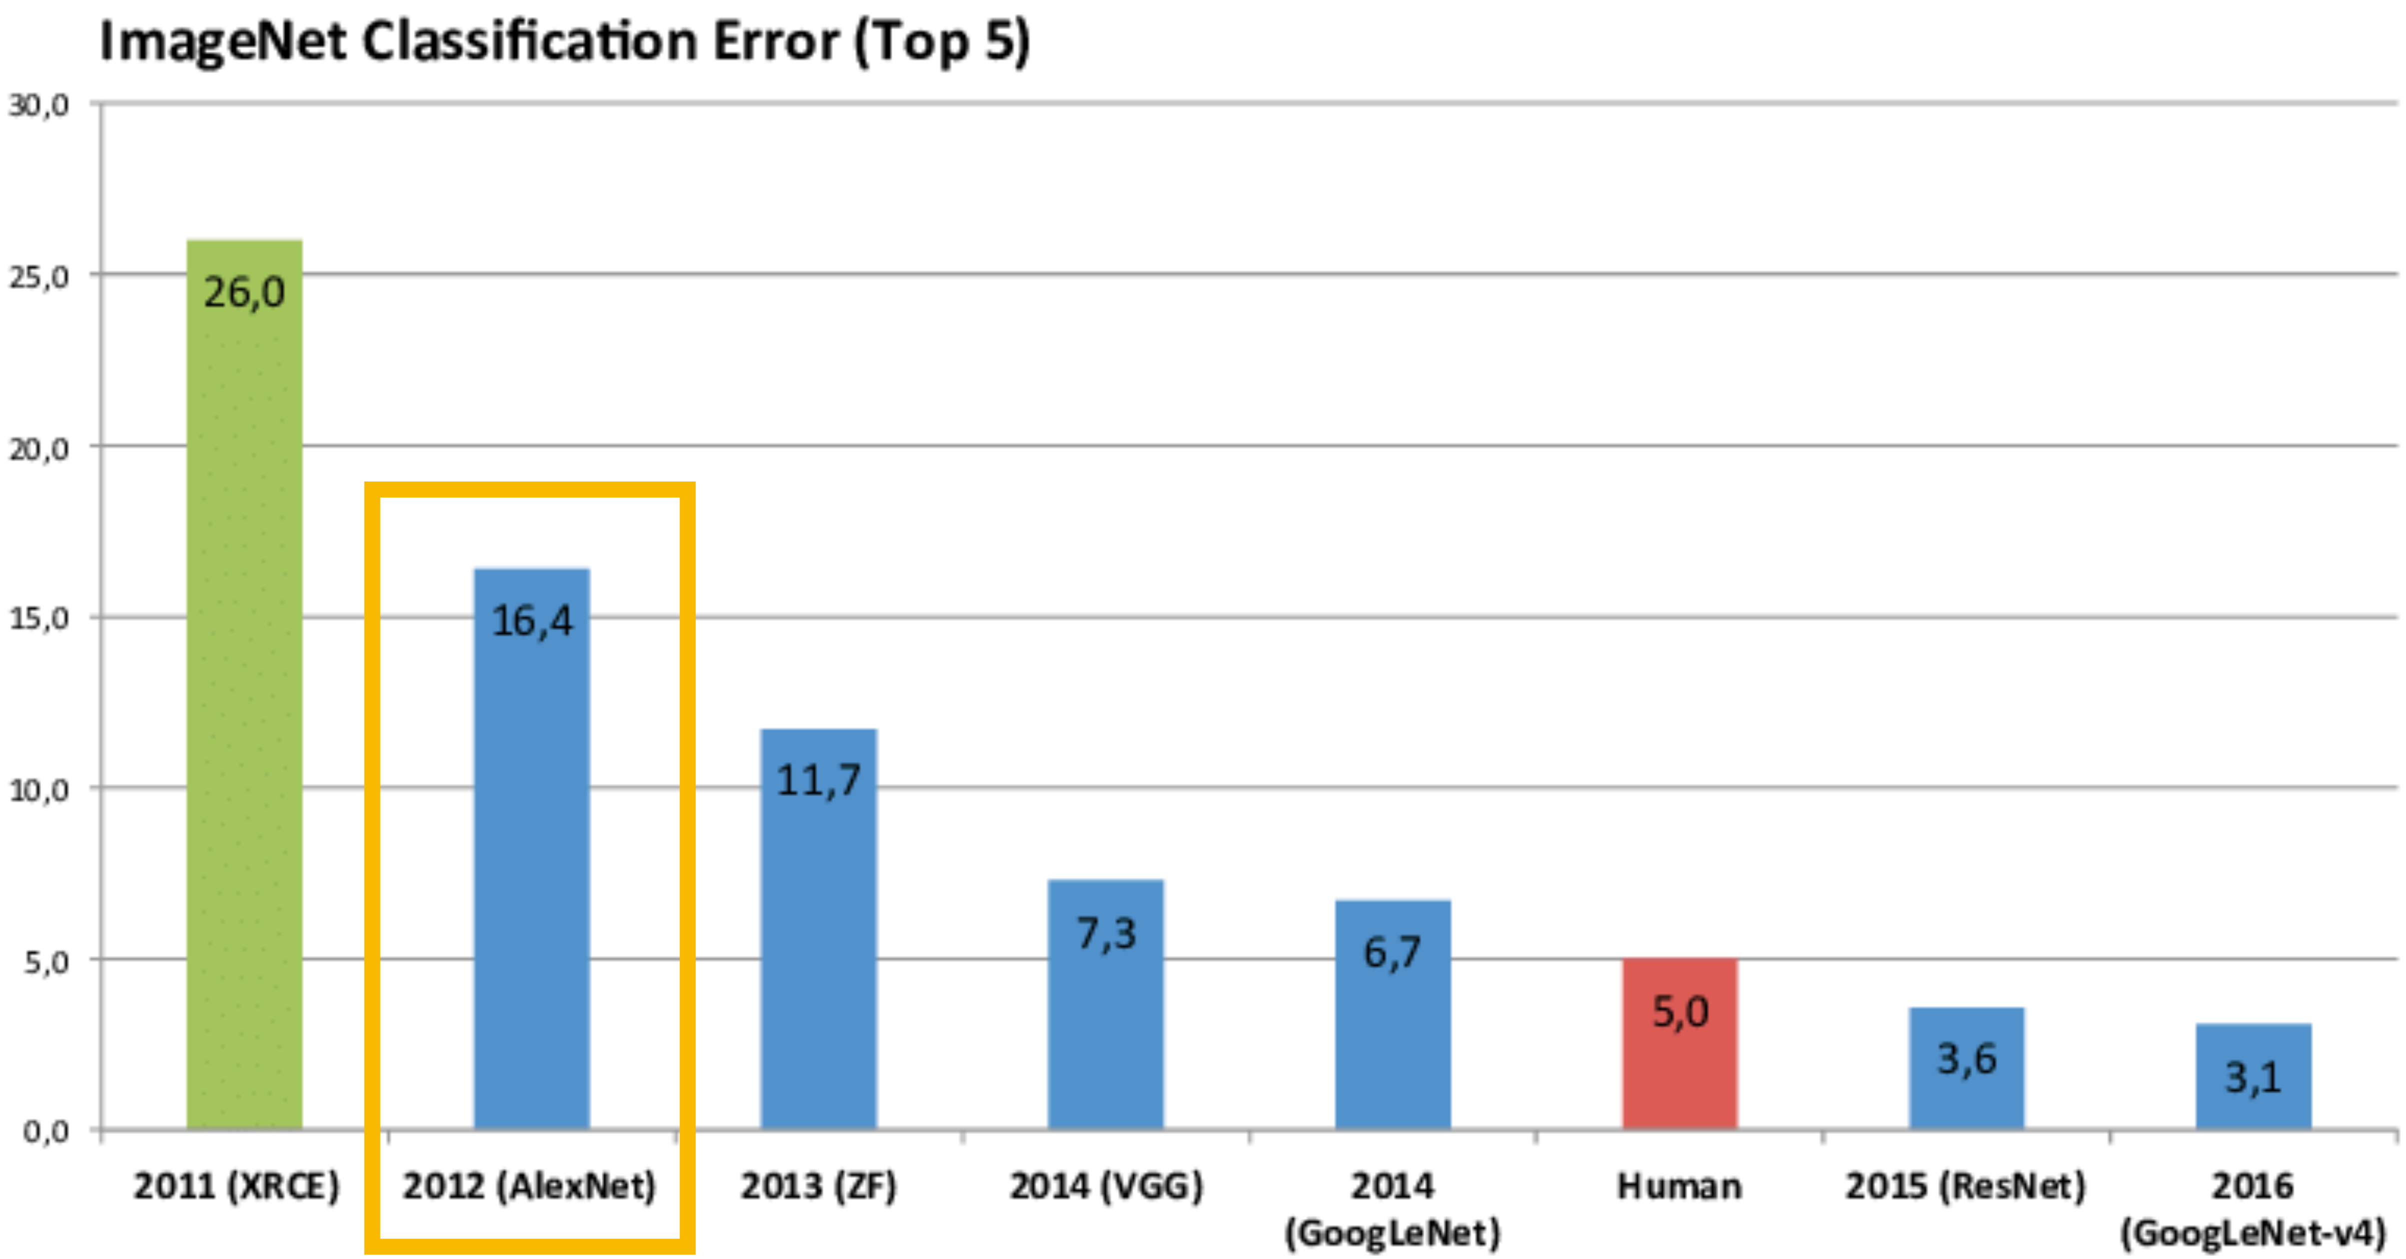

Year (winner)

Figure: Gustav von Zitzewitz

AlexNet [Krizhevsky et al., 2012]

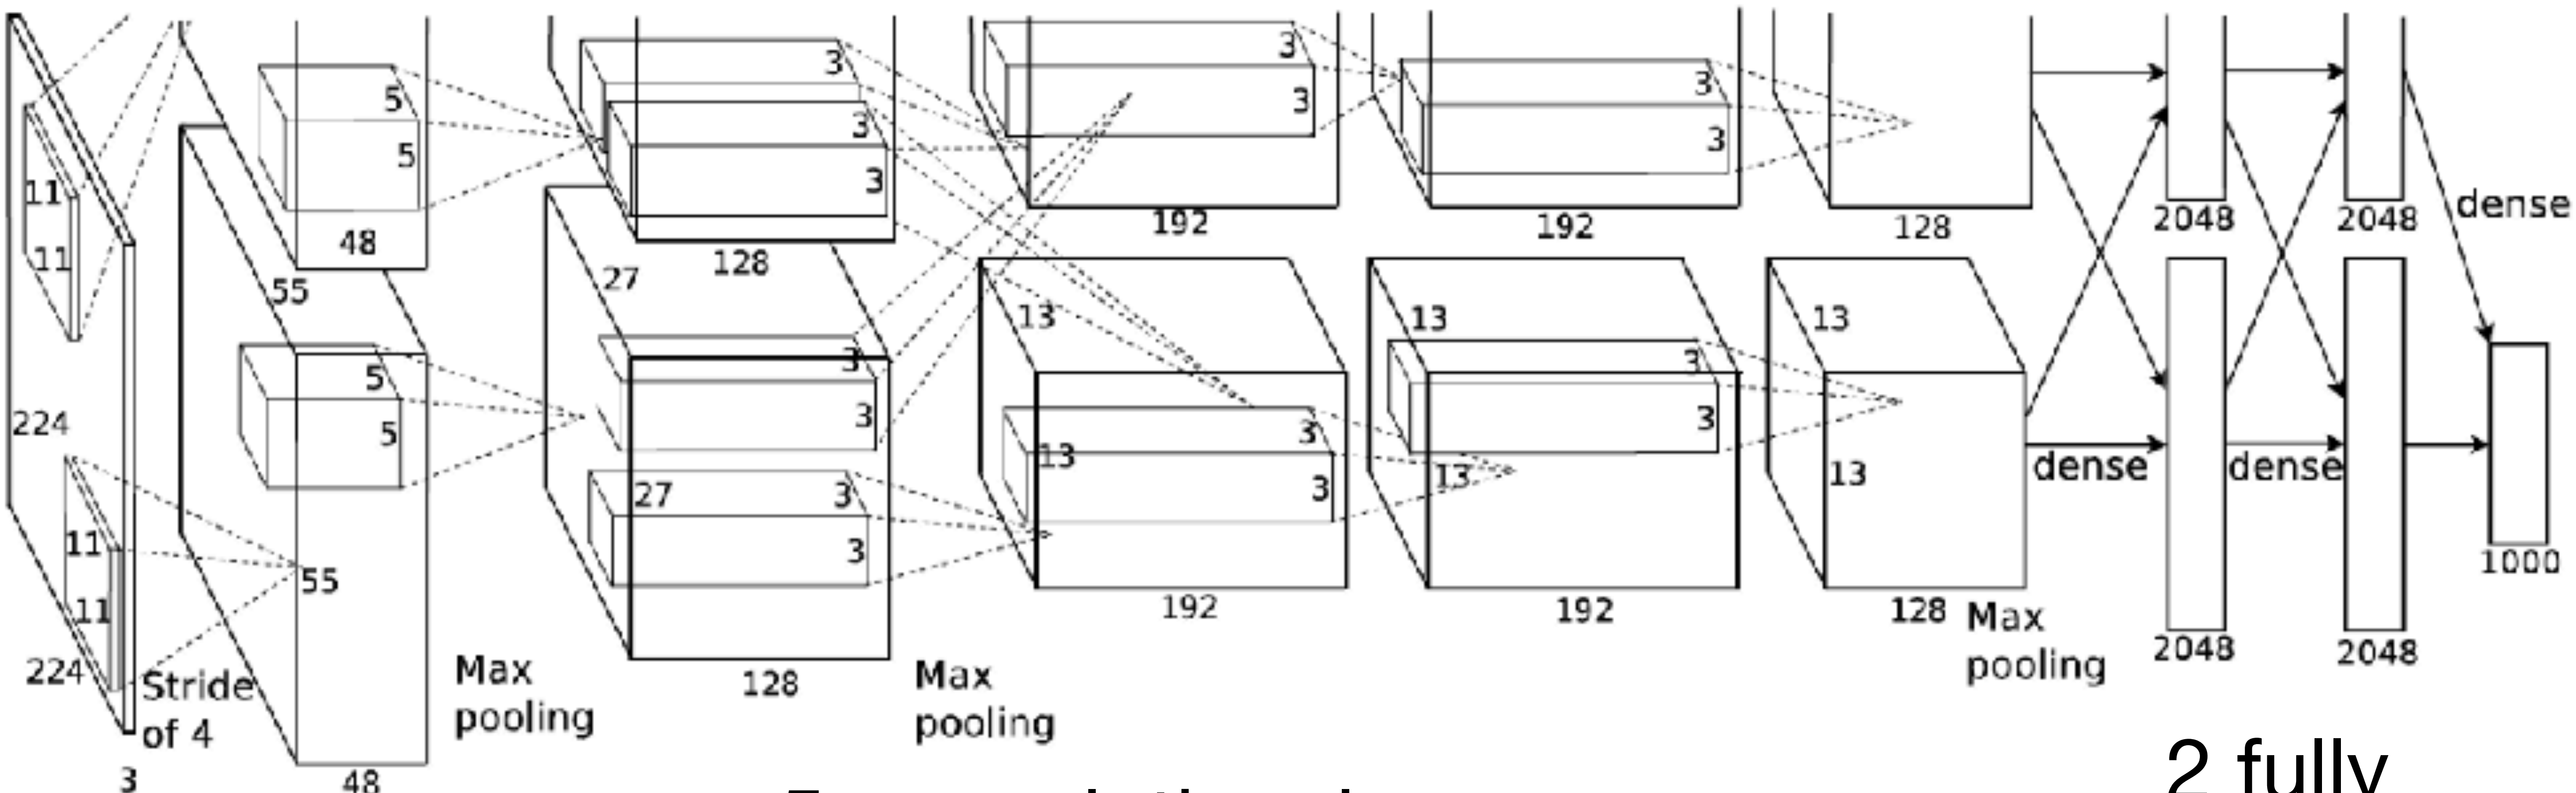

5 convolutional layers

2 fully connected layers

Main Innovations:

- non-smooth ReLU activations
- Used dropout instead of explicit regularization
- Max pooling to reduce the size of the network

ImageNet Challenge Winners

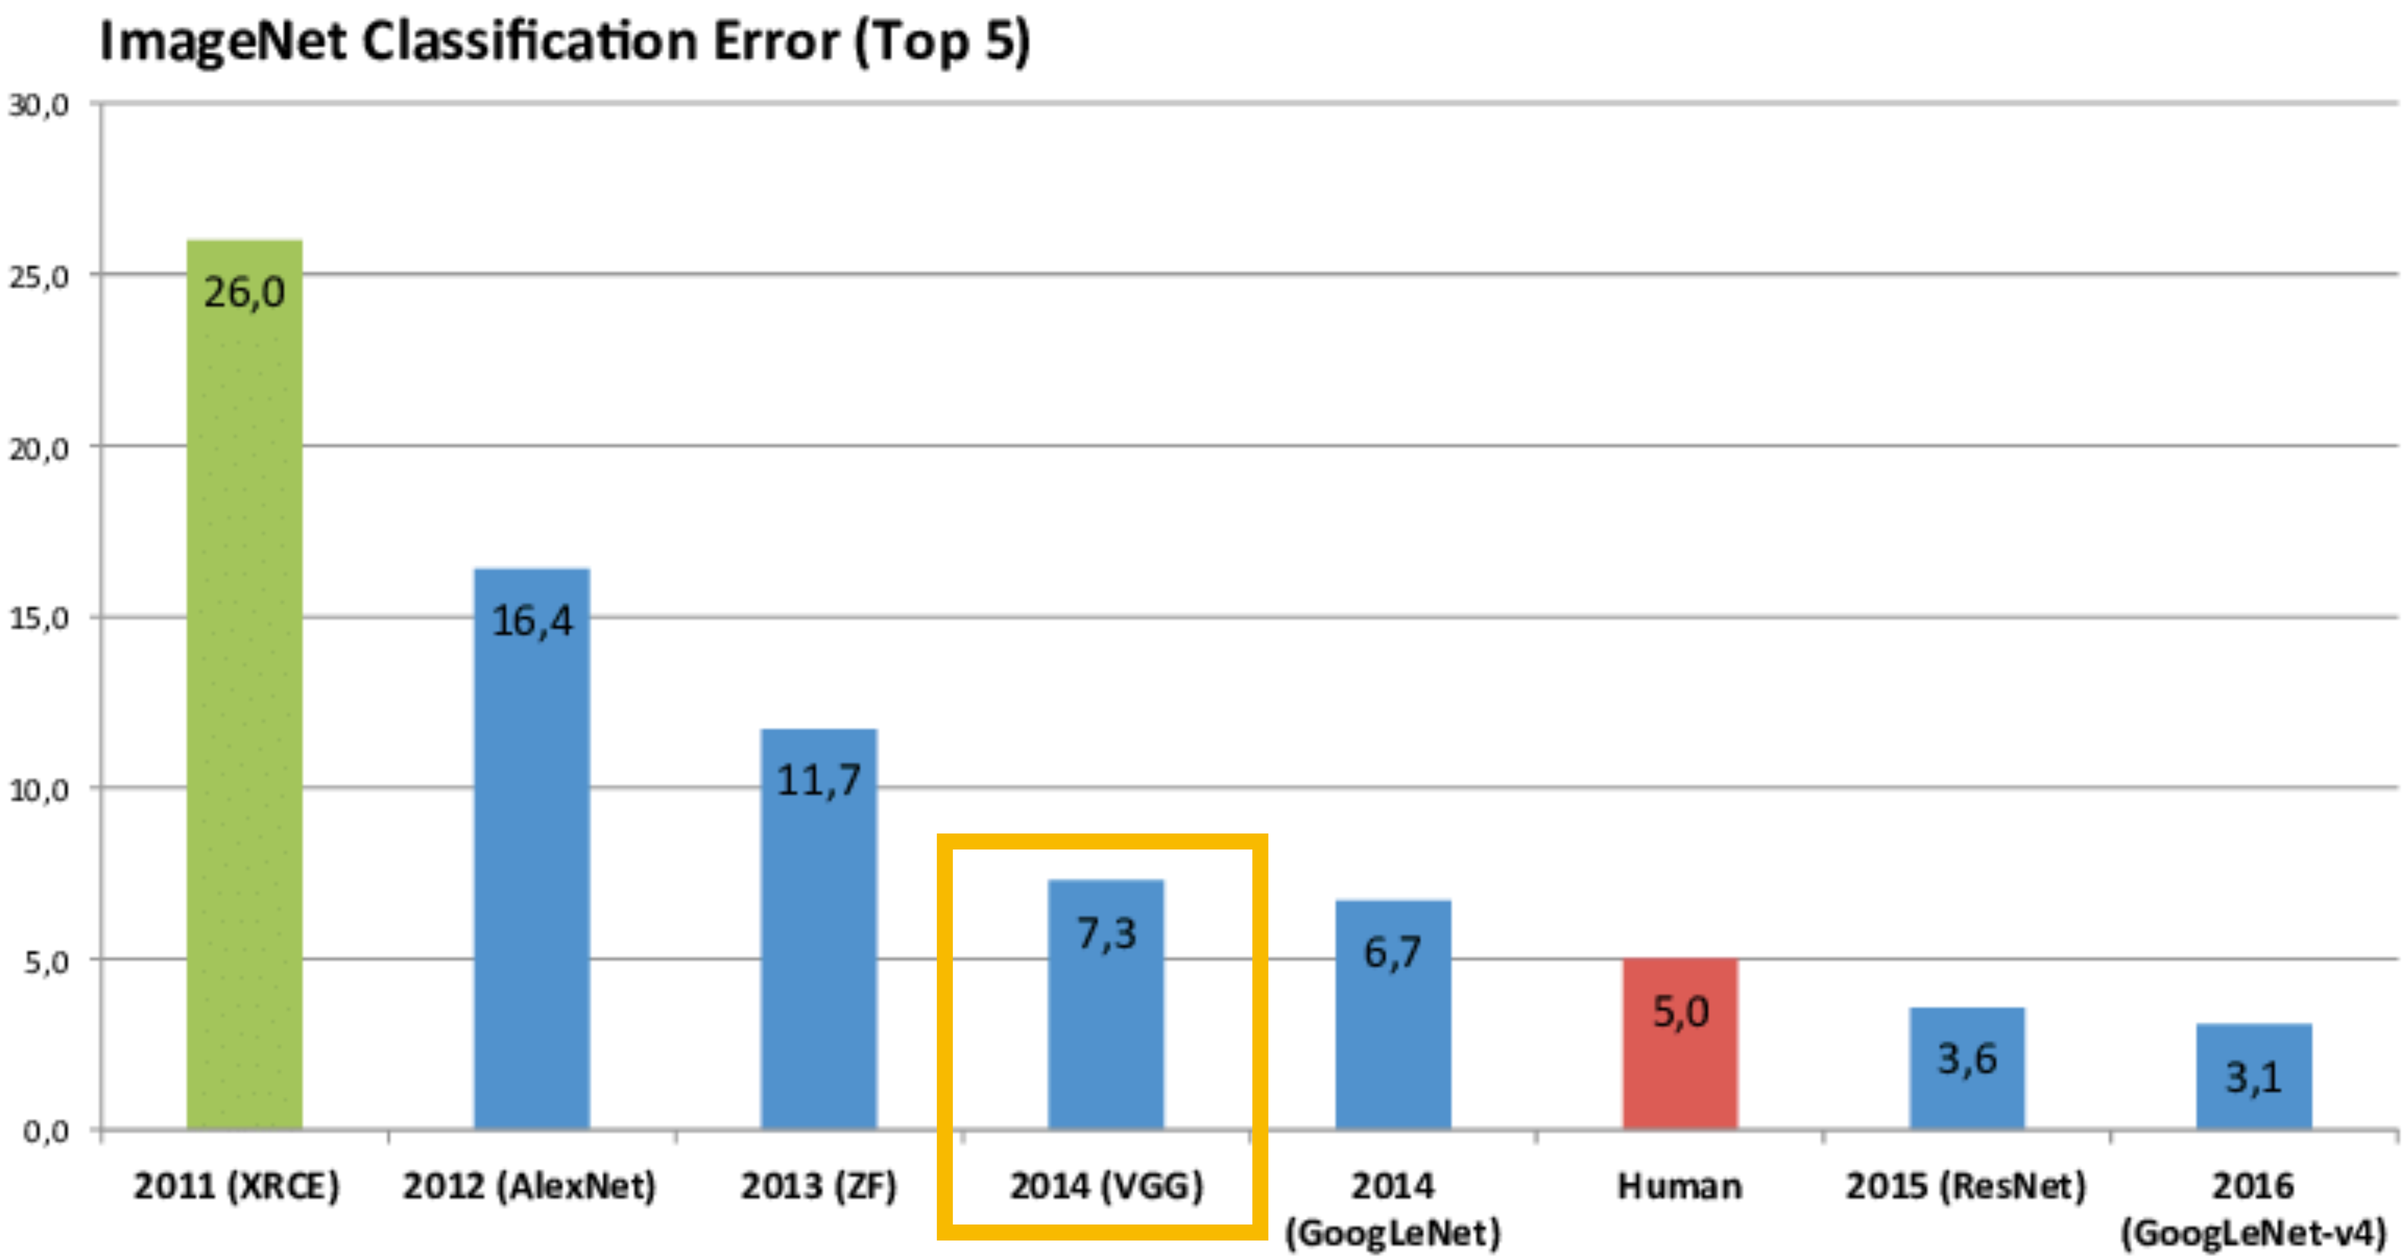

Year (winner)

Figure: Gustav von Zitzewitz

VGGNet [Simonyan & Zisserman , 2012]

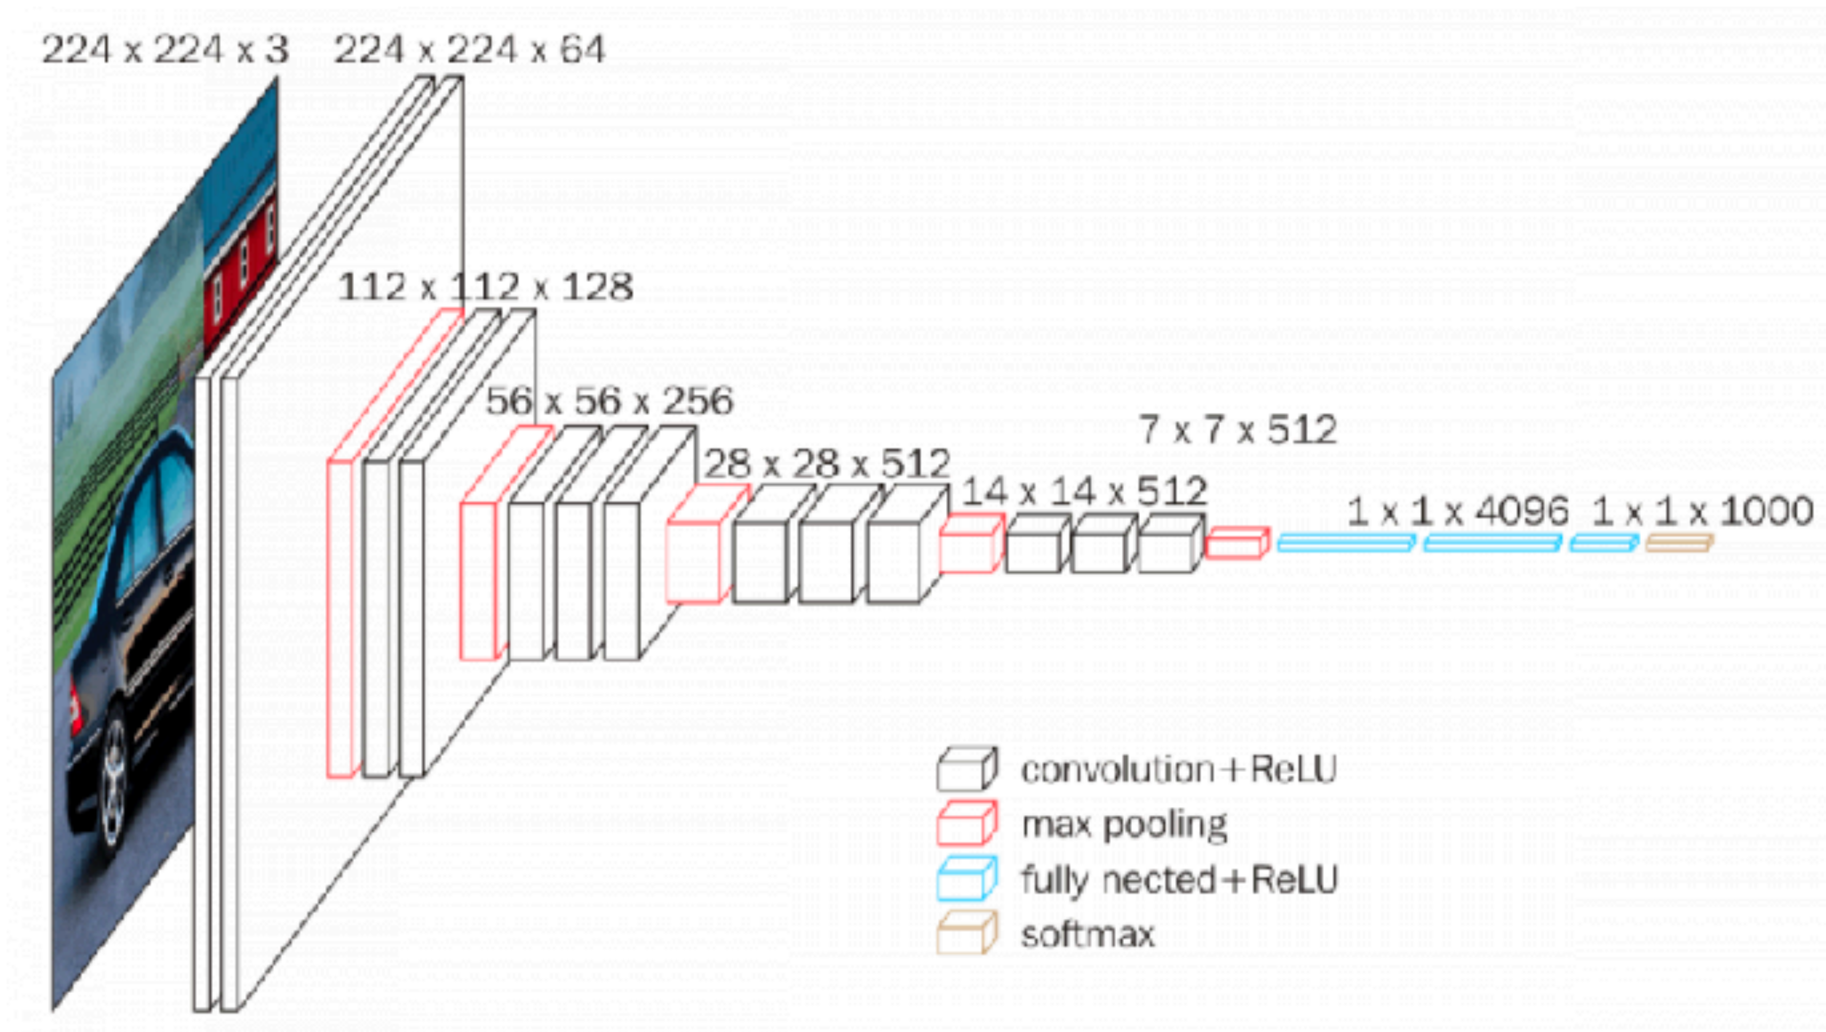

Main Innovations:

- Far deeper: 16-19 convolutional layers
- More & smaller filters per layer
  (e.g., rather than one 7x7 convolution use three 3x3 convolutions)

ImageNet Challenge Winners

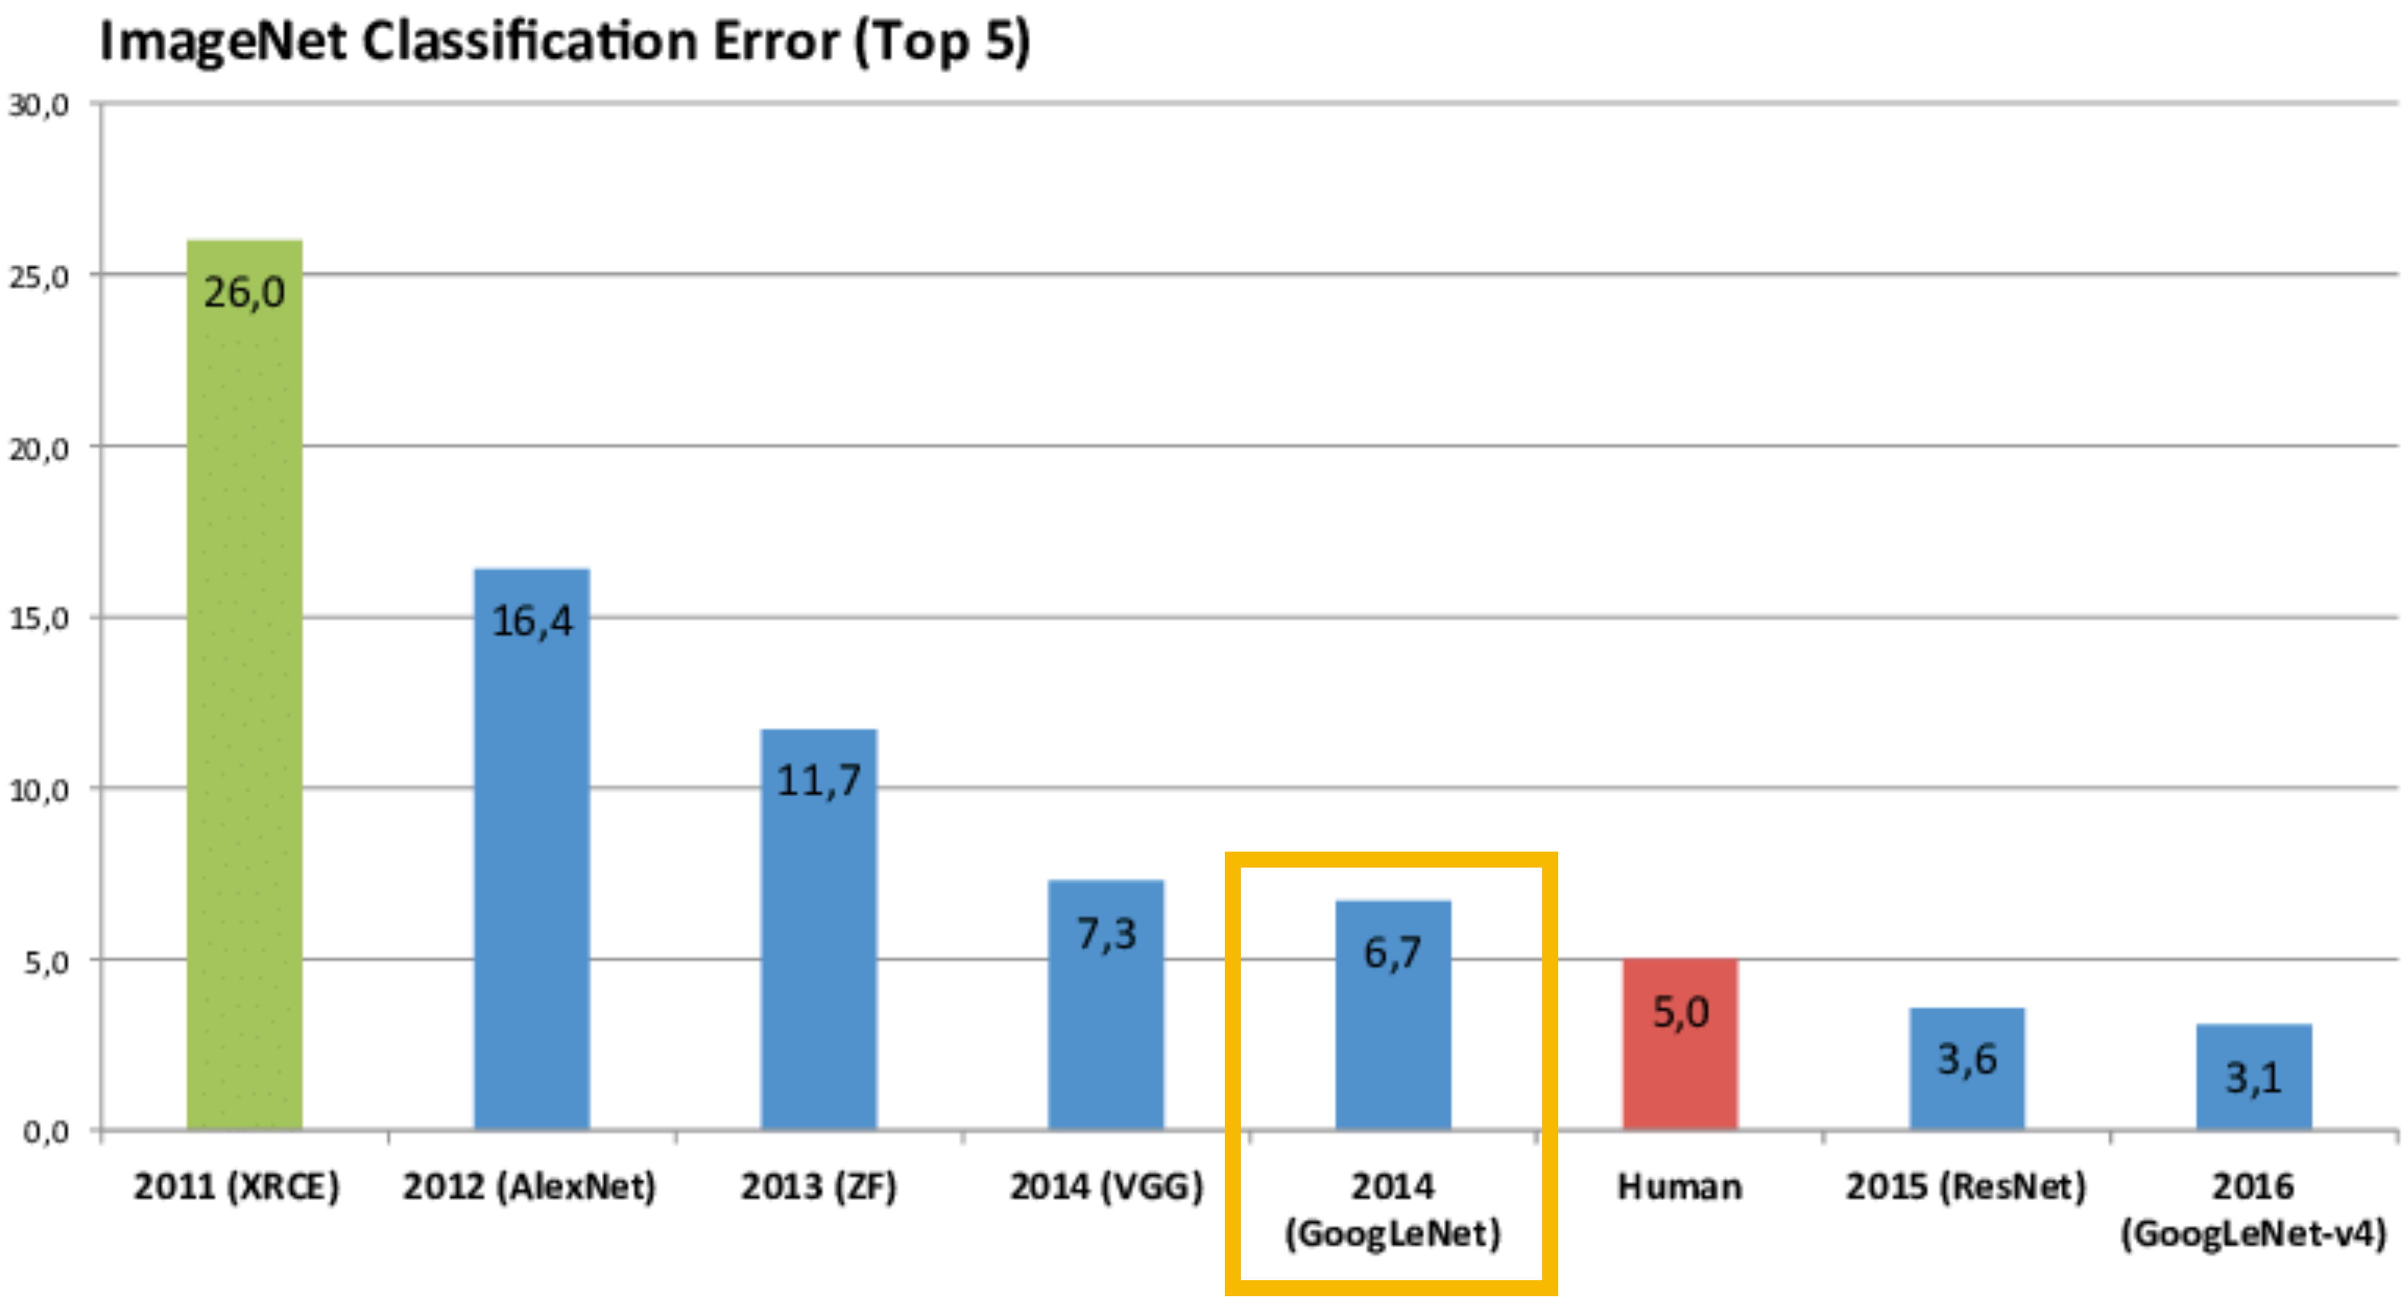

Year (winner)

Figure: Gustav von Zitzewitz

Inception-v1 (a.k.a. GoogleNet) [Szegedy et al., 2014]

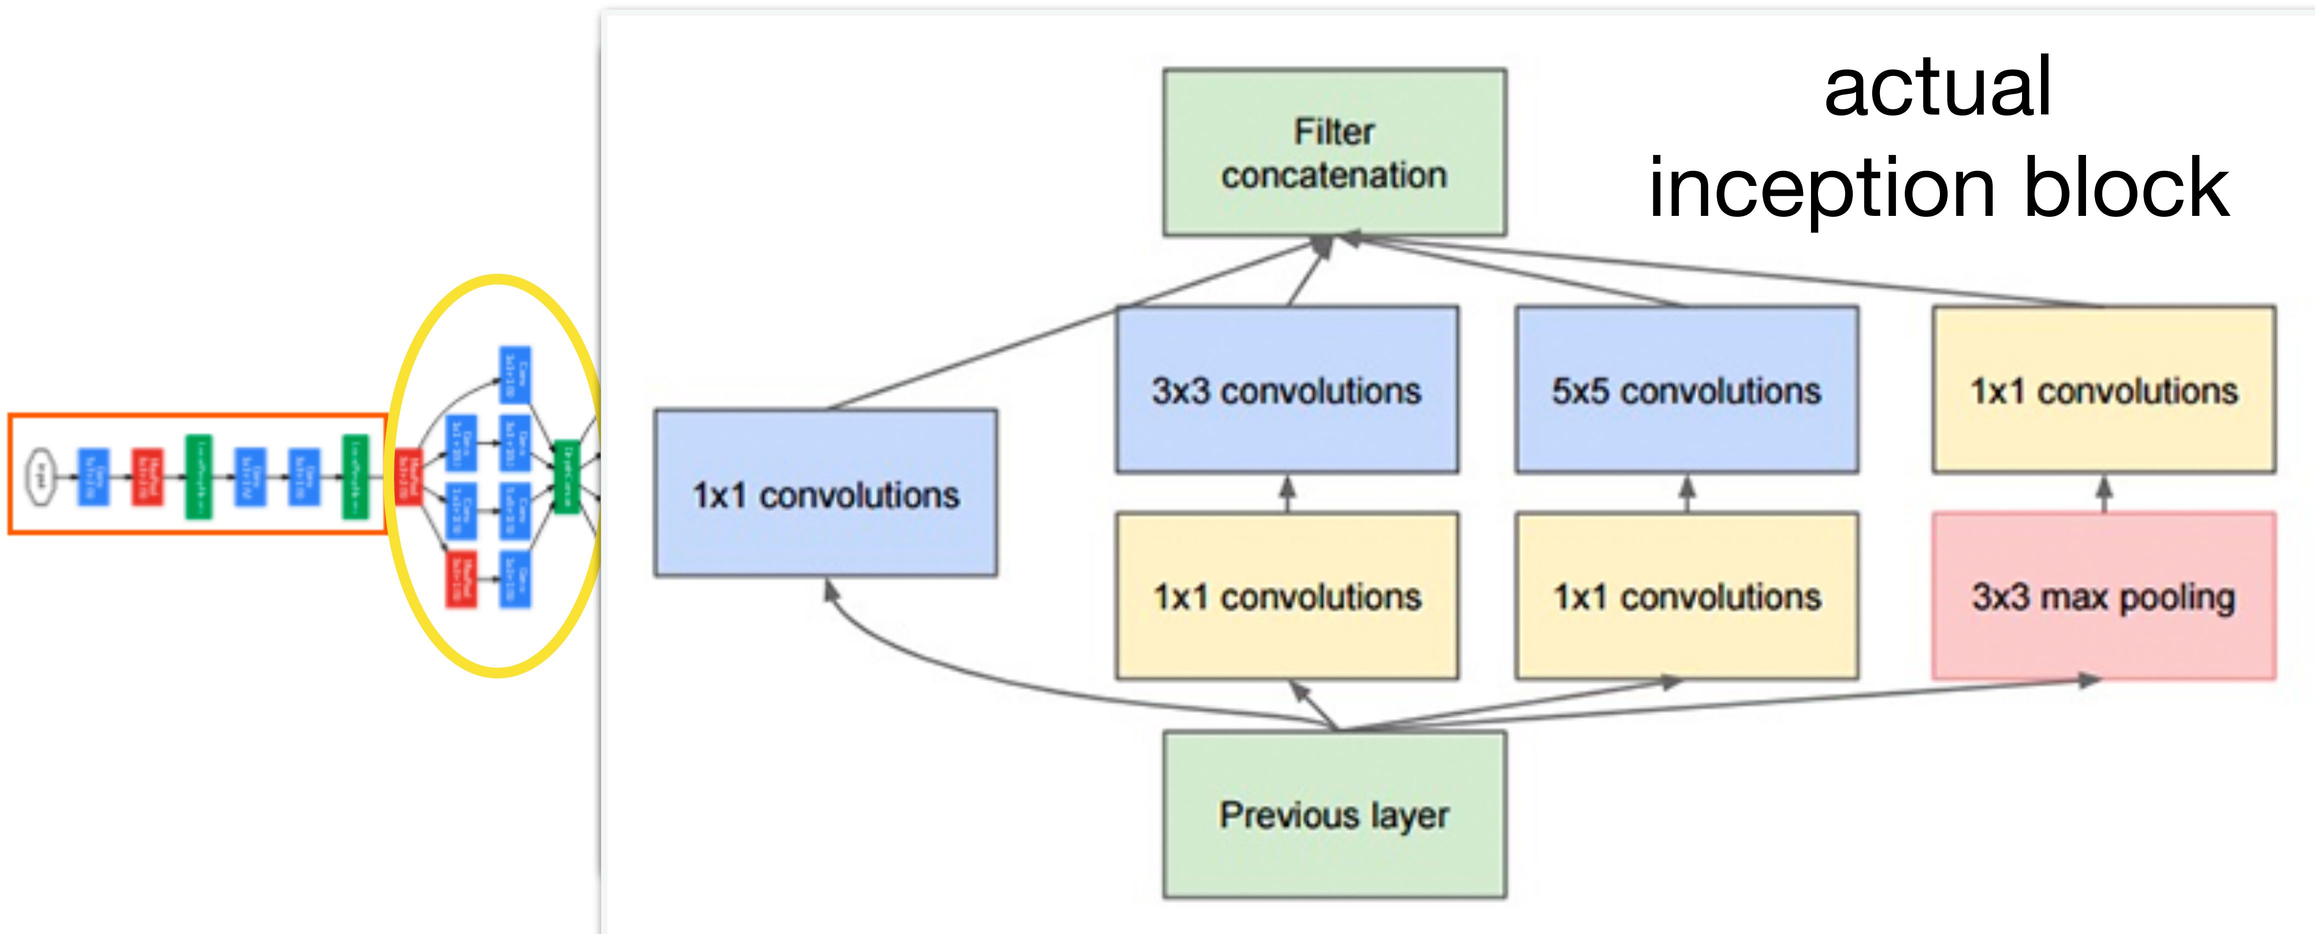

Main Innovations:

- Replace standard convolutional blocks with "inception block"
- Extracts features at multiple scales simultaneously
- No fully connected layers at the end — global average pooling

ImageNet Challenge Winners

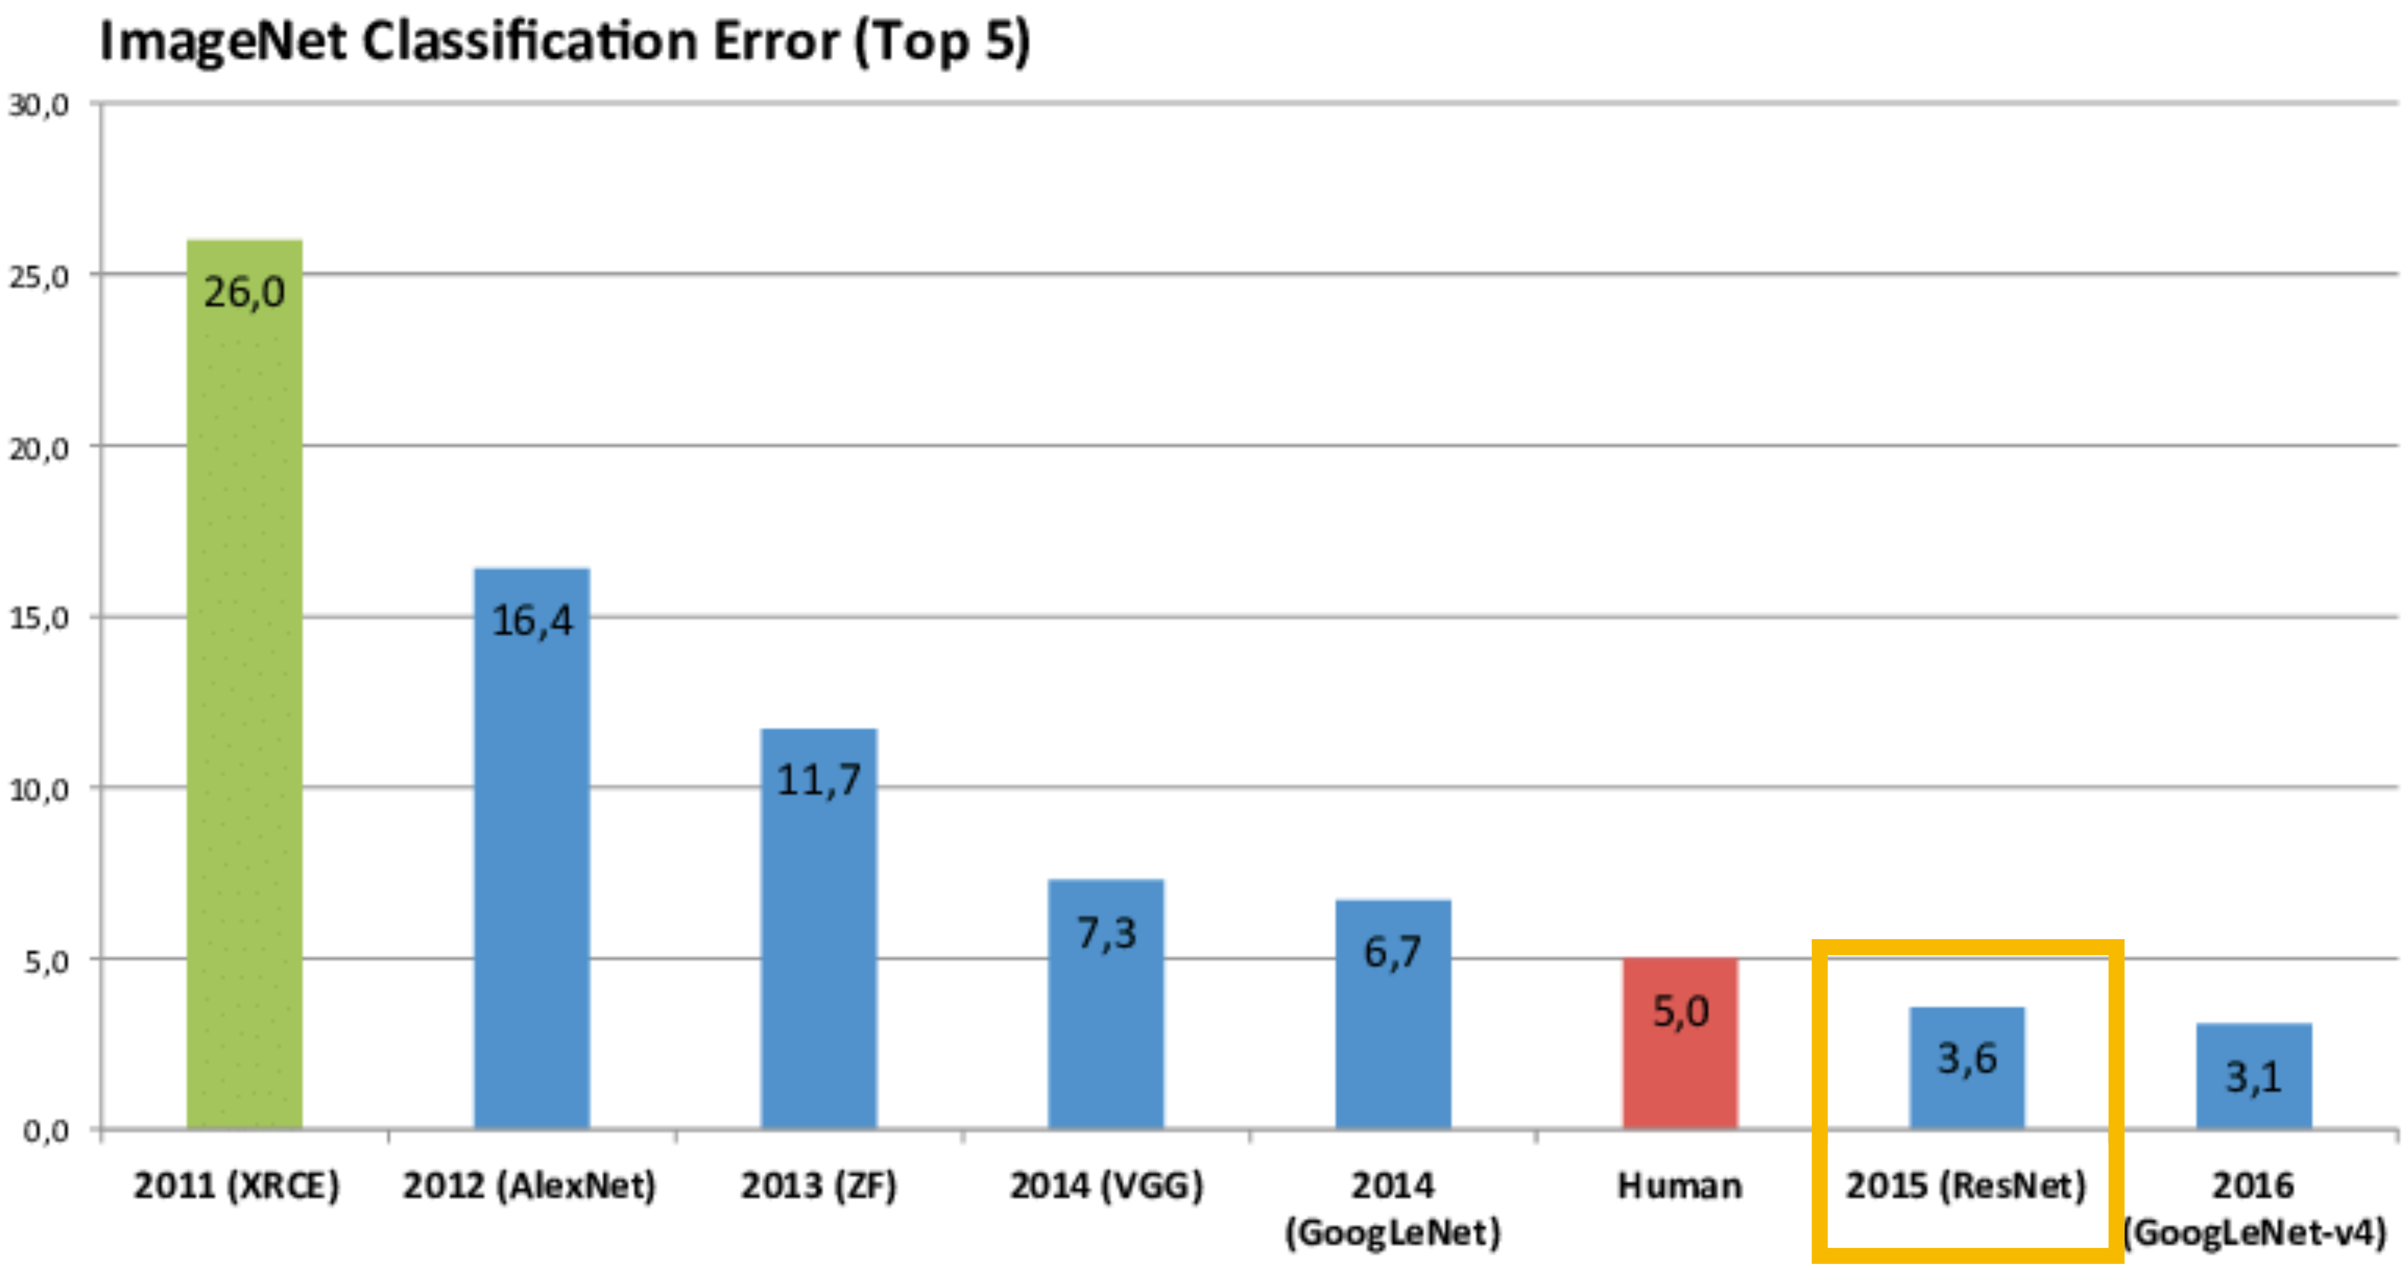

Year (winner)

Figure: Gustav von Zitzewitz

ResNets [He et al., 2014]

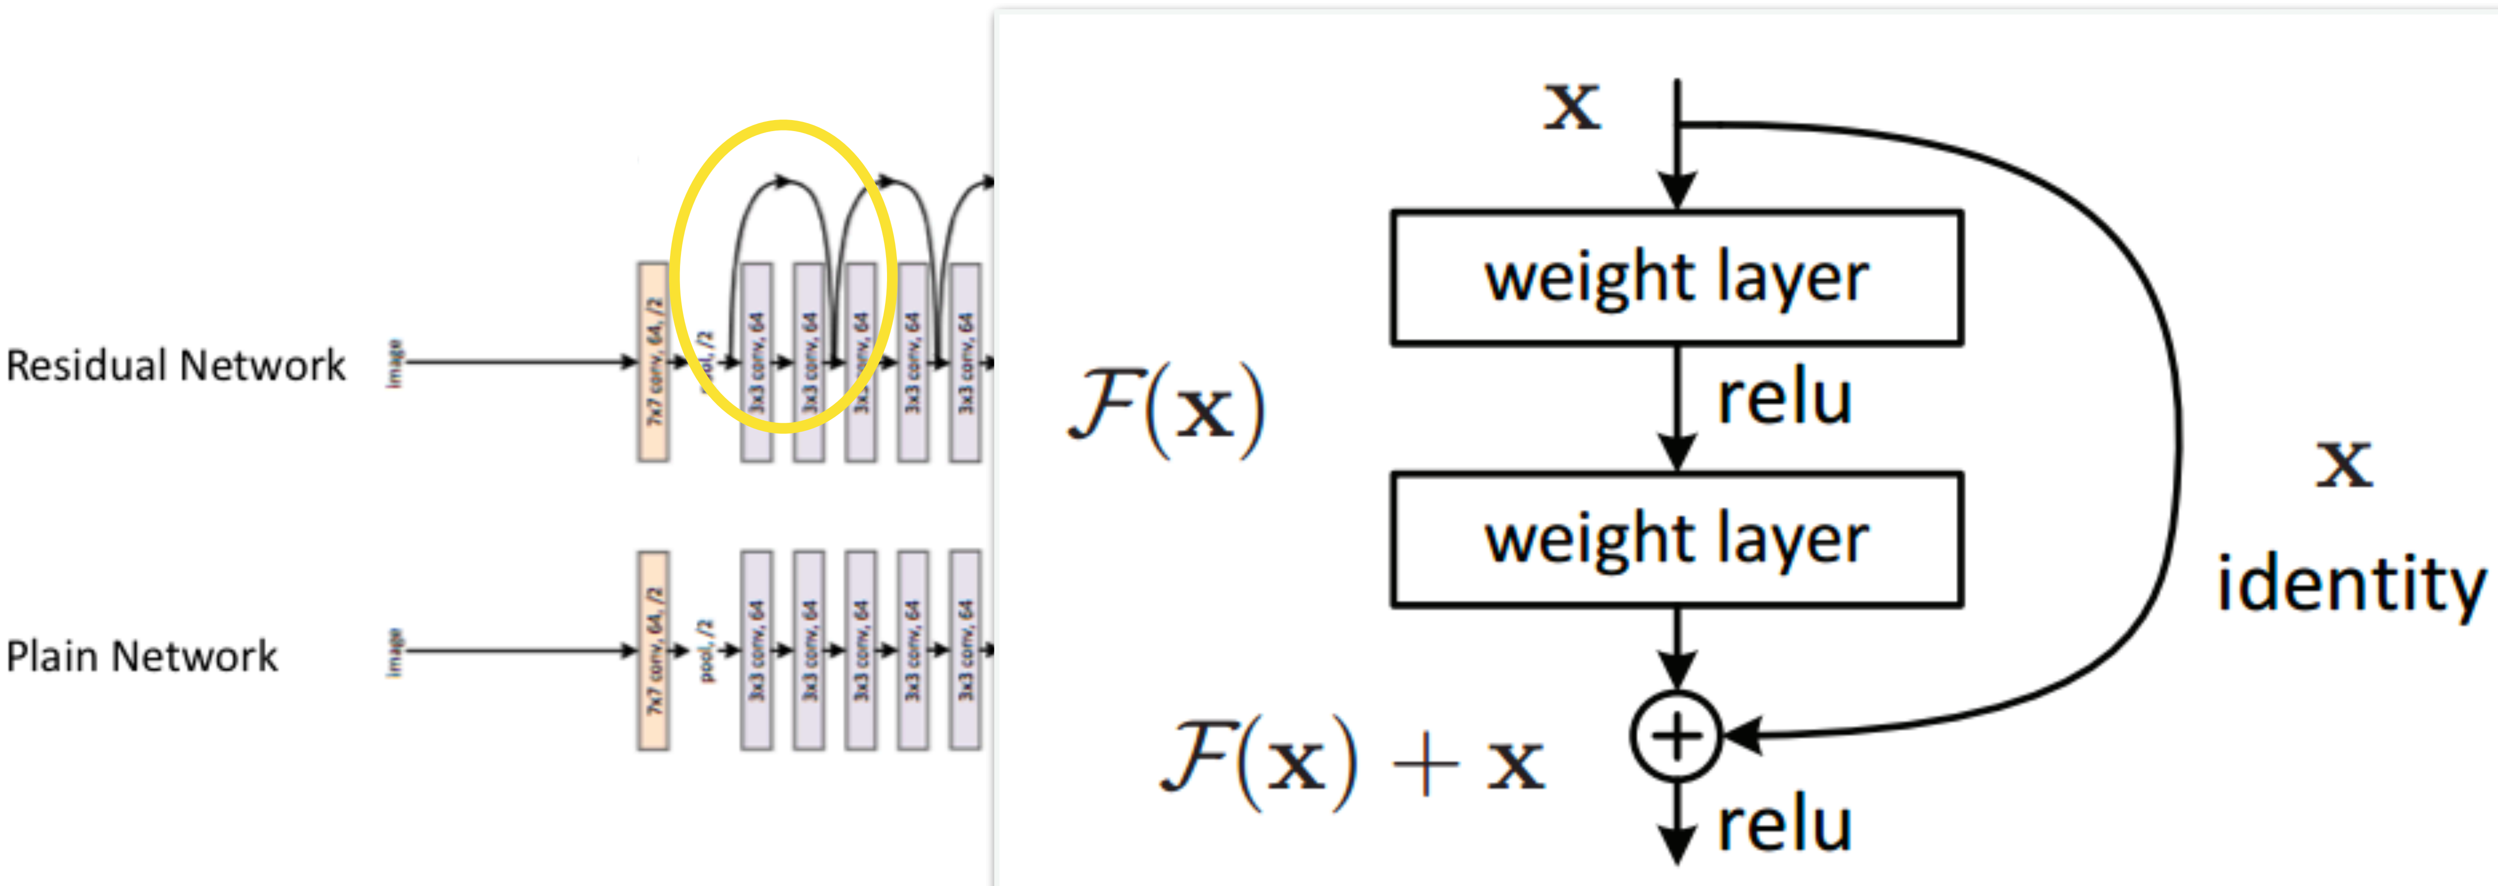

Main Innovations:

- "Skip connections"
- Alleviates "vanishing gradients" issue of deep networks
- Faster training, fewer "hacks" needed (e.g., batch normalization)
- Can train vastly deeper model, e.g., 100+ layers

ResNets—smoother loss landscapes

2-D projections of optimization landscape

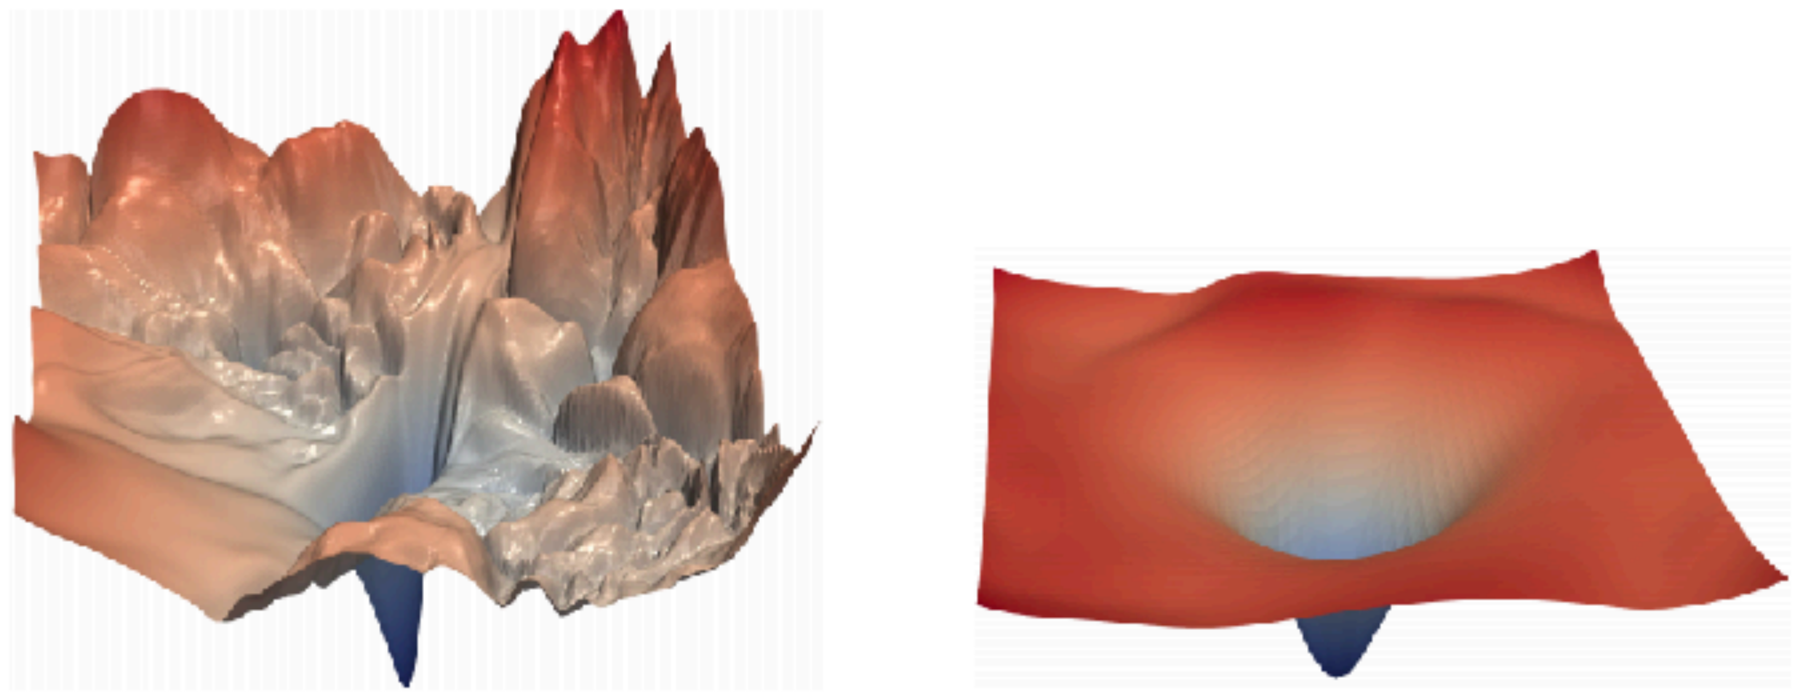

(a) without skip connections (b) with skip connections

Figure 1: The loss surfaces of ResNet-56 with/without skip connections. The proposed filter normalization scheme is used to enable comparisons of sharpness/flatness between the two figures.

Visualizing the Loss Landscape of Neural Nets

Hao Li[1], Zheng Xu[1], Gavin Taylor[2], Christoph Studer[3], Tom Goldstein[1]
[1]University of Maryland, College Park [2]United States Naval Academy [3]Cornell University
{haoli,xuzh,tomg}@cs.umd.edu, taylor@usna.edu, studer@cornell.edu

ImageNet Challenge Winners

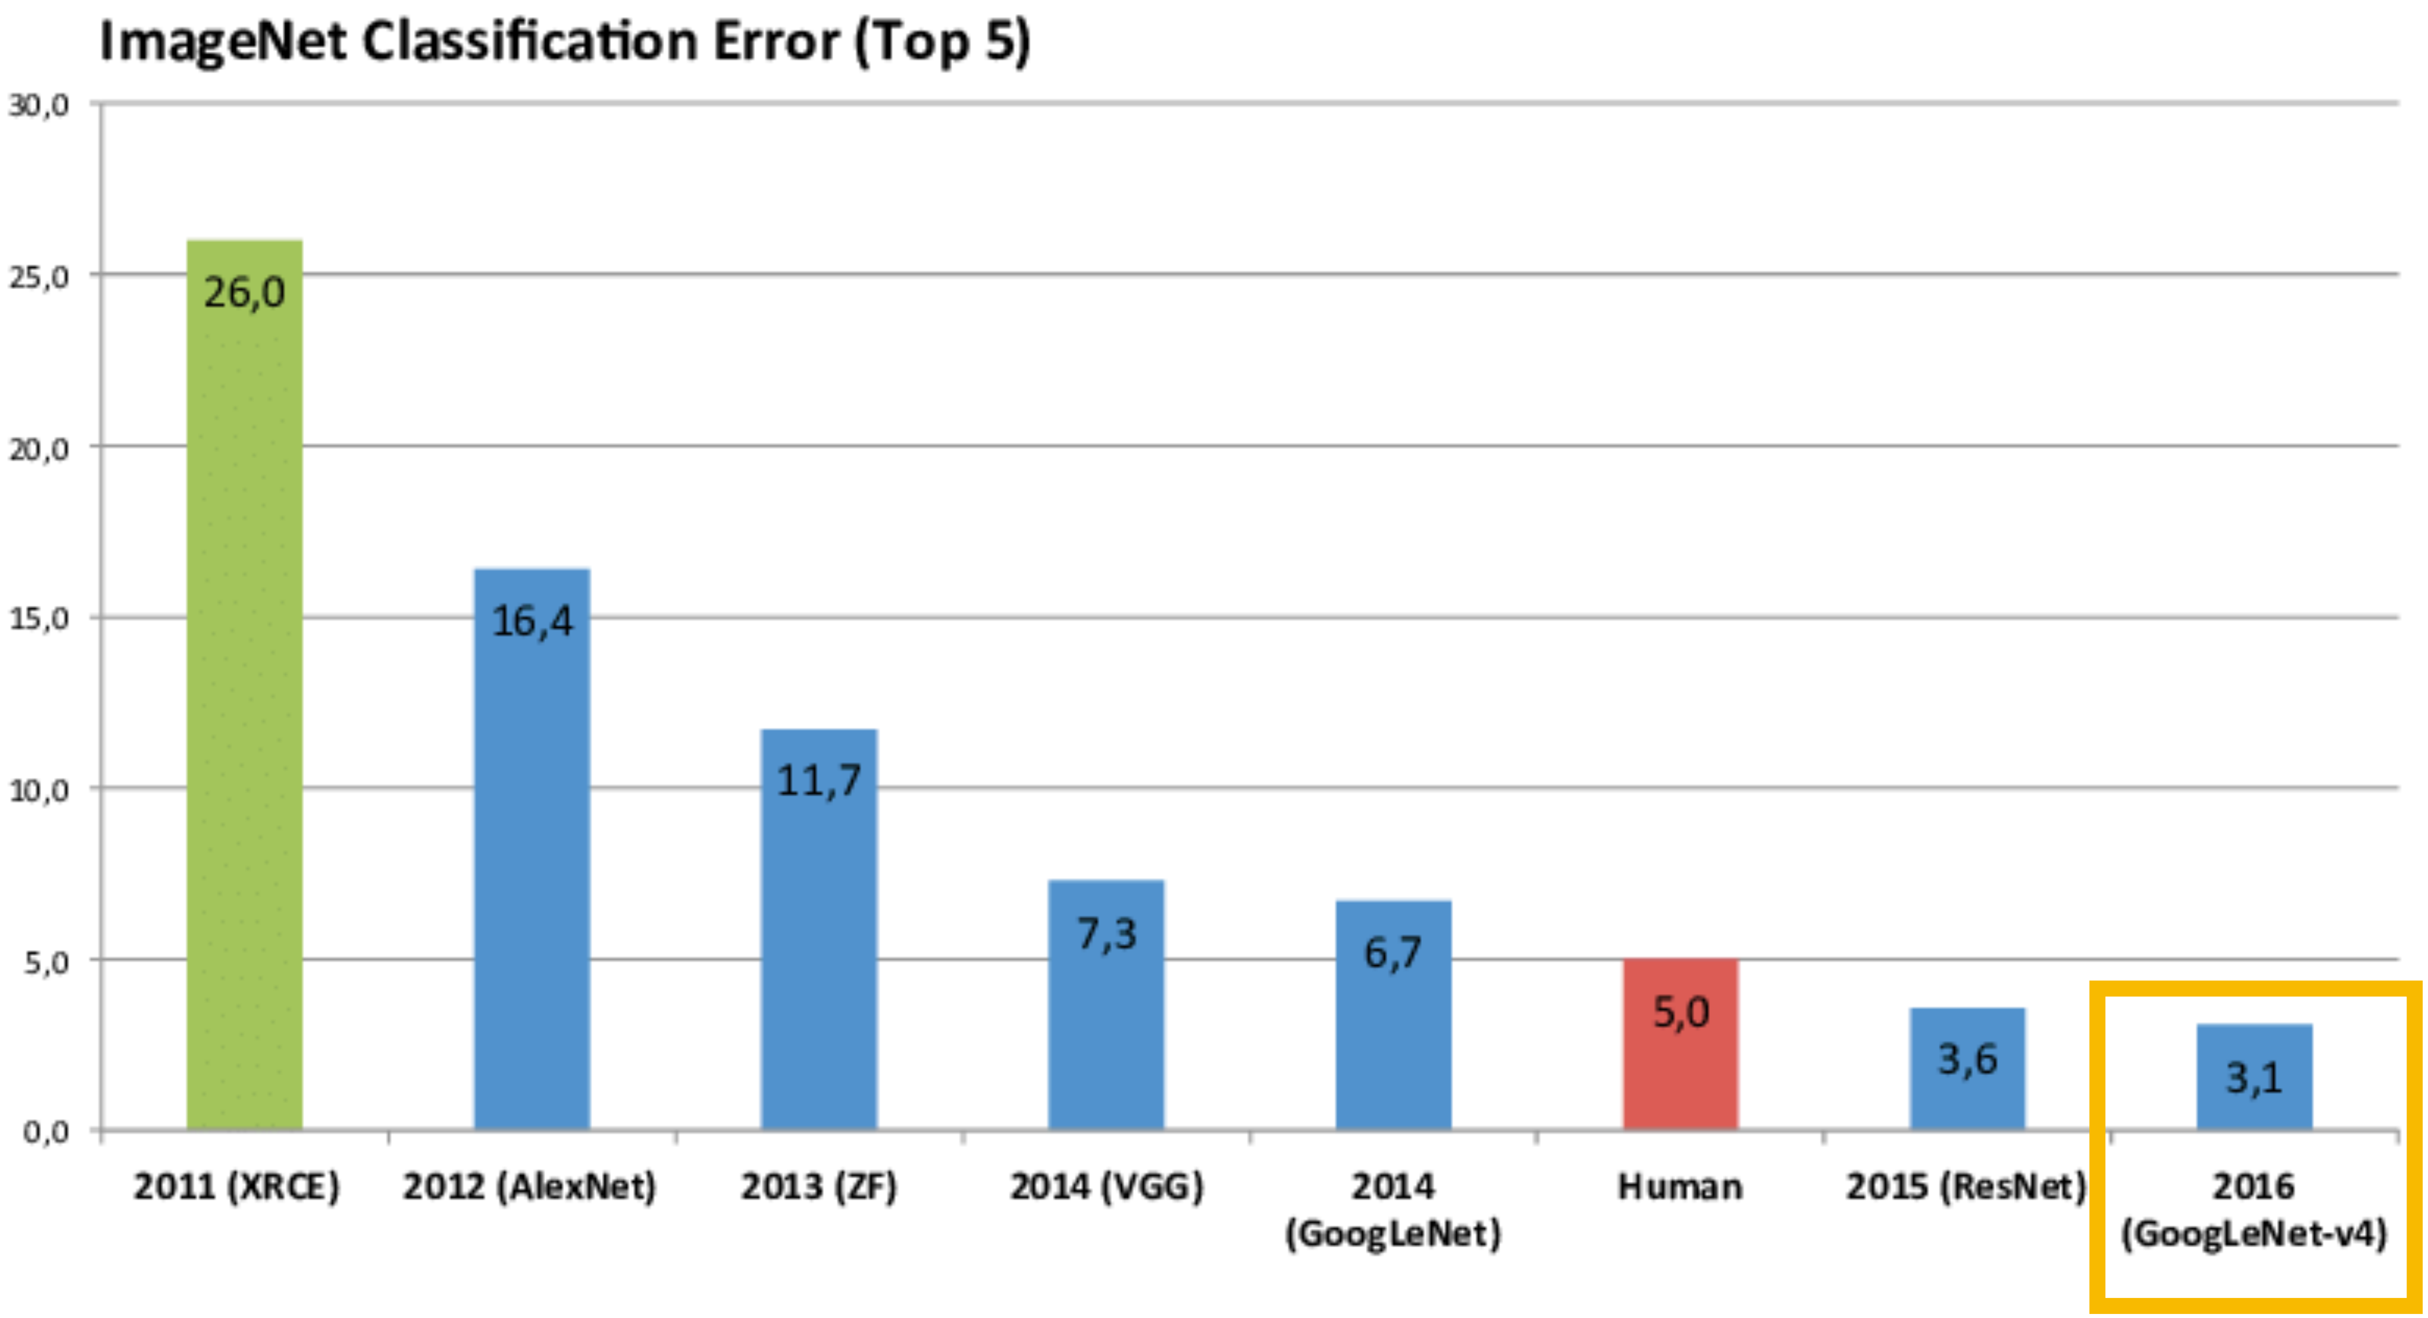

Year (winner)

Figure: Gustav von Zitzewitz

Inception-v3, -v4, -ResNet

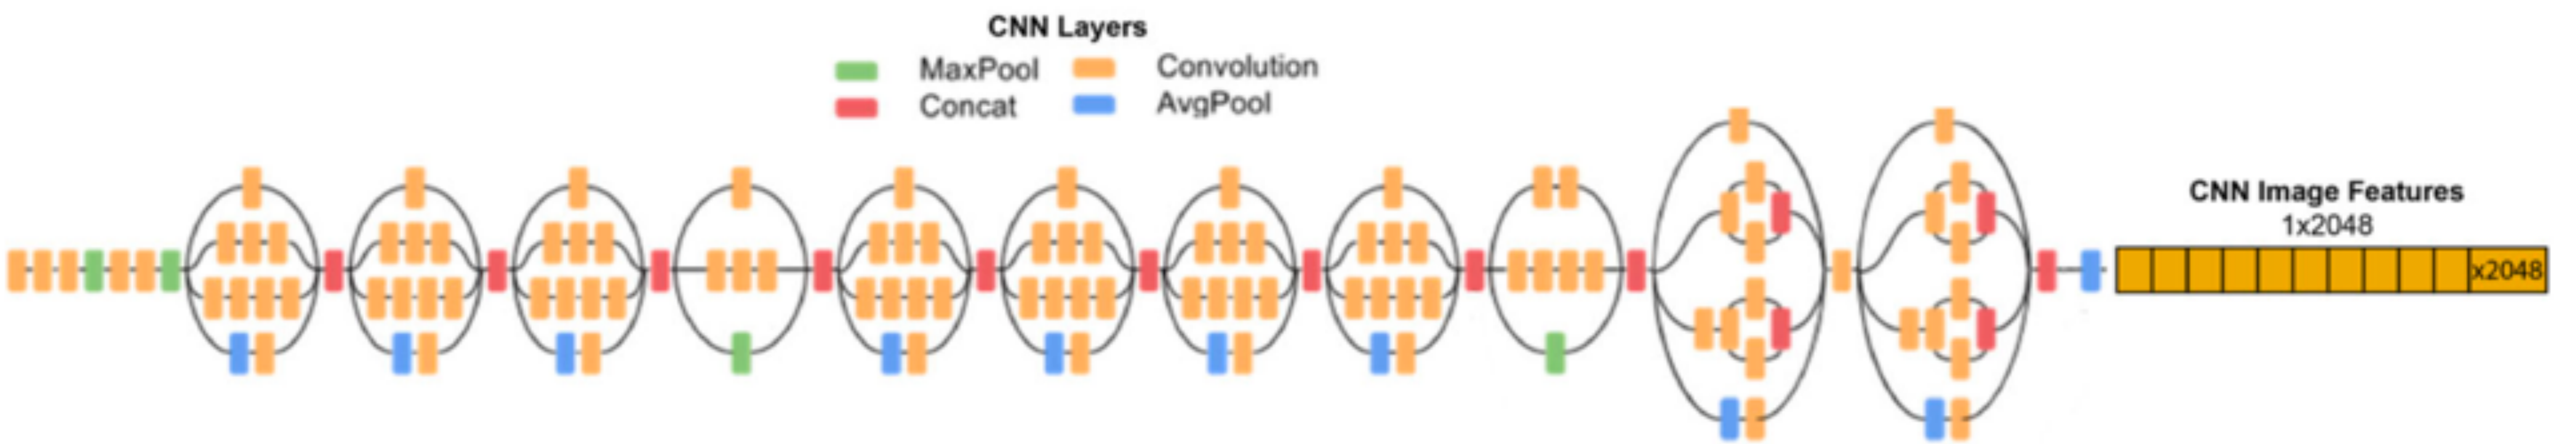

Main Innovations:

- Inception-v3 - Deeper, more efficient inception blocks
- Inception-v4 - “”
- Inception-ResNet - adds skip connections to inception blocks

…and many others

DenseNets
[Huang et al., 2016]

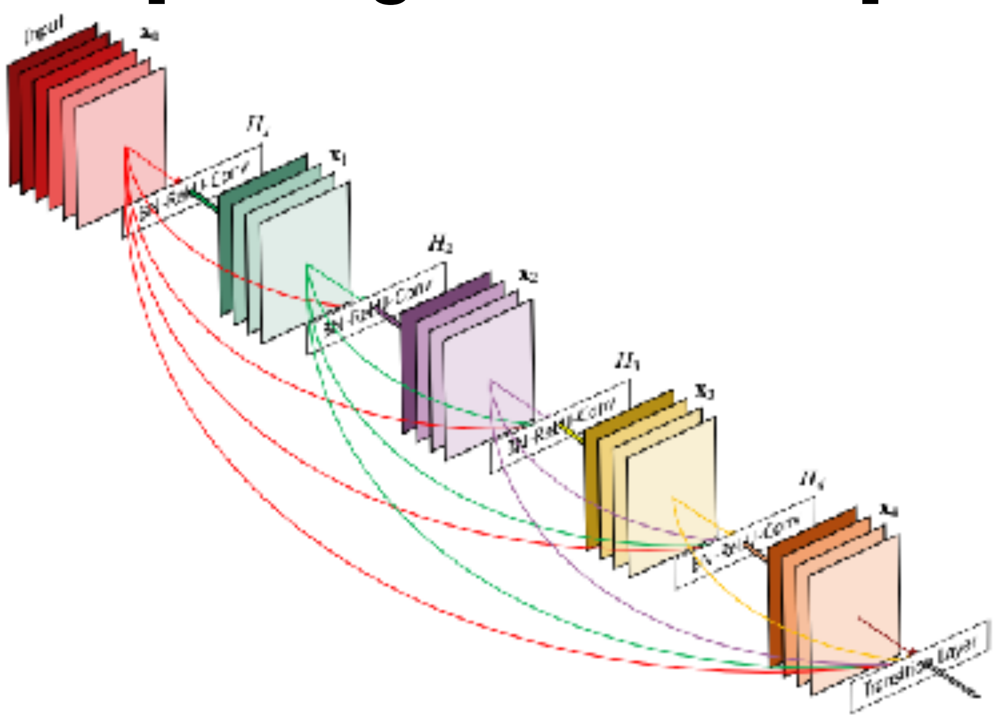

FractalNets
[Larsson et al., 2016]

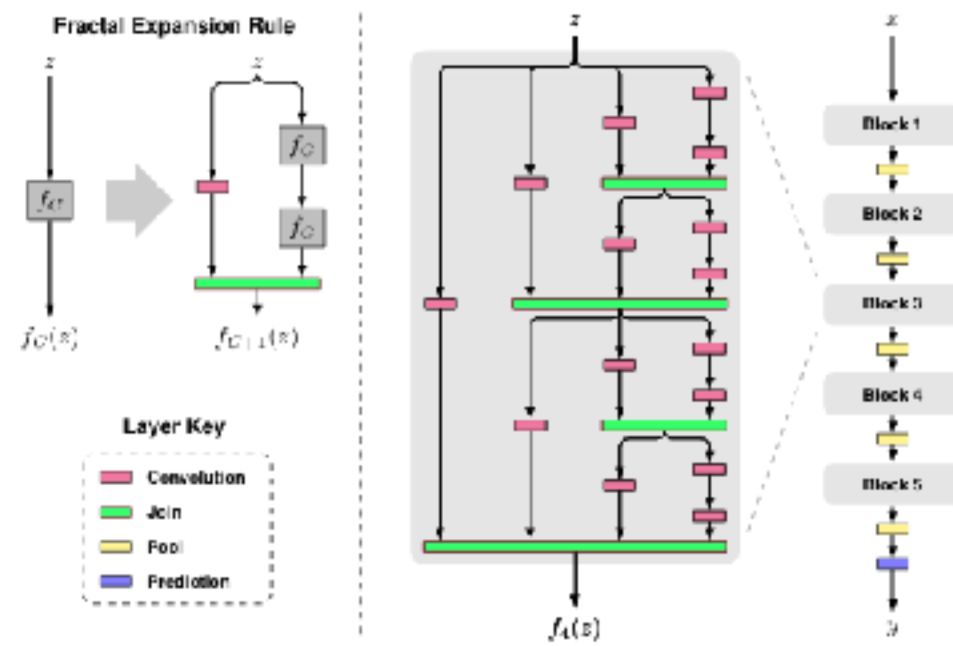

TOYOTA TECHNOLOGICAL INSTITUTE AT CHICAGO

Squeeze-and-Excitation Nets
[Hu et al., 2017]

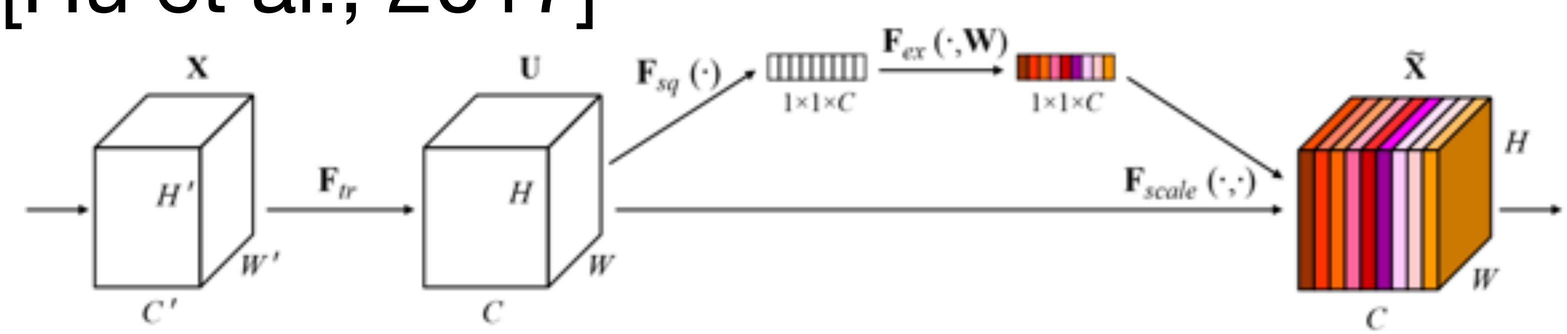

Fig. 1. A Squeeze-and-Excitation block.

What network architecture should I use?

- Short answer: It doesn't matter too much!
- Most state-of-the-art networks are available in standard deep learning toolboxes
- Task will dictate architecture
- Constraints:
  - Memory
  - Size of training set
  - Deployment

Challenges adapting deep CNN’s to biomedical imaging problems

- **Challenge 1: Limited Training Data**
  - 1M+ training examples ImageNet, biomedical imaging 100-10k typical
  - How do we train a deep CNN without overfitting?
- **Challenge 2: Complex Input Formats**
  - 3D volumes are commonplace in medical imaging
  - multi-stream or multi-modal data (e.g., CT + MRI scans, text + image)
  - measurement domain data (e.g., raw data from MRI scanner)
- **Challenge 3: Tasks Beyond Classification**
  - Task is not classification/regression (or is inefficiently represented as such)
    - Segmentation
    - Image restoration/reconstruction

Biomedical image classification/detection

Application: Detecting skin cancer by classification of lesions in photographs

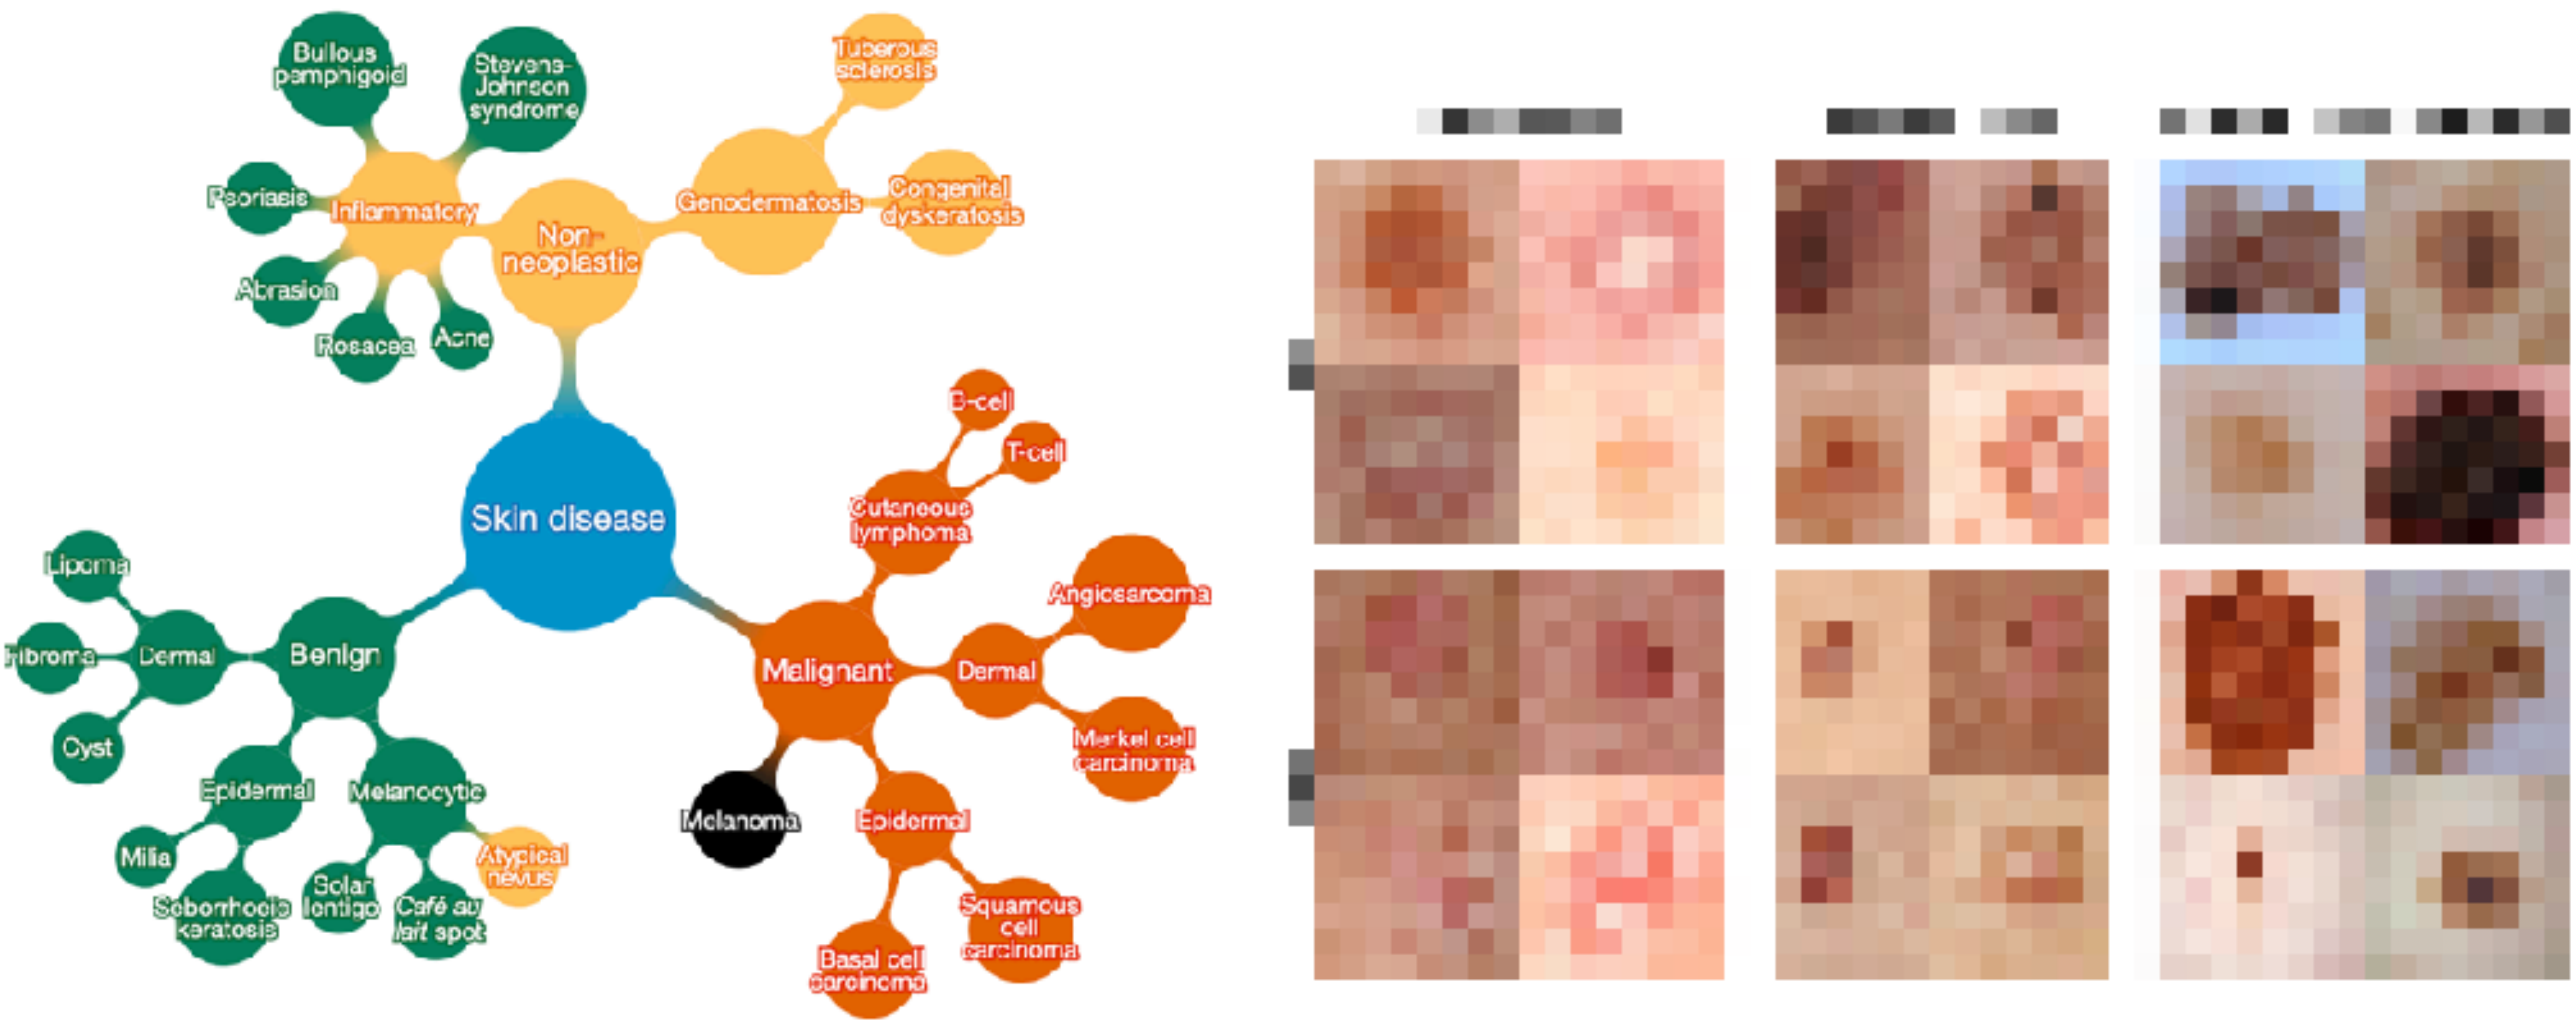

Figure: [Esteva et al., 2017]

- Nature paper [Esteva et al., 2017]
- Dataset of 129,450 clinical images
- 2,032 different diseases.

Use Inception-v3 network

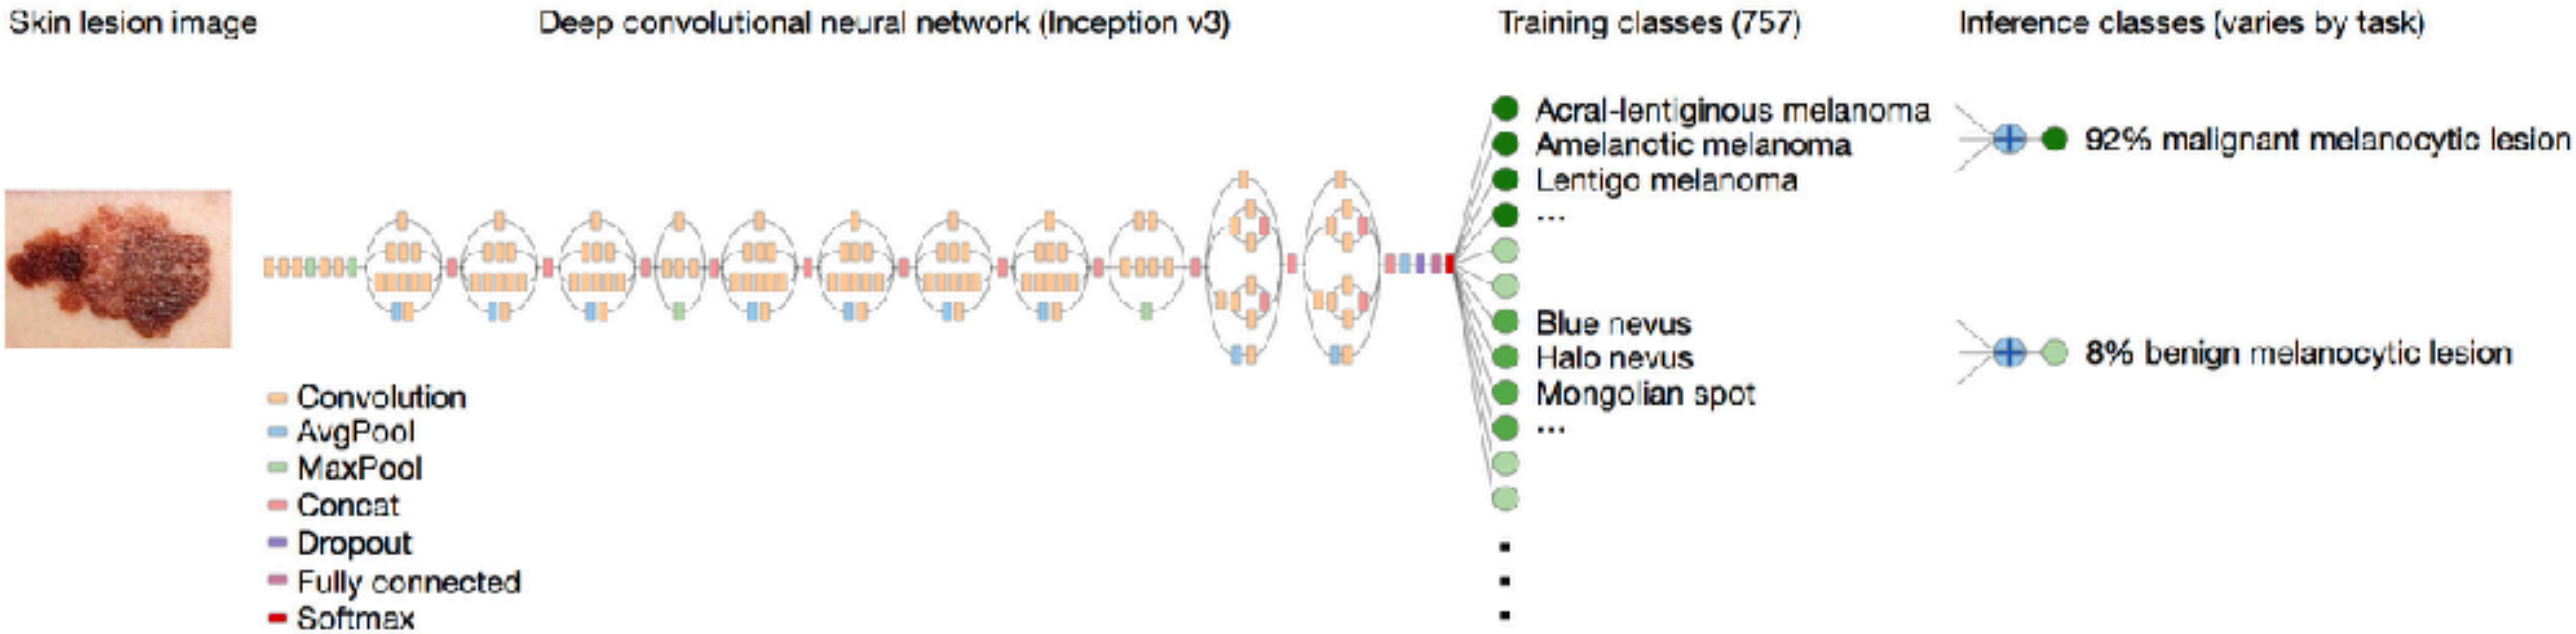

CNN outperforms dermatologists!

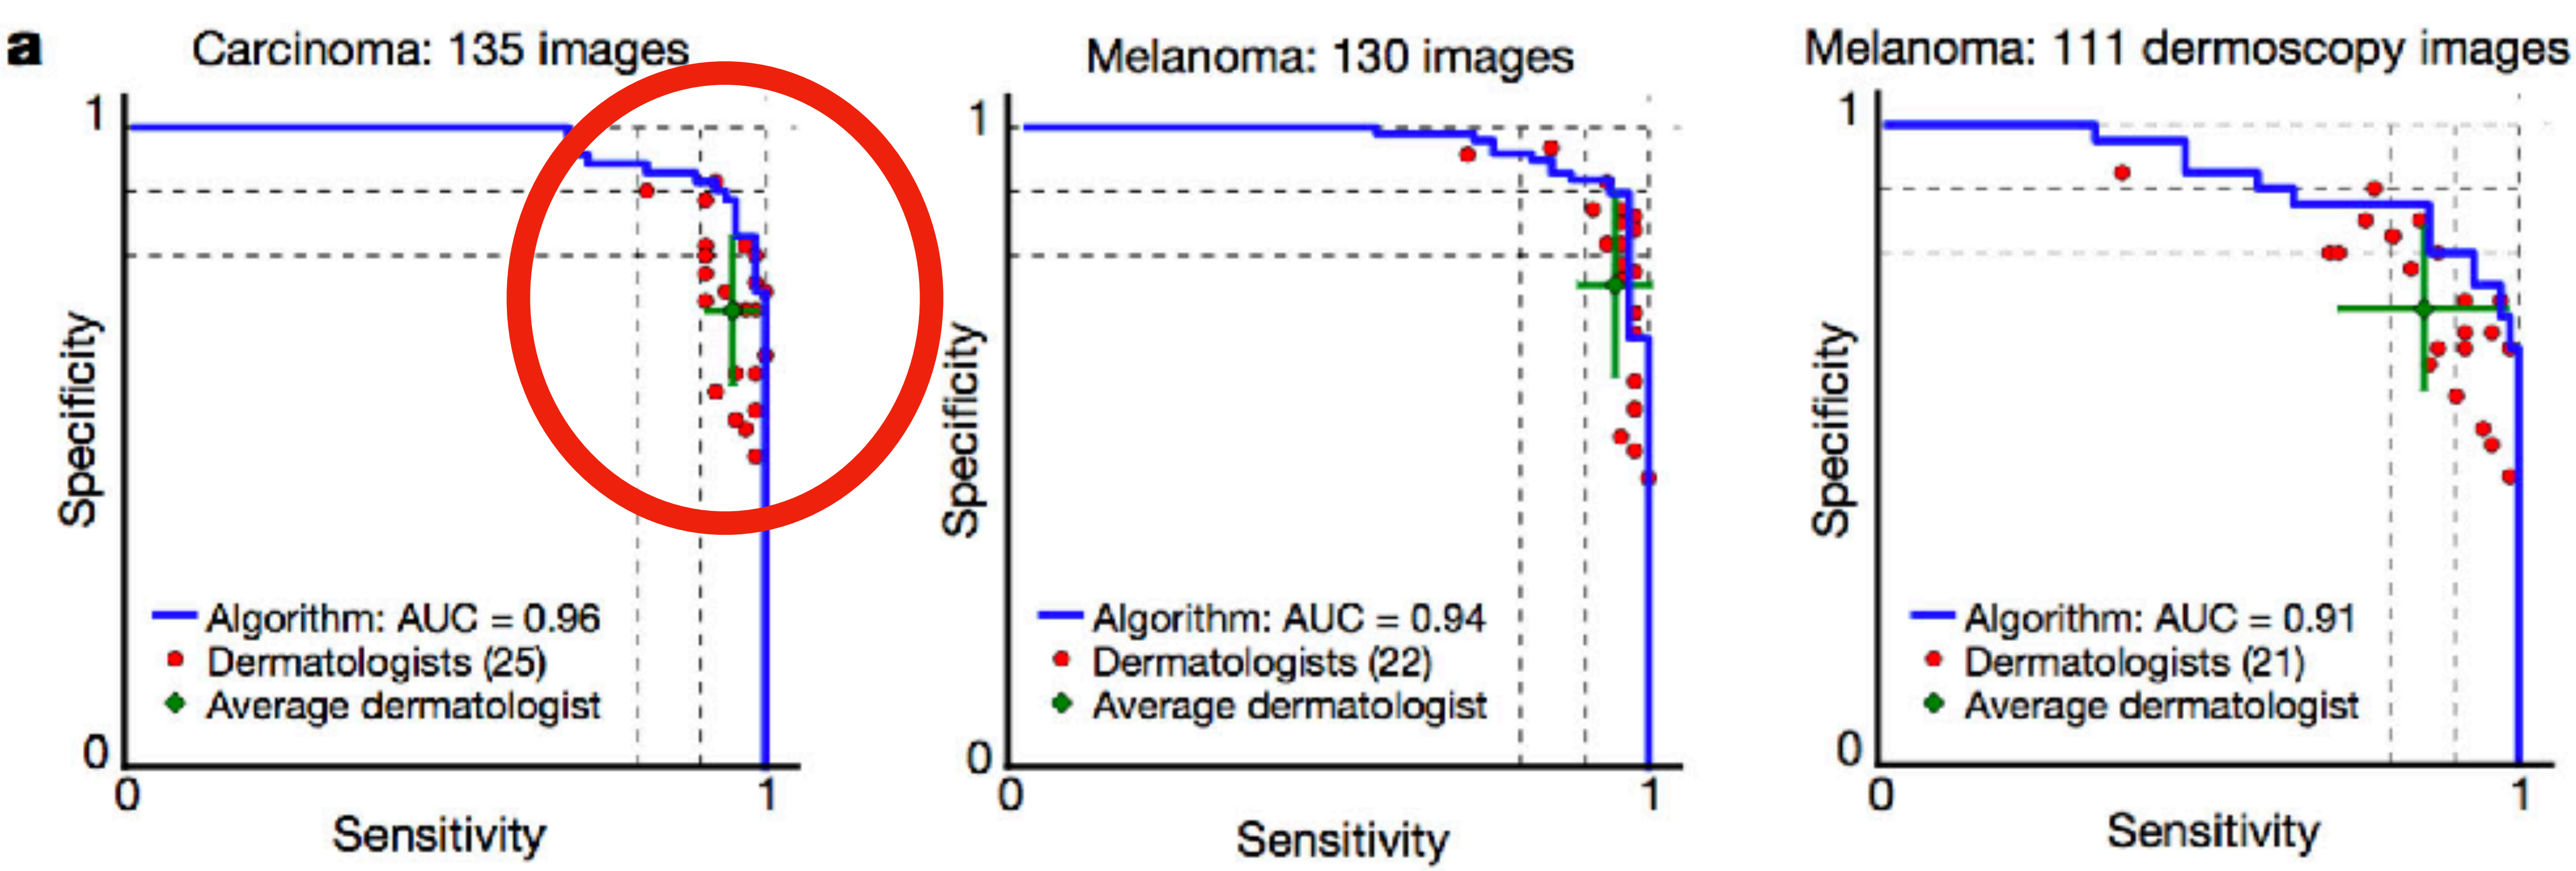

Transfer Learning

- (Esteva et al., 2017), and nearly every other biomedical image classification approach makes use of *transfer learning*
- **Idea:** Pre-train the network on ImageNet, then fine-tune by retraining on your own data.
    - retrain only final layers
    - retrain end-to-end

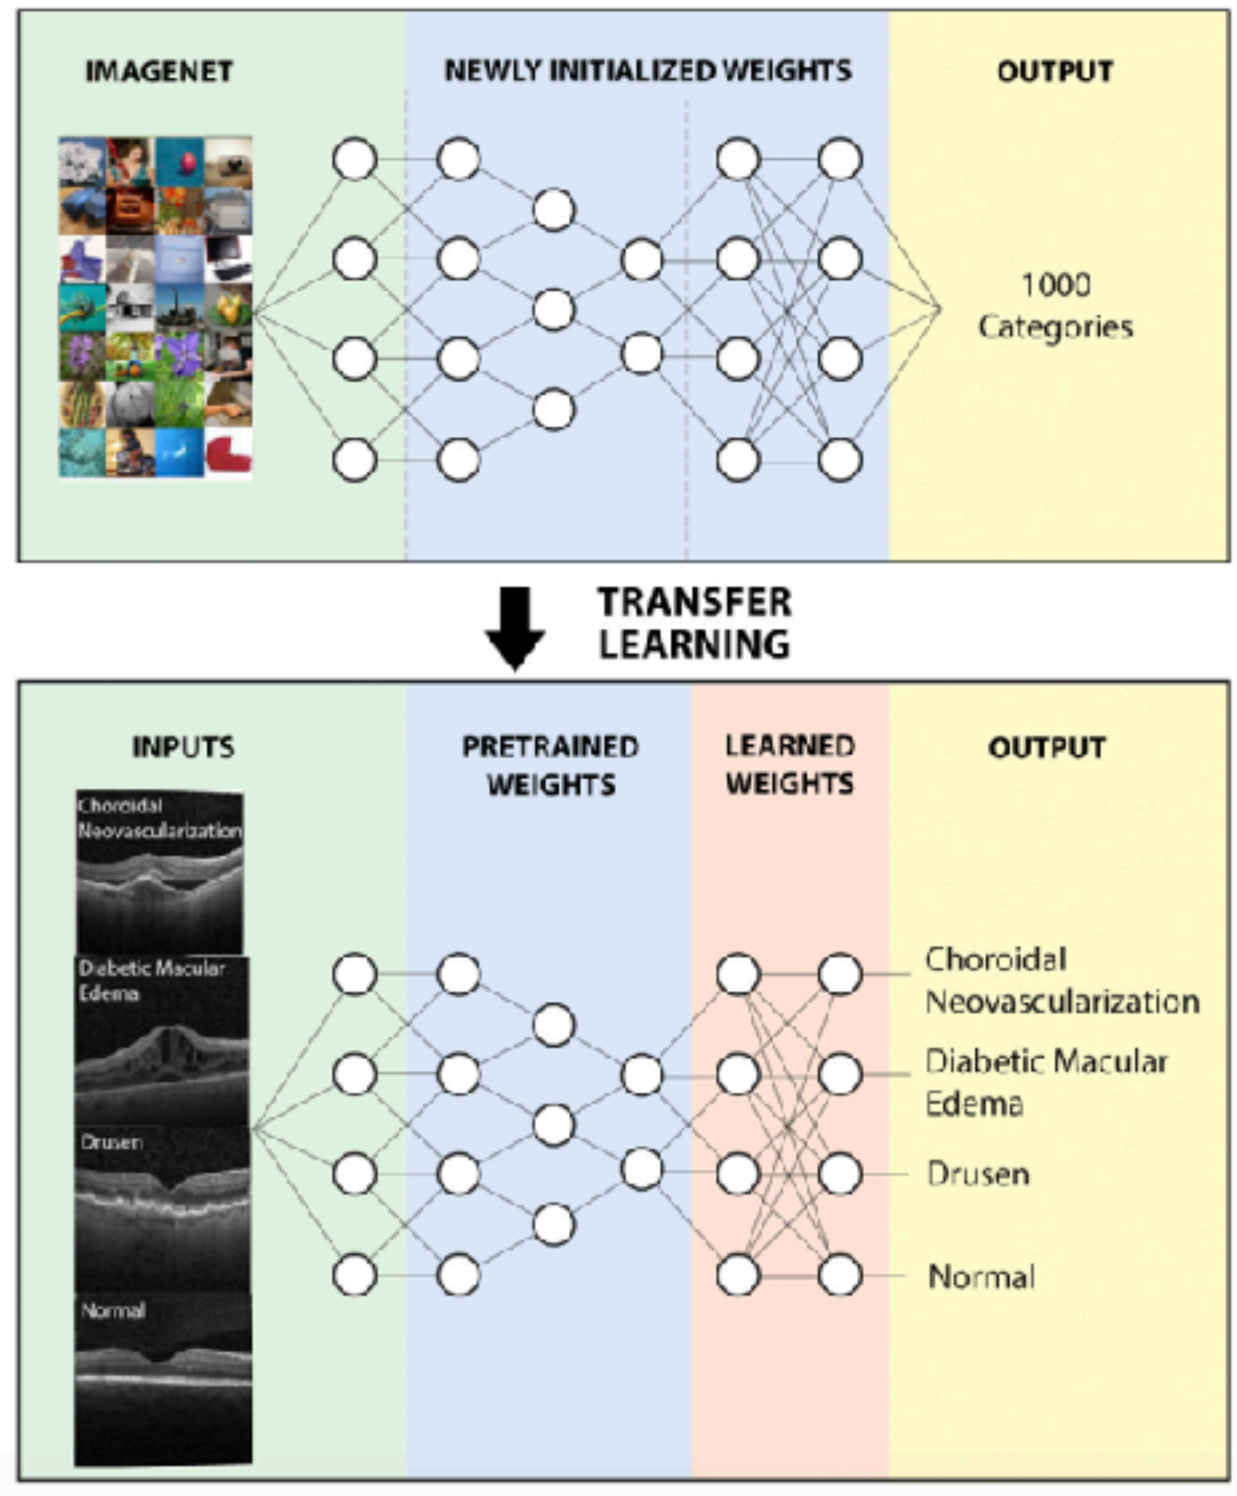

Figure: (Kermany et al., 2018)

Application:
Lung nodule detection in chest CT scans

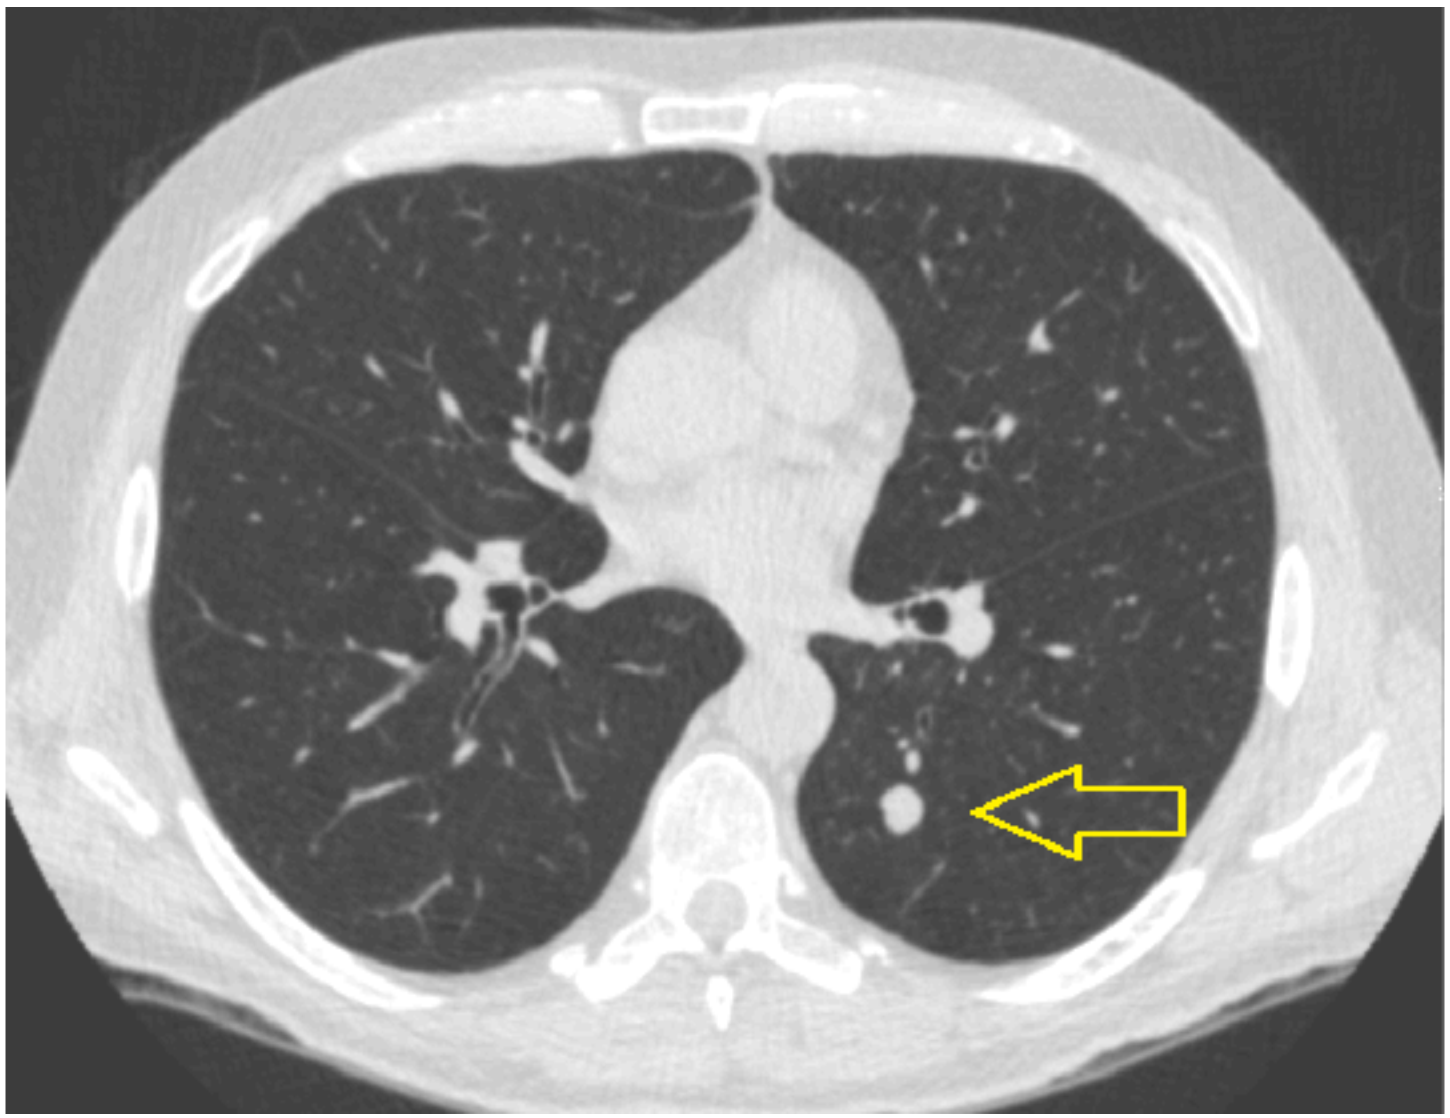

Figure: http://www.diagnijmegen.nl/index.php/Lung_Cancer

- Early stage lung cancers detectable via low-dose CT scans
- Manifest as small pulmonary nodules
- Demanding task for radiologists: ~200-400 axial slices per scan

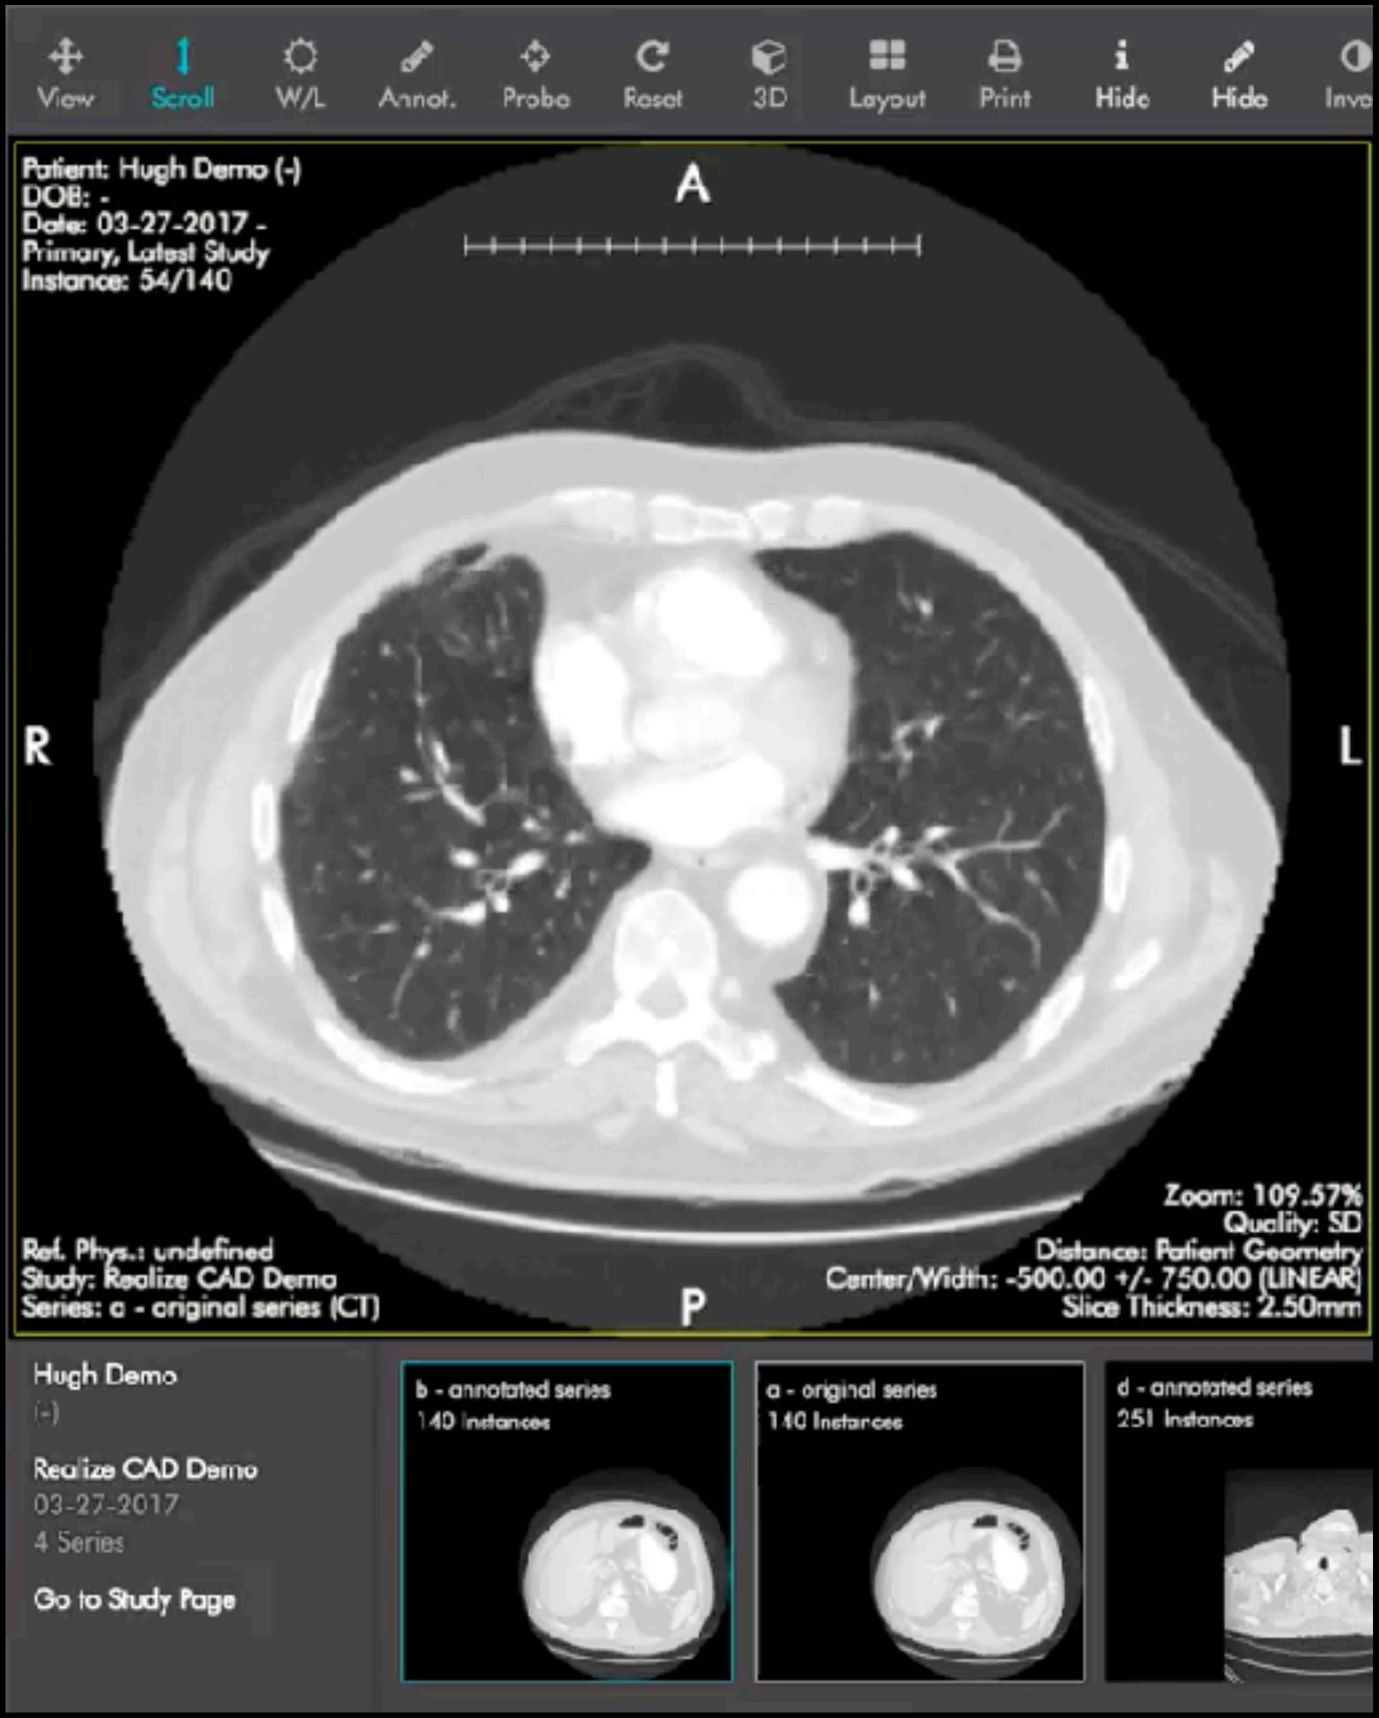

Realize AI: https://ambrahealth.com/directory/realize-ai/
Video: https://www.youtube.com/watch?v=X_8bpuL0G3Q

Potential for CNN's

IEEE TRANSACTIONS ON MEDICAL IMAGING, VOL. 14, NO. 4, DECEMBER 1995 711

Artificial Convolution Neural Network Techniques and Applications for Lung Nodule Detection

Shih-Chung B. Lo, Shyh-Liang A. Lou, *Member, IEEE*, Jyh-Shyan Lin, Matthew T. Freedman, Minze V. Chien, *Member, IEEE*, and Seong K. Mun, *Member, IEEE*

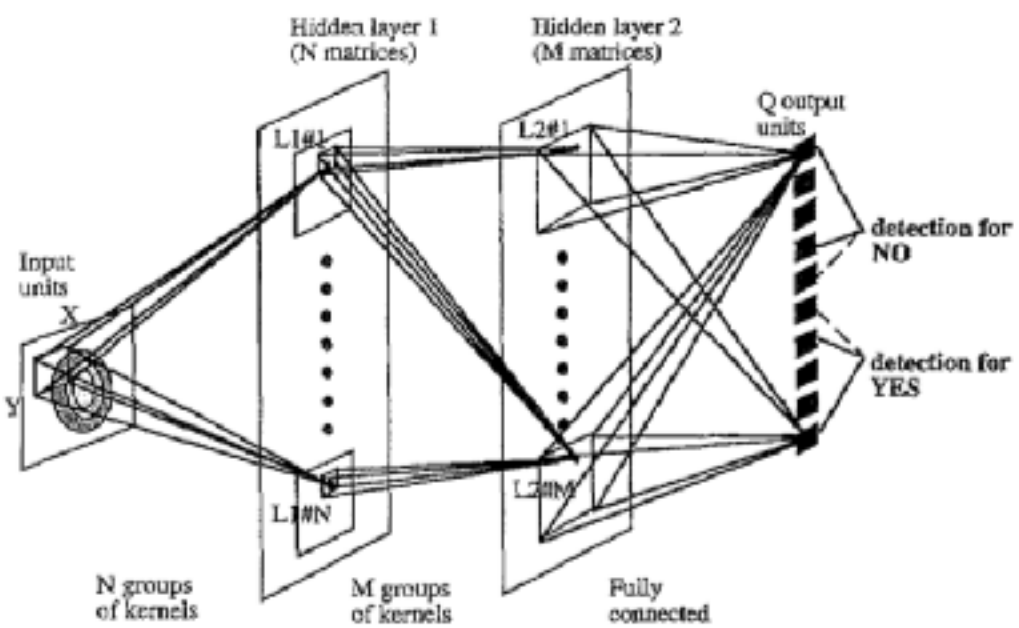

Fig. 3. Artificial convolution neural network for detection of lung nodule.

Modern approaches

- Multi-scale approach [Shen et al., 2015]
- Trains 3 CNN's simultaneously on patches at different scales

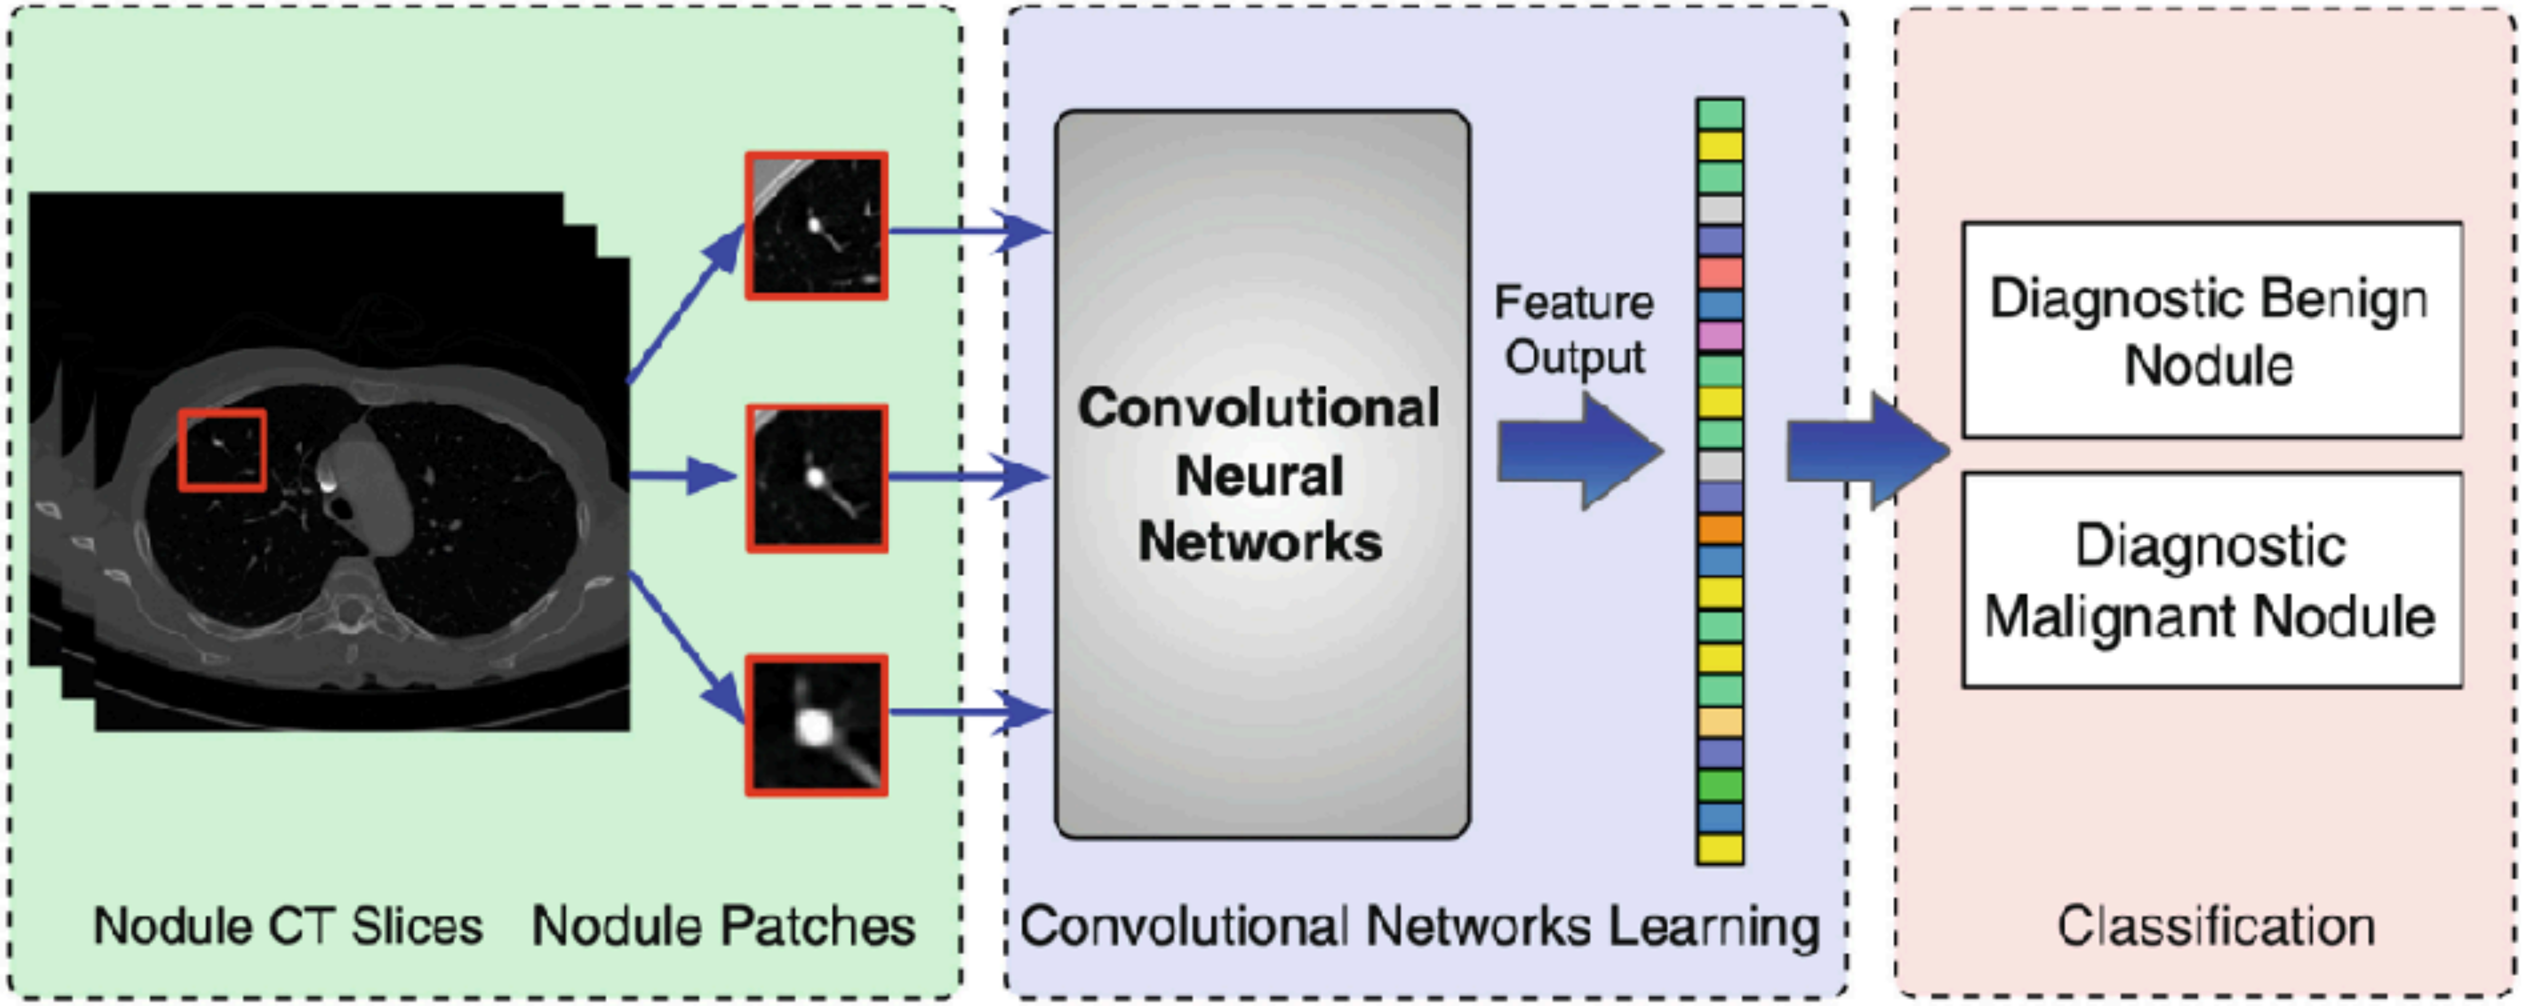

- Uses domain specific knowledge:
  nodule sizes vary from < 3 mm to >30 mm

Modern approaches

- Multi-view approach [Setio et al., 2016]
- Trains 9 CNN’s simultaneously for 9 different views of nodule

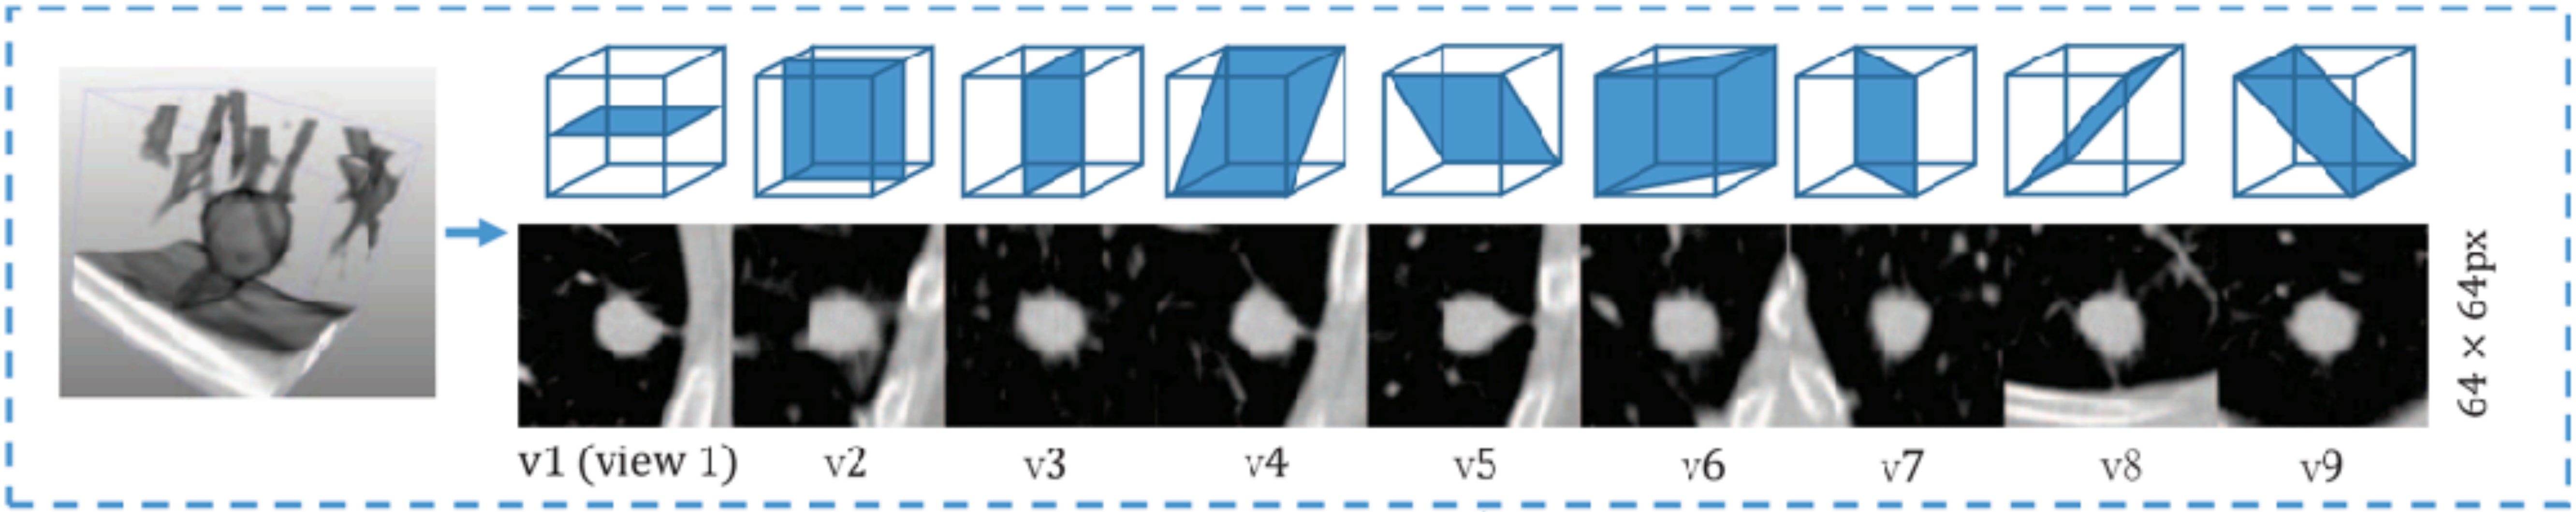

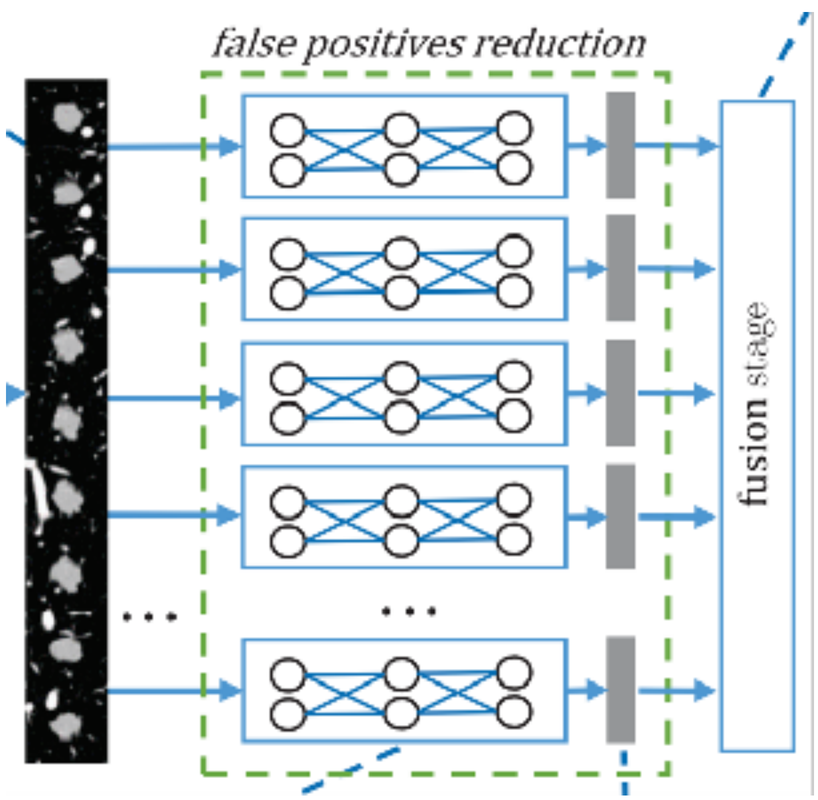

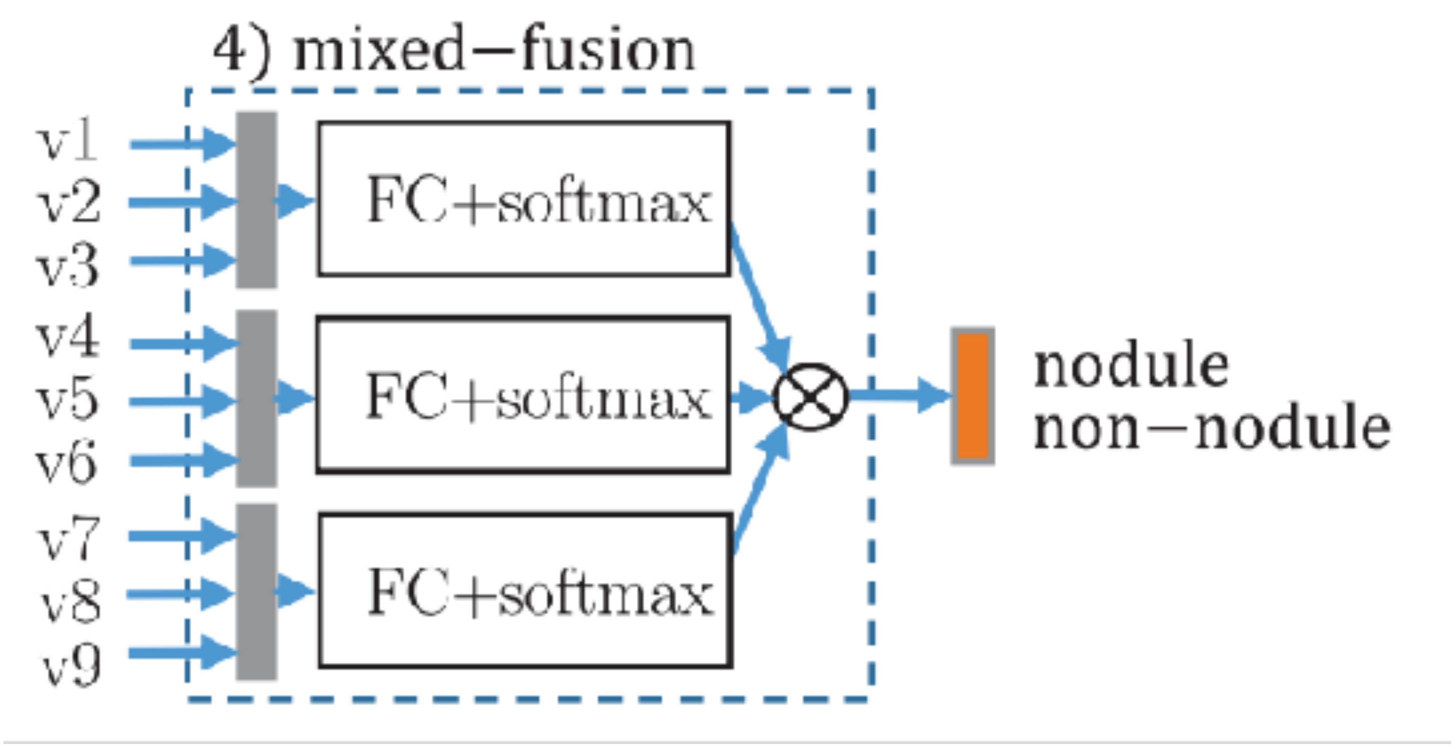

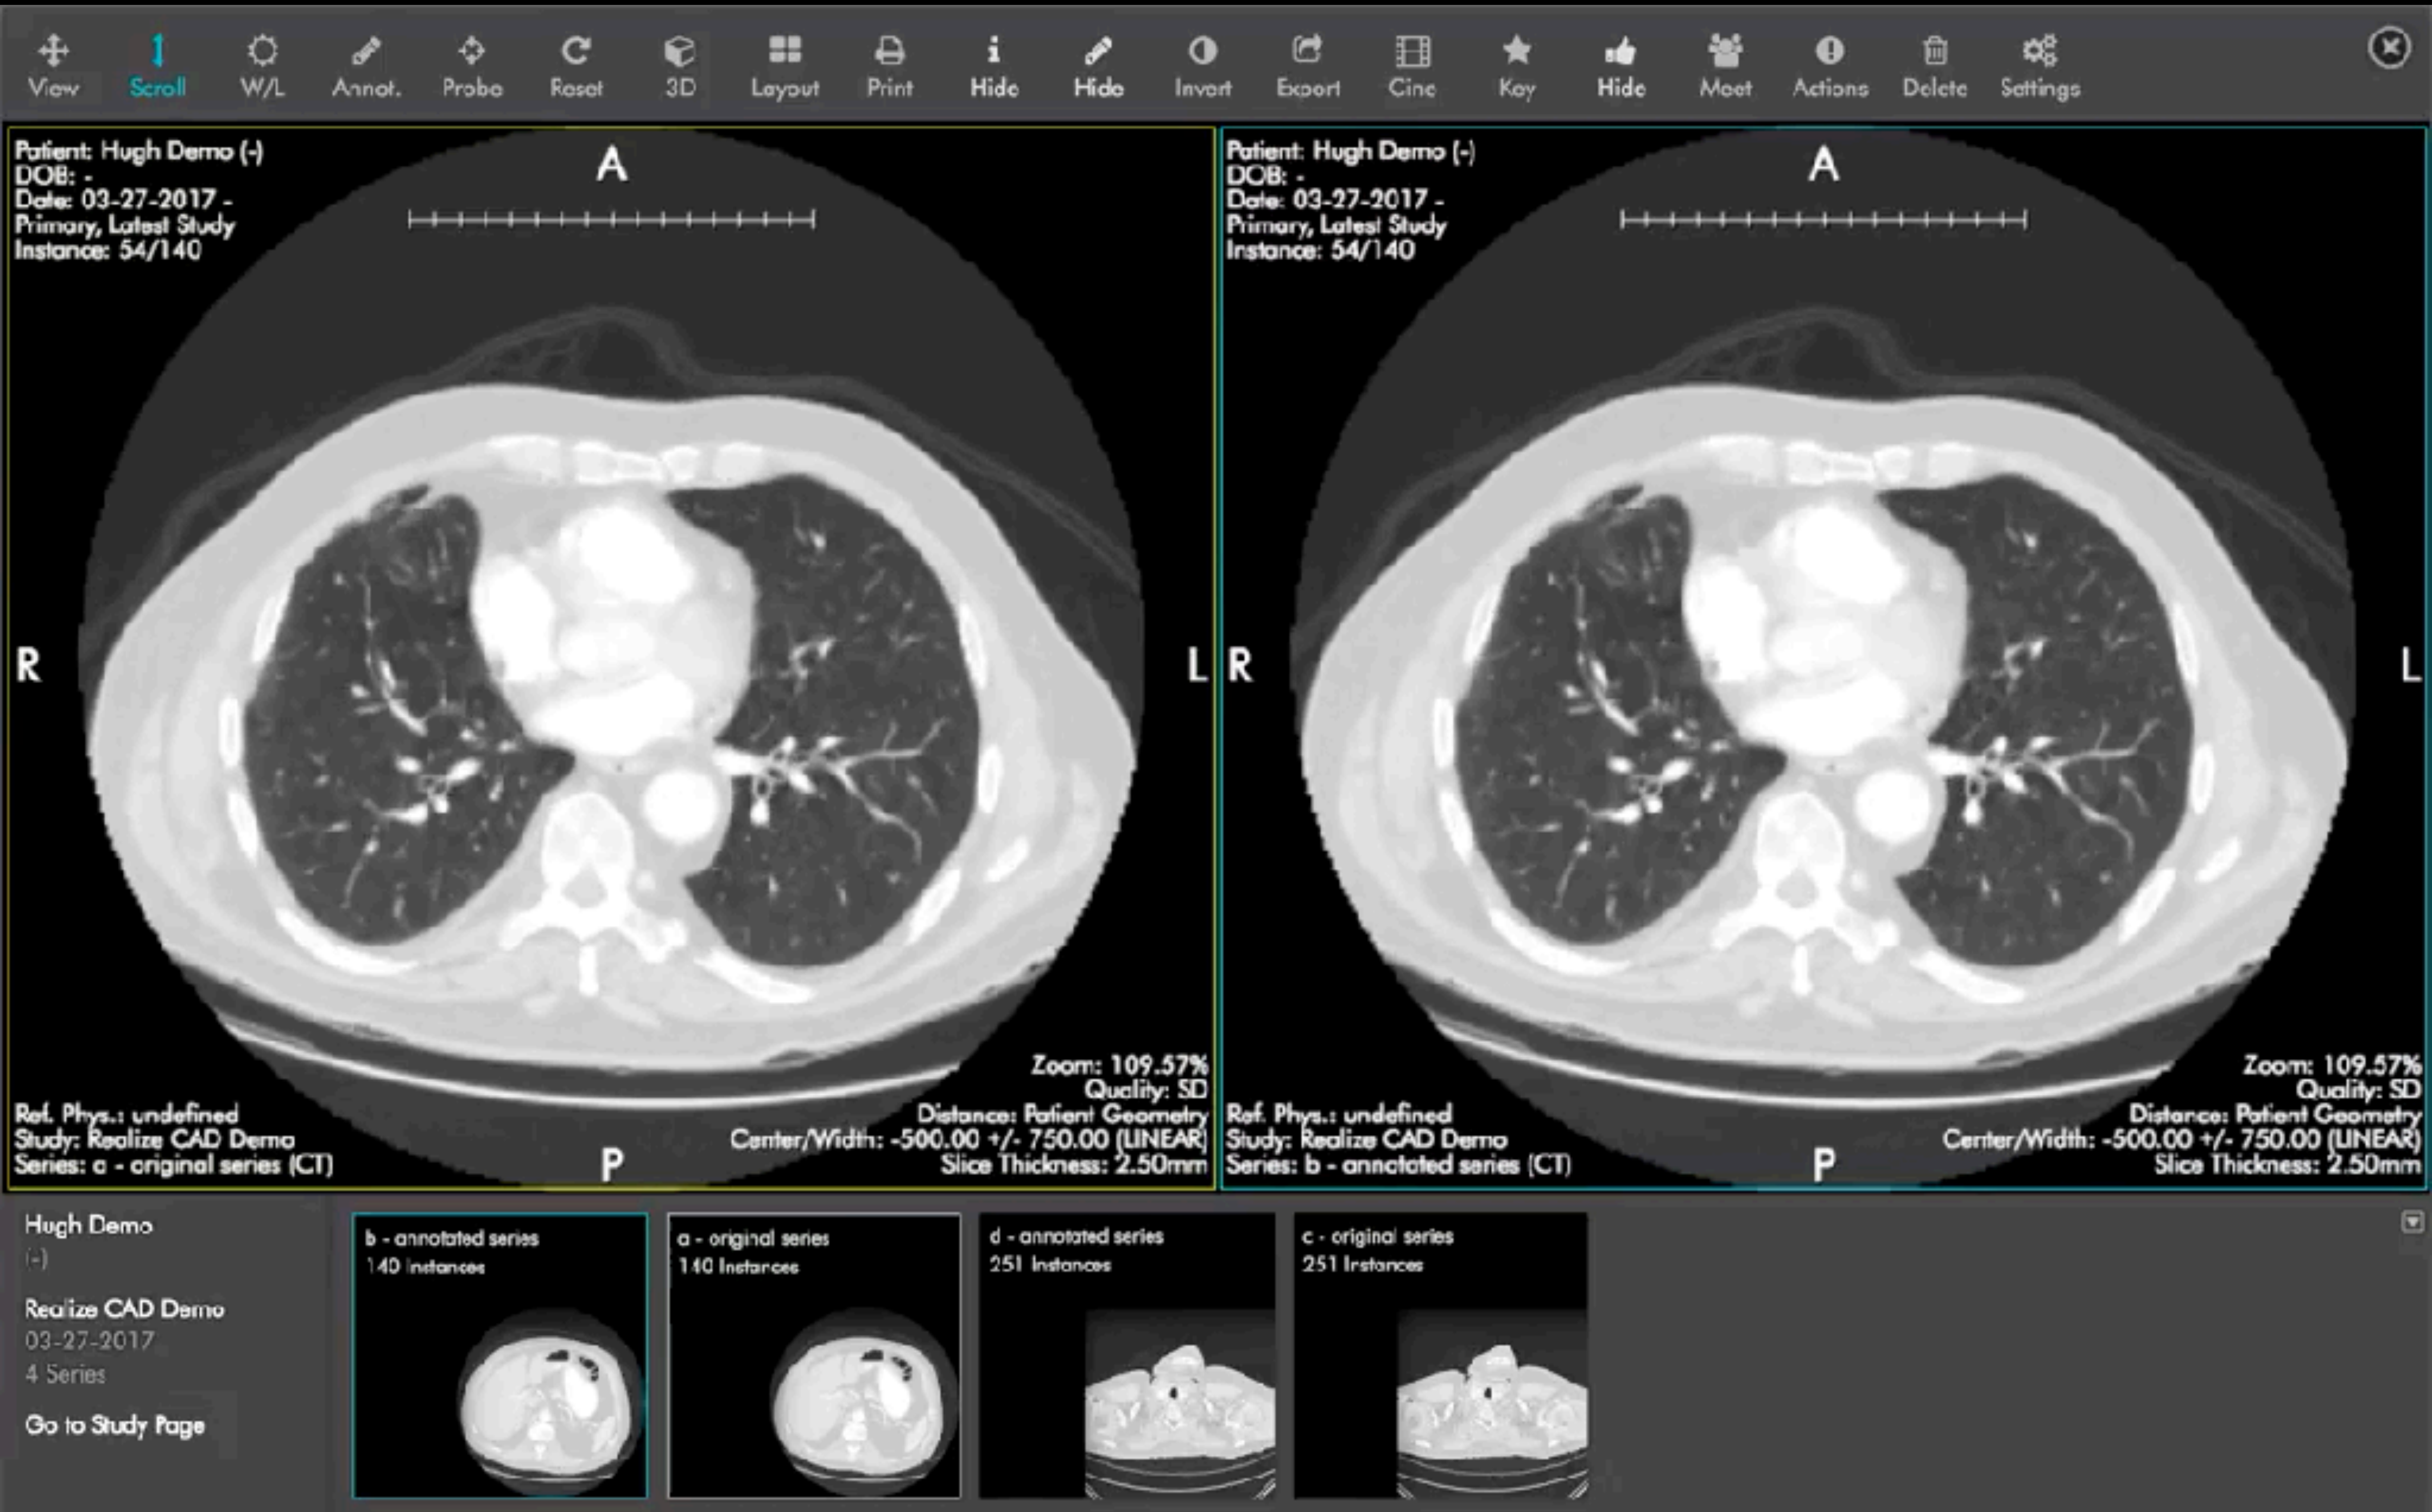

Realize AI: https://ambrahealth.com/directory/realize-ai/
Video: https://www.youtube.com/watch?v=X_8bpuL0G3Q

Takeaway

Main innovations in classifying biomedical images with CNNs are in "meta-architectures" that make use of domain specific knowledge

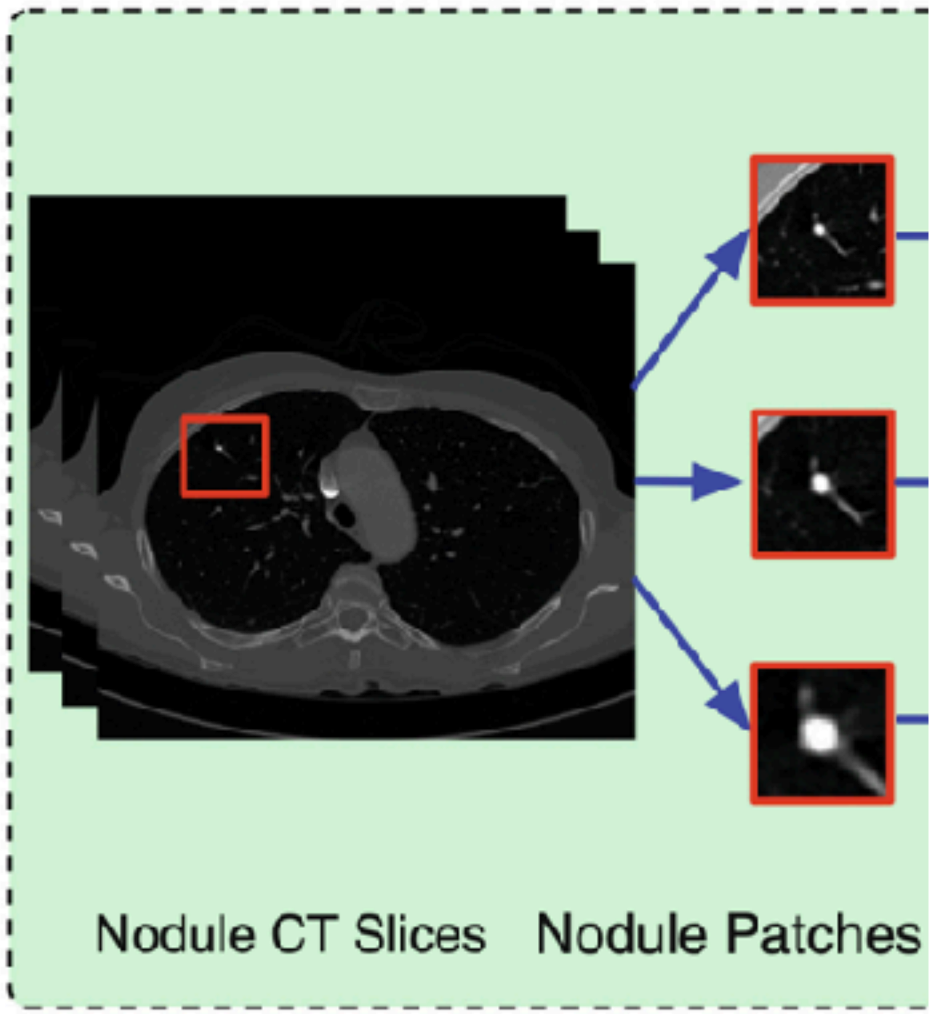

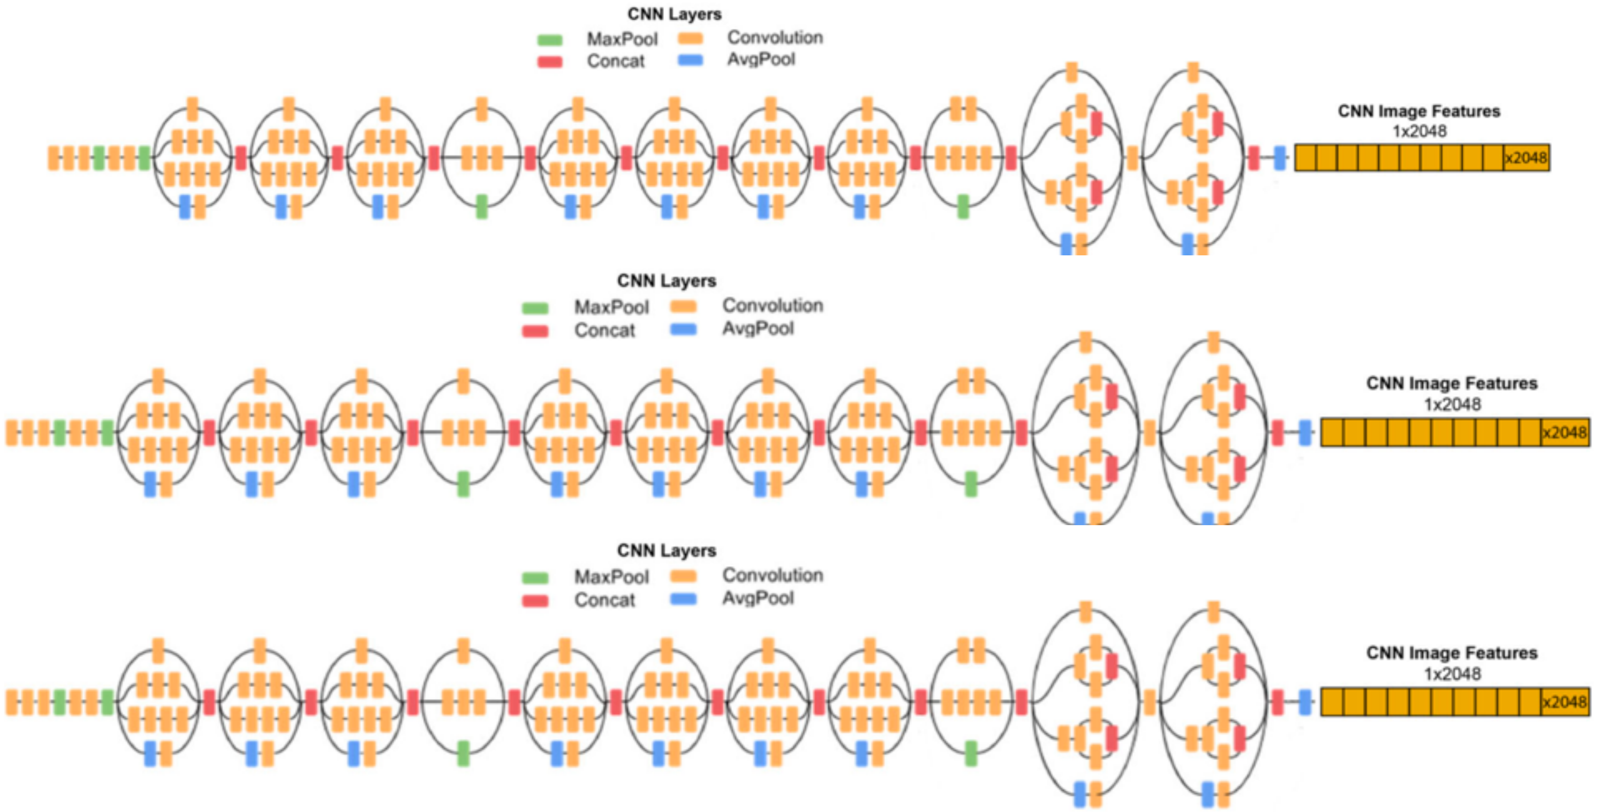

Biomedical Image Segmentation with the U-net

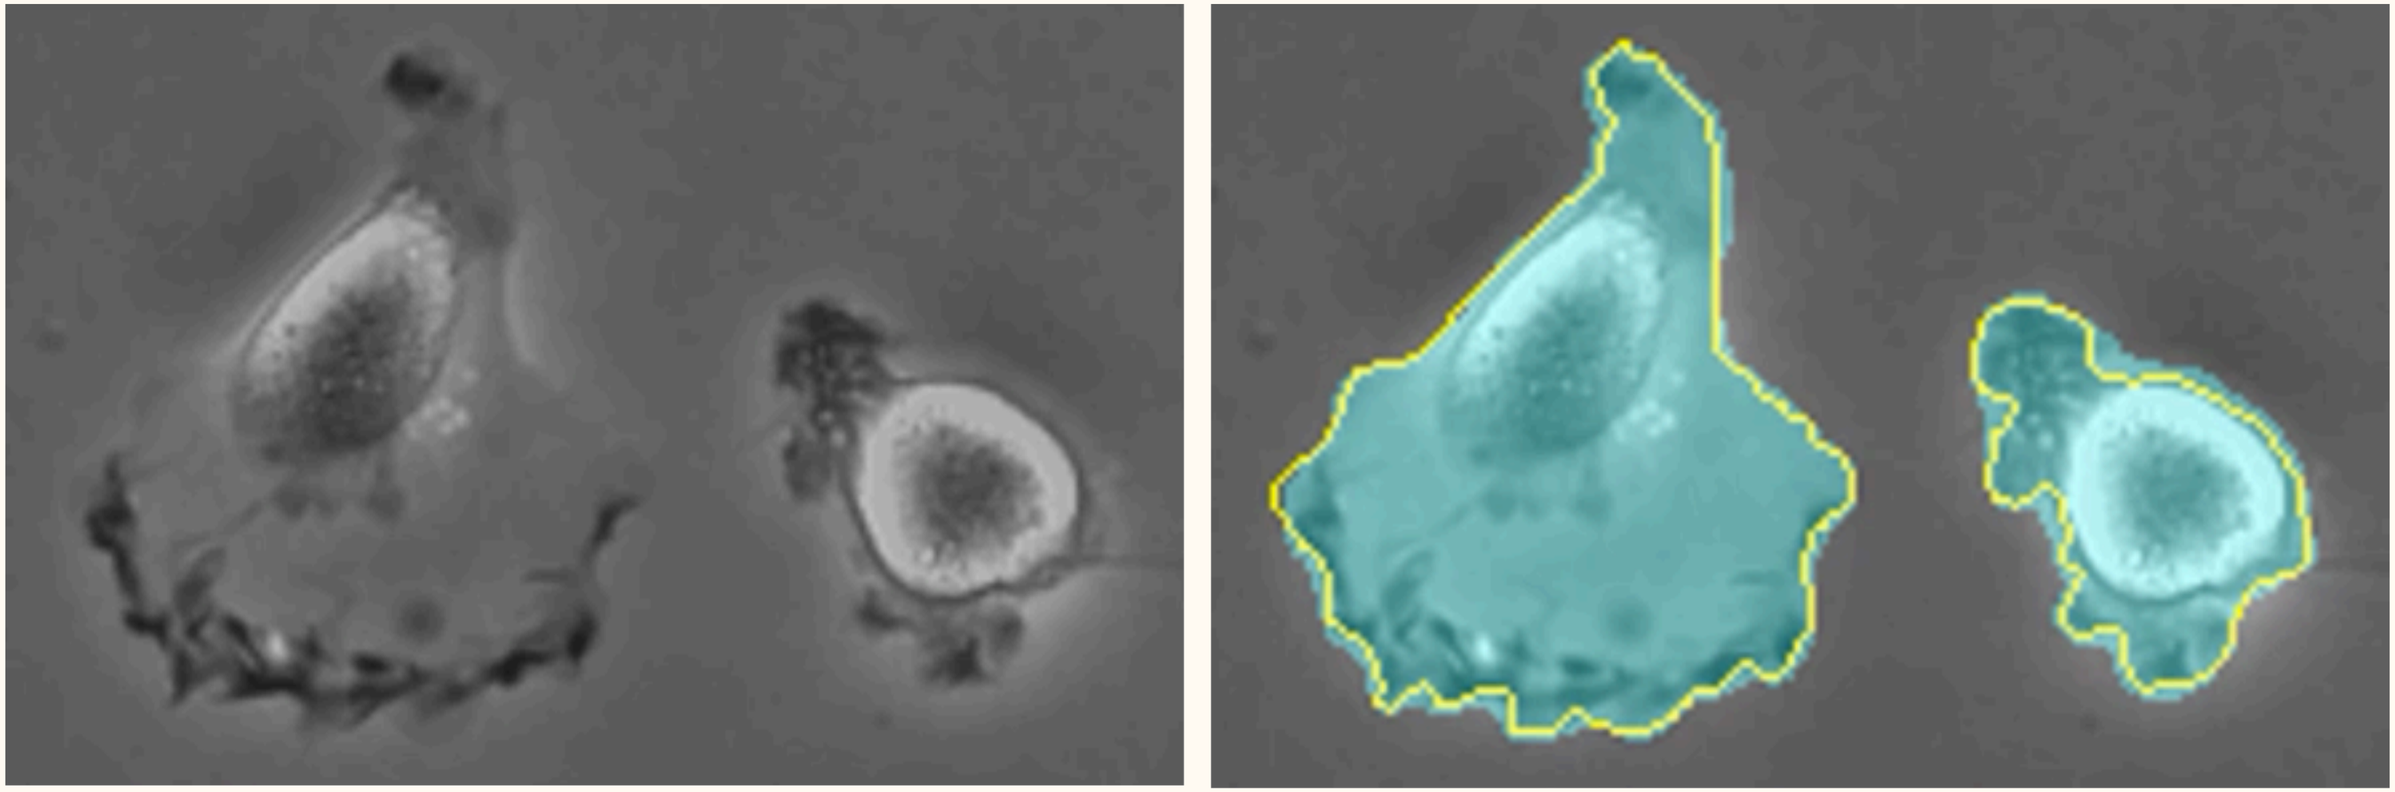

What is segmentation?

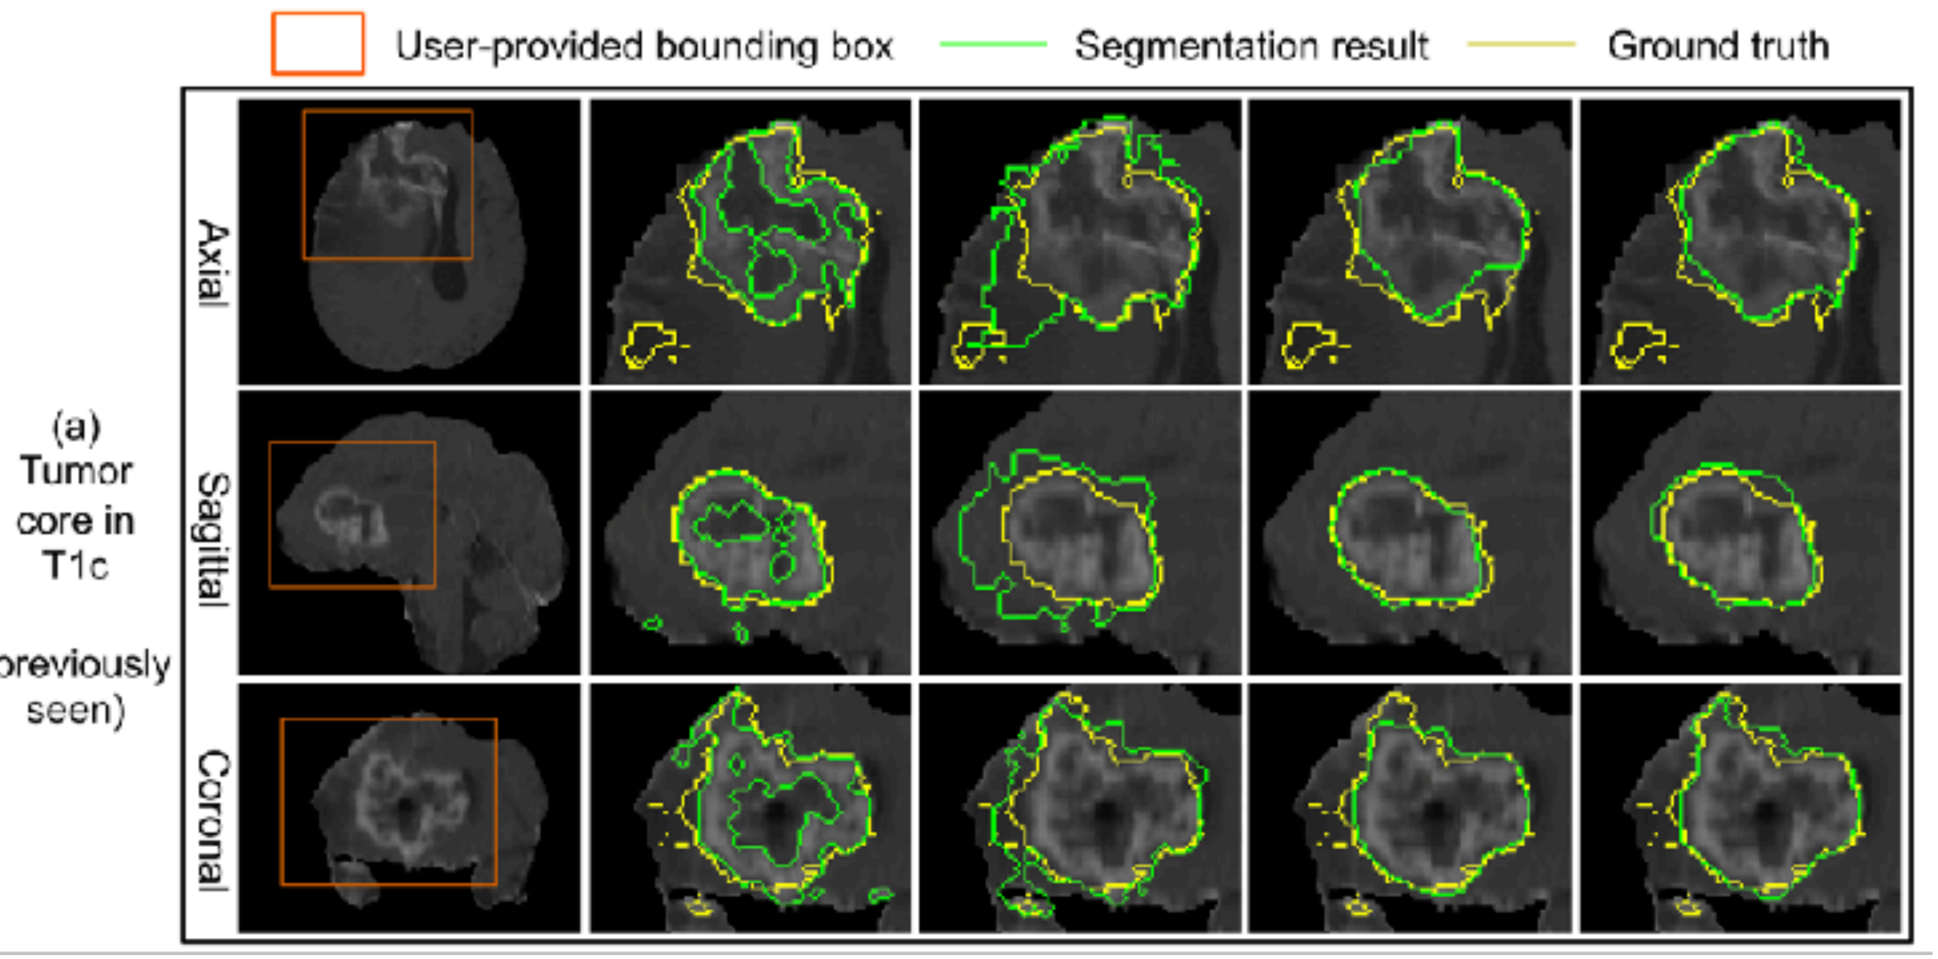

Example: Tumor segmentation in MRI brain scans

Goal: partition image into multiple regions that share attributes for *localization* or *quantification*

Application: Segmentation of neuronal structures in electron microscope stacks

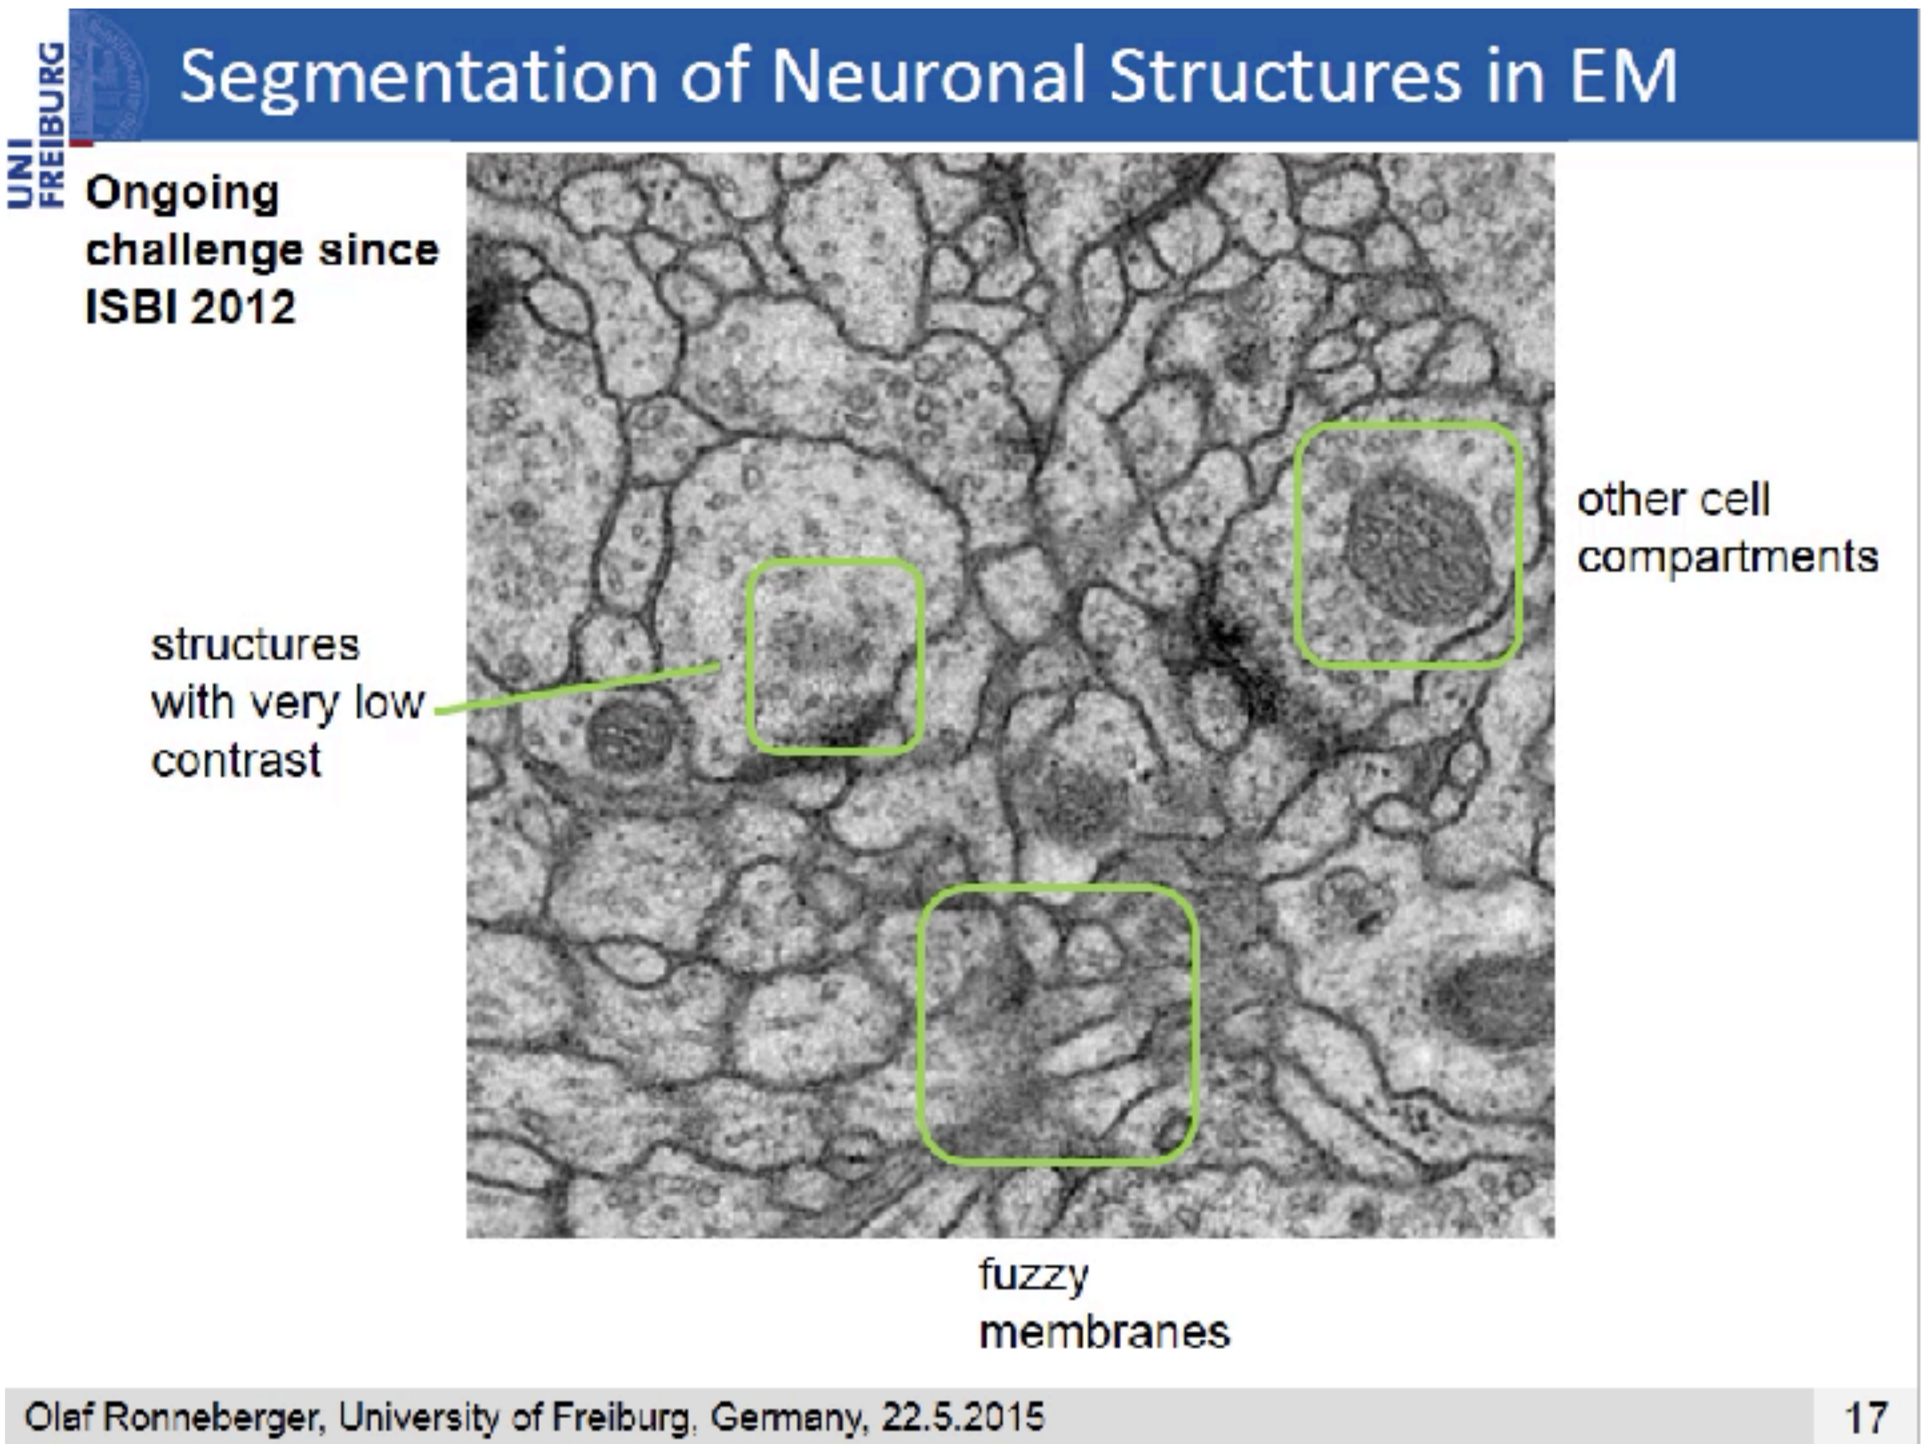

https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/

Segmentation = pixel-wise classification?

- Classify pixel-wise with deep CNN classifier
- Use a “sliding window” approach

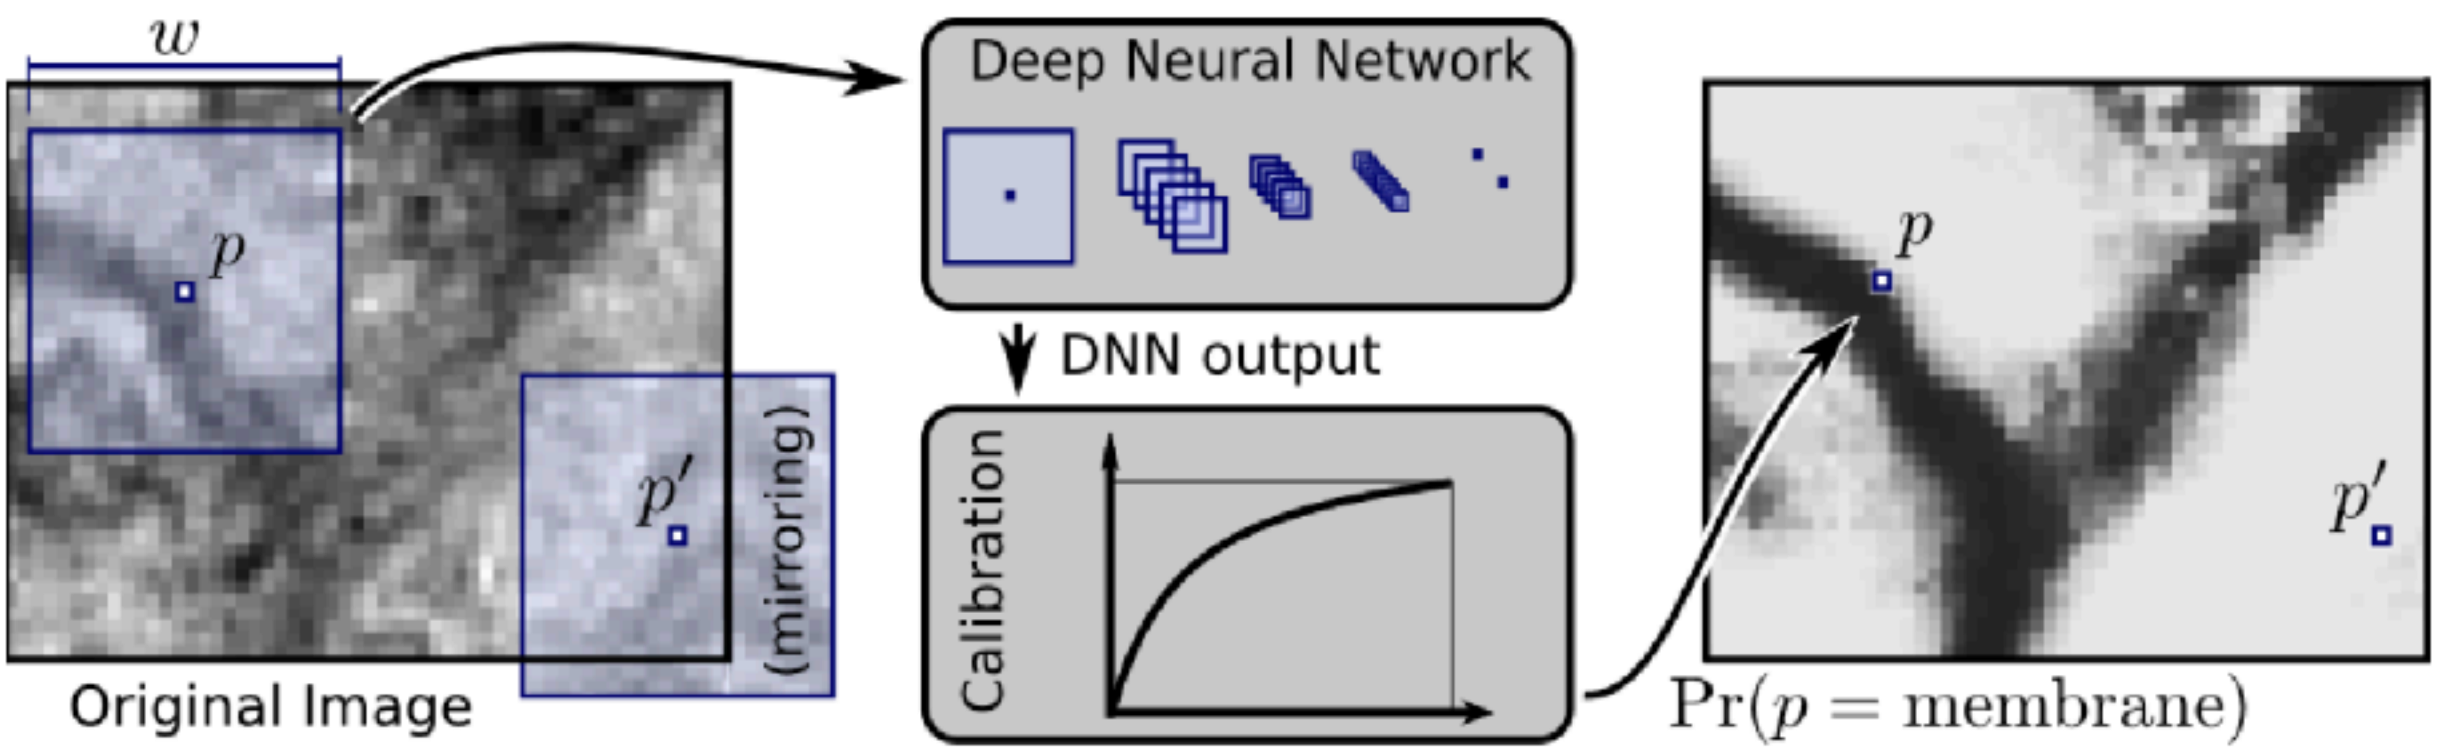

[Ciresan et al., 2012]

- Drawbacks: Inefficient to scale to large images
- Only uses local information

Fully convolutional neural networks (fCNN)

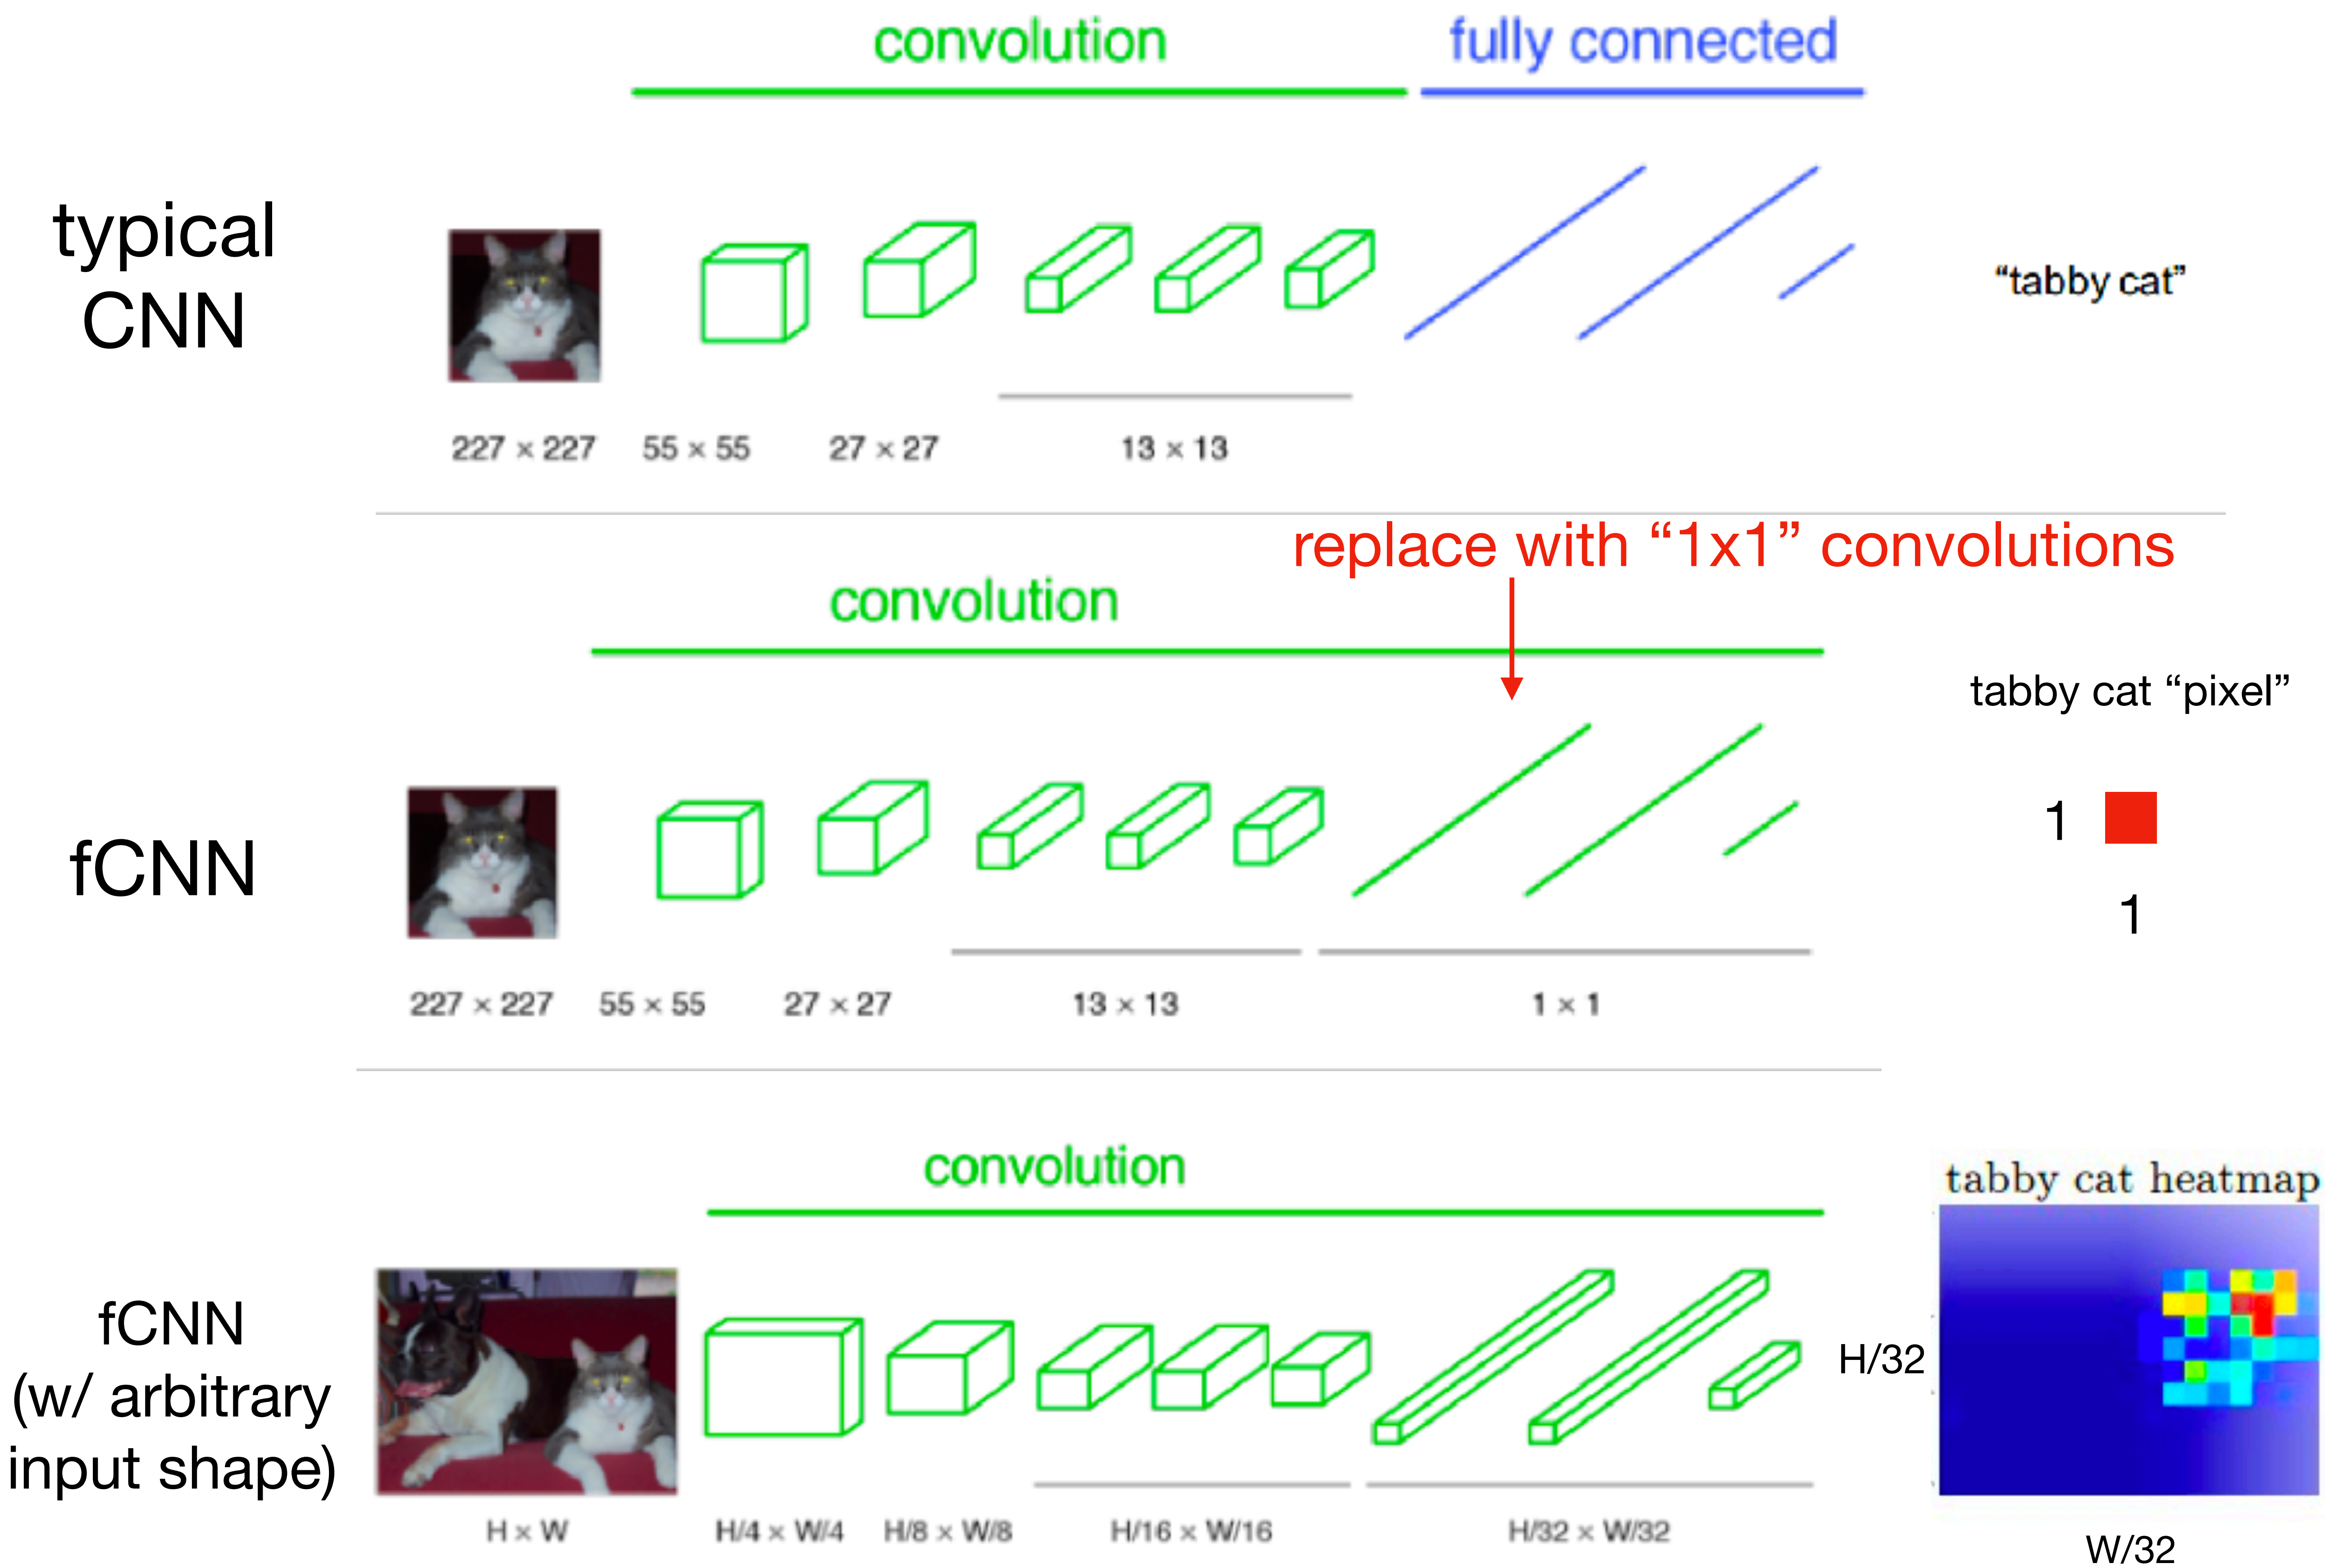

Figures: https://leonardoaraujosantos.gitbooks.io/artificial-inteligence/content/image_segmentation.html

fCNN Segmentation Network [Long et al., 2014]

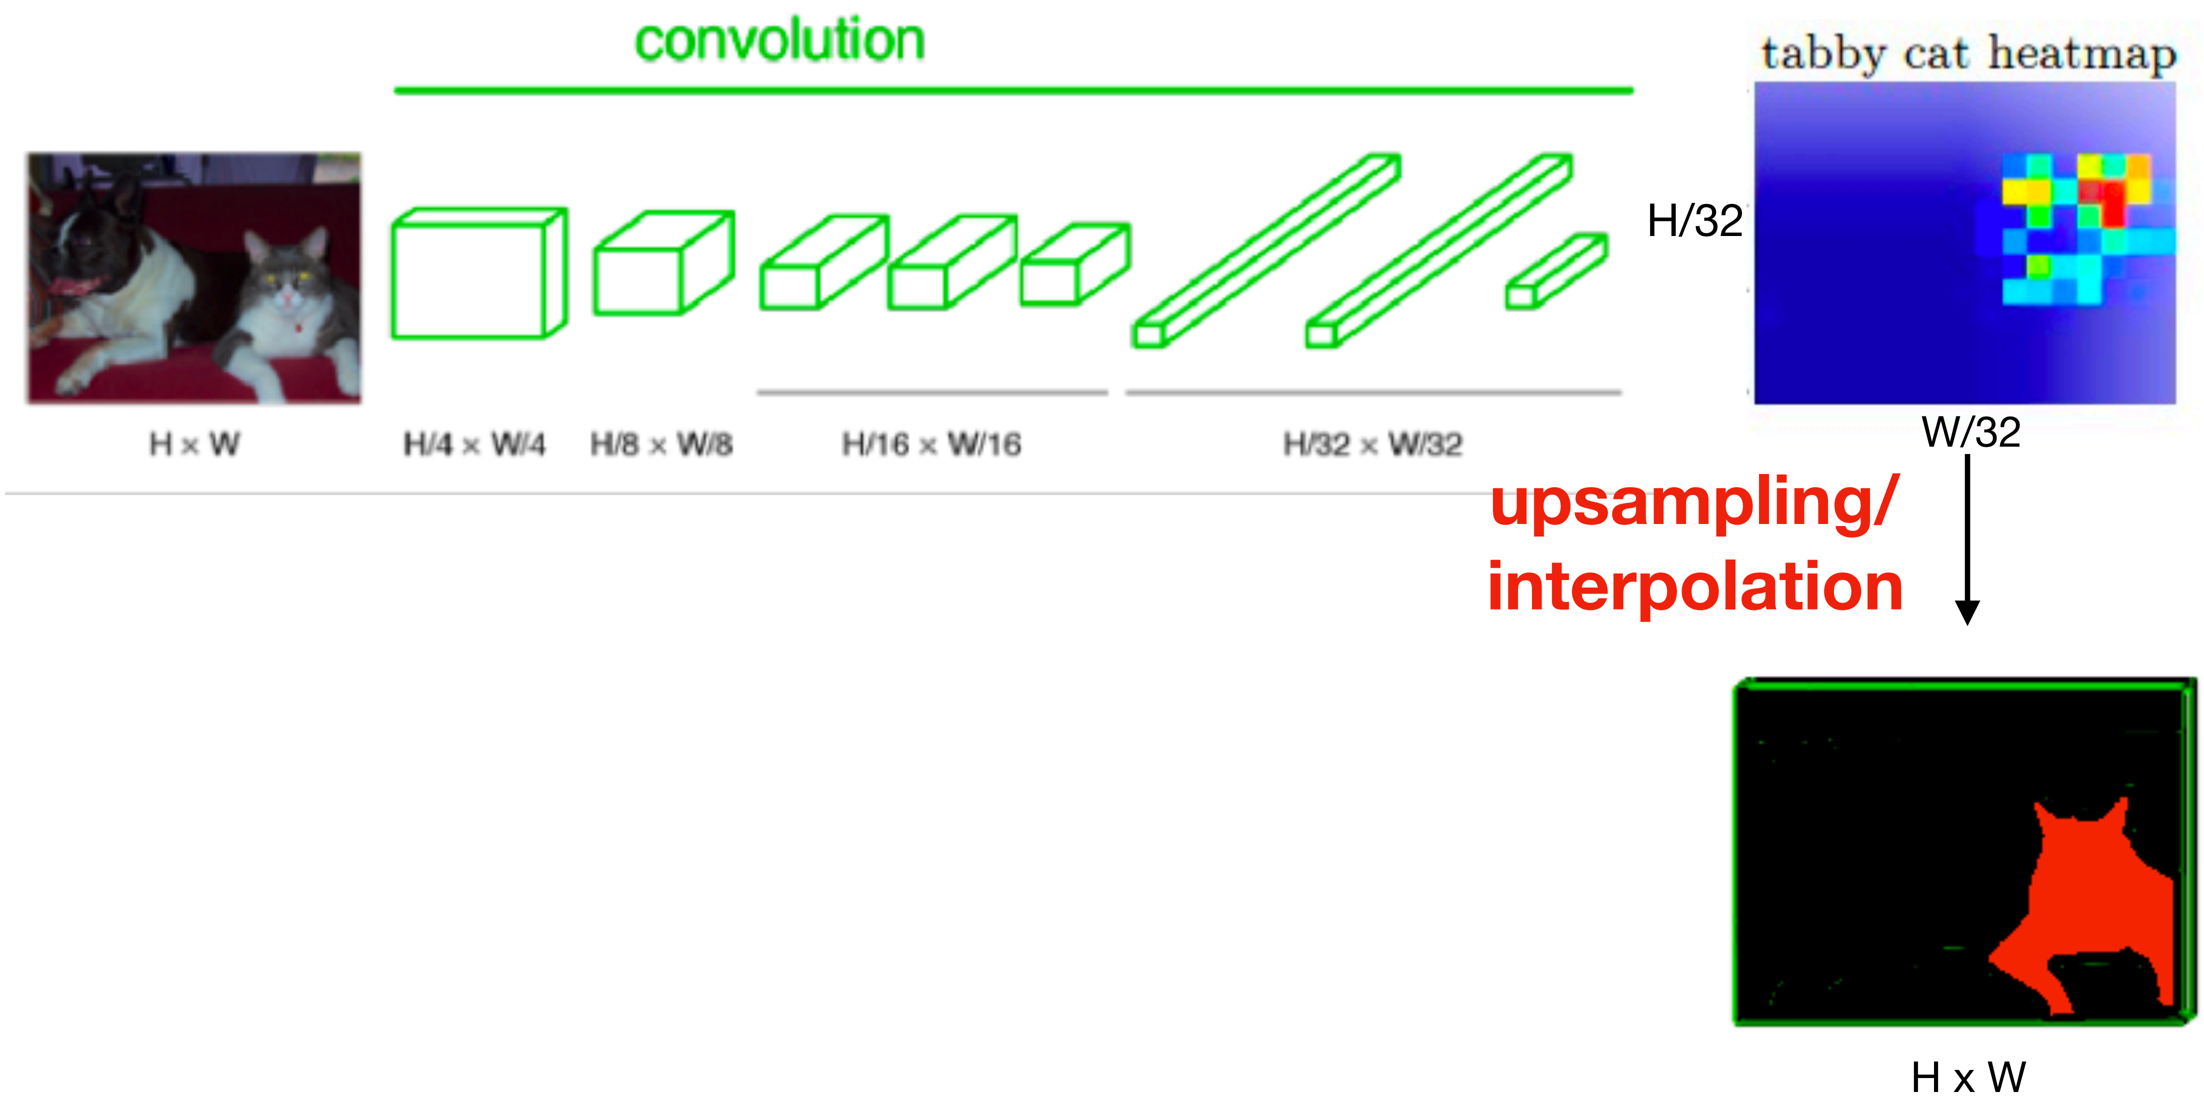

pixelwise output+loss

Figures: https://leonardoaraujosantos.gitbooks.io/artificial-inteligence/content/image_segmentation.html

Skip connections in fCNNs

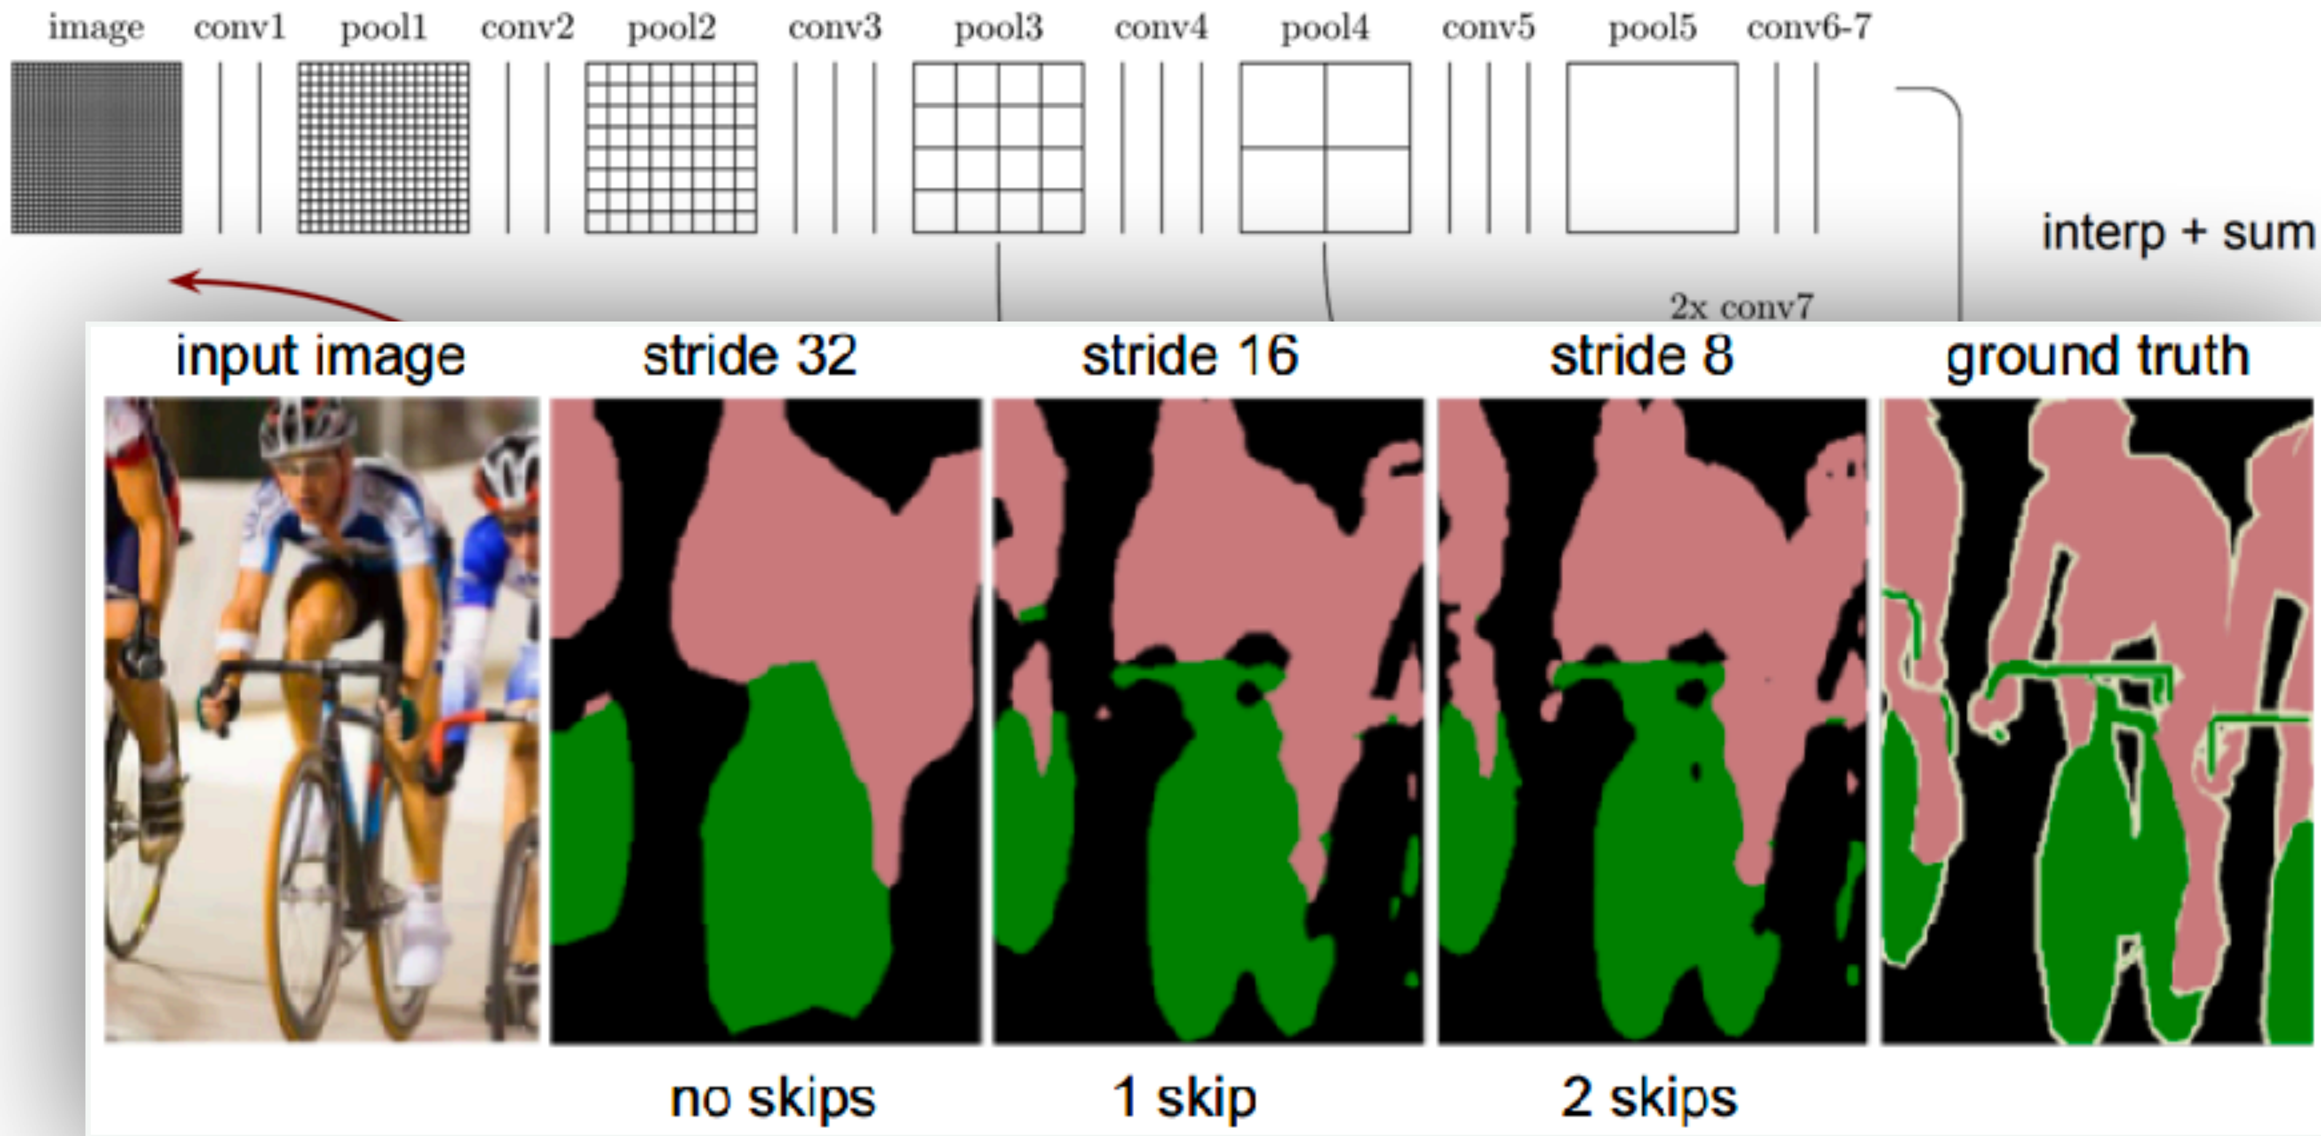

Figure: [Long et al., 2014]

U-net architecture [Ronneberger et al. 2015]

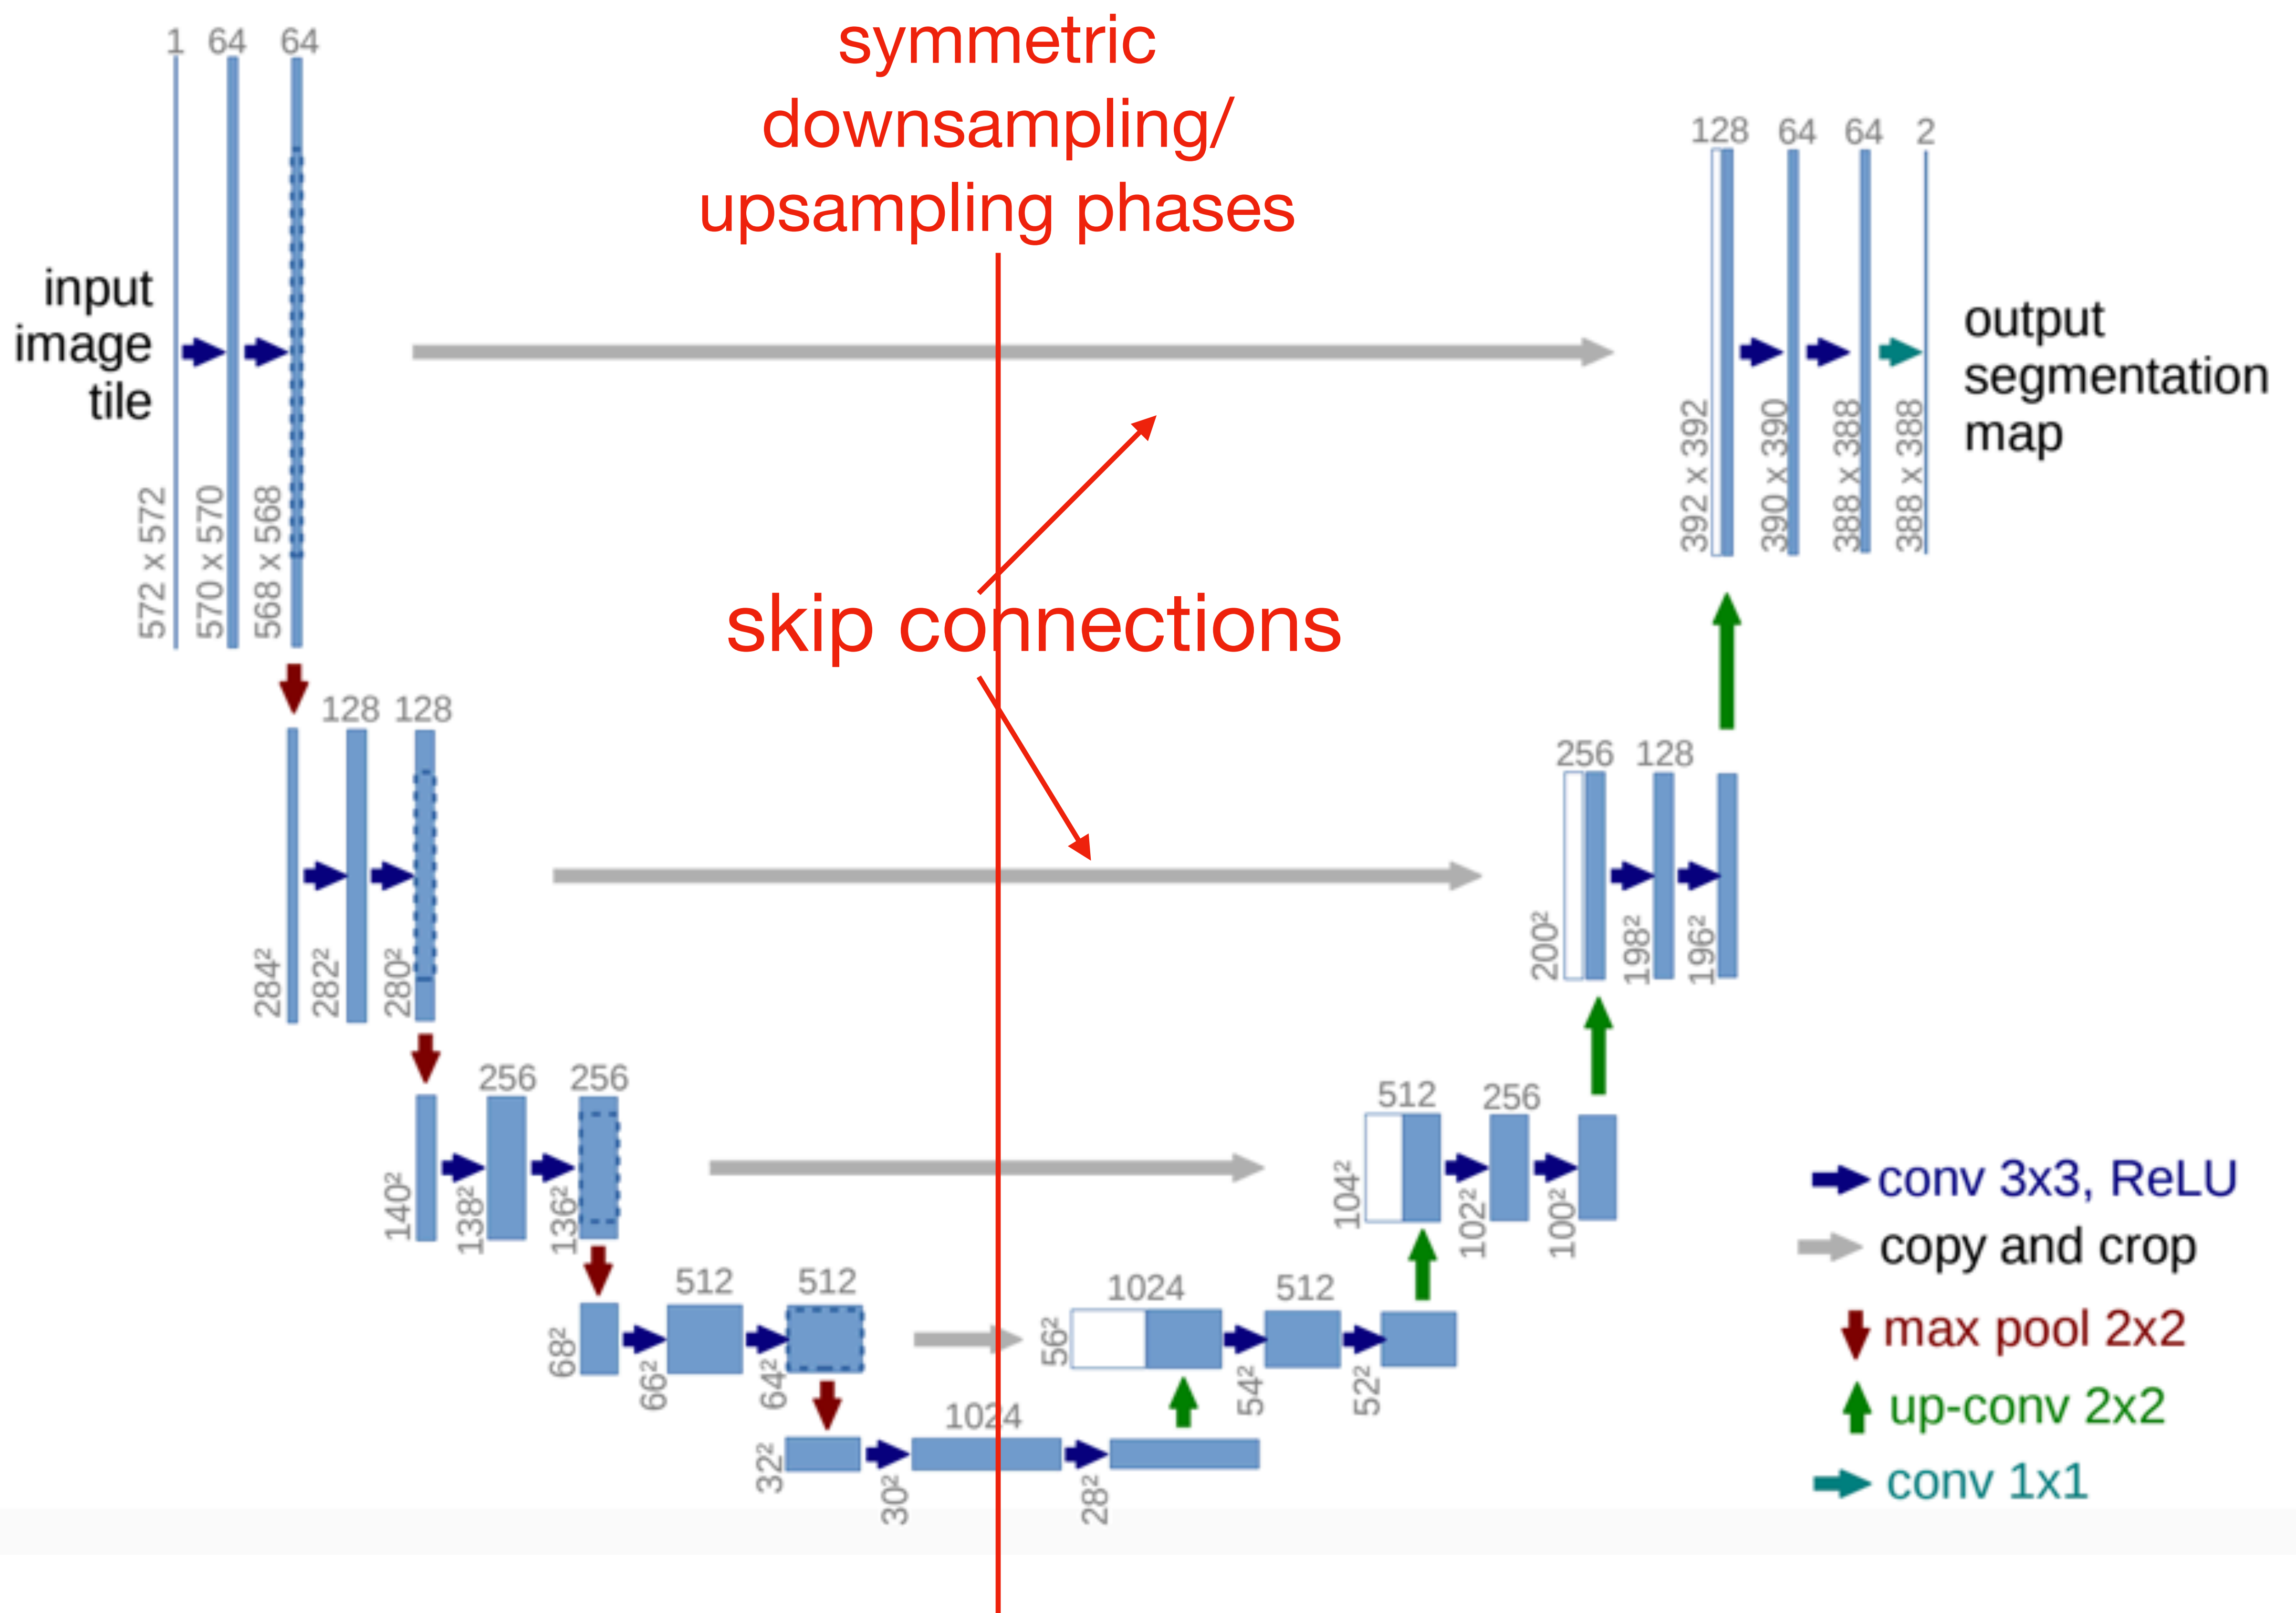

Figure: [Ronneberger et al. 2015]

Application: Segmentation of neuronal structures in electron microscope stacks

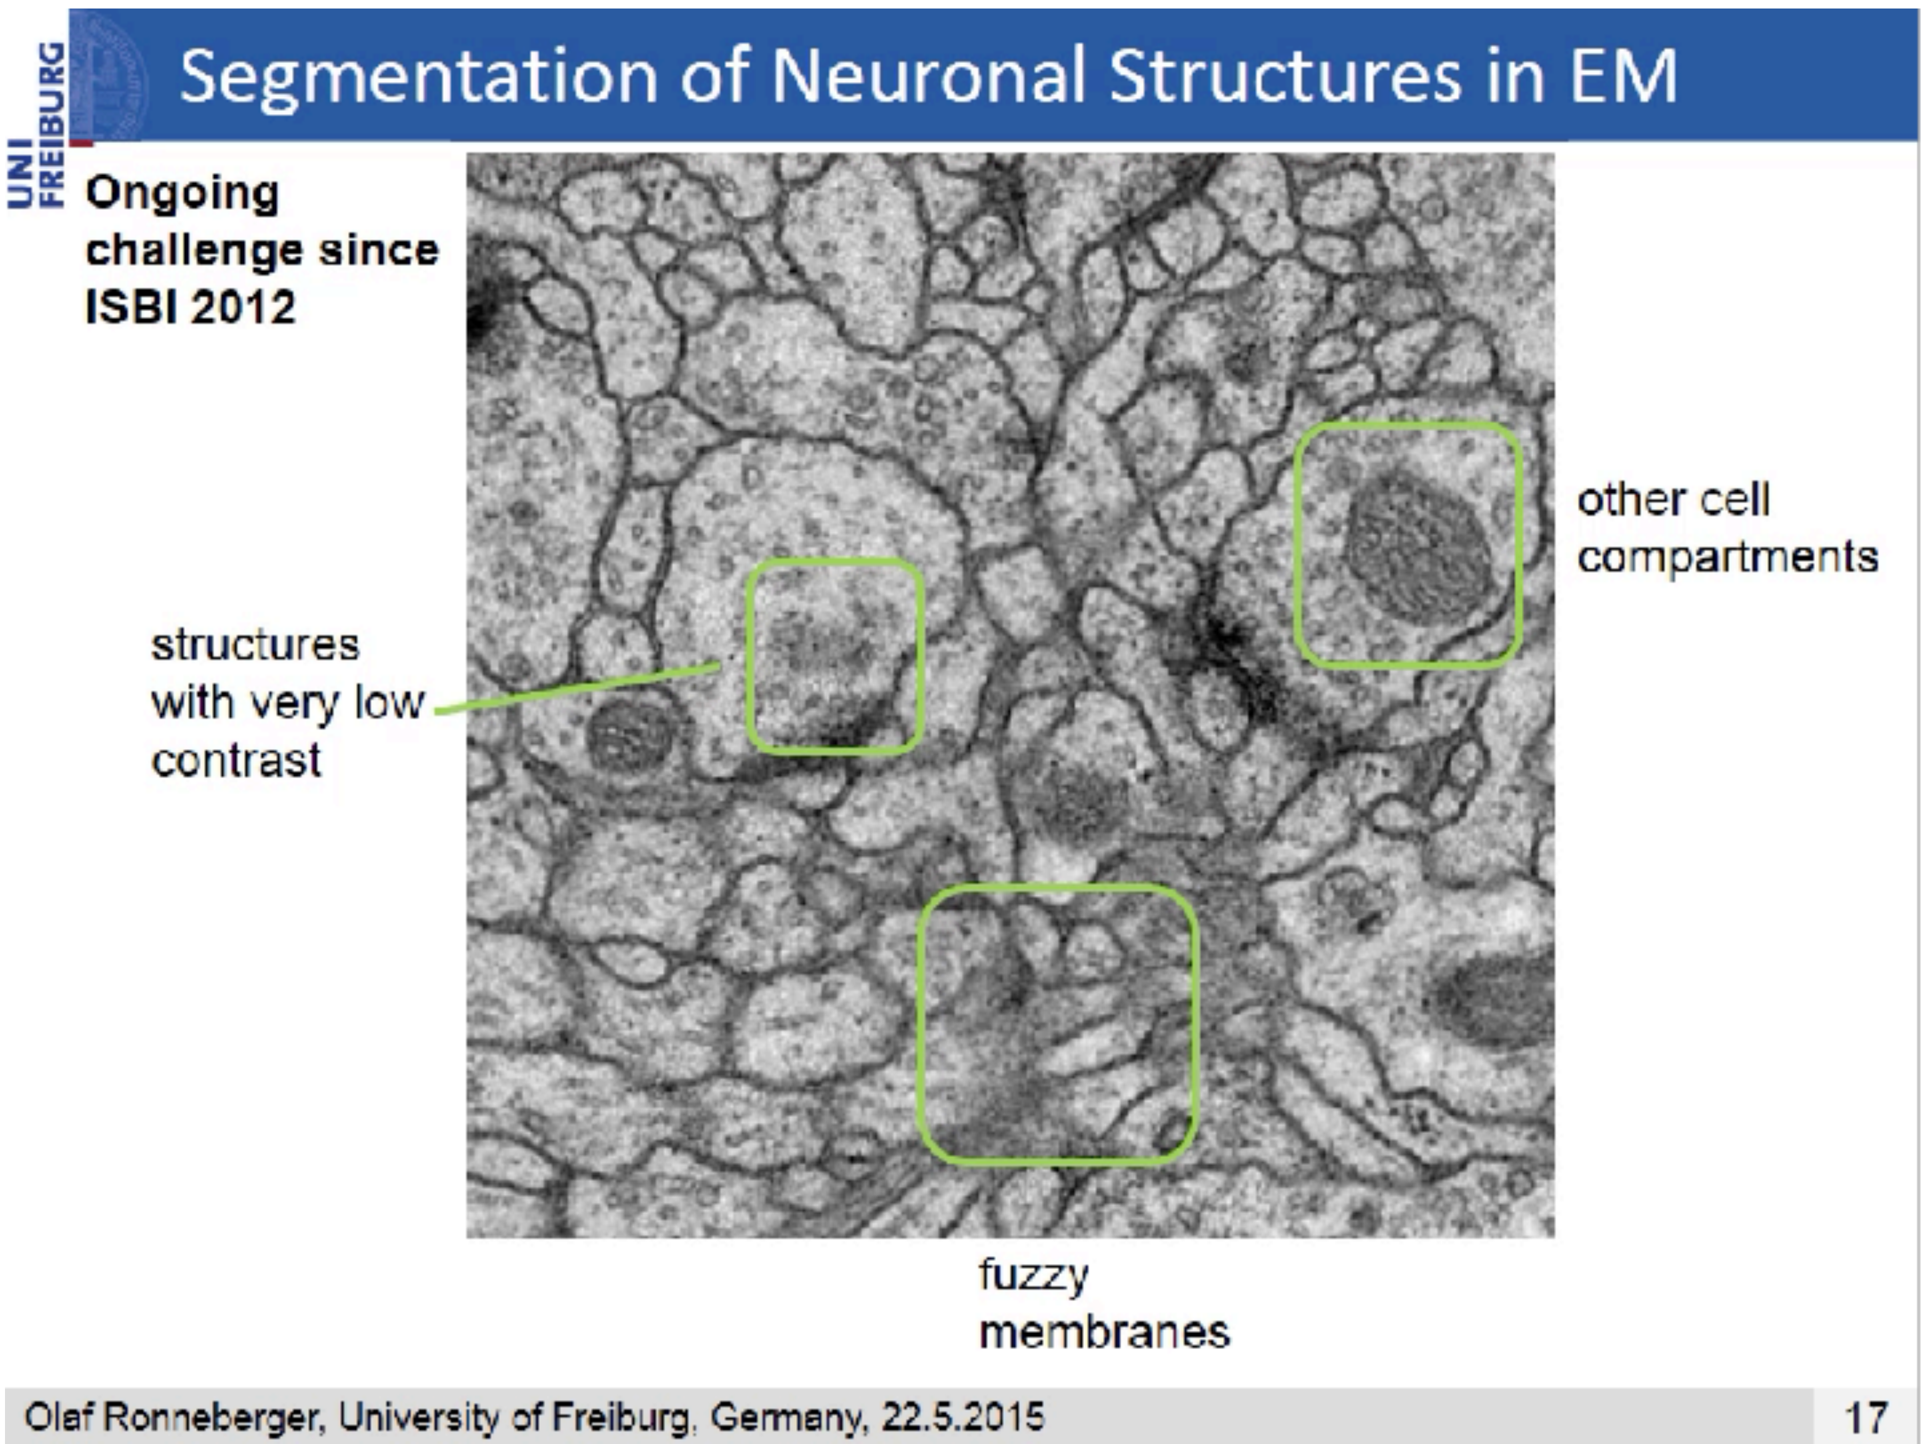

https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/

Issue: Very Little Training Data

- Hand-labelled segmentations difficult to obtain
  - e.g., ISBI 2012 challenge has only 30 training images!

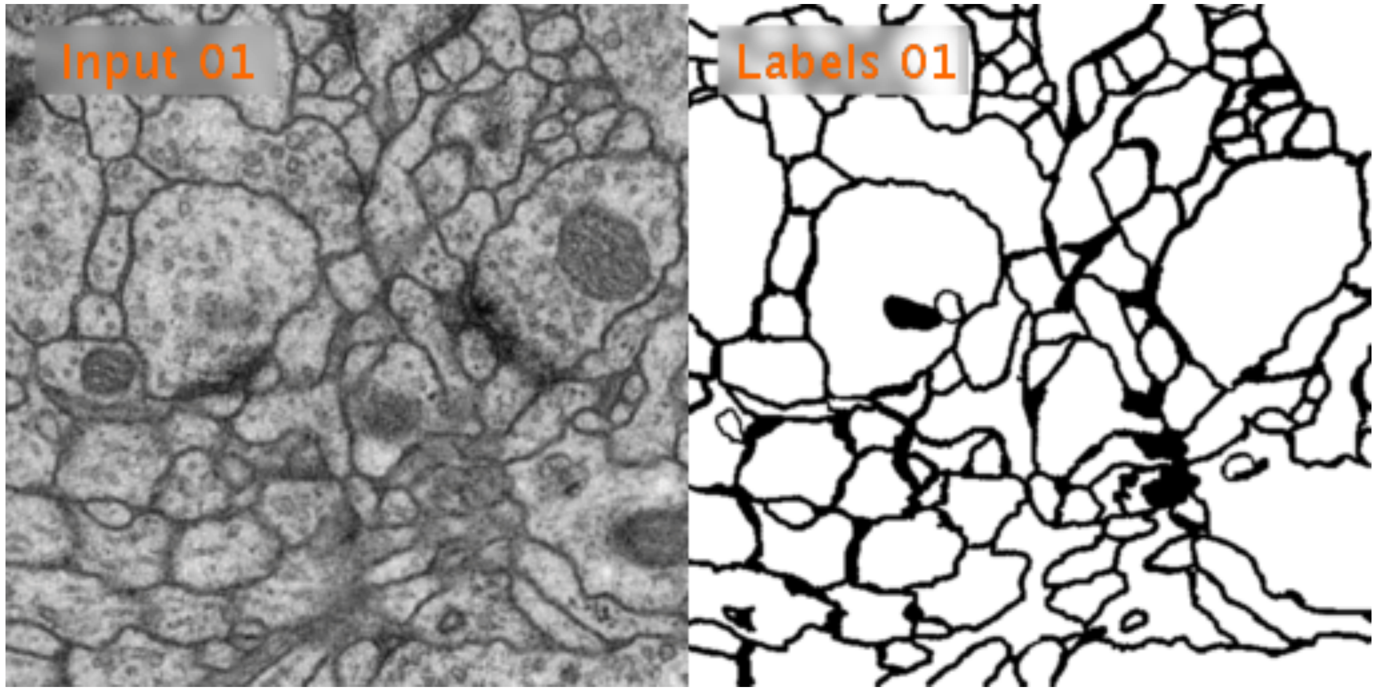

- Transfer learning less useful in segmentation context

http://brainiac2.mit.edu/isbi_challenge/home

Solution: Data Augmentation

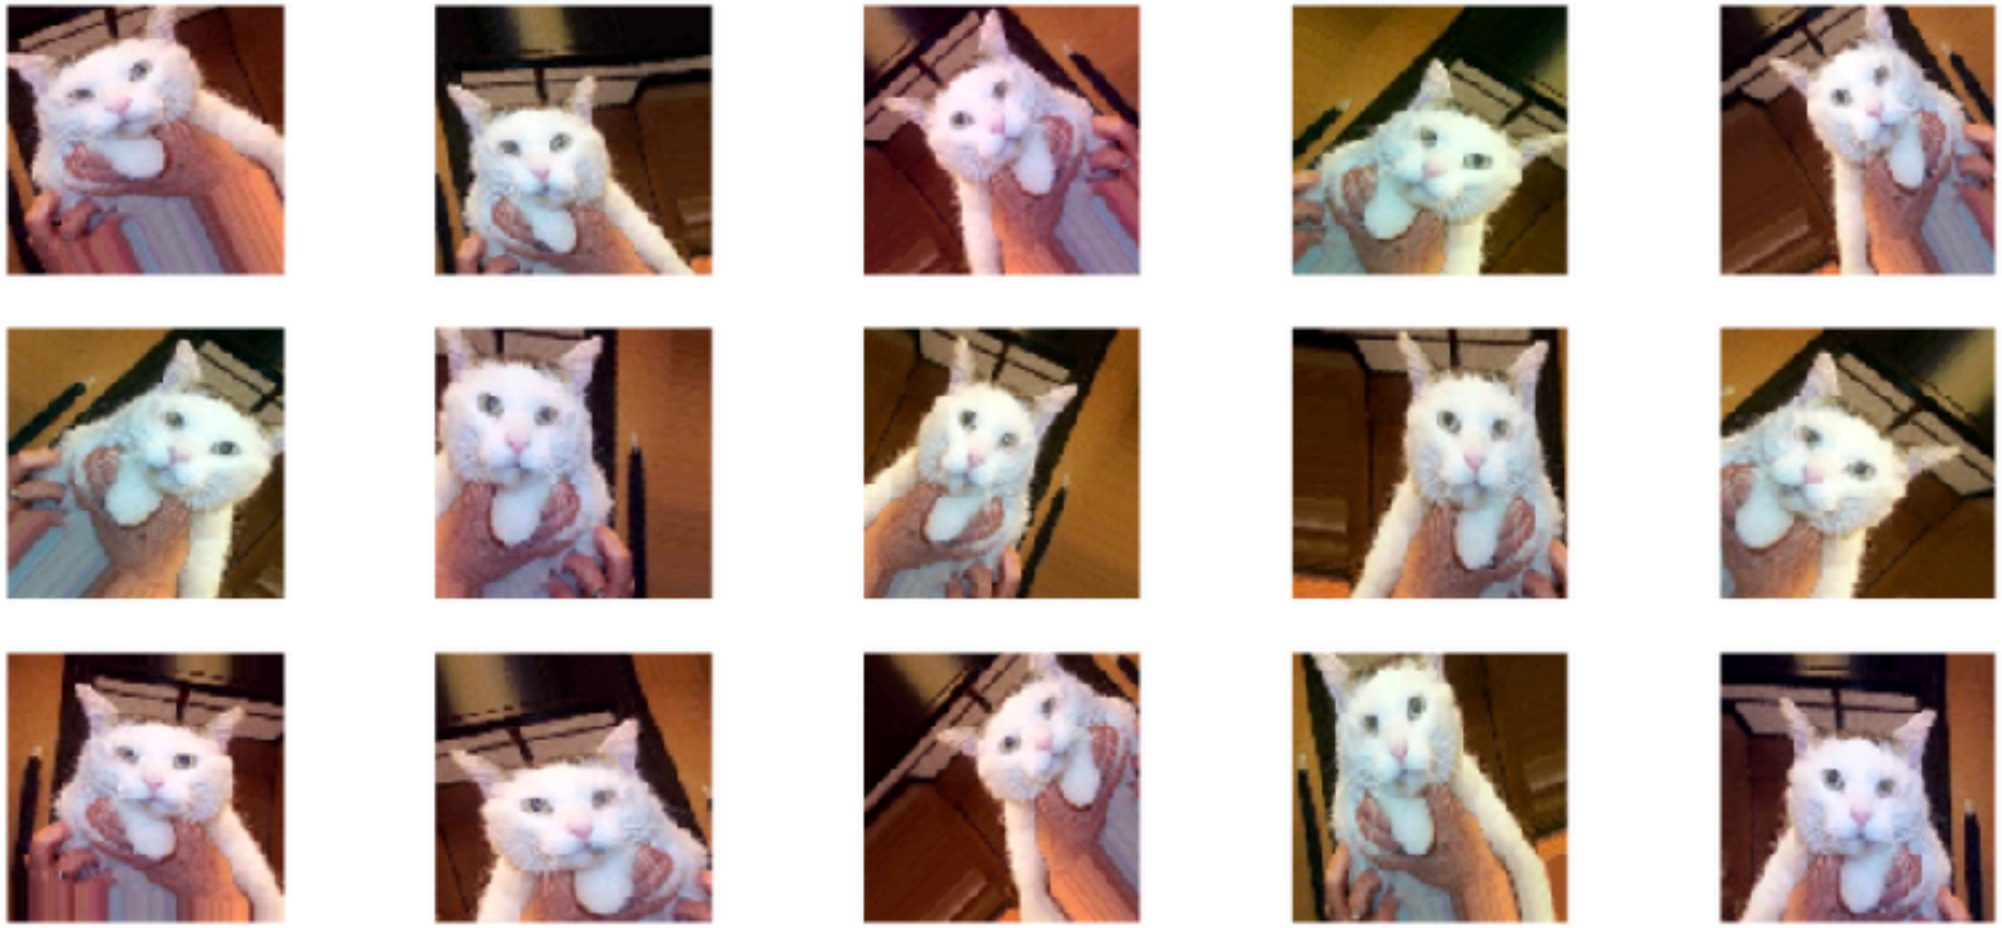

Figure: https://m2dsupsdlclass.github.io/lectures-labs/slides/04_conv_nets/index.html#82

Solution: Data Augmentation

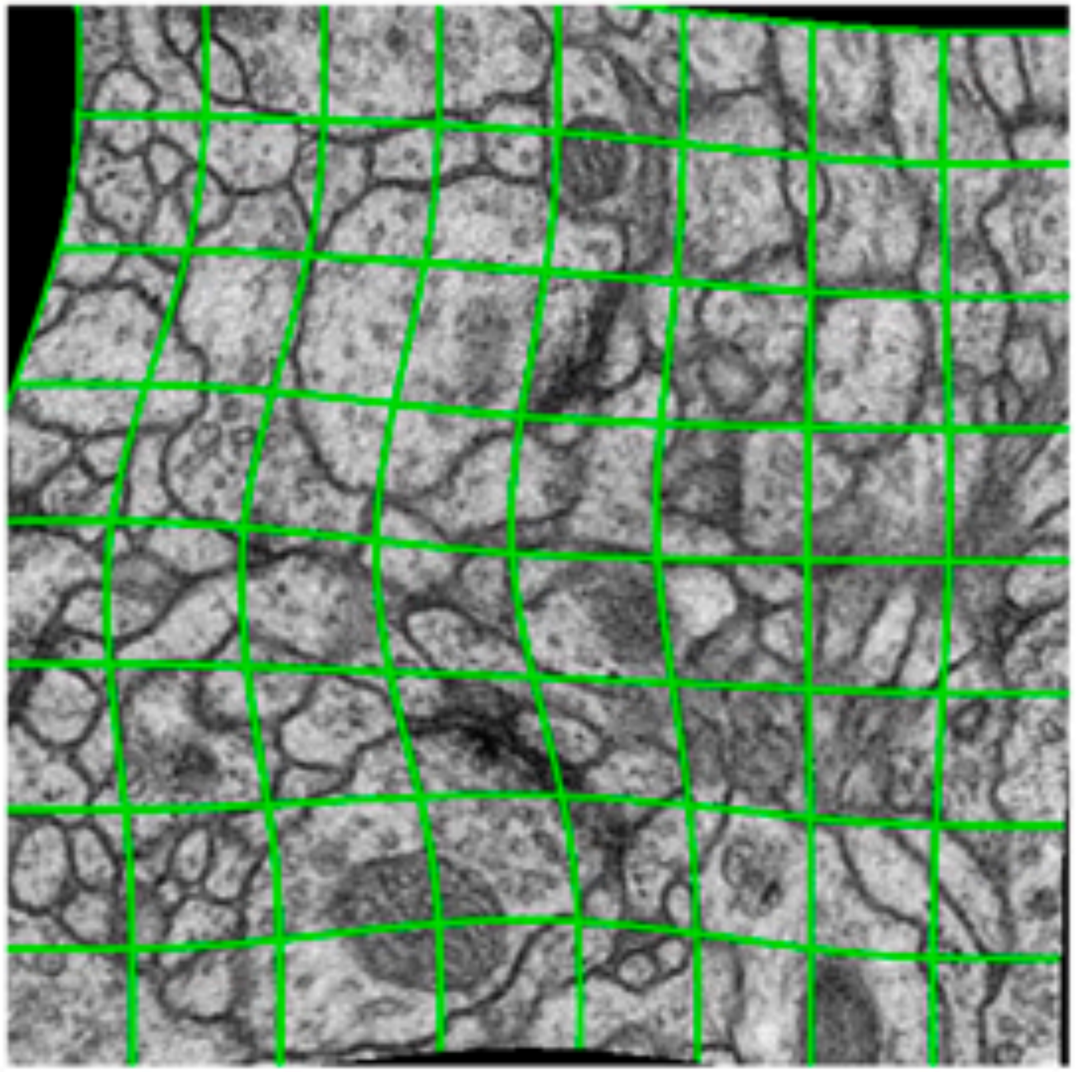

resulting deformed image
(for visualization: no rotation, no shift, no extrapolation)

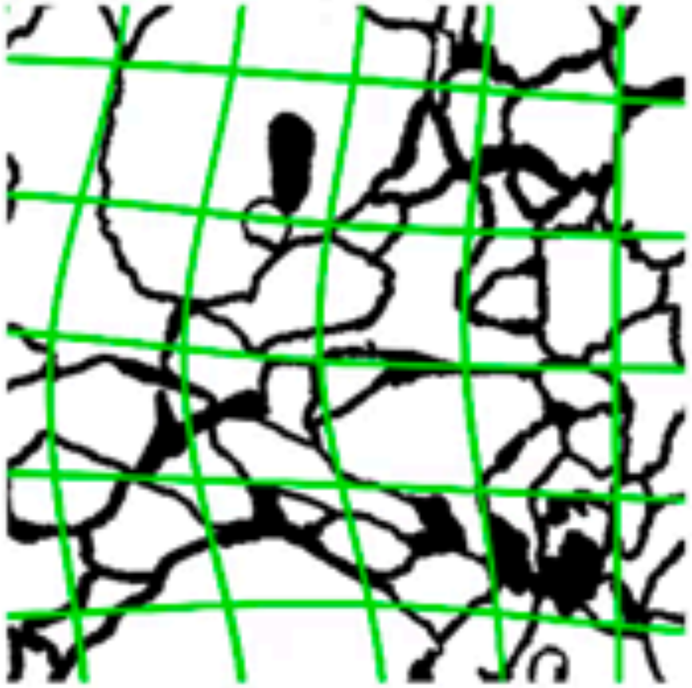

correspondingly deformed
manual labels

Other Improvements: Task specific loss functions

binary cross-entropy: $$L_{bce} = \sum_i y_i \log o_i + (1 - y_i) \log (1 - o_i)$$

y_i =true labels
o_i =predictions

"Dice" loss:
(common metric used in segmentation)

$$L_{Dice} = -\frac{2 \sum_i o_i y_i}{\sum_i o_i + \sum_i y_i}$$

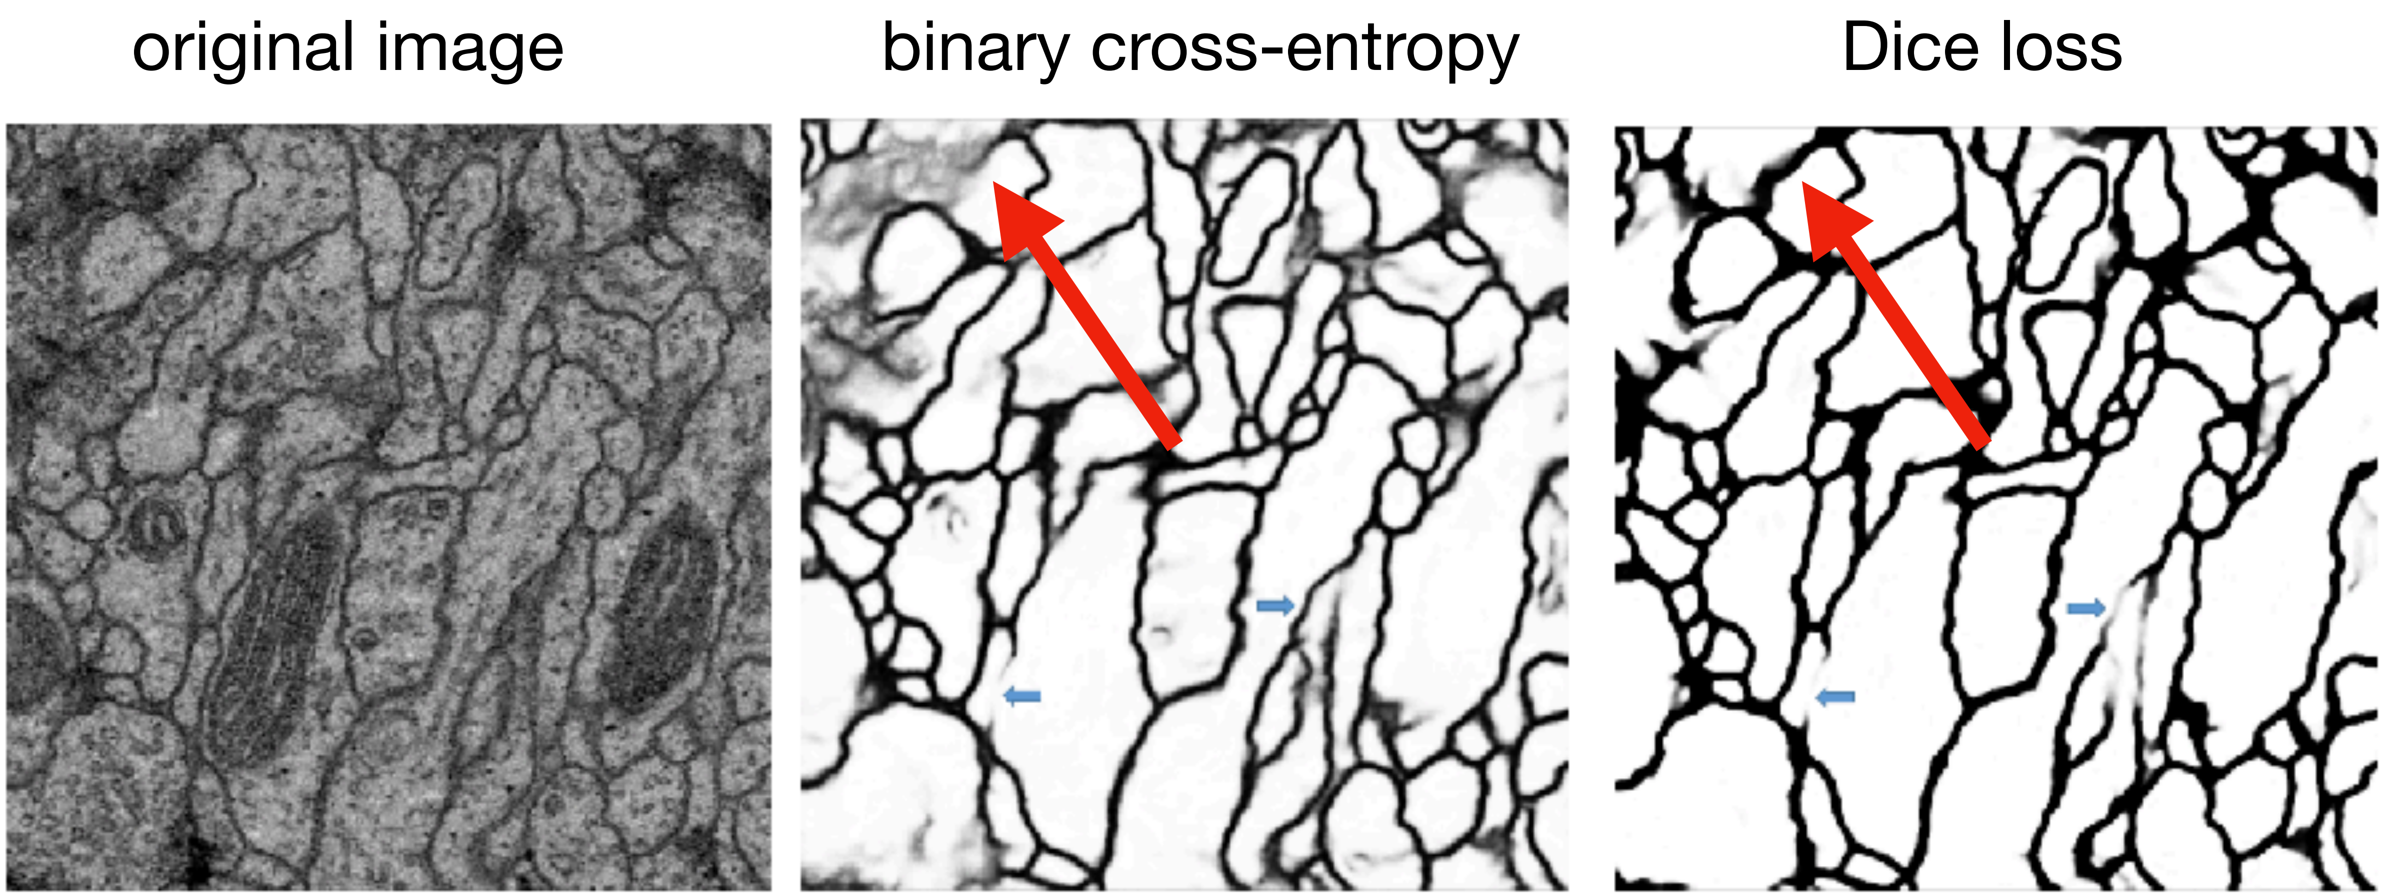

Figure: [Drozdzal et al. 2016]

Other applications: Cell segmentation in light microscopy images

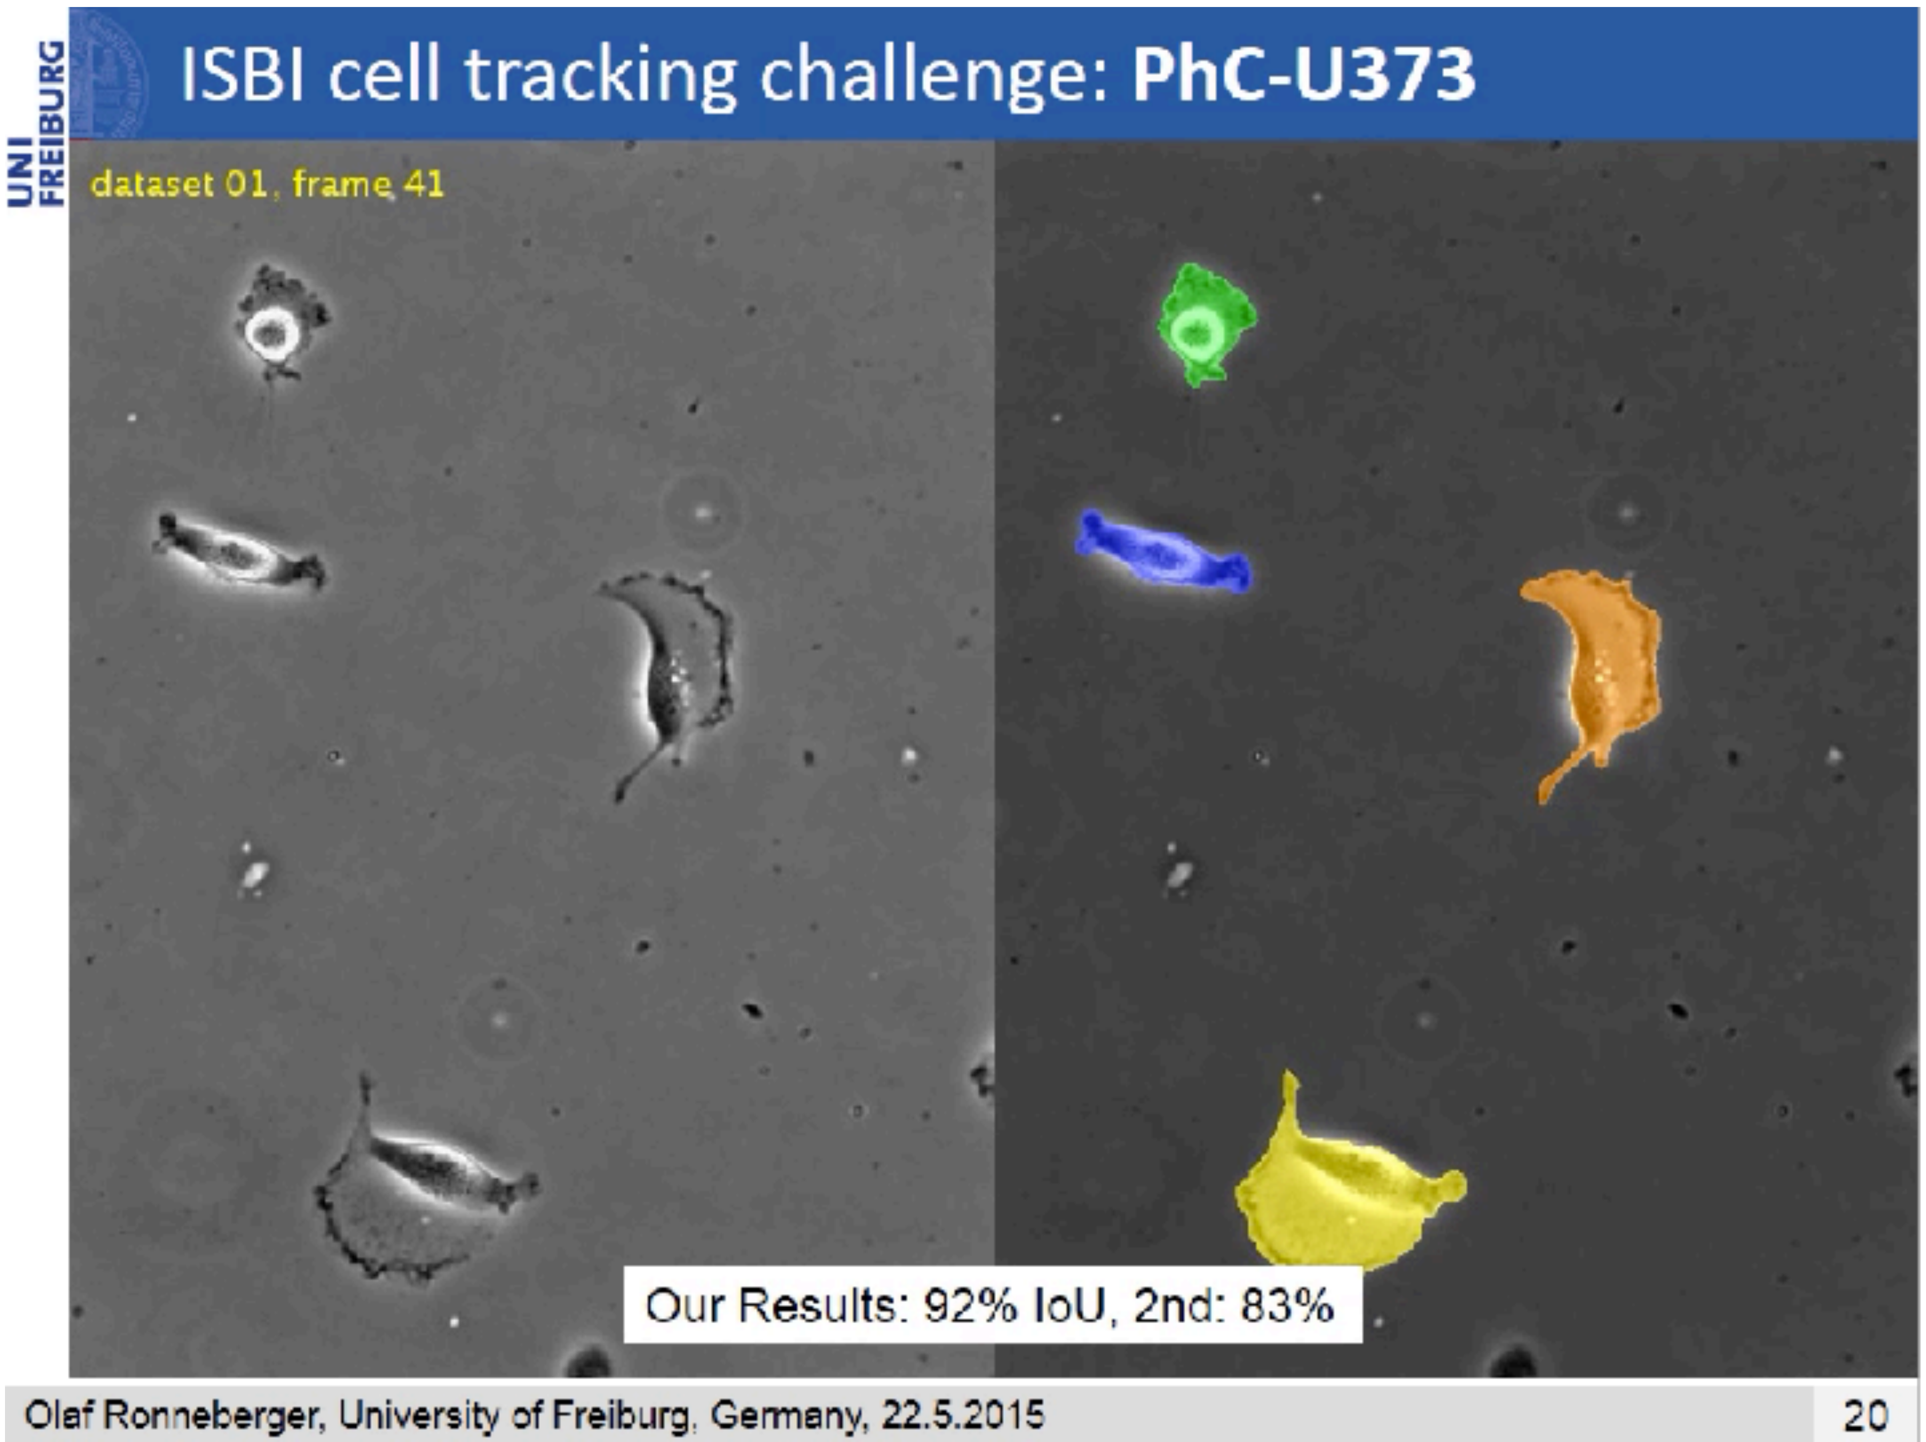

Other applications: Segmentation of prostate in 3D MRI scans

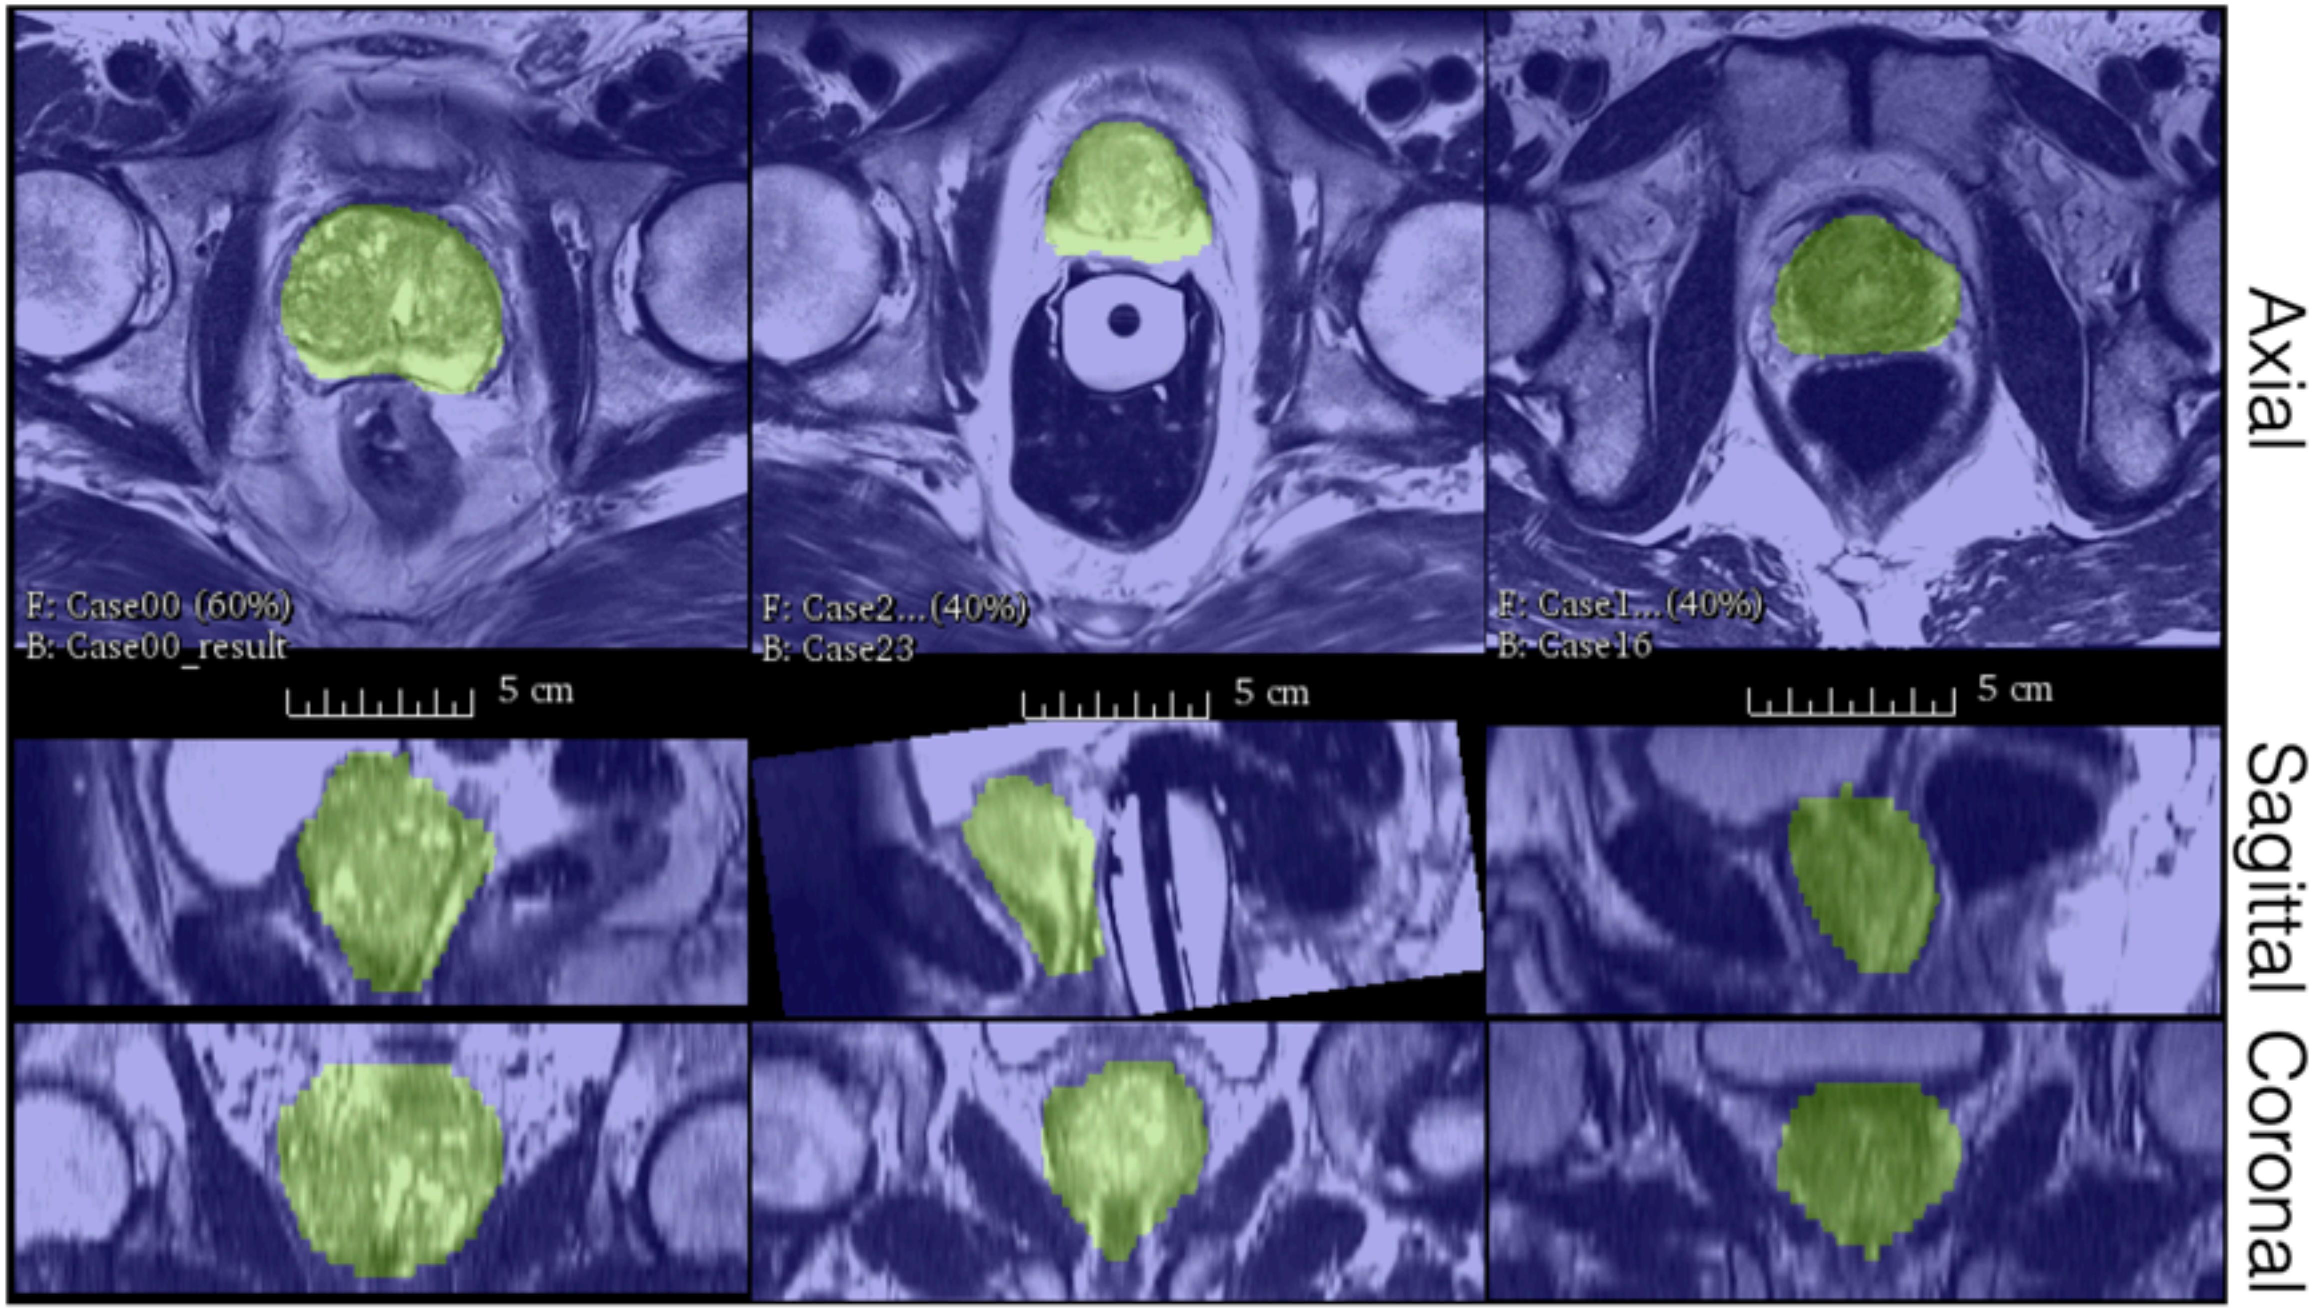

V-net [Milletari et al., 2016]

Takeaway

- Standard CNN classification architectures are inefficient/poor choices for segmenting biomedical images.
- High-quality segmentation of biomedical images is made possible with fully connected neural networks (such as the U-net)
- Domain specific knowledge (data augmentation & custom loss functions) yields more improvements.

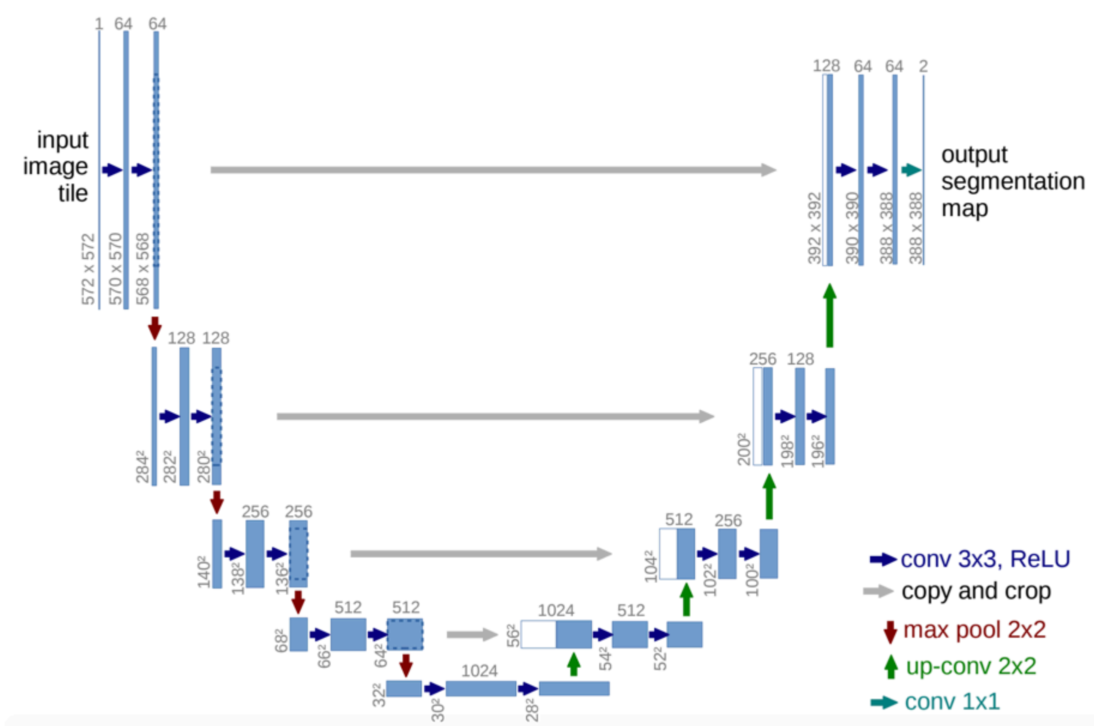

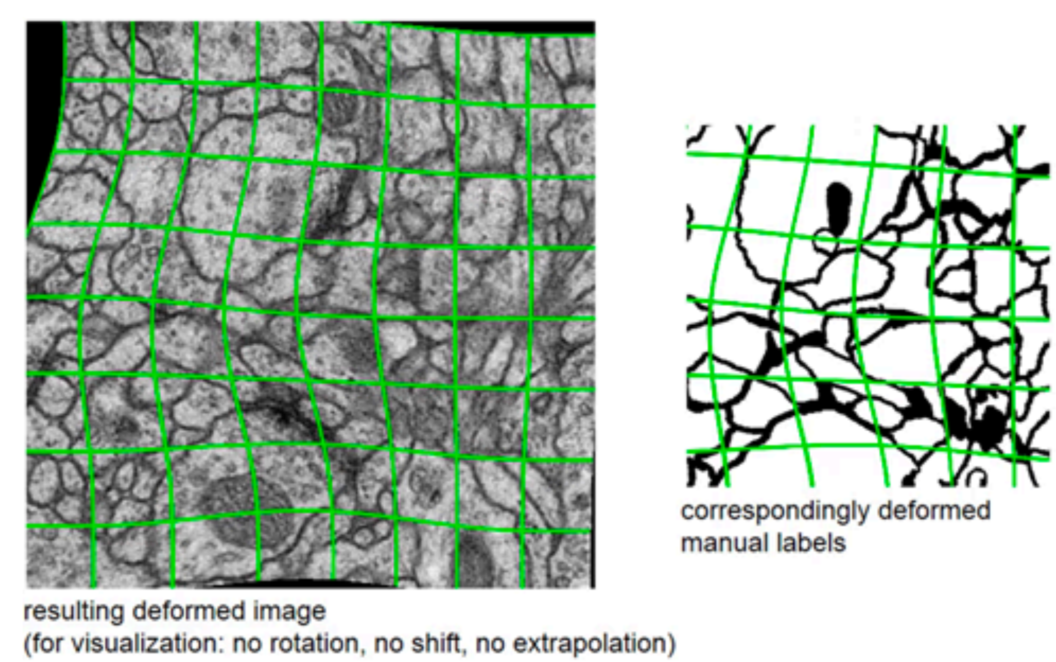

Recap and Outlook

Successful applications of deep learning in biomedical image analysis

- **Classification/Detection**
  - Skin lesion classification from photographs for skin cancer detection
  - Lung nodule classification in CT images for lung cancer detection
- **Segmentation**
  - Segmentation of neuronal structures in electron microscope stacks
  - Cell tracking in light microscopy images
  - Prostate segmentation in 3D MRI images
- **& Many, many more —**
  Hundreds of new publications and patents every year

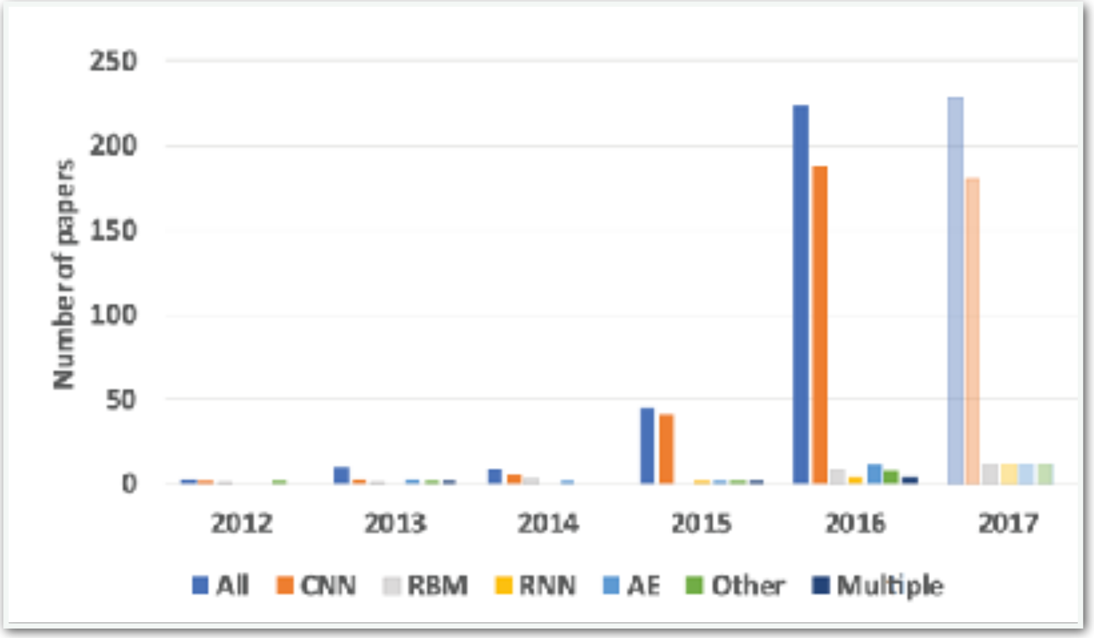

[Litjens et al., 2017]

Challenges in deep learning for biomedical imaging

- **Challenge 1: Limited Training Data**
  - Transfer Learning — pre-train on ImageNet
  - Data Augmentation — shifts, rotations, warps of data
  - Not talked about: Generative models, few-shot learning
- **Challenge 2: Complex Input Formats**
  - Multi-scale/multi-view concatenations of CNN's
  - Not talked about: 3-D CNN's, incorporating semantic information
- **Challenge 3: Tasks Beyond Classification**
  - Fully convolutional neural nets for segmentation
  - Modified loss functions - Dice loss in place of cross-entropy
  - Not talked about: **image restoration/reconstruction problems (next)**

Part II:
Deep learning for
biomedical image reconstruction

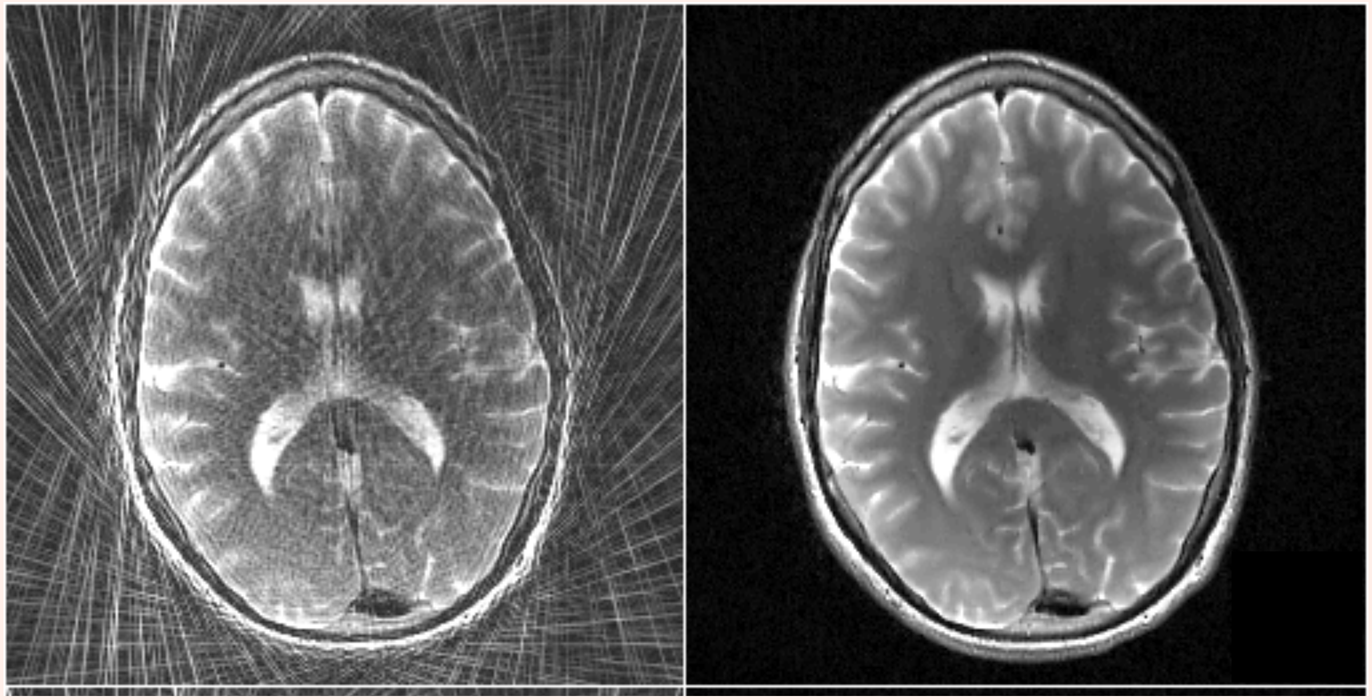

Biomedical imaging pipeline

Aquire Data → Reconstruct Images → Process Images → Analyze Diagnose Quantify Interpret Intervene

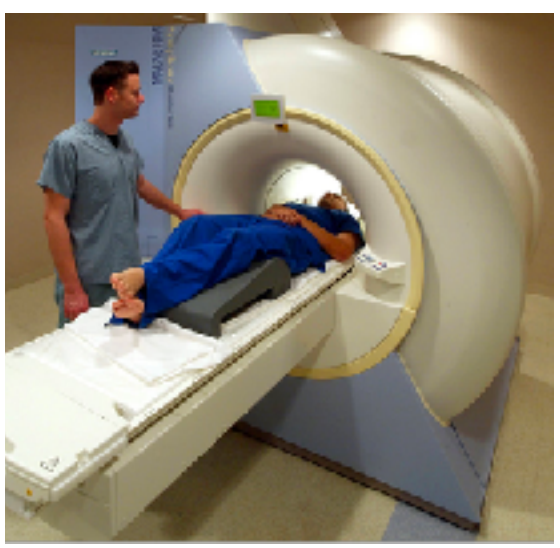

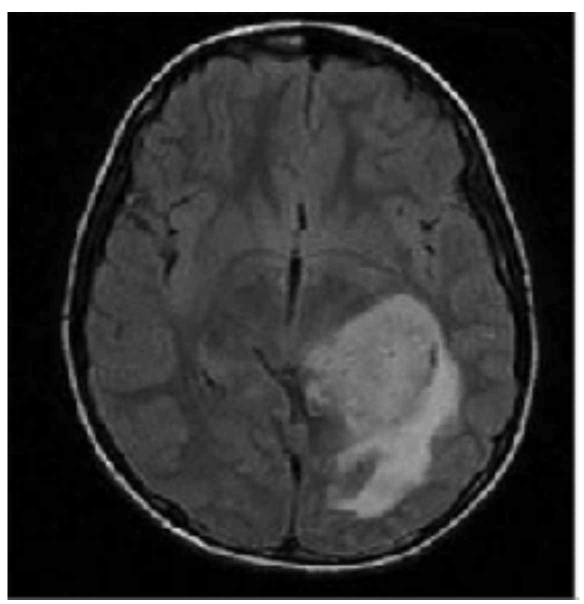

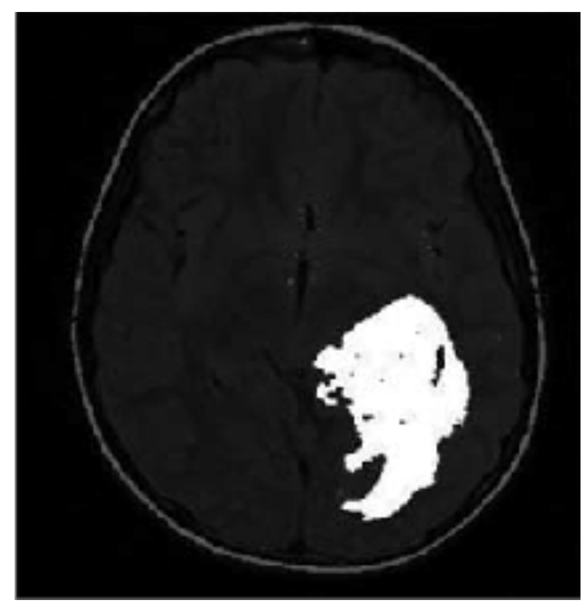

Why bother? In MRI…

Magnetic Resonance Imaging (MRI)

- Long scan-time (30-90 minutes)
- Physical limits to how fast one can take measurements
- Could take fewer measurements, but at the expense of a noisy/ lower-resolution image

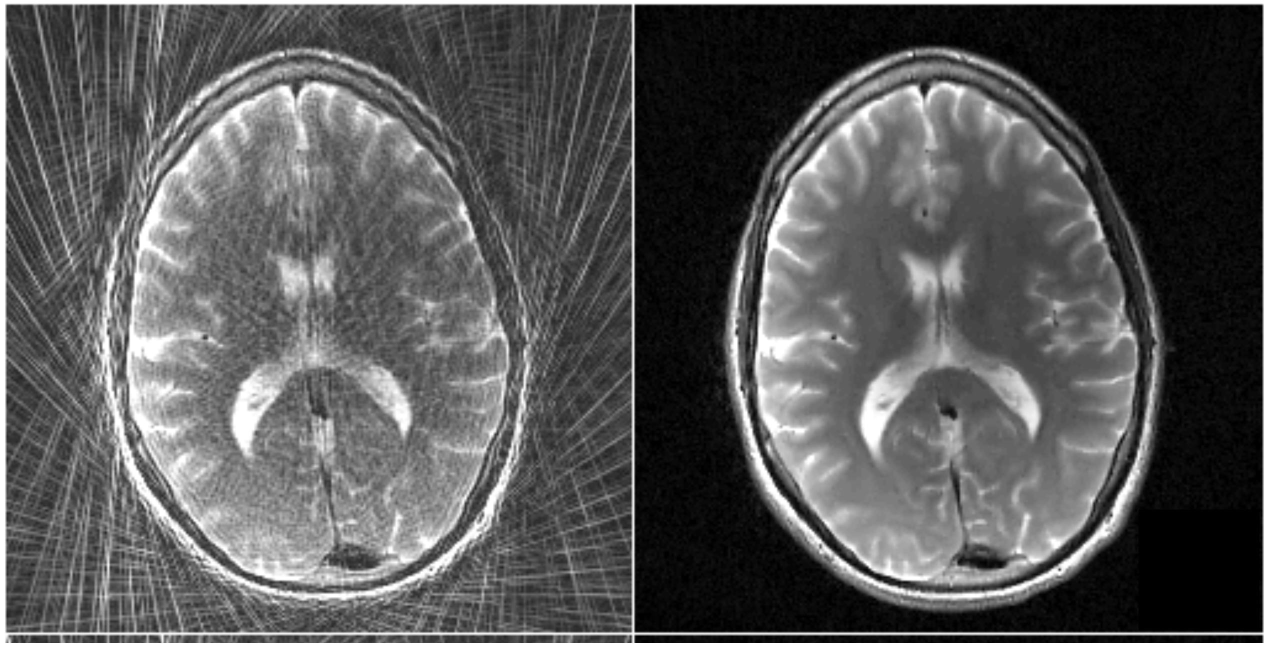

Goals:

faster scans: accelerate MRI acquisition (take fewer measurements)

faster recons: reduce computational cost of reconstruction

better images: improve spatio-temporal resolution (e.g., dynamic MRI)

Why bother? In CT…

X-ray Computed Tomography (CT, aka a CAT scan)

Uses ionizing radiation — potentially harmful to patient

ORIGINAL INVESTIGATION

Projected Cancer Risks From Computed Tomographic Scans Performed in the United States in 2007

Amy Berrington de González, DPhil; Mahadevappa Mahesh, MS, PhD; Kwang-Pyo Kim, PhD; Mythreyi Bhargavan, PhD; Rebecca Lewis, MPH; Fred Mettler, MD; Charles Land, PhD

"Overall, we estimated that approximately 29,000 future cancers could be related to CT scans performed in the US in 2007."

Goals:

lower dose: try to use lower radiation doses, yet achieve same image quality

faster recons: reduce computational cost of reconstruction

Other applications: Medical imaging

Positron Emission Tomography (PET)

200x Low-dose PET Reconstruction using Deep Learning

Junshen Xu[†], Enhao Gong[†], John Pauly and Greg Zaharchuk[*]

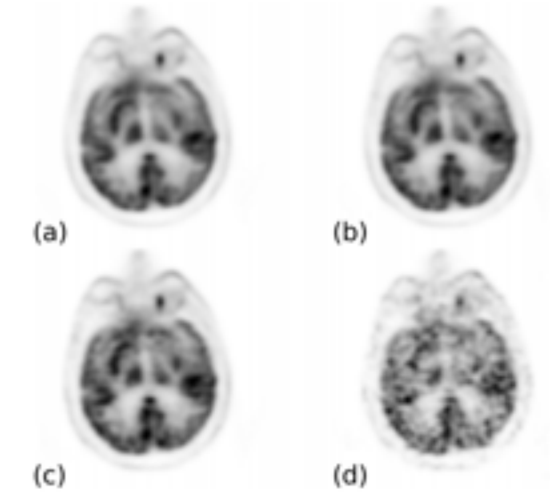

Fig. 1. PET images with normal dose and different levels of dose reduction. (a) standard-dose, (b) quarter-dose, (c) twentieth-dose, and (d) two-hundredth-dose.

Ultrasound Imaging

TOWARDS CT-QUALITY ULTRASOUND IMAGING USING DEEP LEARNING

Sanketh Vedula[*,†] *Ortal Senouf*[*,†] *Alex M. Bronstein*[†] *Oleg V. Michailovich*[‡] *Michael Zibulevsky*[†]

[†] Technion – Israel Institute of Technology
[‡] Electrical and Computer Engineering, University of Waterloo, Canada

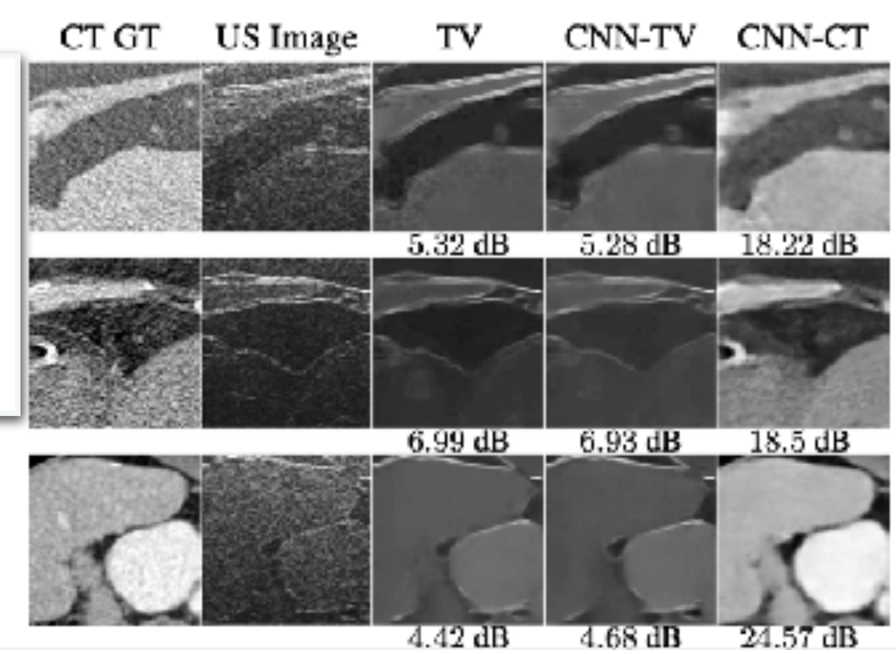

Other applications: Biological imaging

Cell imaging

super-resolution localization fluorescence microscopy

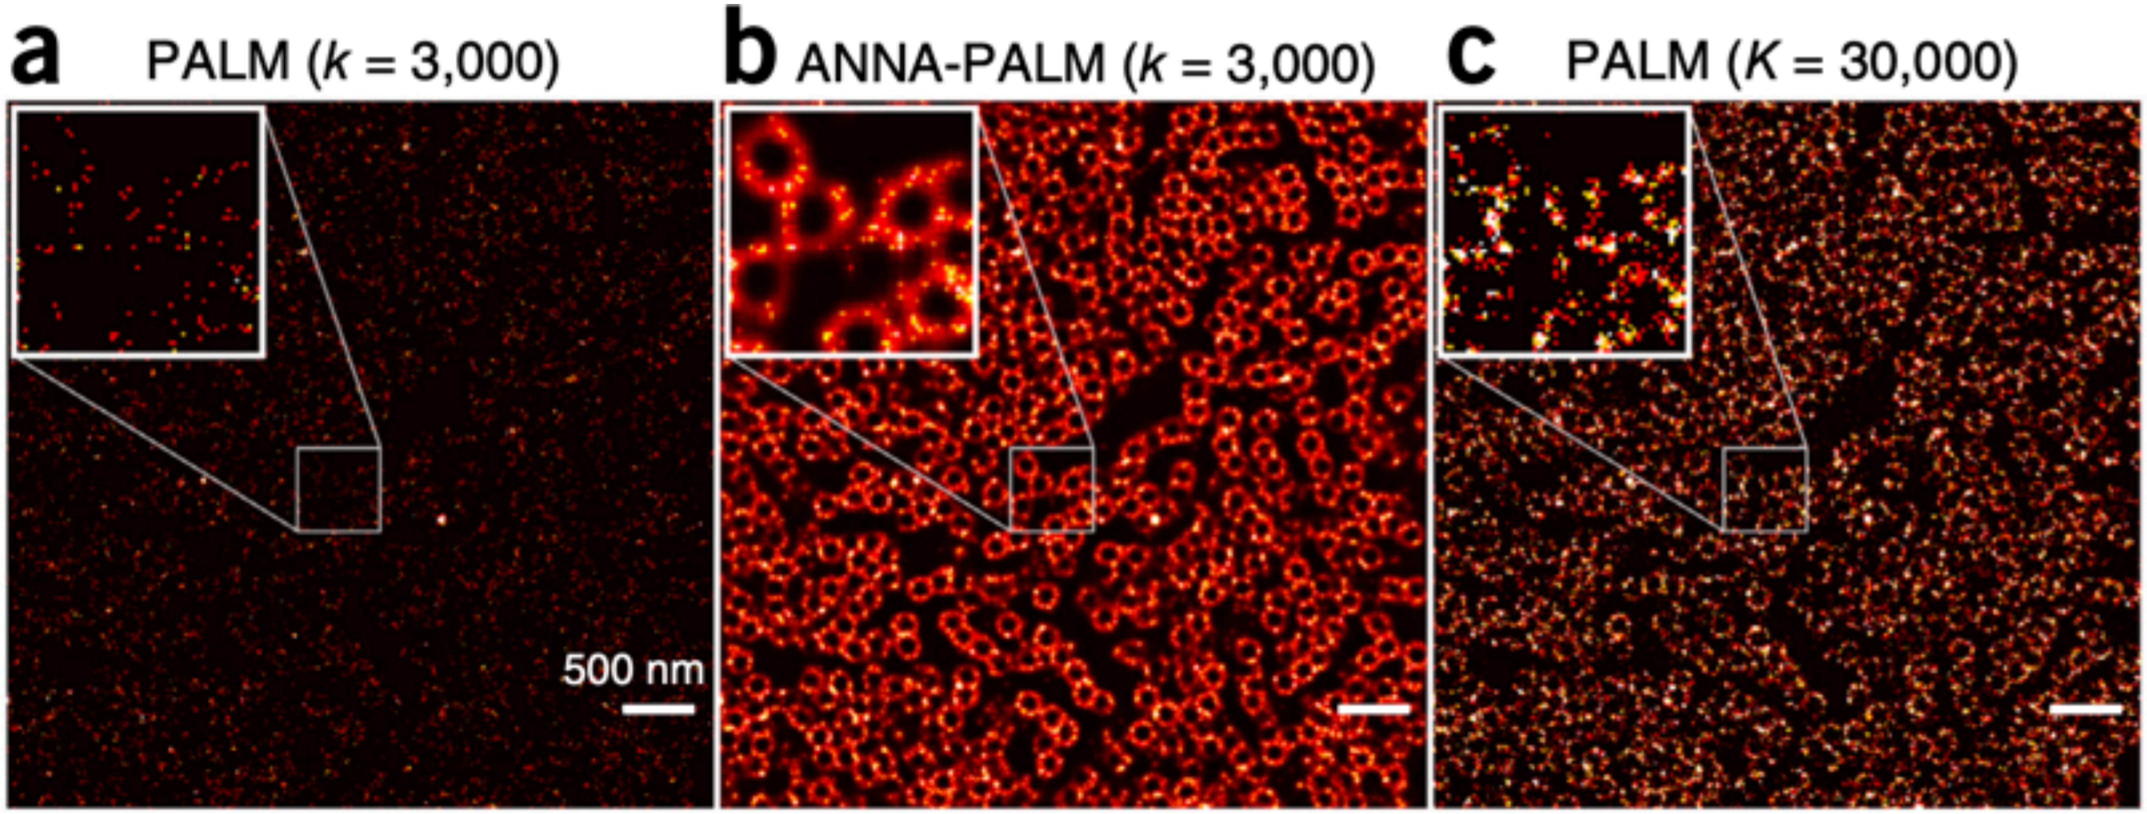

conventional | with deep learning | conventional x10 more data

Deep learning massively accelerates super-resolution localization microscopy

Wei Ouyang[1–3], Andrey Aristov[1–3], Mickaël Lelek[1–3], Xian Hao[1–3] & Christophe Zimmer[1–3]

Medical image reconstruction basics

raw data

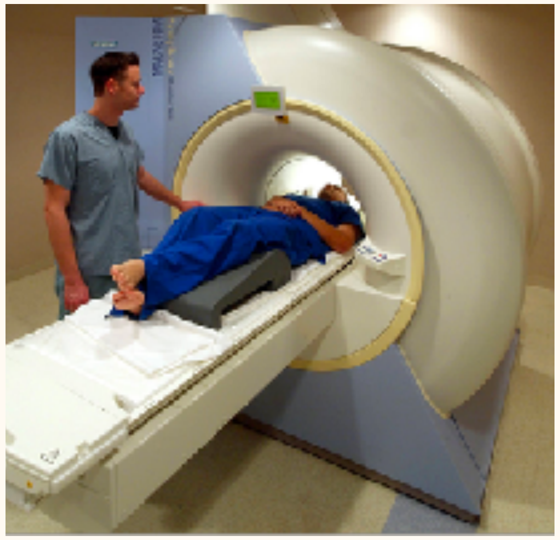

?

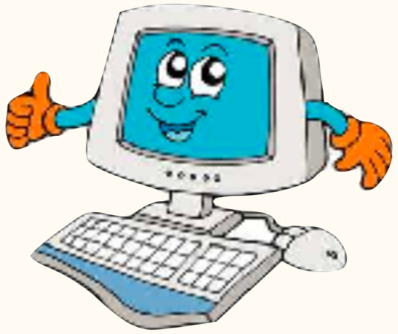

images

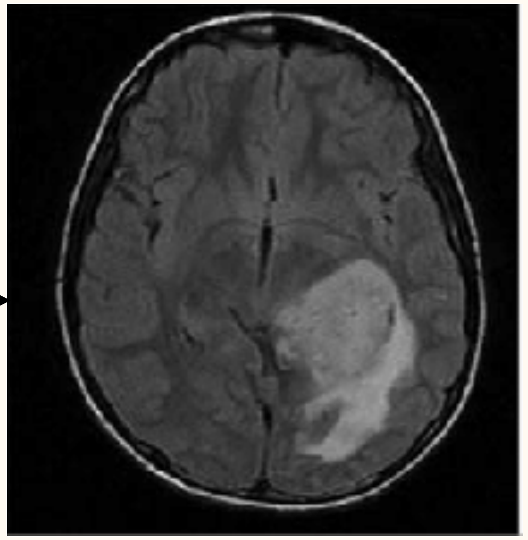

Background: MRI Acquisition

MRI: Data is acquired in spatial frequency domain (k-space)

k-space | image domain

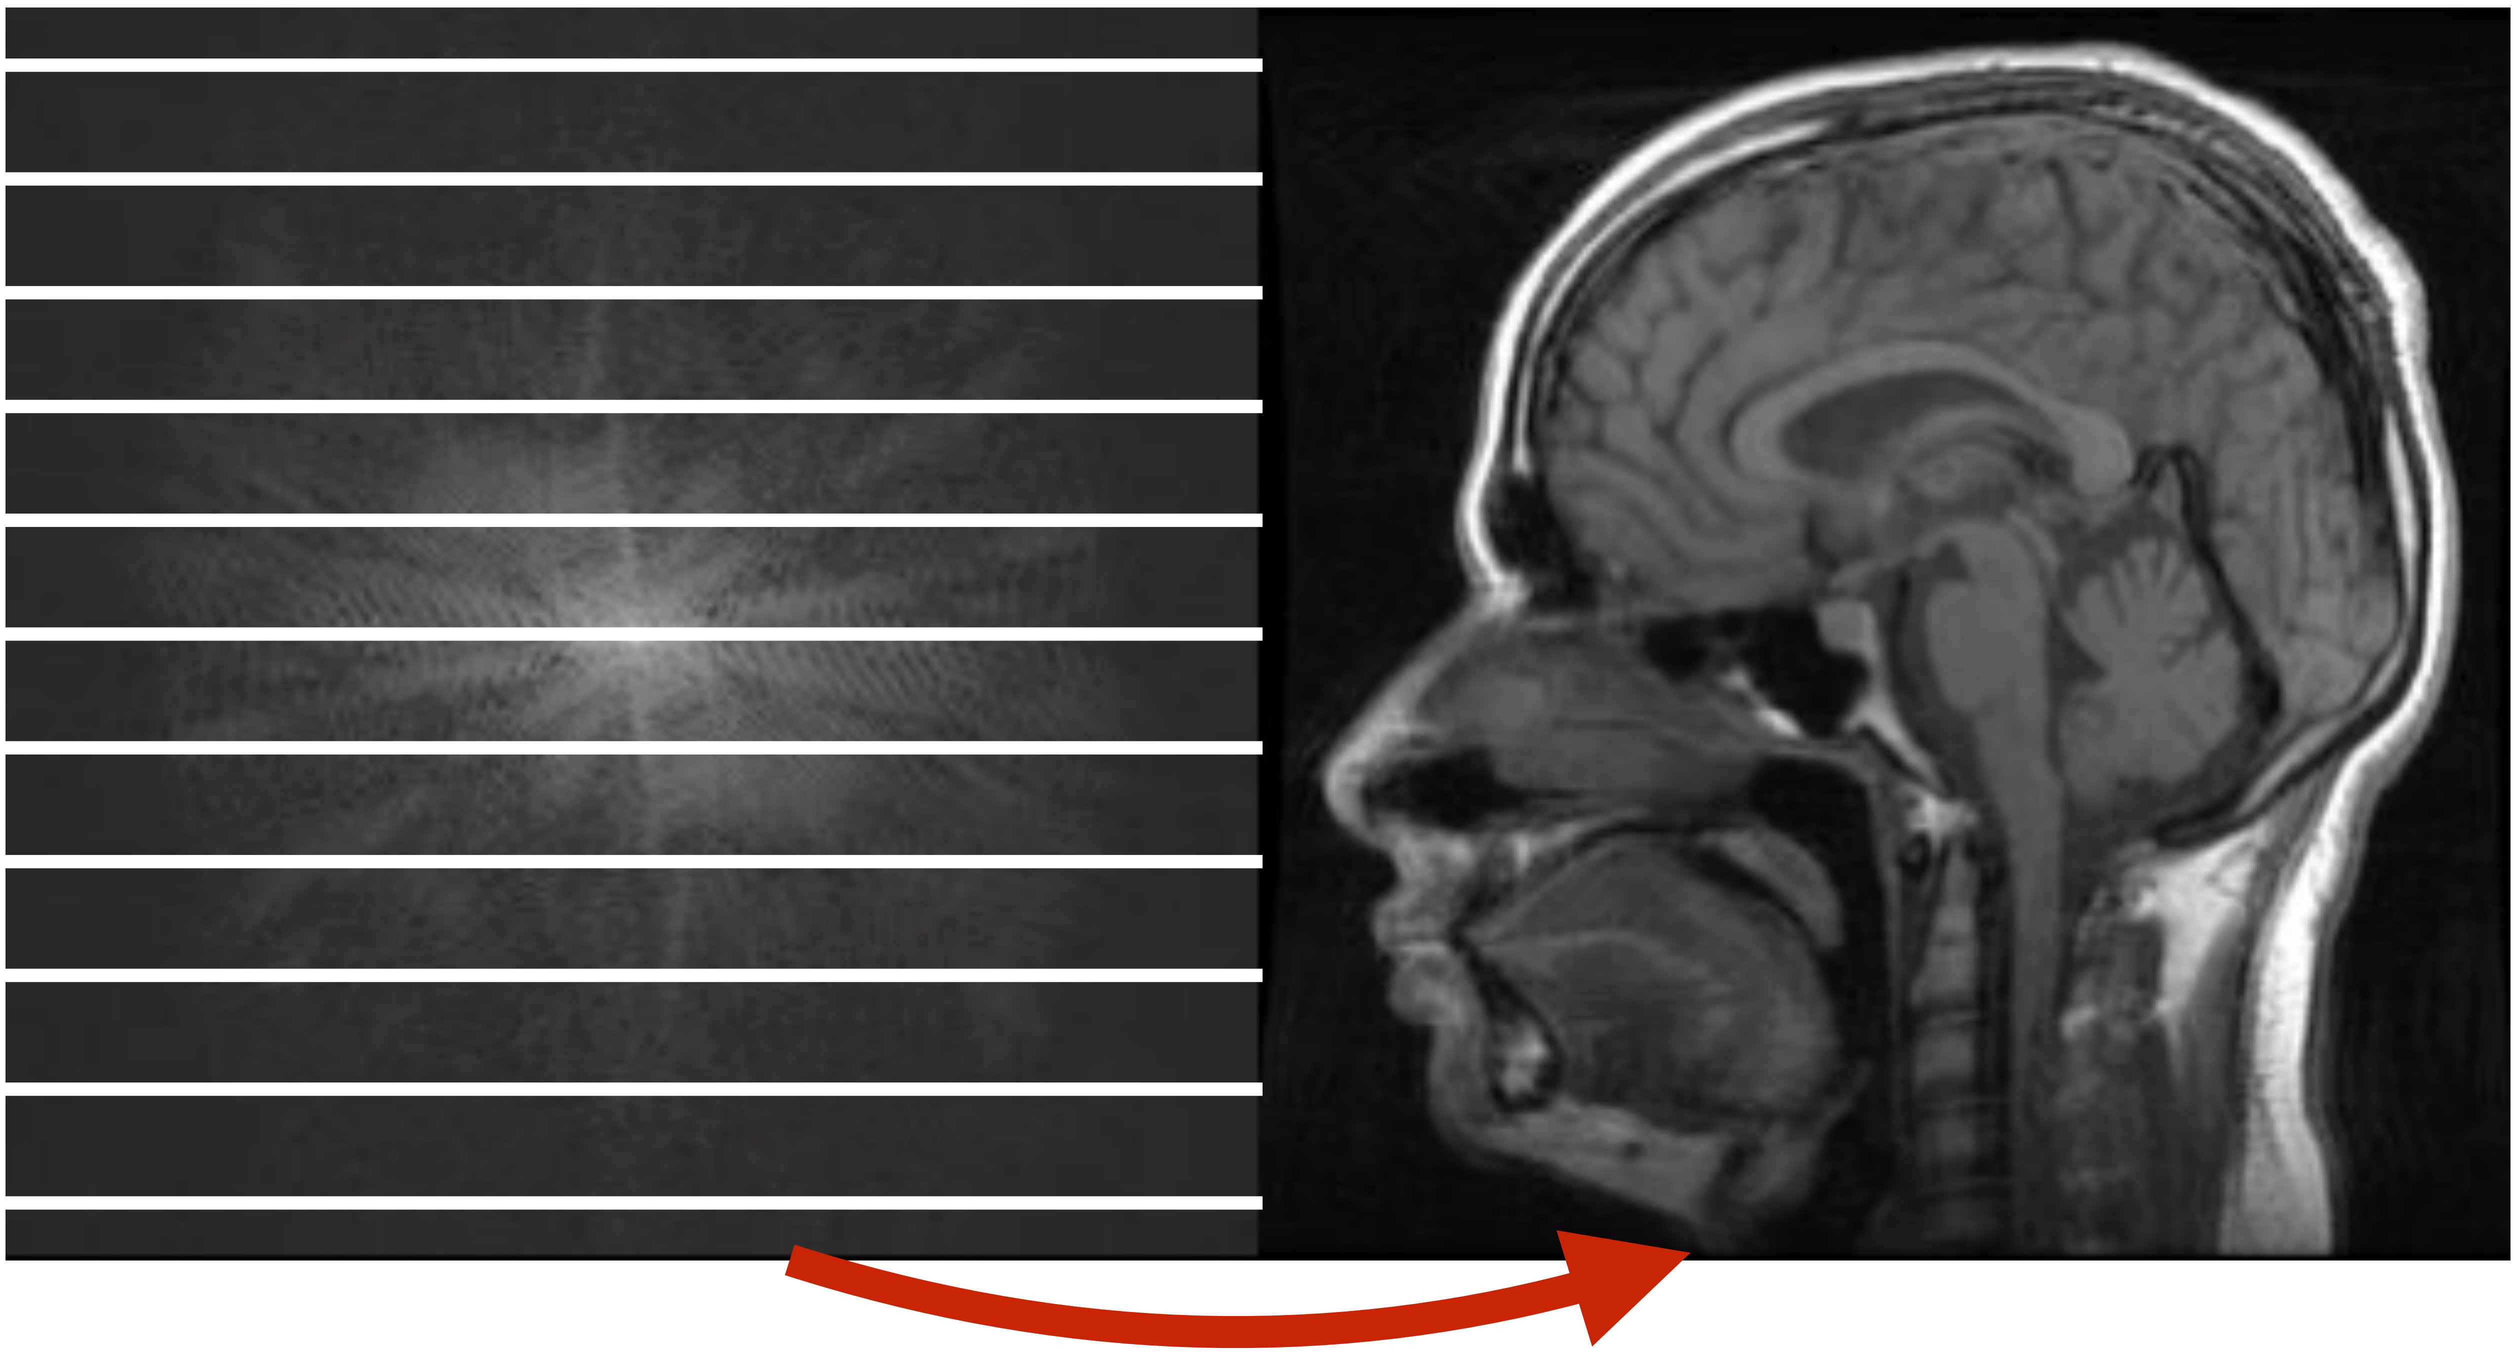

Inverse Fourier transform

Background: Computed Tomography

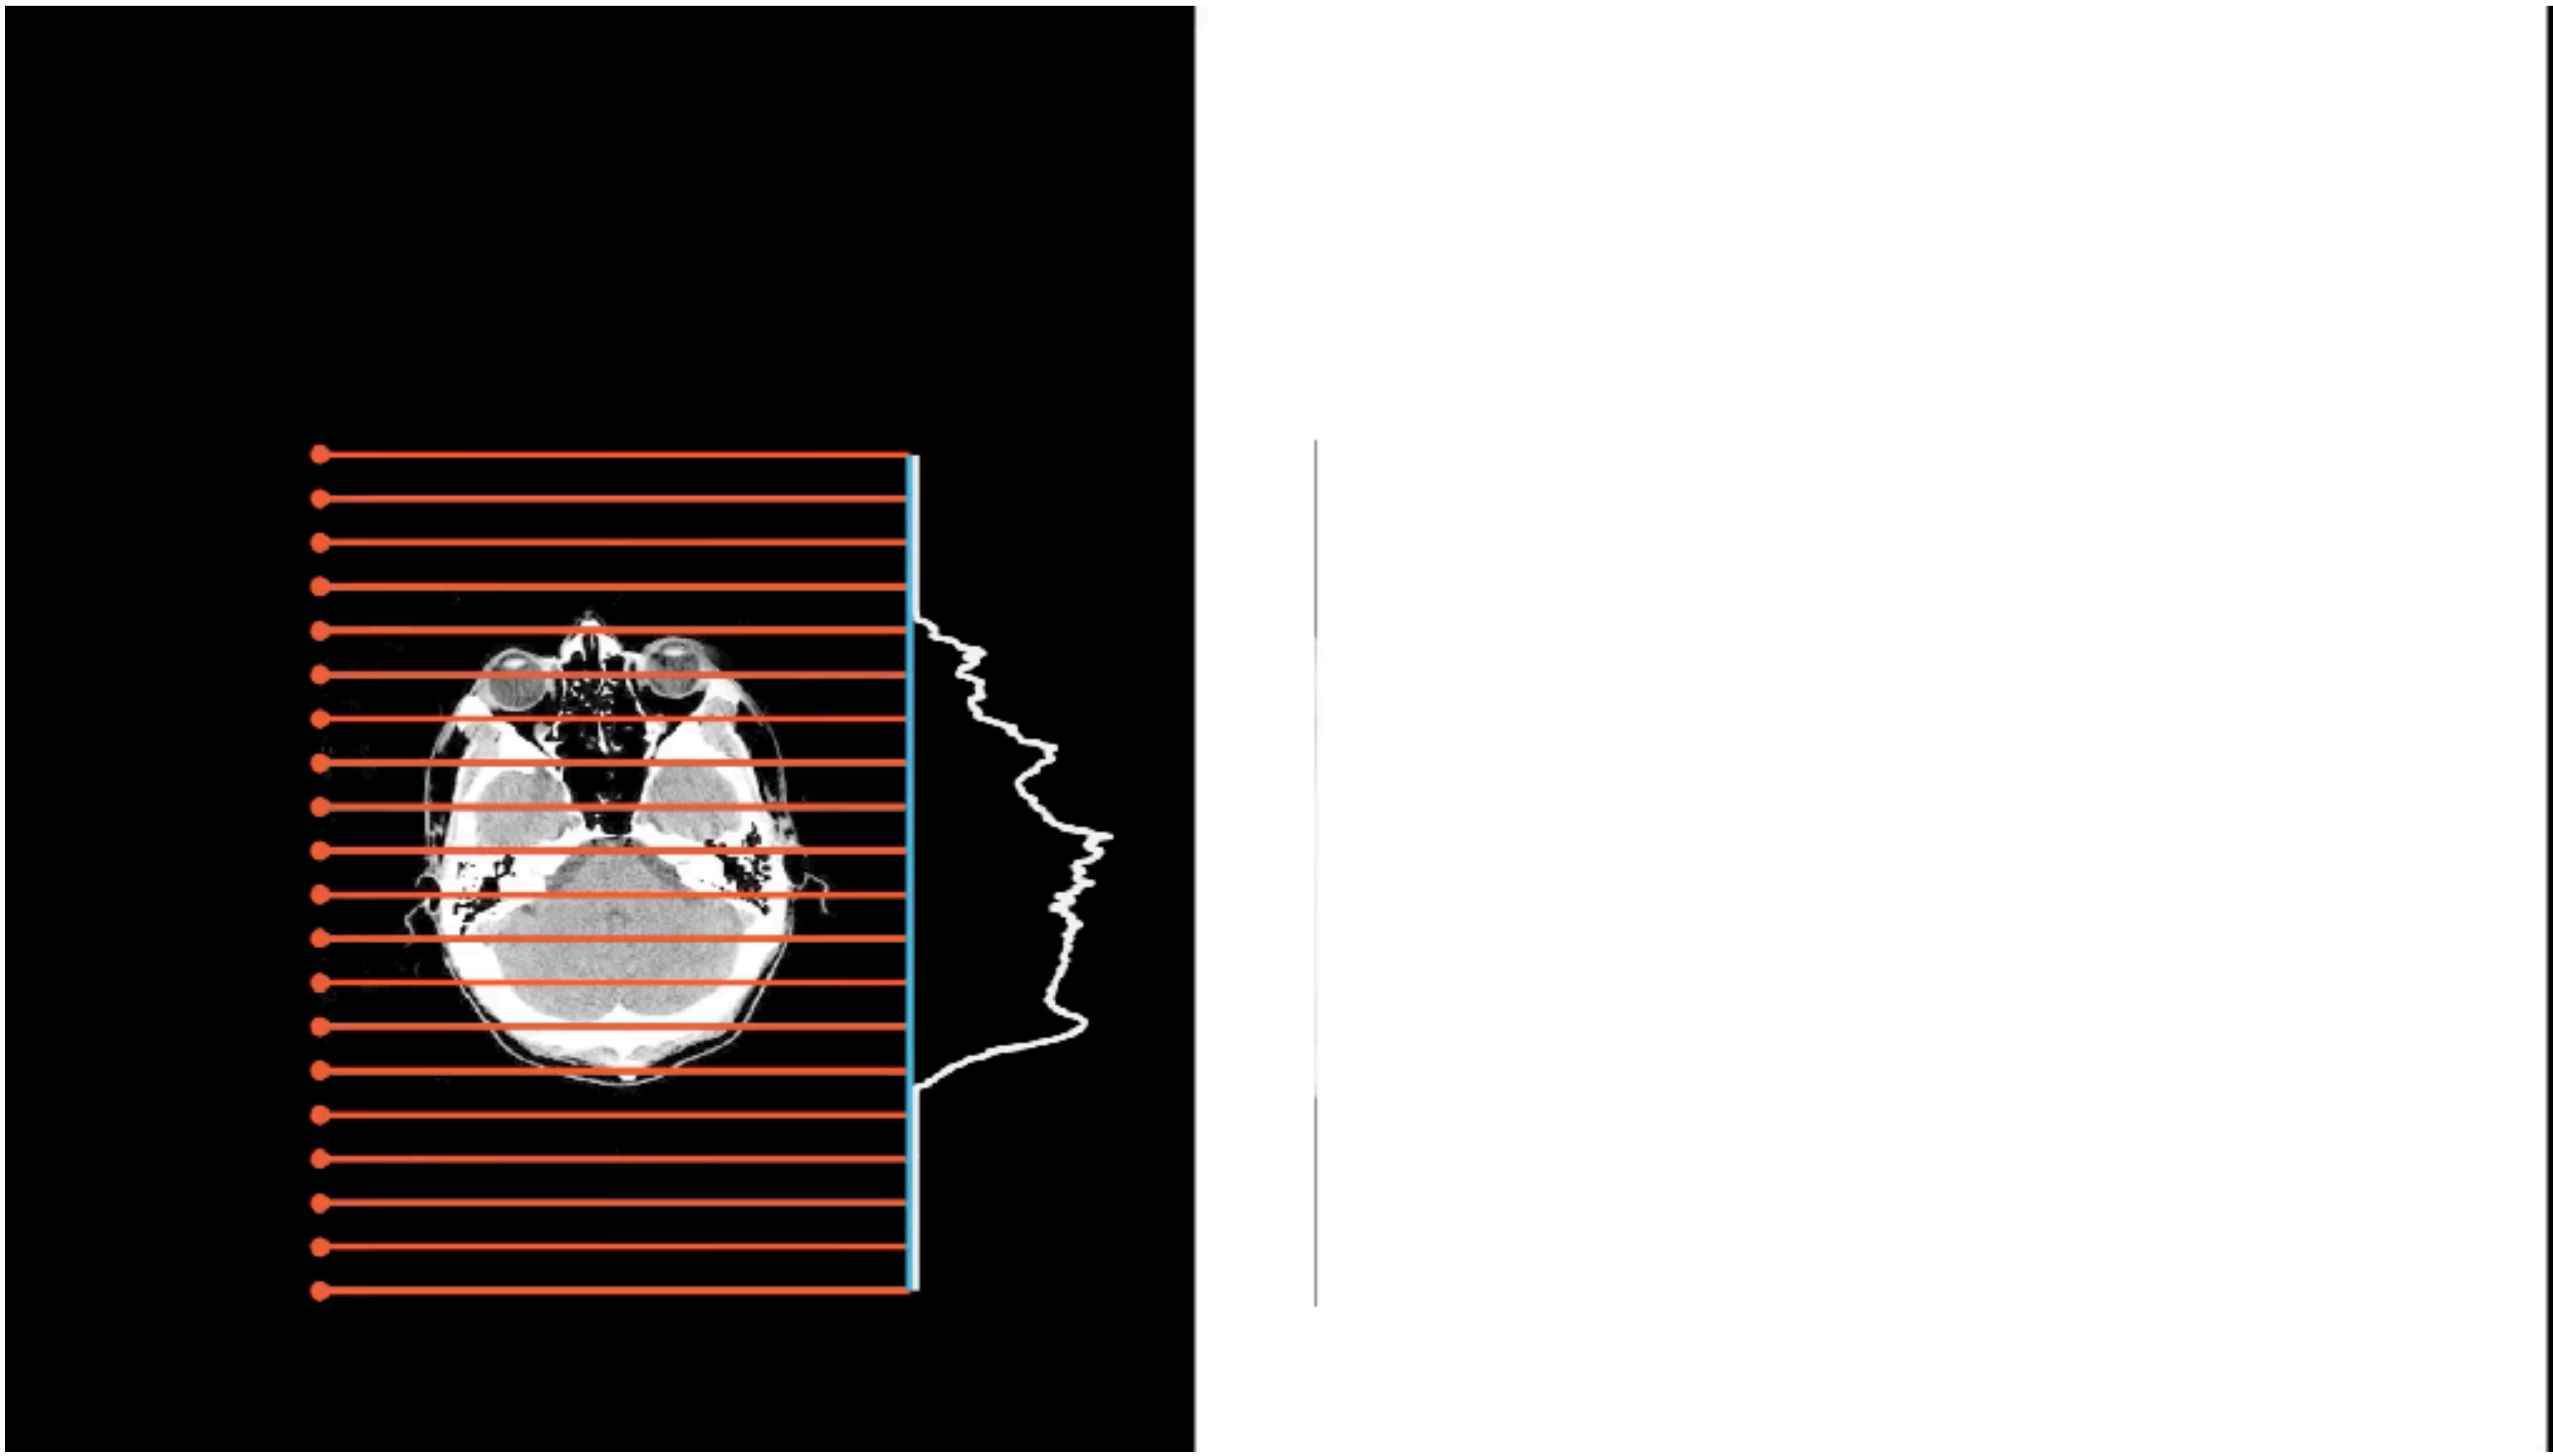

Video credit: Samuli Siltanen **https://www.youtube.com/watch?v=q7Rt_OY_7tU**

Abstraction: Linear inverse problem

linear measurement operator

$$y = H(x)$$

measurements **image**

Abstraction: Linear inverse problem

linear measurement operator

$$y = H(x)$$

measurements **image**

Example: MRI

k-space measurements

MRI image (tissue density)

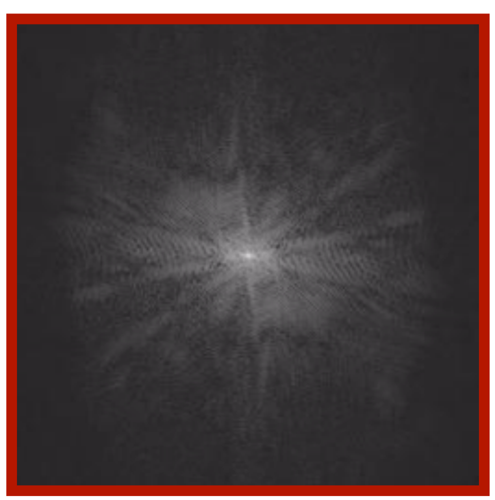

$$= H_{\mathrm{MRI}}\Big($$

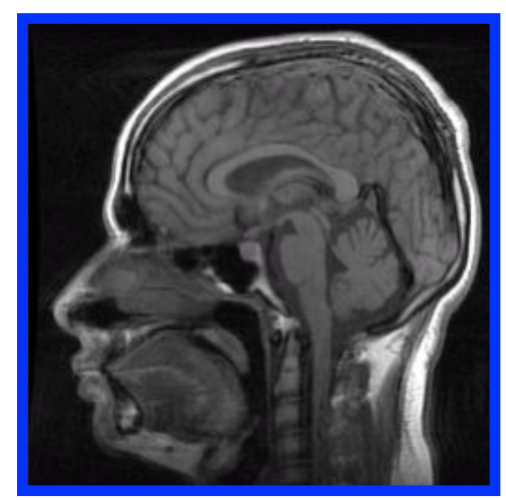

$$\Big)$$

Fourier Transform

Abstraction: Linear inverse problem

linear measurement operator

$$y = H(x)$$

measurements **image**

Example: CT imaging

sinogram

$$= H_{CT}($$

CT image (attenuation)

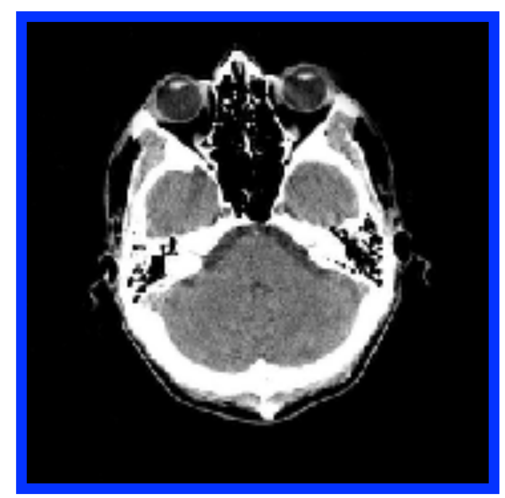

$$)$$

Radon/X-ray Transform

Abstraction: Linear inverse problem

linear measurement operator

$$y = H(x)$$

measurements **image**

Write as matrix equation:

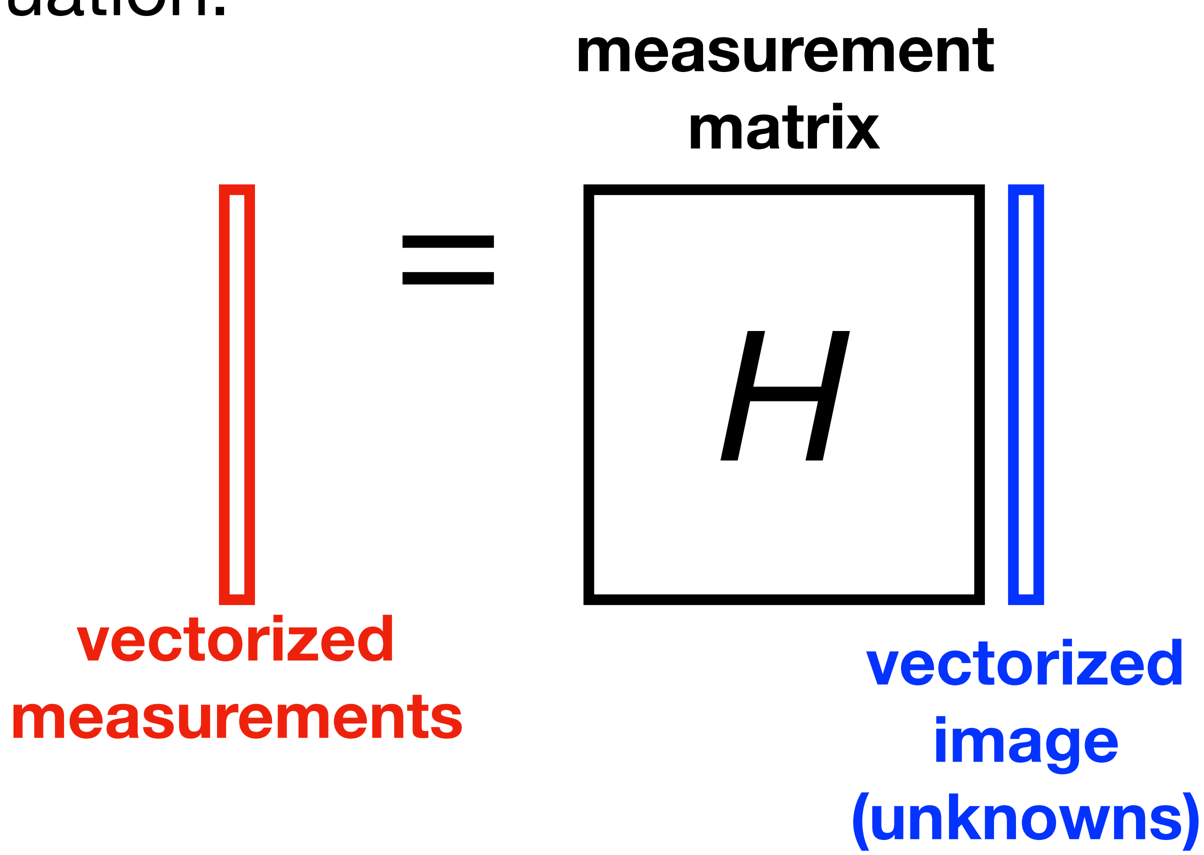

Abstraction: Linear inverse problem

linear measurement operator

$$y = H(x)$$

measurements **image**

Write as matrix equation:

matrix inverse

Too big to invert exactly— Find approximate solution

$$\boxed{H^{-1}} \; [\ ] = [\ ]$$

vectorized measurements

vectorized image (unknowns)

Conventional reconstructions

approximate inverse

$$\tilde{H}^{-1}(y) = \tilde{x}$$

measurements

reconstructed image

Example: MRI imaging

y

"k-space"

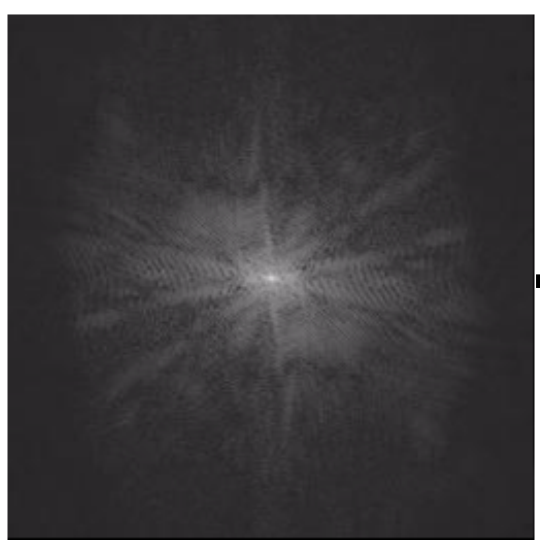

$\tilde{H}^{-1}$

$\tilde{x}$

"inverse discrete Fourier transform"

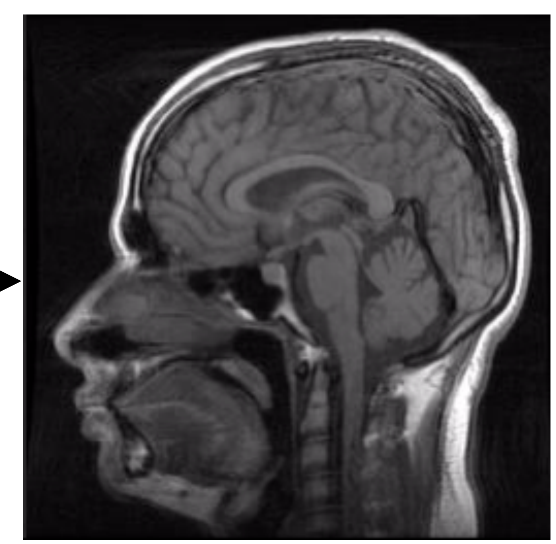

Matlab, Numpy: `ifftn()`

Image source: https://sites.google.com/site/fanyangspace111/cone-beam-tomographyct-reconstruction-using-filtered-back-projection

Conventional reconstructions

approximate inverse

$$\tilde{H}^{-1}(y) = \tilde{x}$$

measurements **reconstructed image**

Example: CT imaging

y

“sinogram”

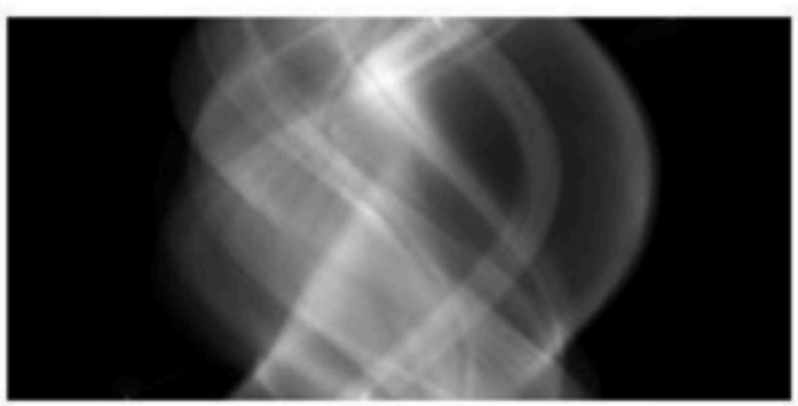

$\tilde{x}_0 = \tilde{H}_0^{-1}(y)$

“back-projection”

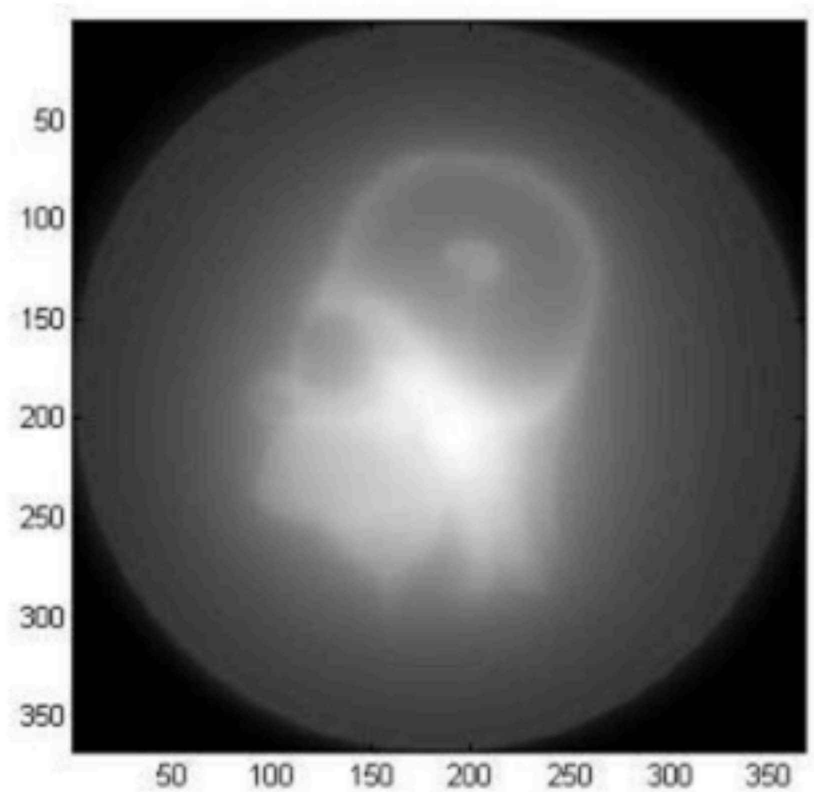

$\tilde{x} = \tilde{H}^{-1}(y)$

“filtered back-projection”

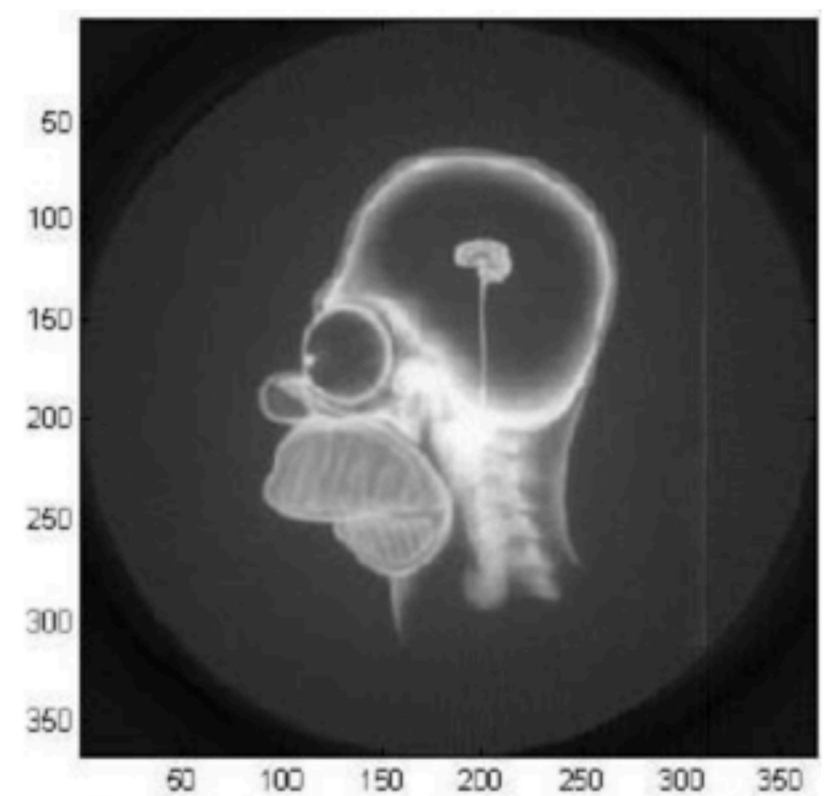

Image source: https://sites.google.com/site/fanyangspace111/cone-beam-tomographyct-reconstruction-using-filtered-back-projection

Conventional reconstructions

approximate inverse

$$\tilde{H}^{-1}(y) \quad = \quad \tilde{x}$$

measurements **reconstructed image**

- Need fully-sampled data to get (#measurements = #pixels)
- **Goal:** take fewer measurements (undersample) to speed up acquisition (in MRI), reduce dose (in CT)
- What happens to conventional reconstructions when we undersample?

Undersampling in MRI

sample locations in k-space | under-sampled recon | fully-sampled recon

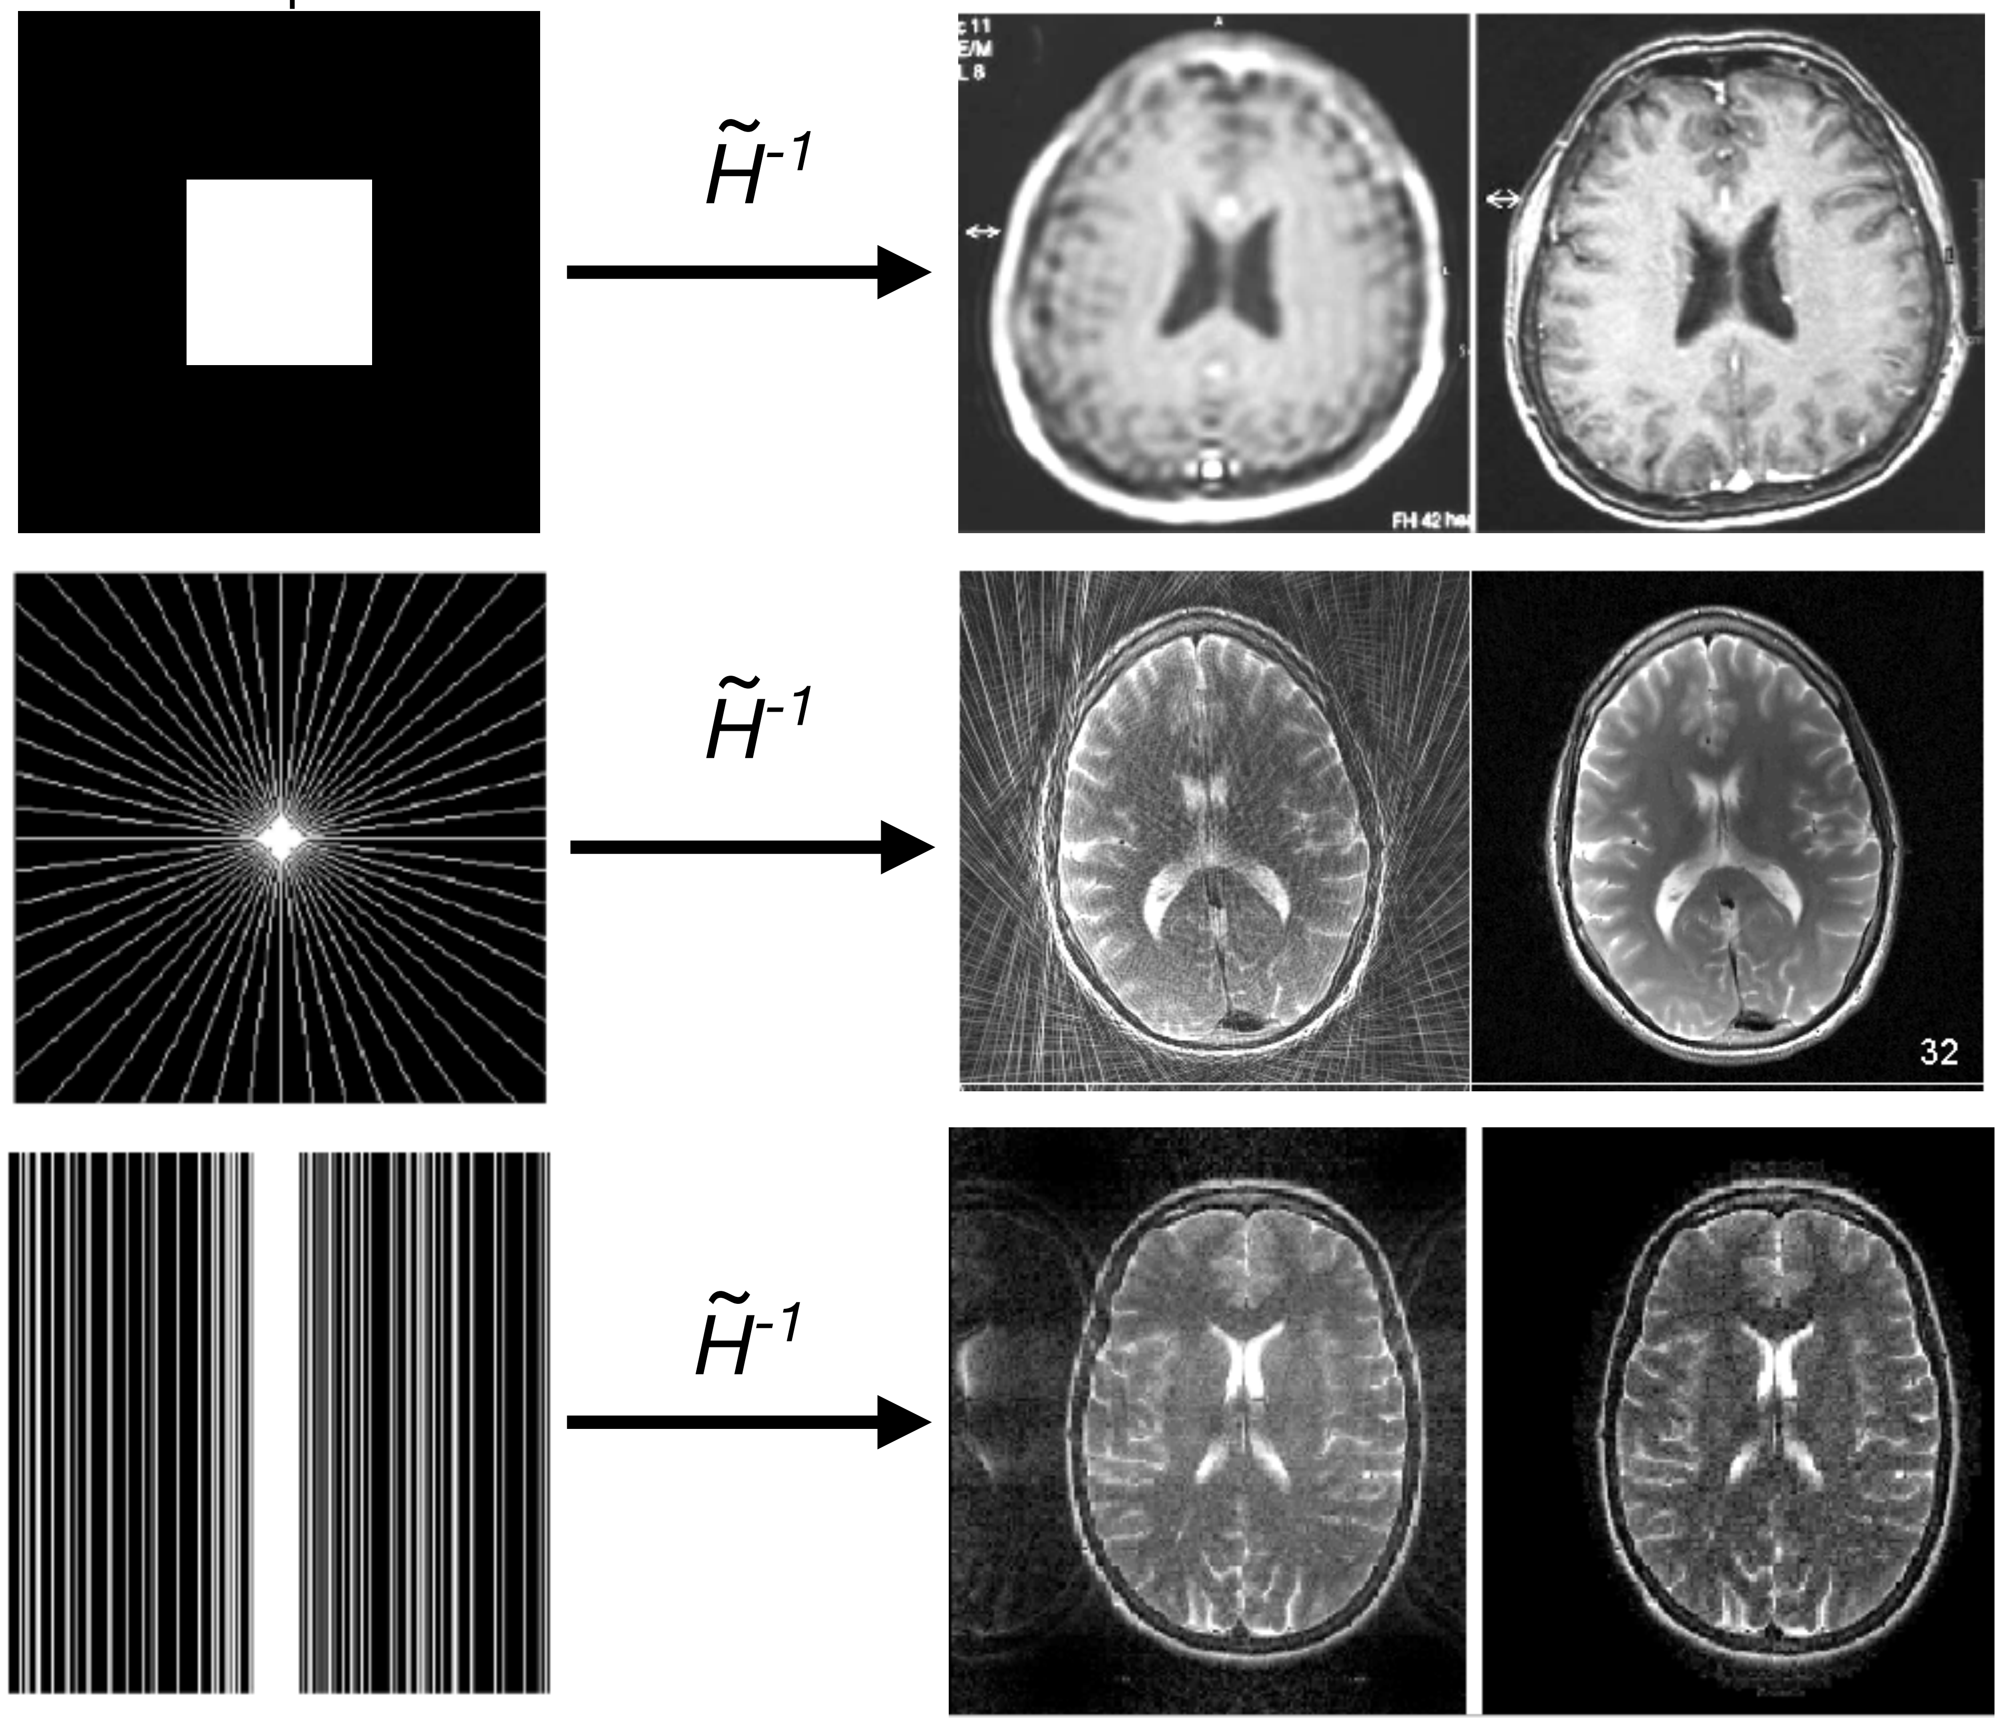

Undersampling in CT

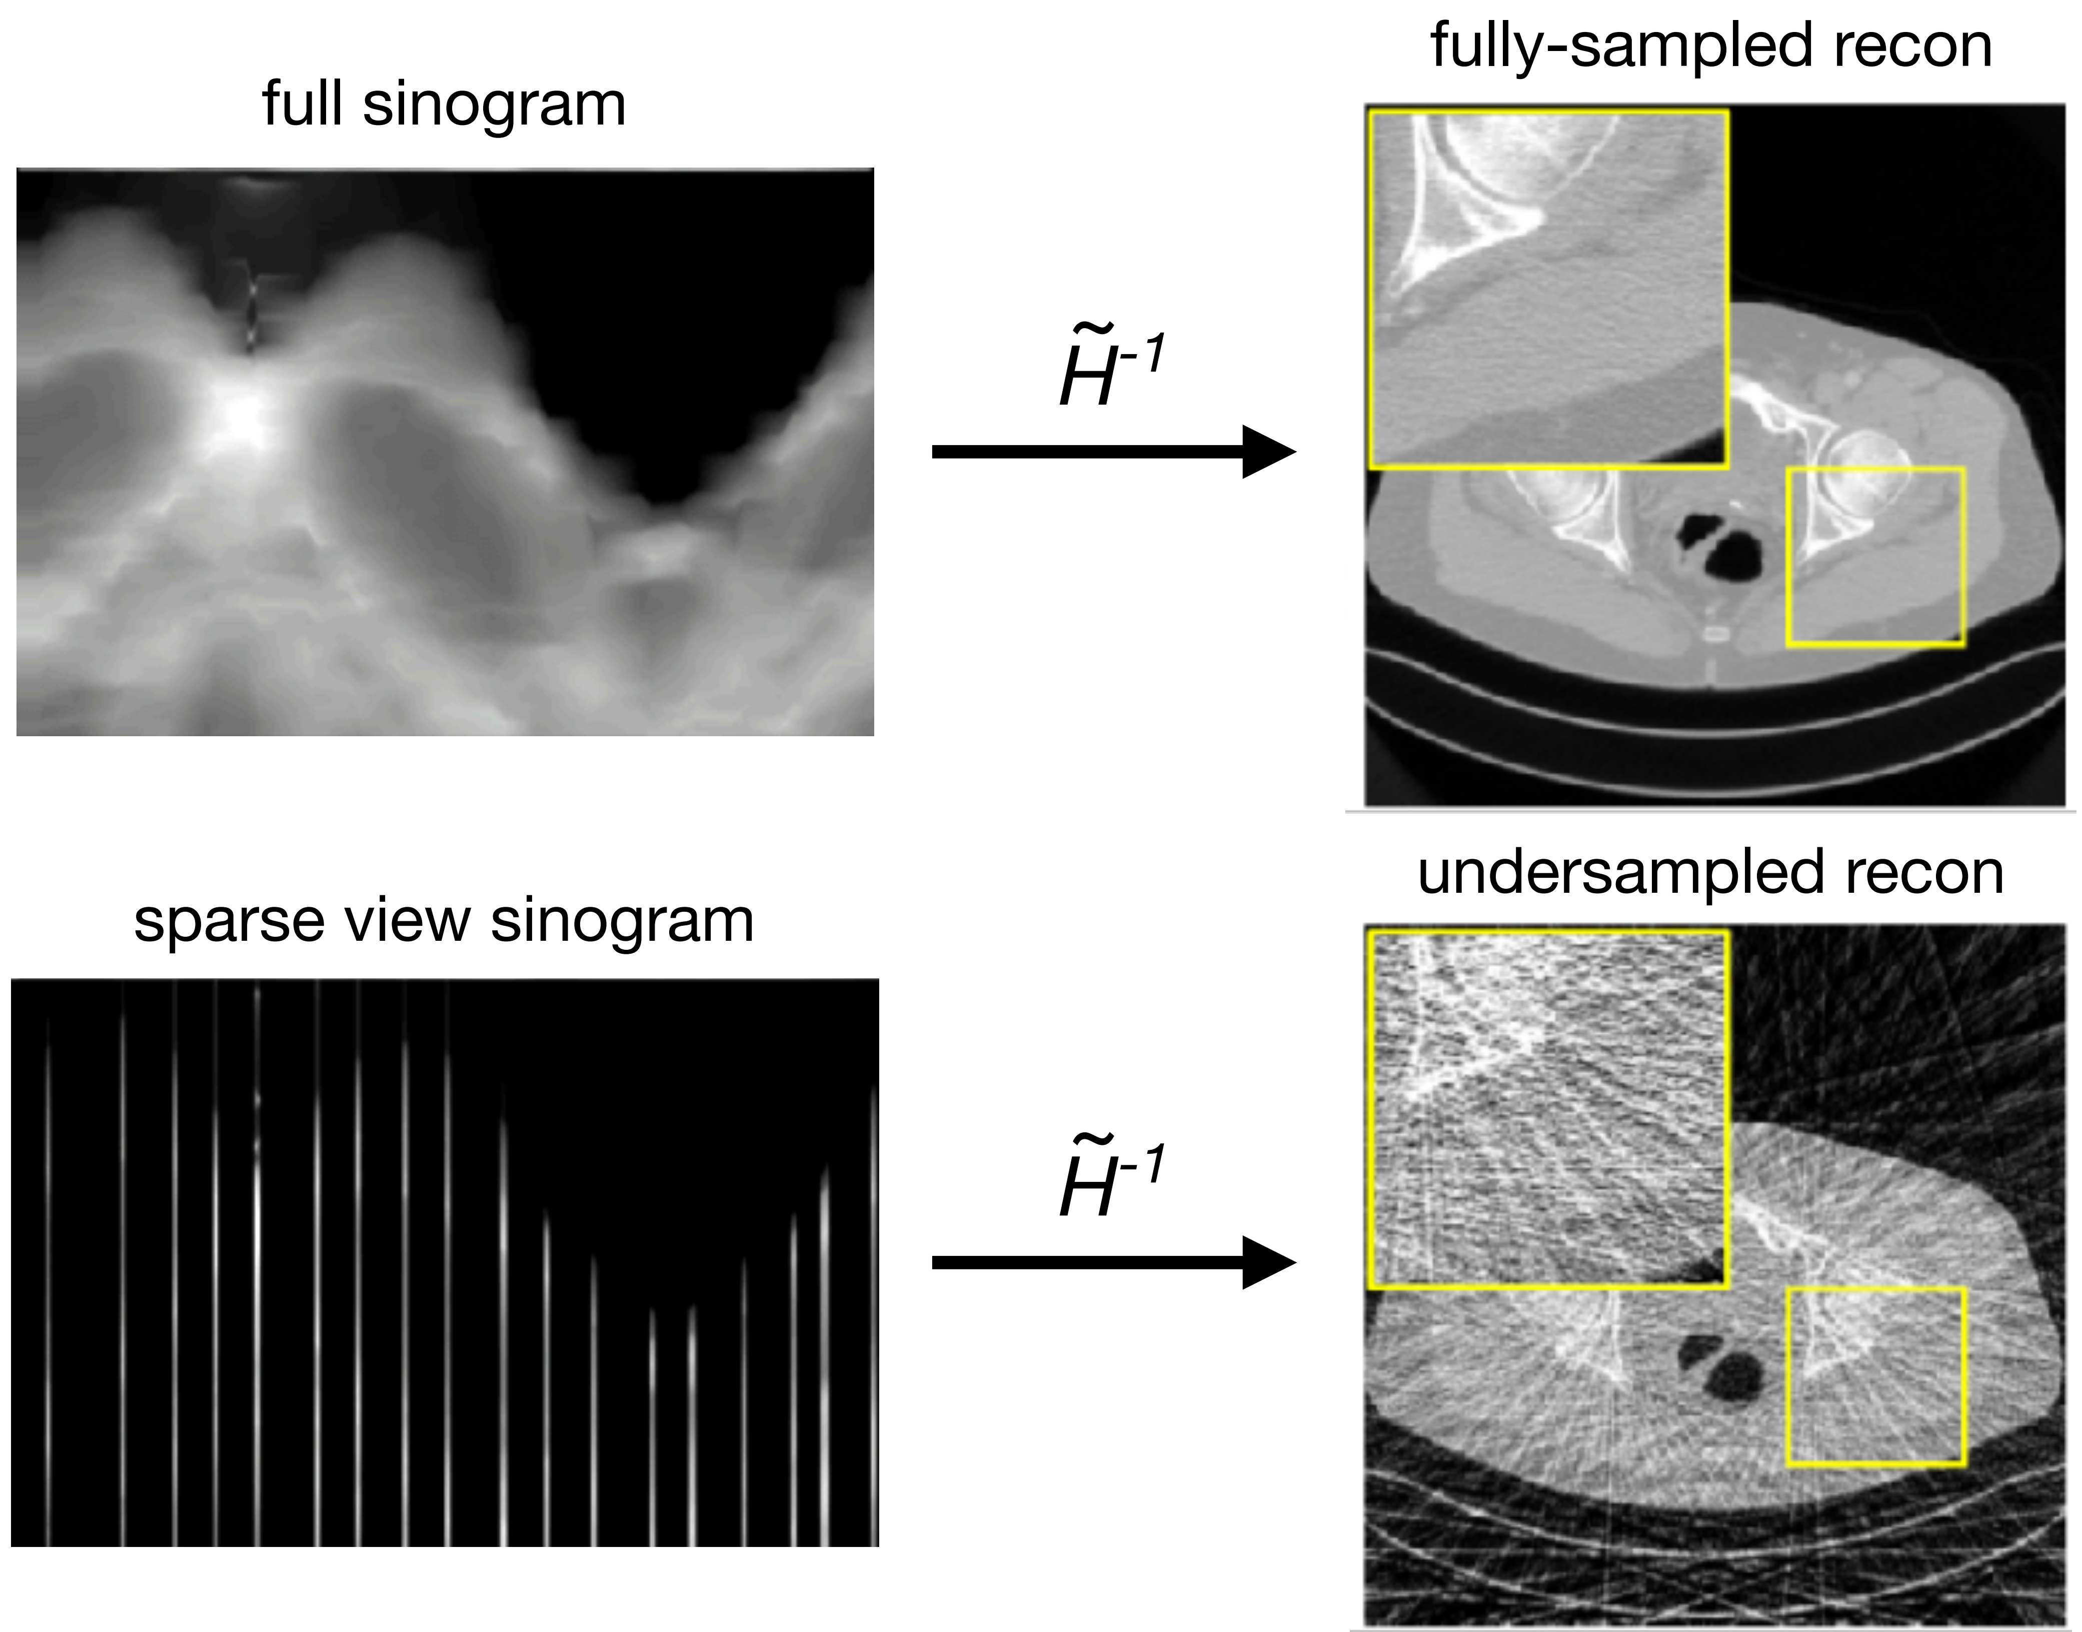

Undersampling

linear measurement operator

$$y = H(x)$$

measurements **image**

Write as matrix equation:

measurement matrix

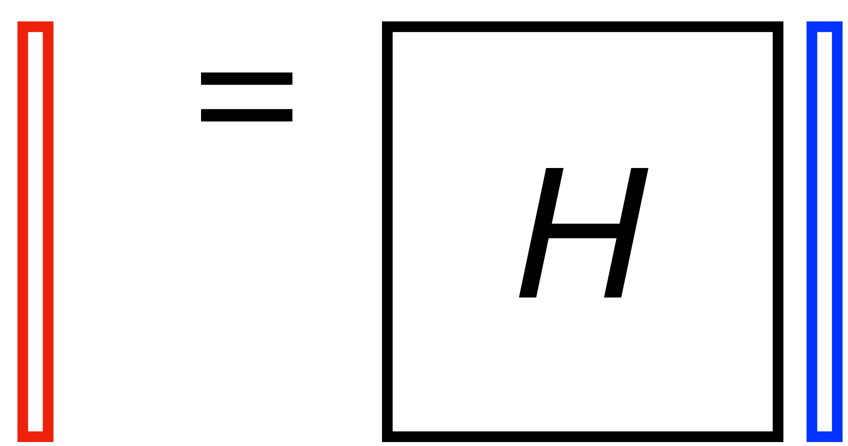

vectorized measurements

vectorized image (unknowns)

Fully-sampled

Same # of equations as unknowns

Undersampling

linear measurement operator

$$y = H(x)$$

measurements **image**

Write as matrix equation:

measurement matrix

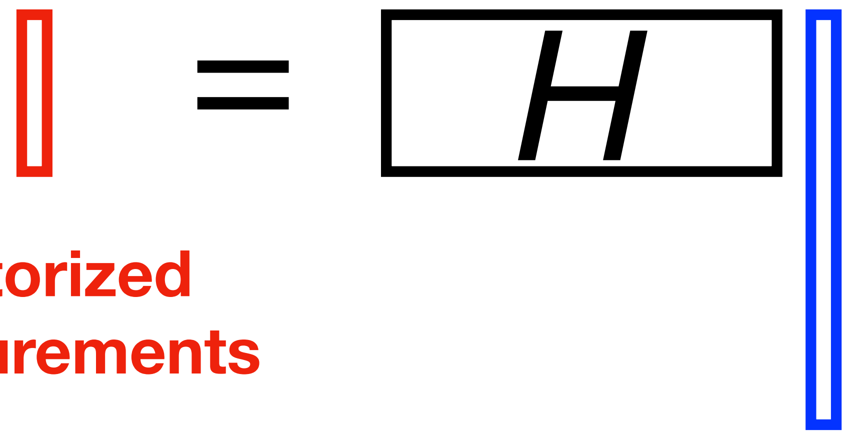

vectorized measurements

vectorized image (unknowns)

Undersampled

fewer equations than unknowns

Infinitely many solutions!

Compressed sensing - 2006

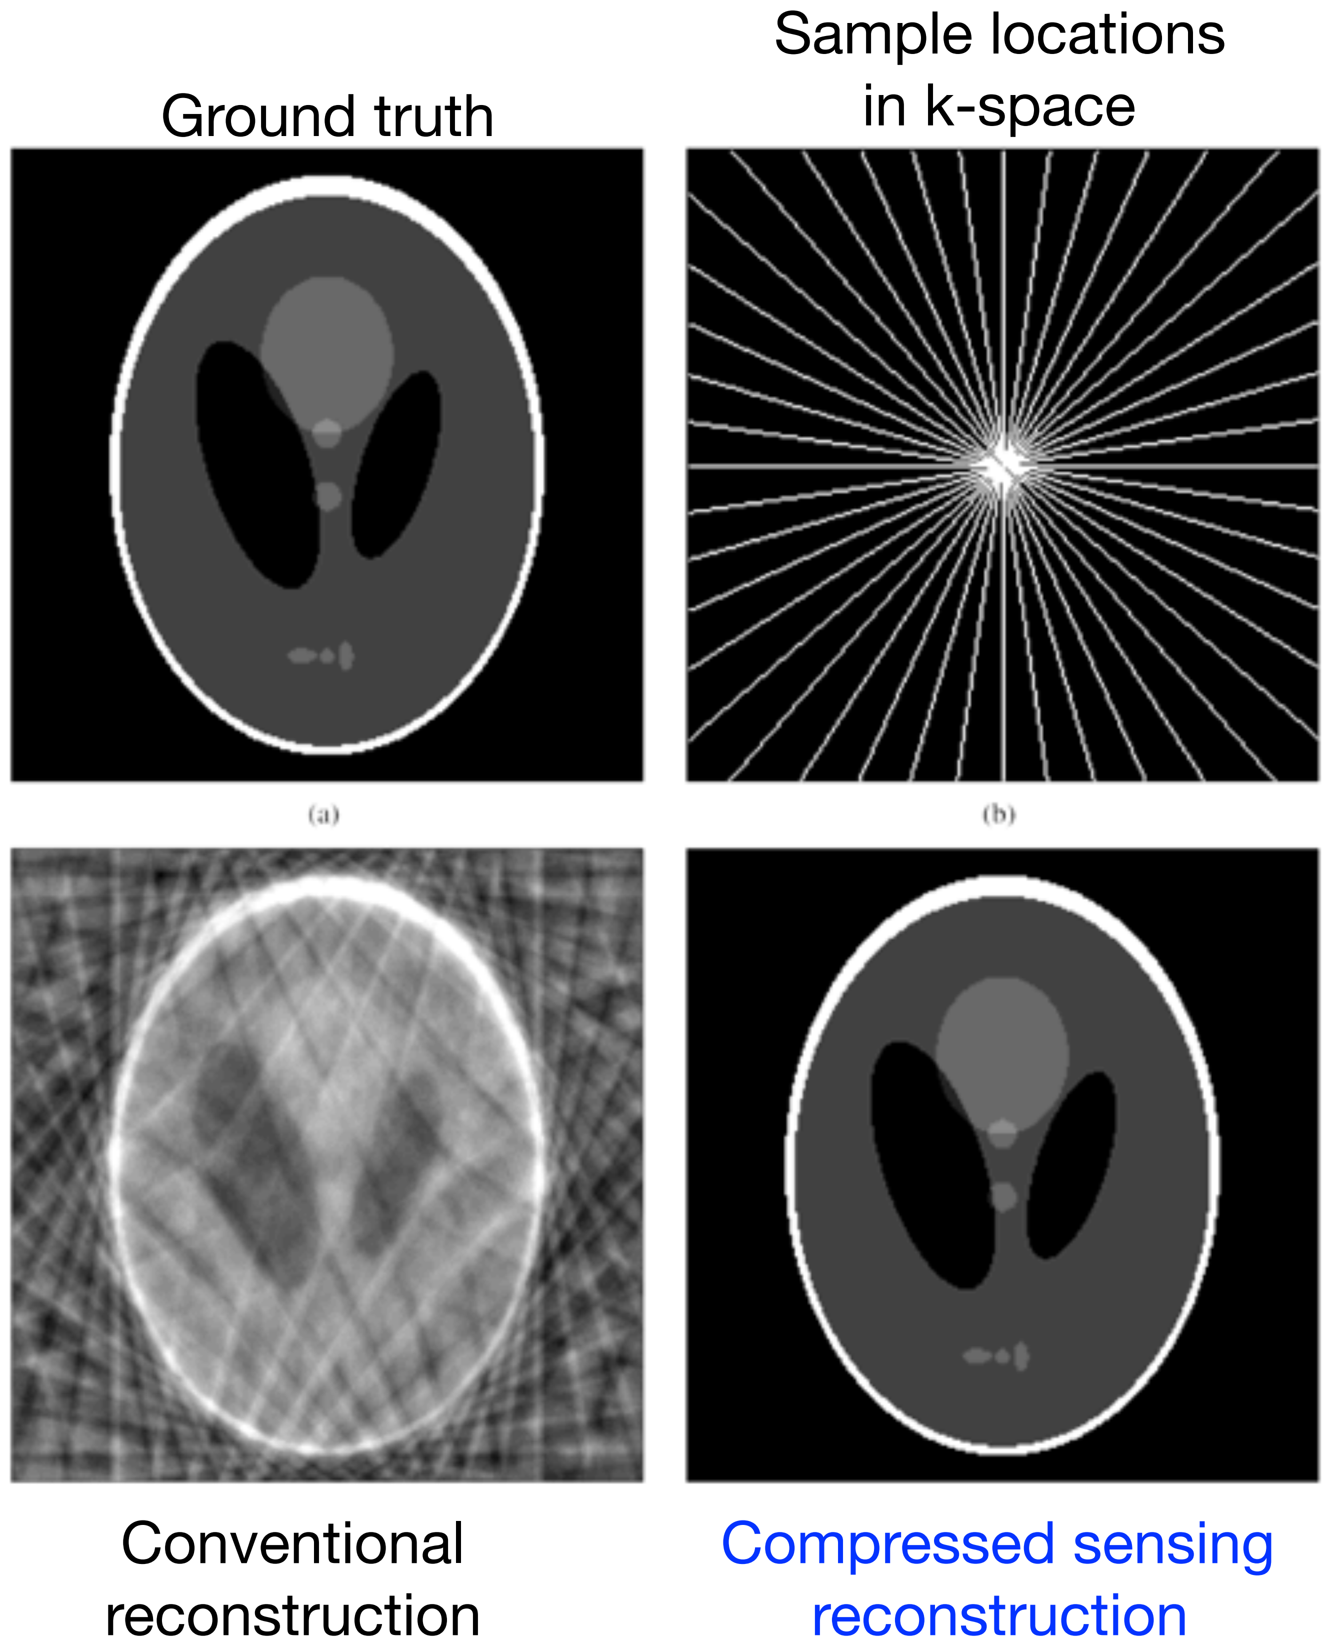

Exact reconstruction!

Candes, Romberg, Tao 2006

Compressed Sensing Reconstruction

Pose reconstruction as an optimization problem:

$$\underset{x}{\text{minimize}} \underbrace{\|Hx - y\|^2}_{\text{data-fit term}} + \underbrace{r(x)}_{\text{regularizer}}$$

Typically ***r(x)*** is chosen to promote sparsity of the image in some domain

e.g.,
Wavelet
sparsity

Wx →

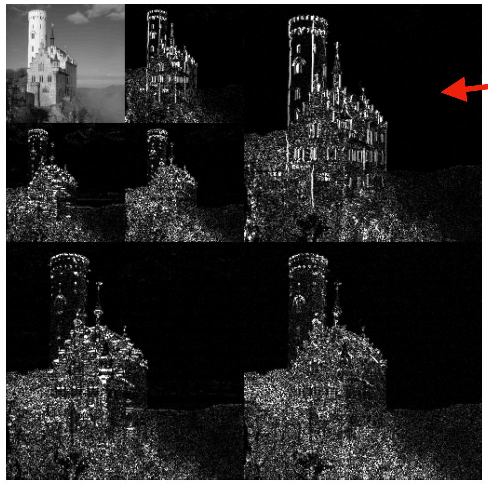

coefficients
mostly
zero

r(x) = ||Wx||$_1$

Figure by Alessio Damato, https://en.wikipedia.org/wiki/Wavelet_transform

Solve by an iterative method, e.g., **gradient descent**

Computationally costly: ~100x slower than conventional reconstruction

Compressed Sensing Dynamic MRI

Free-Breathing Liver Perfusion Imaging

- Retrospective selection of resolution
- Reconstruction with different timing possible

- Example: 13 spokes → 2 sec resolution
- Perfusion imaging during free breathing

- Here: 384 x 384 x 30 matrix
- Spatial resolution 1.0 x 1.0 x 3.0 mm^3
- Temporal resolution 1.5 sec

Recon time: ~6 hours

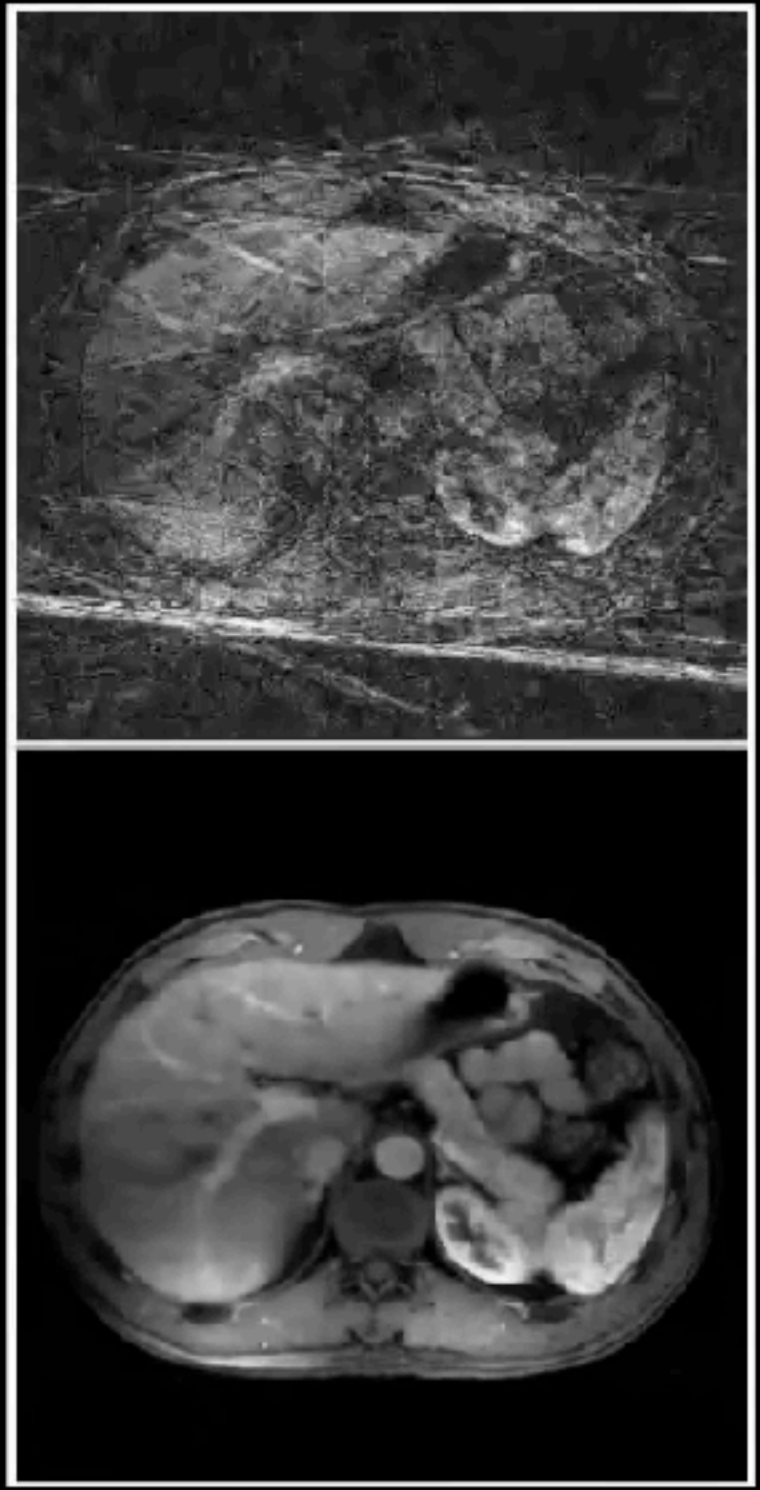

Top: Gridding
Bottom: GRASP

NYU Langone MEDICAL CENTER

Chandarana et al, ISMRM 2012: 5529

The Truth About Compressed Sensing

"In the literature, a lack of translation to final users is presently discernable: while there are over 120 papers about compressed sensing in MRI published in Magnetic Resonance in Medicine, there are only 8 papers in Radiology.

…

it is essential for the radiographer to get image feedback within seconds of the scan terminating for accelerated imaging to be practically useful."

Quote from [Hollingsworth, 2015]

- Hollingsworth, K.G., 2015. Reducing acquisition time in clinical MRI by data undersampling and compressed sensing reconstruction. Phys. Med, Biol. 60(21), p.R297.

ML to the rescue

Optimization algorithm

$$y \longrightarrow \boxed{\arg\min_{\beta} \|y - X\beta\|_2^2 + r(\beta)} \longrightarrow \hat{\beta}$$

Feed-forward deep neural network

$y \longrightarrow$

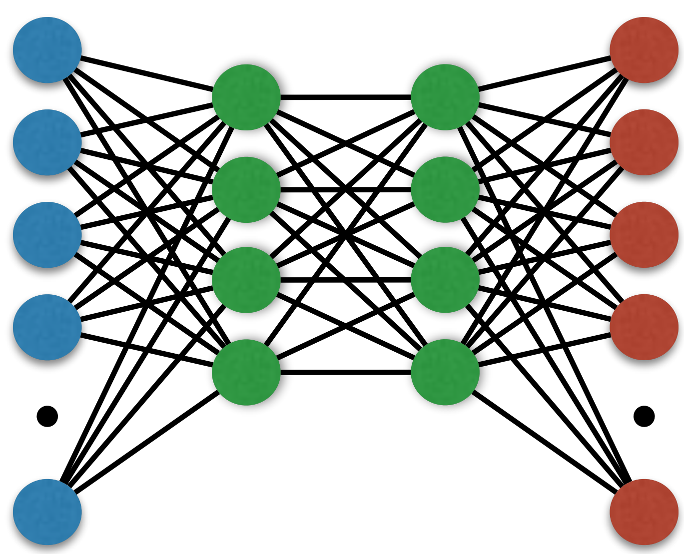

$\longrightarrow \hat{\beta}$

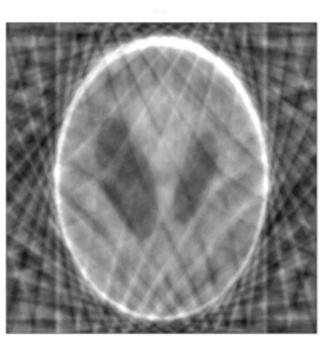

(c)

Learn from training pairs

(d)

Deep learning
for image reconstruction

Approach 1: Learn to “enhance” traditional reconstructions
Approach 2: Train a generative model
Approach 3: Unrolling of optimization algorithms

Approach 1: Learn to "enhance" traditional reconstructions

“Enhancing” with Deep Learning

Single Image Super-resolution

Low-resolution Input

Deep Neural Network Output

Ground Truth

https://webdav.tue.mpg.de/pixel/enhancenet/

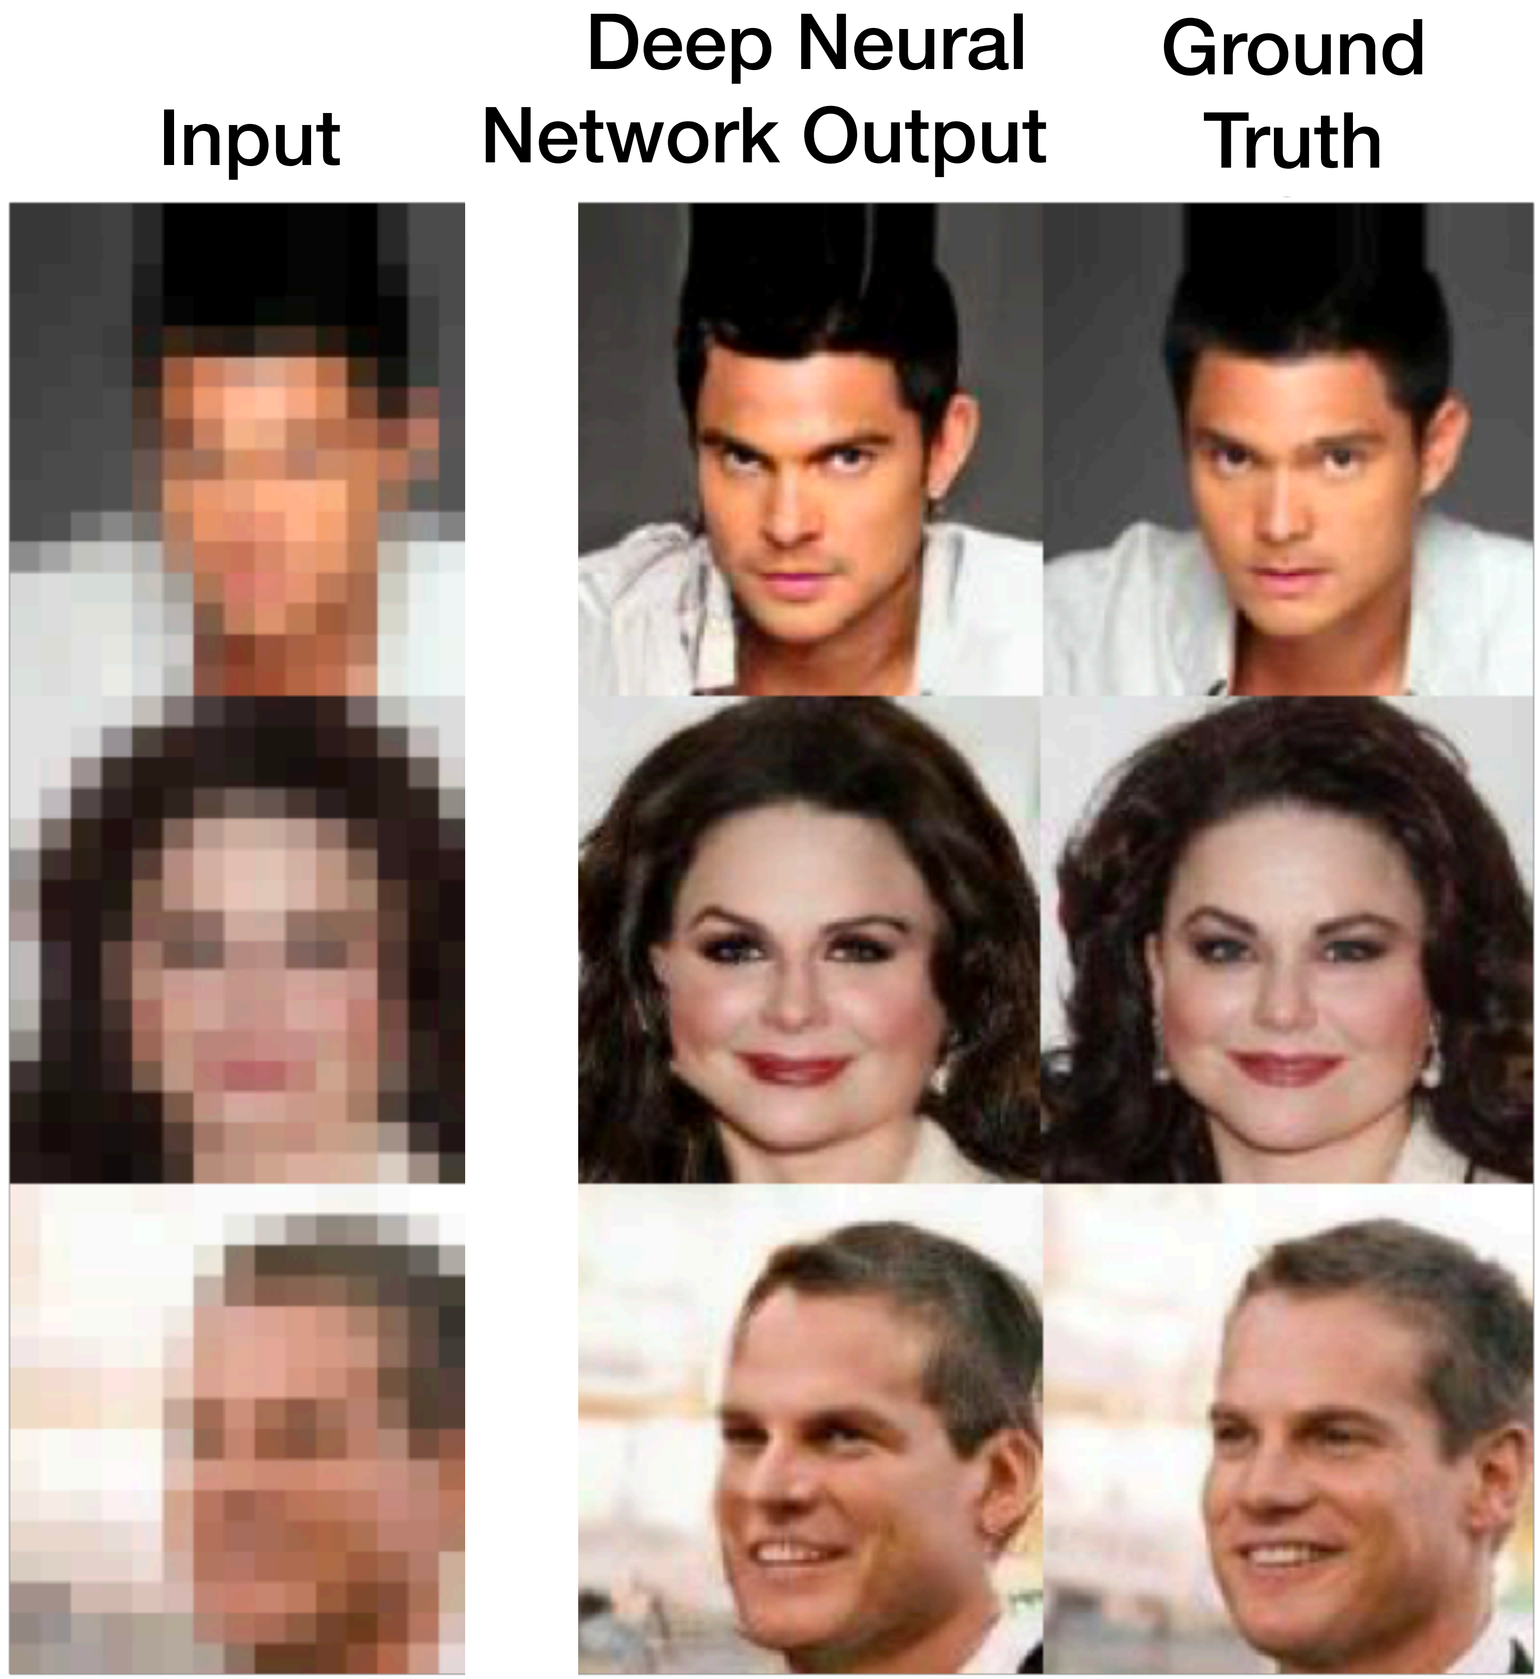

Progressive Face Super-Resolution via Attention to Facial Landmark

Deokyun Kim, Minseon Kim, Gihyun Kwon, Dae-Shik Kim

(Submitted on 22 Aug 2019)

How to use CNN's to "enhance" images?

Most existing CNNs are designed for classification tasks, use "max pooling" layers

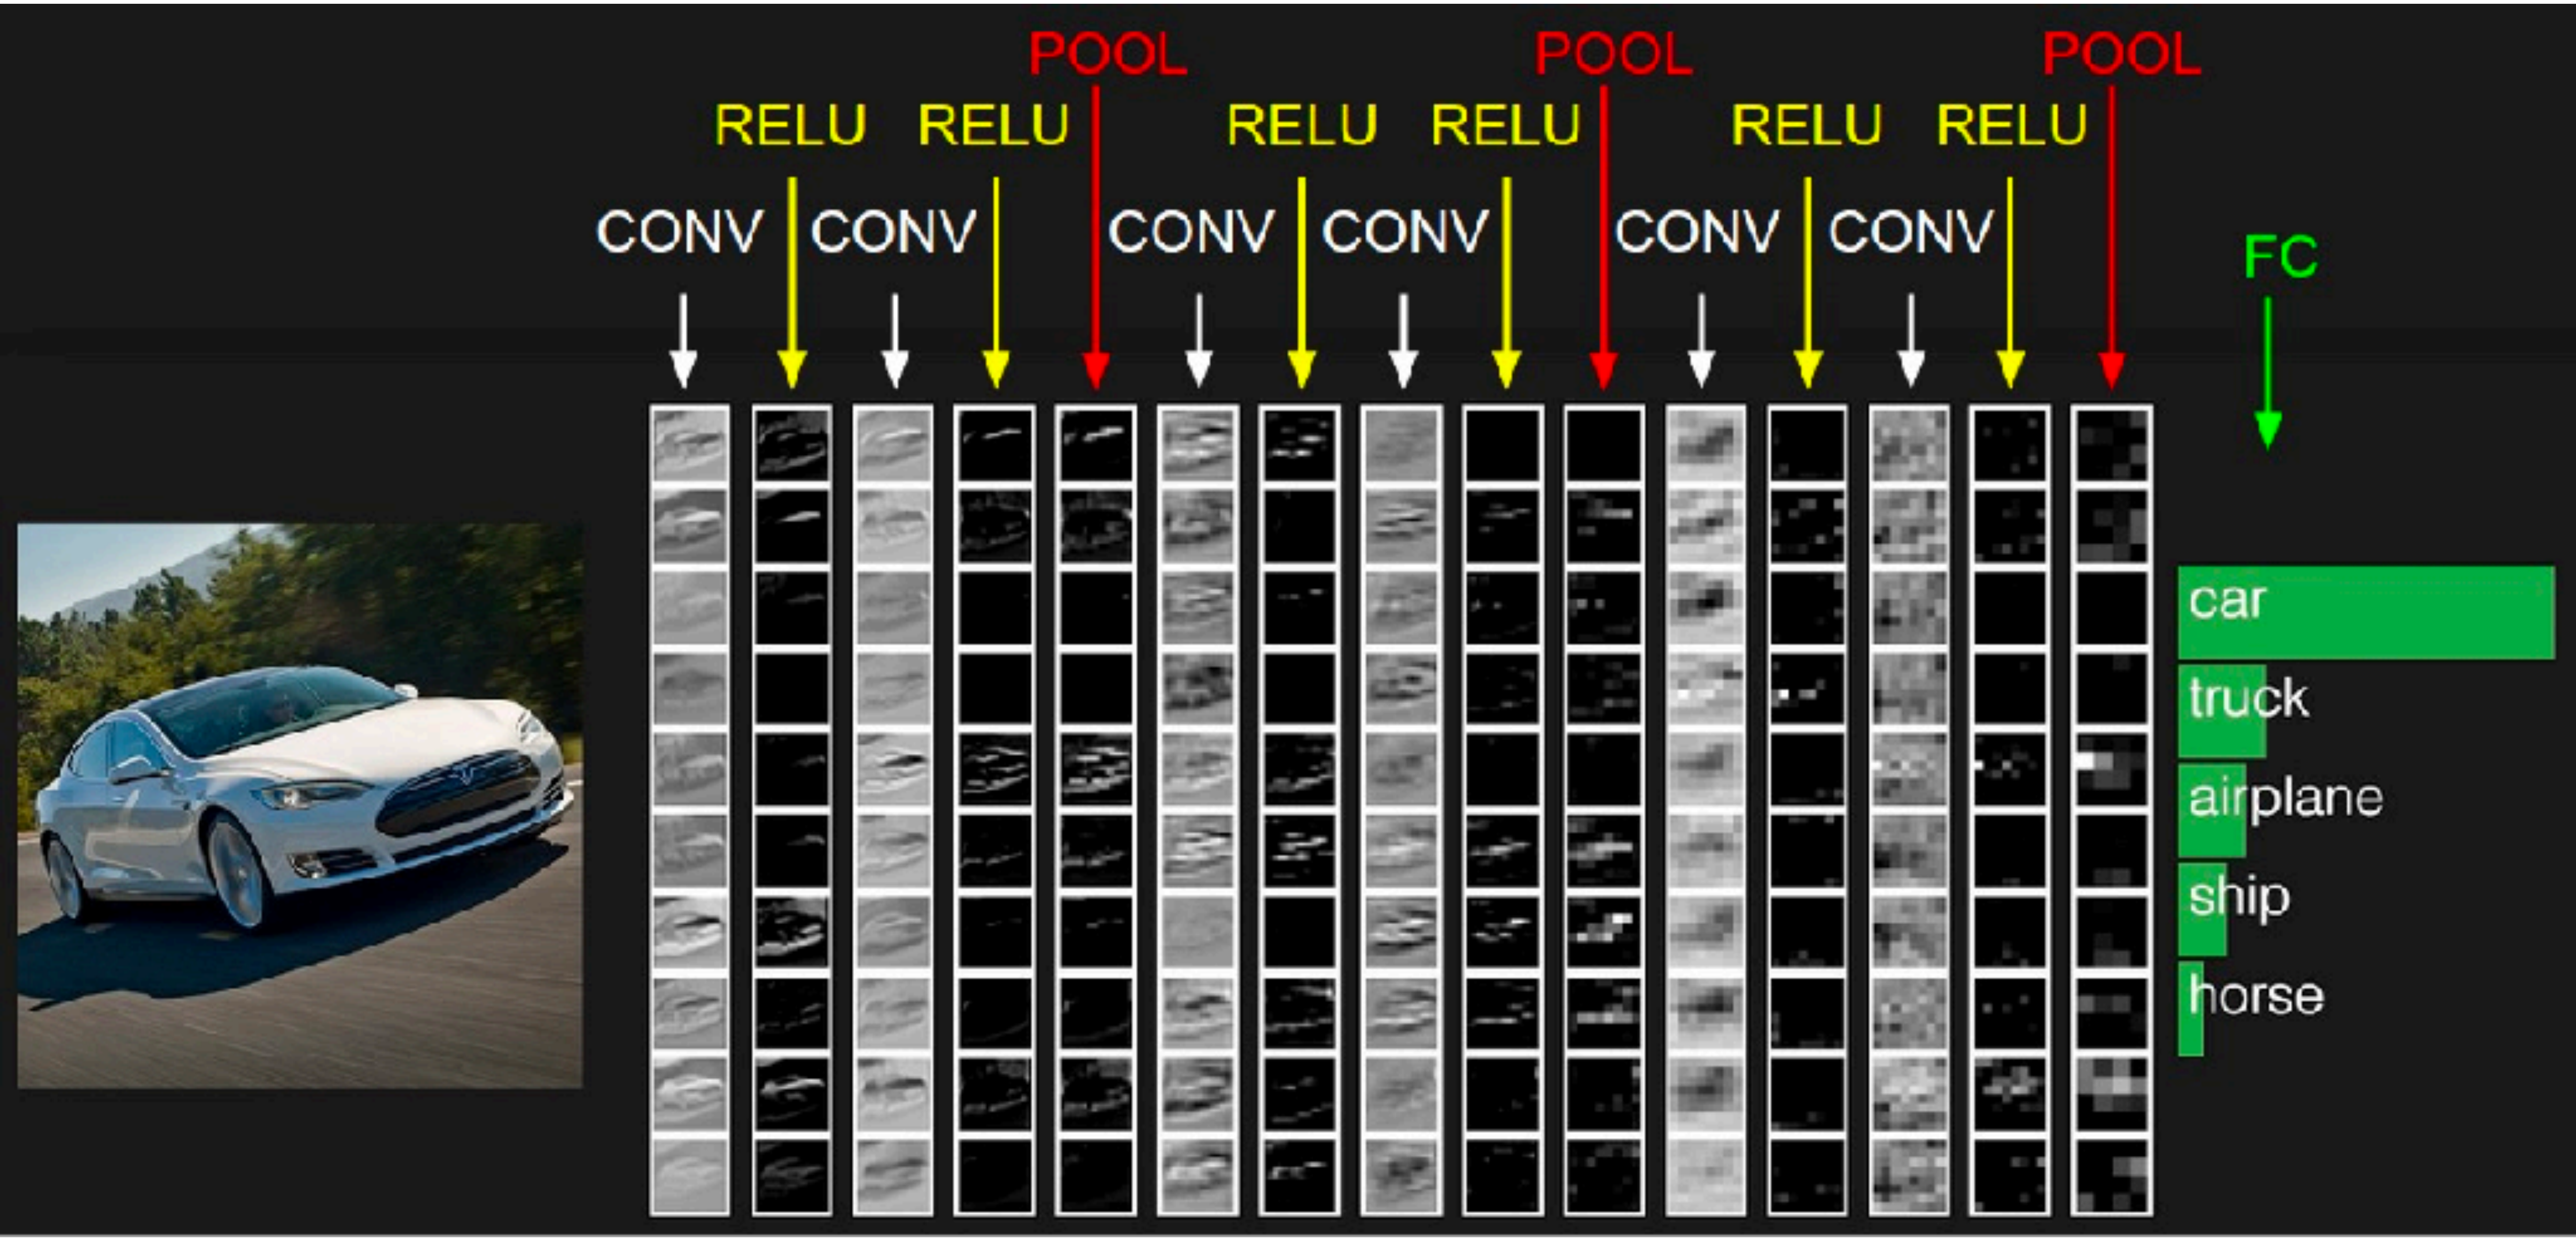

Super-resolution with a CNN

IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, VOL. 38, NO. 2, FEBRUARY 2016 295

Image Super-Resolution Using Deep Convolutional Networks

Chao Dong, Chen Change Loy, *Member, IEEE*, Kaiming He, *Member, IEEE*, and Xiaoou Tang, *Fellow, IEEE*

3 Layer CNN

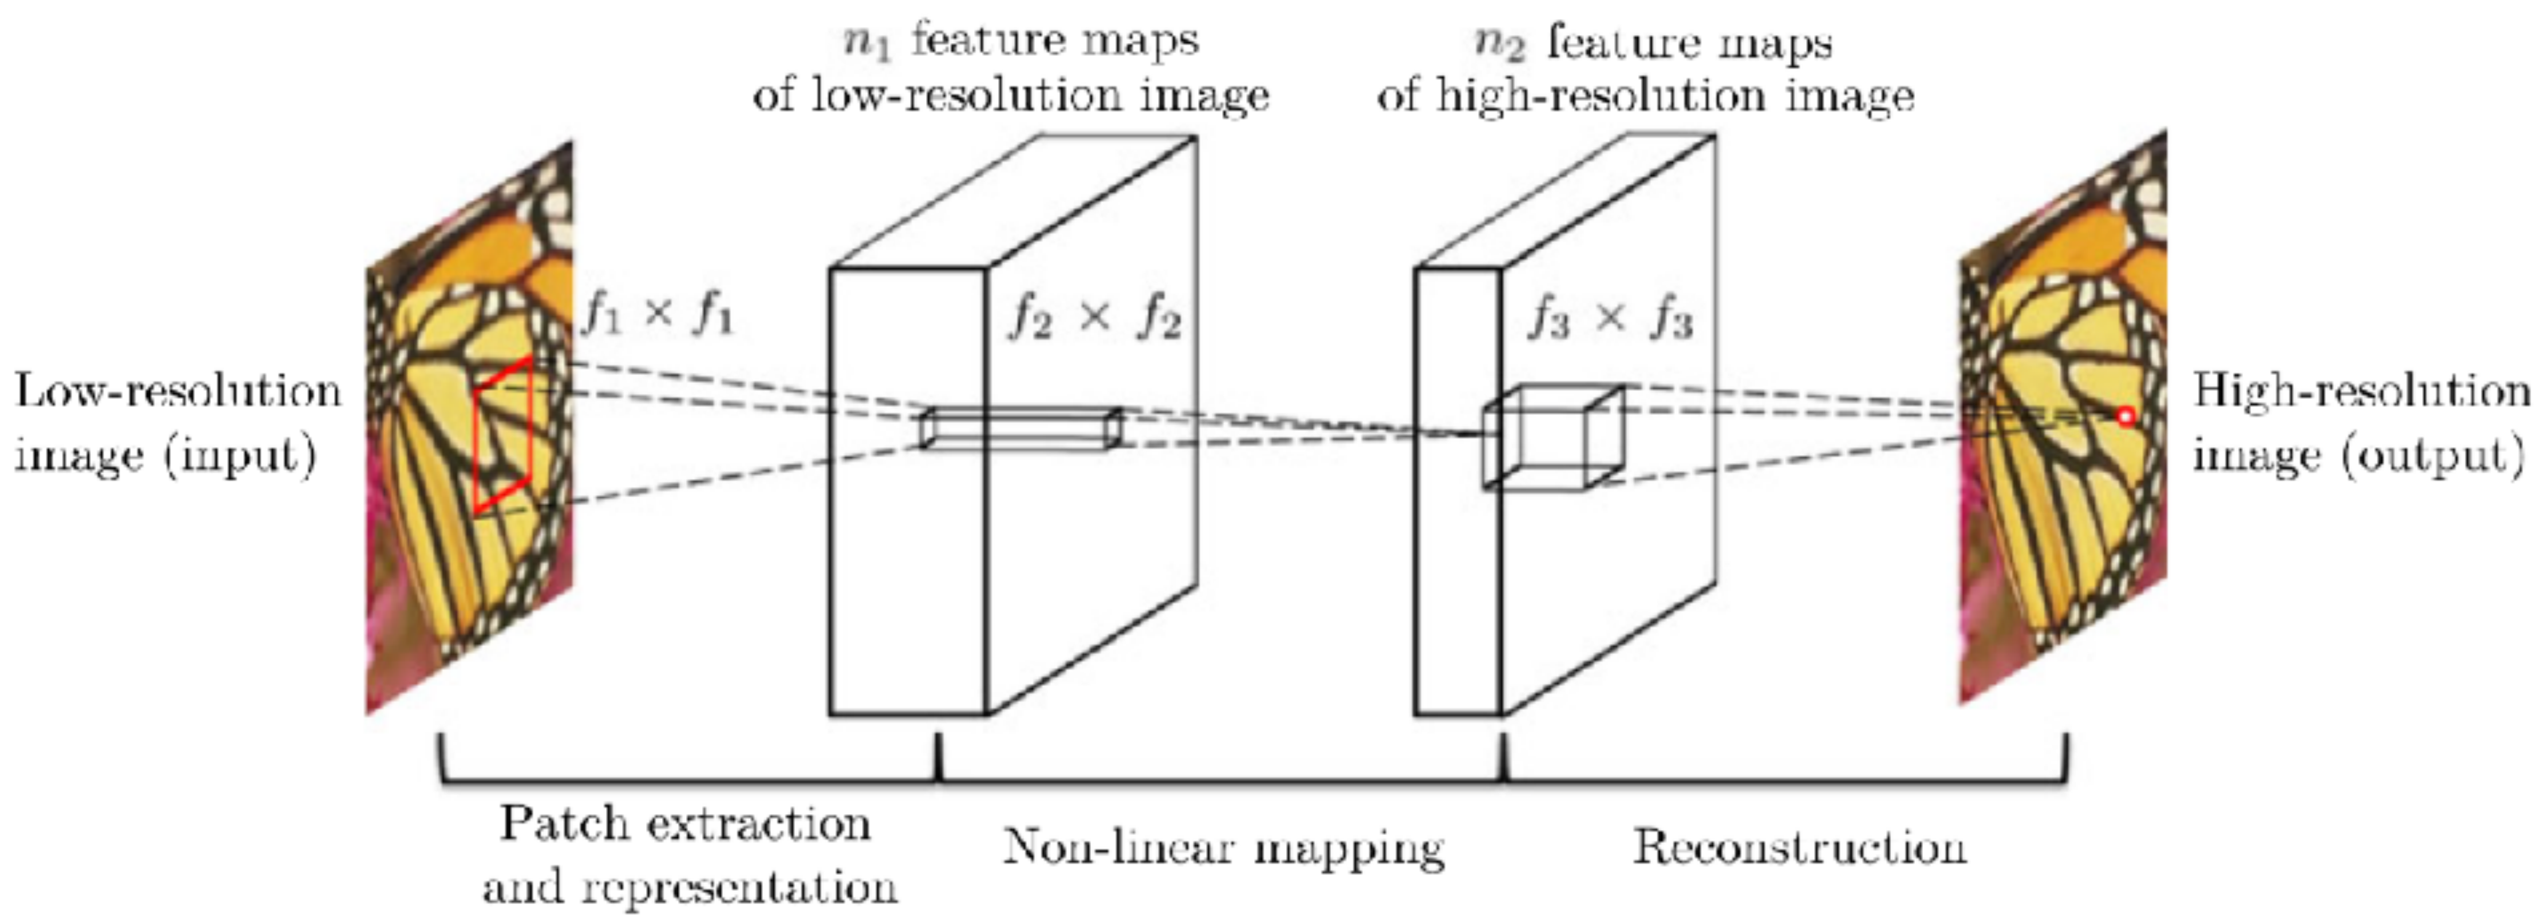

Super-resolution with CNNs

raw data (low resolution image)

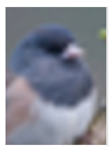

$\tilde{H}^{-1}$

bicubic interpolation

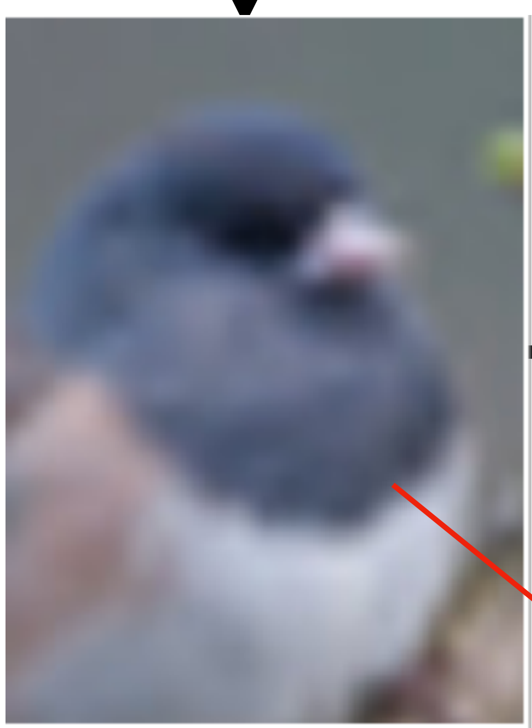

approximation of high-resolution image

blurry/blocky artifacts due to re-scaling

Train deep CNN to remove artifacts

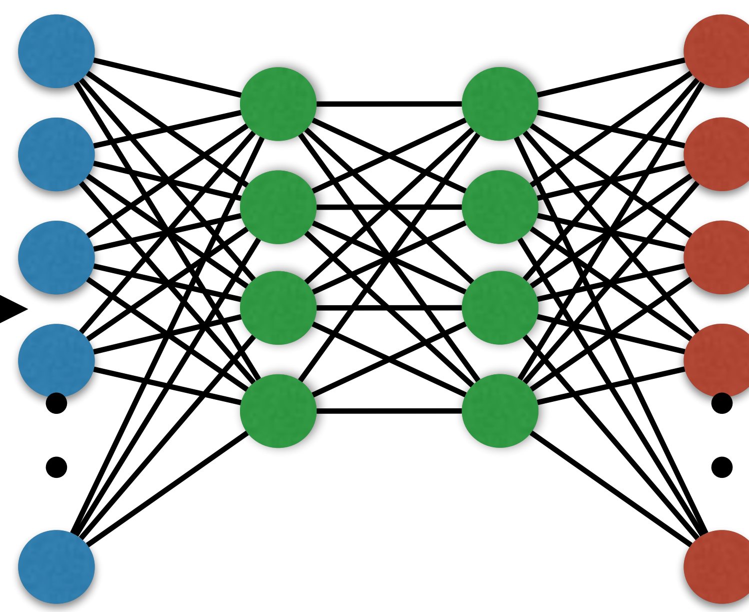

reconstruction

Pictures from: http://webdav.tuebingen.mpg.de/pixel/enhancenet/

Extension to medical imaging: CT reconstruction

raw data y
(sparse view sinogram)

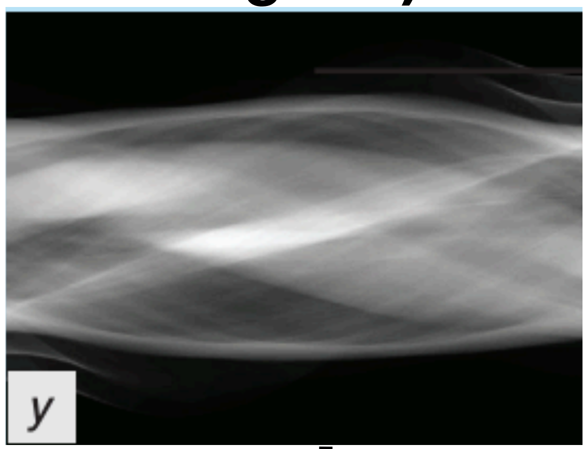

$\tilde{H}^{-1}$

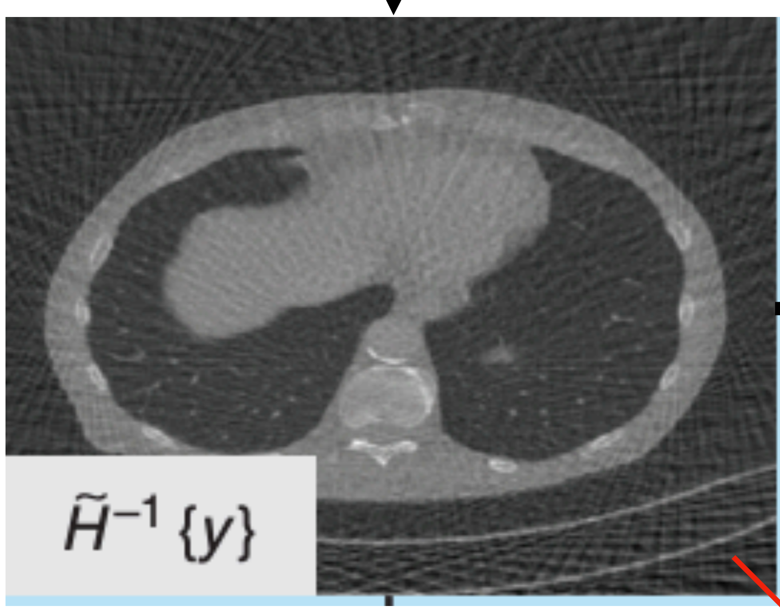

conventional reconstruction
(filtered back-projection)

streaking artifacts due to undersampling

Train deep CNN
to remove artifacts

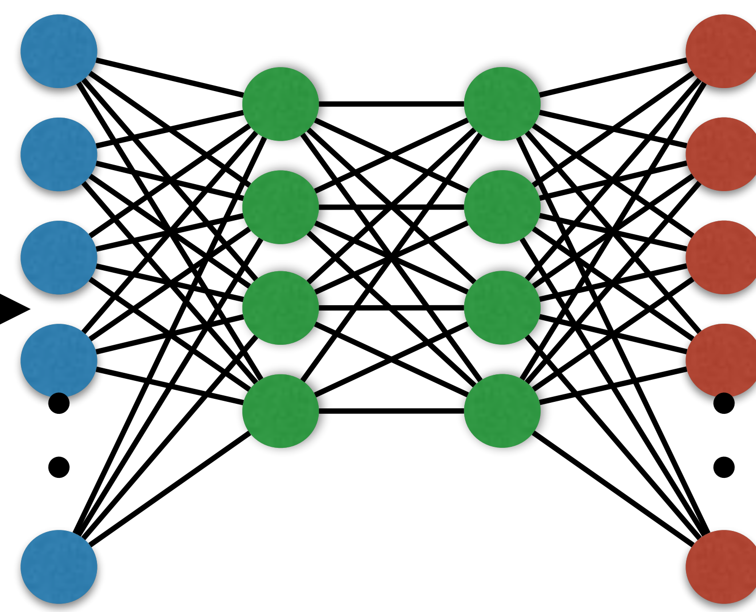

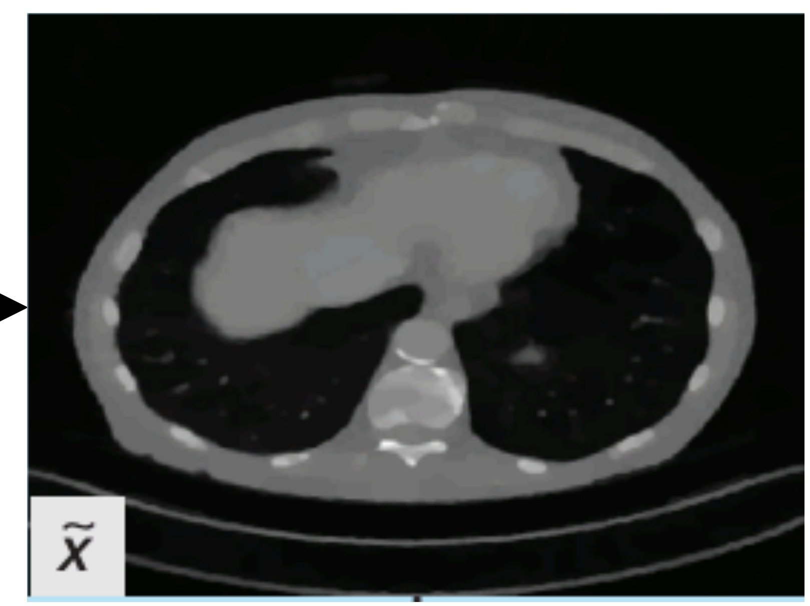

reconstruction

McCann, M. T., Jin, K. H., & Unser, M. (2017). Convolutional neural networks for inverse problems in imaging: A review. IEEE Signal Processing Magazine, 34(6), 85-95.

Example “Deep” CT Reconstruction

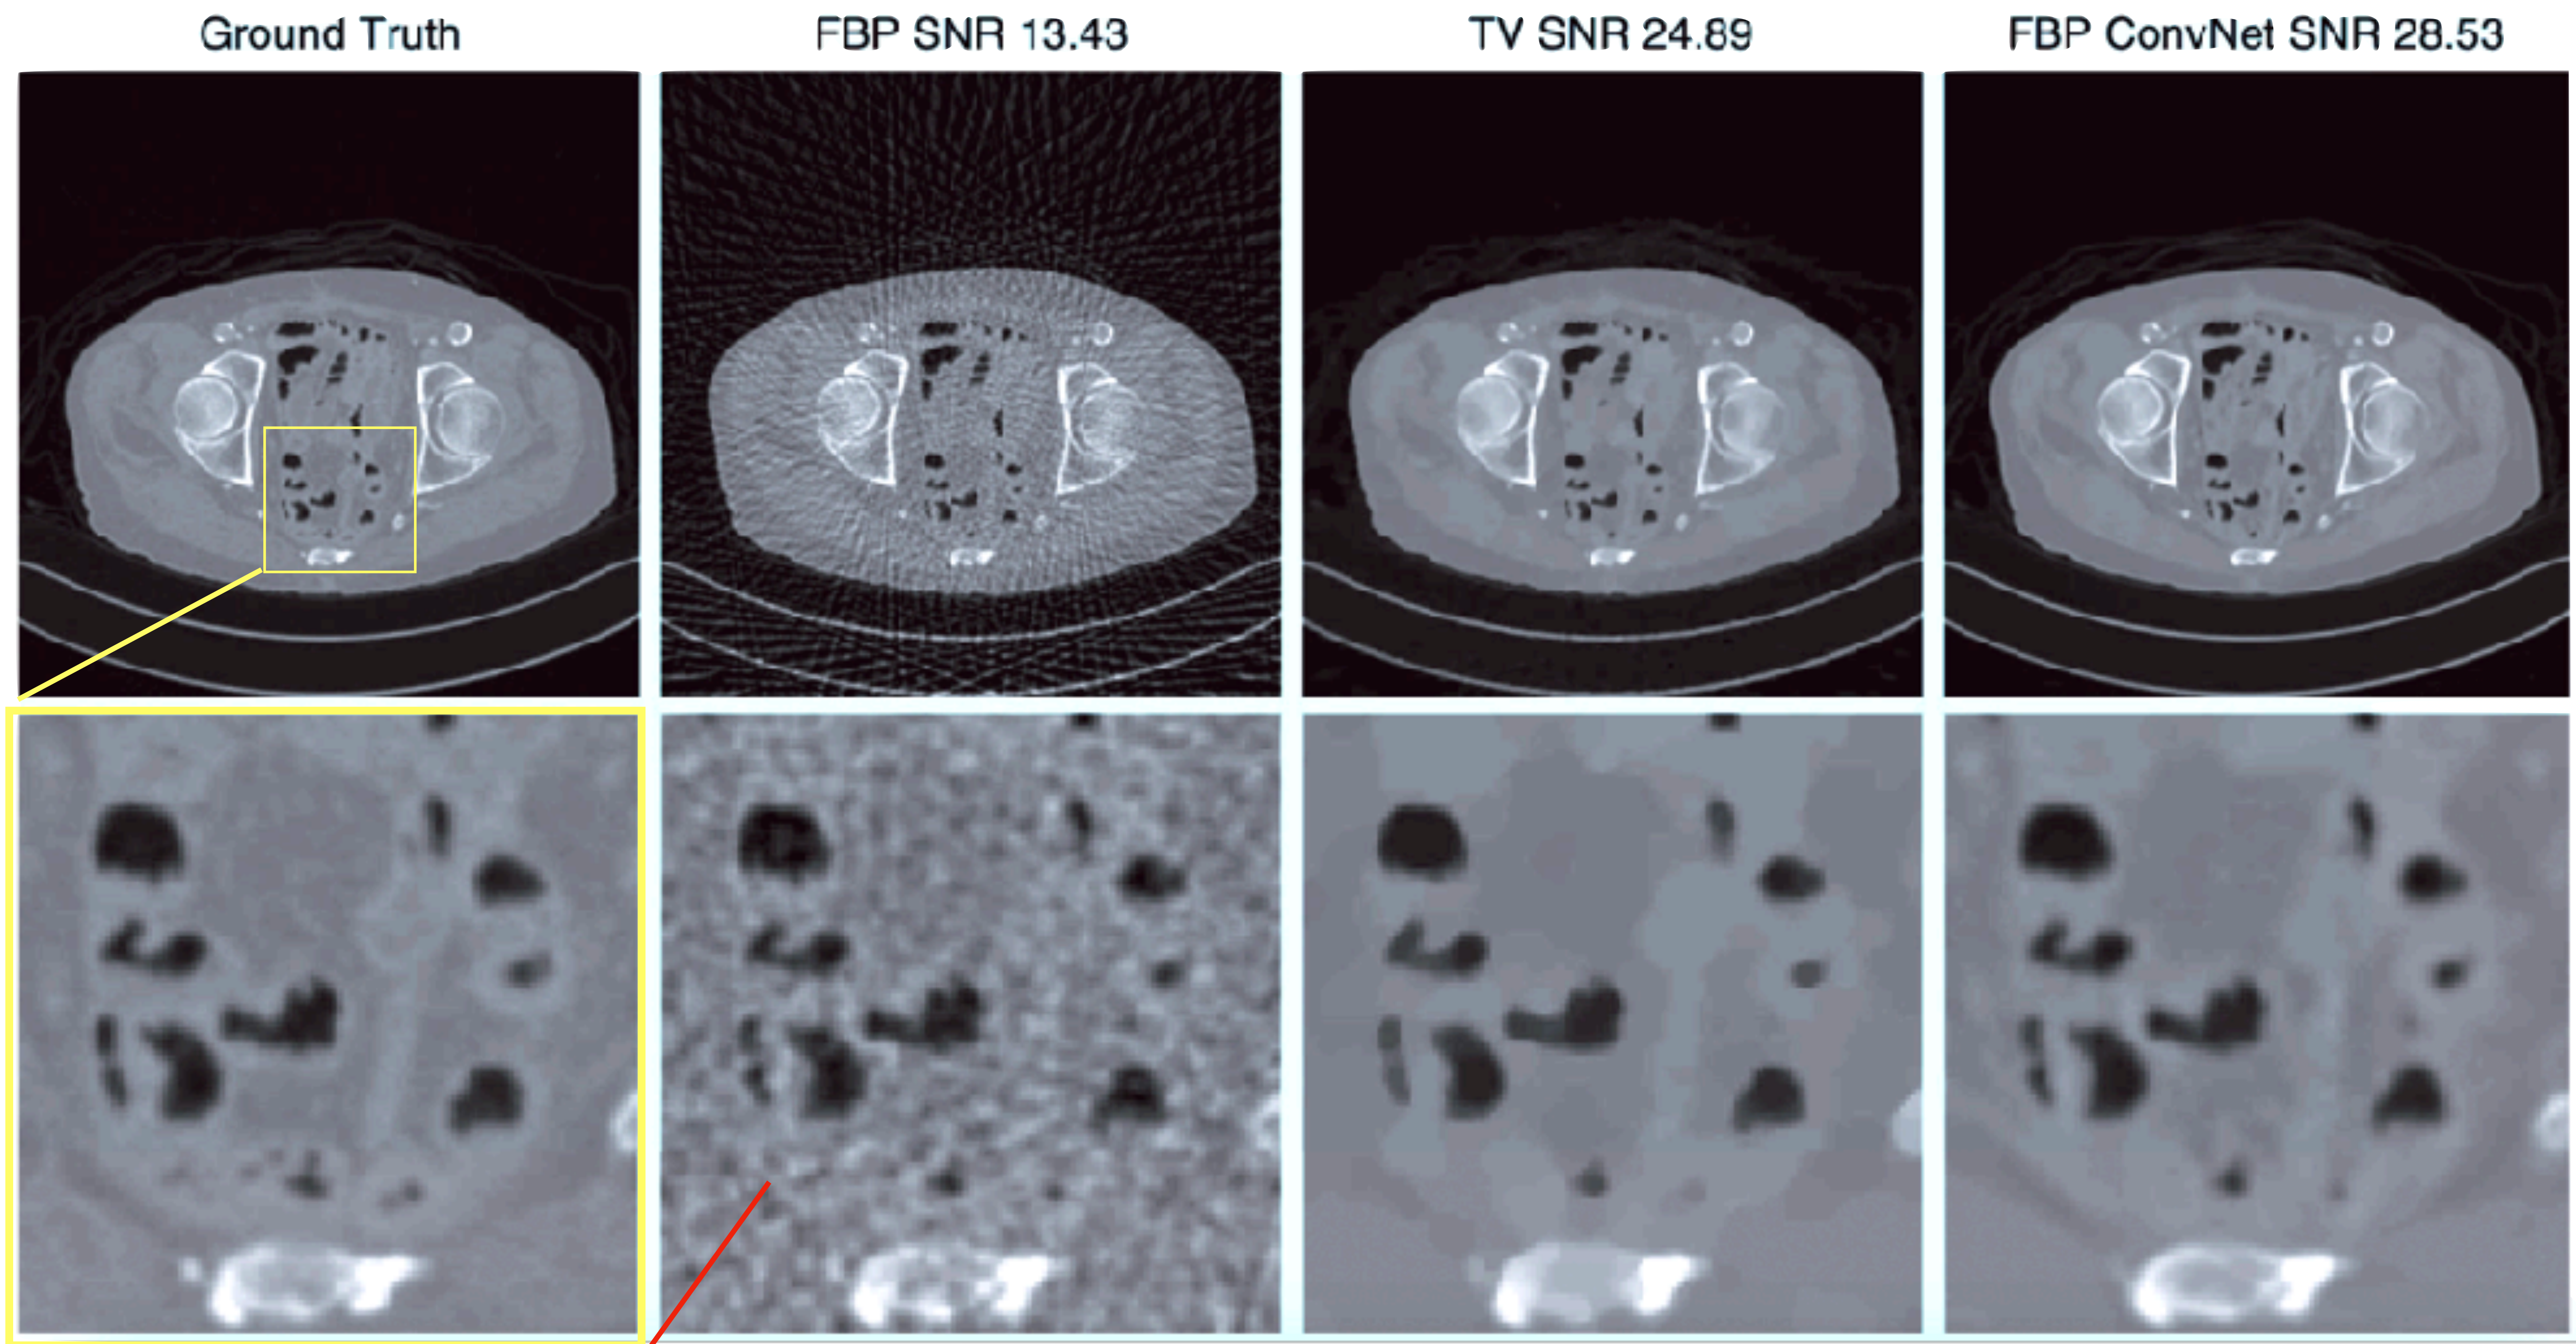

artifacts due to undersampling

conventional reconstruction

compressed sensing reconstruction (~60 s)

deep CNN reconstruction (~0.1 s)

McCann, M. T., Jin, K. H., & Unser, M. (2017). Convolutional neural networks for inverse problems in imaging: A review. IEEE Signal Processing Magazine, 34(6), 85-95.

Example “Deep” MRI Reconstruction

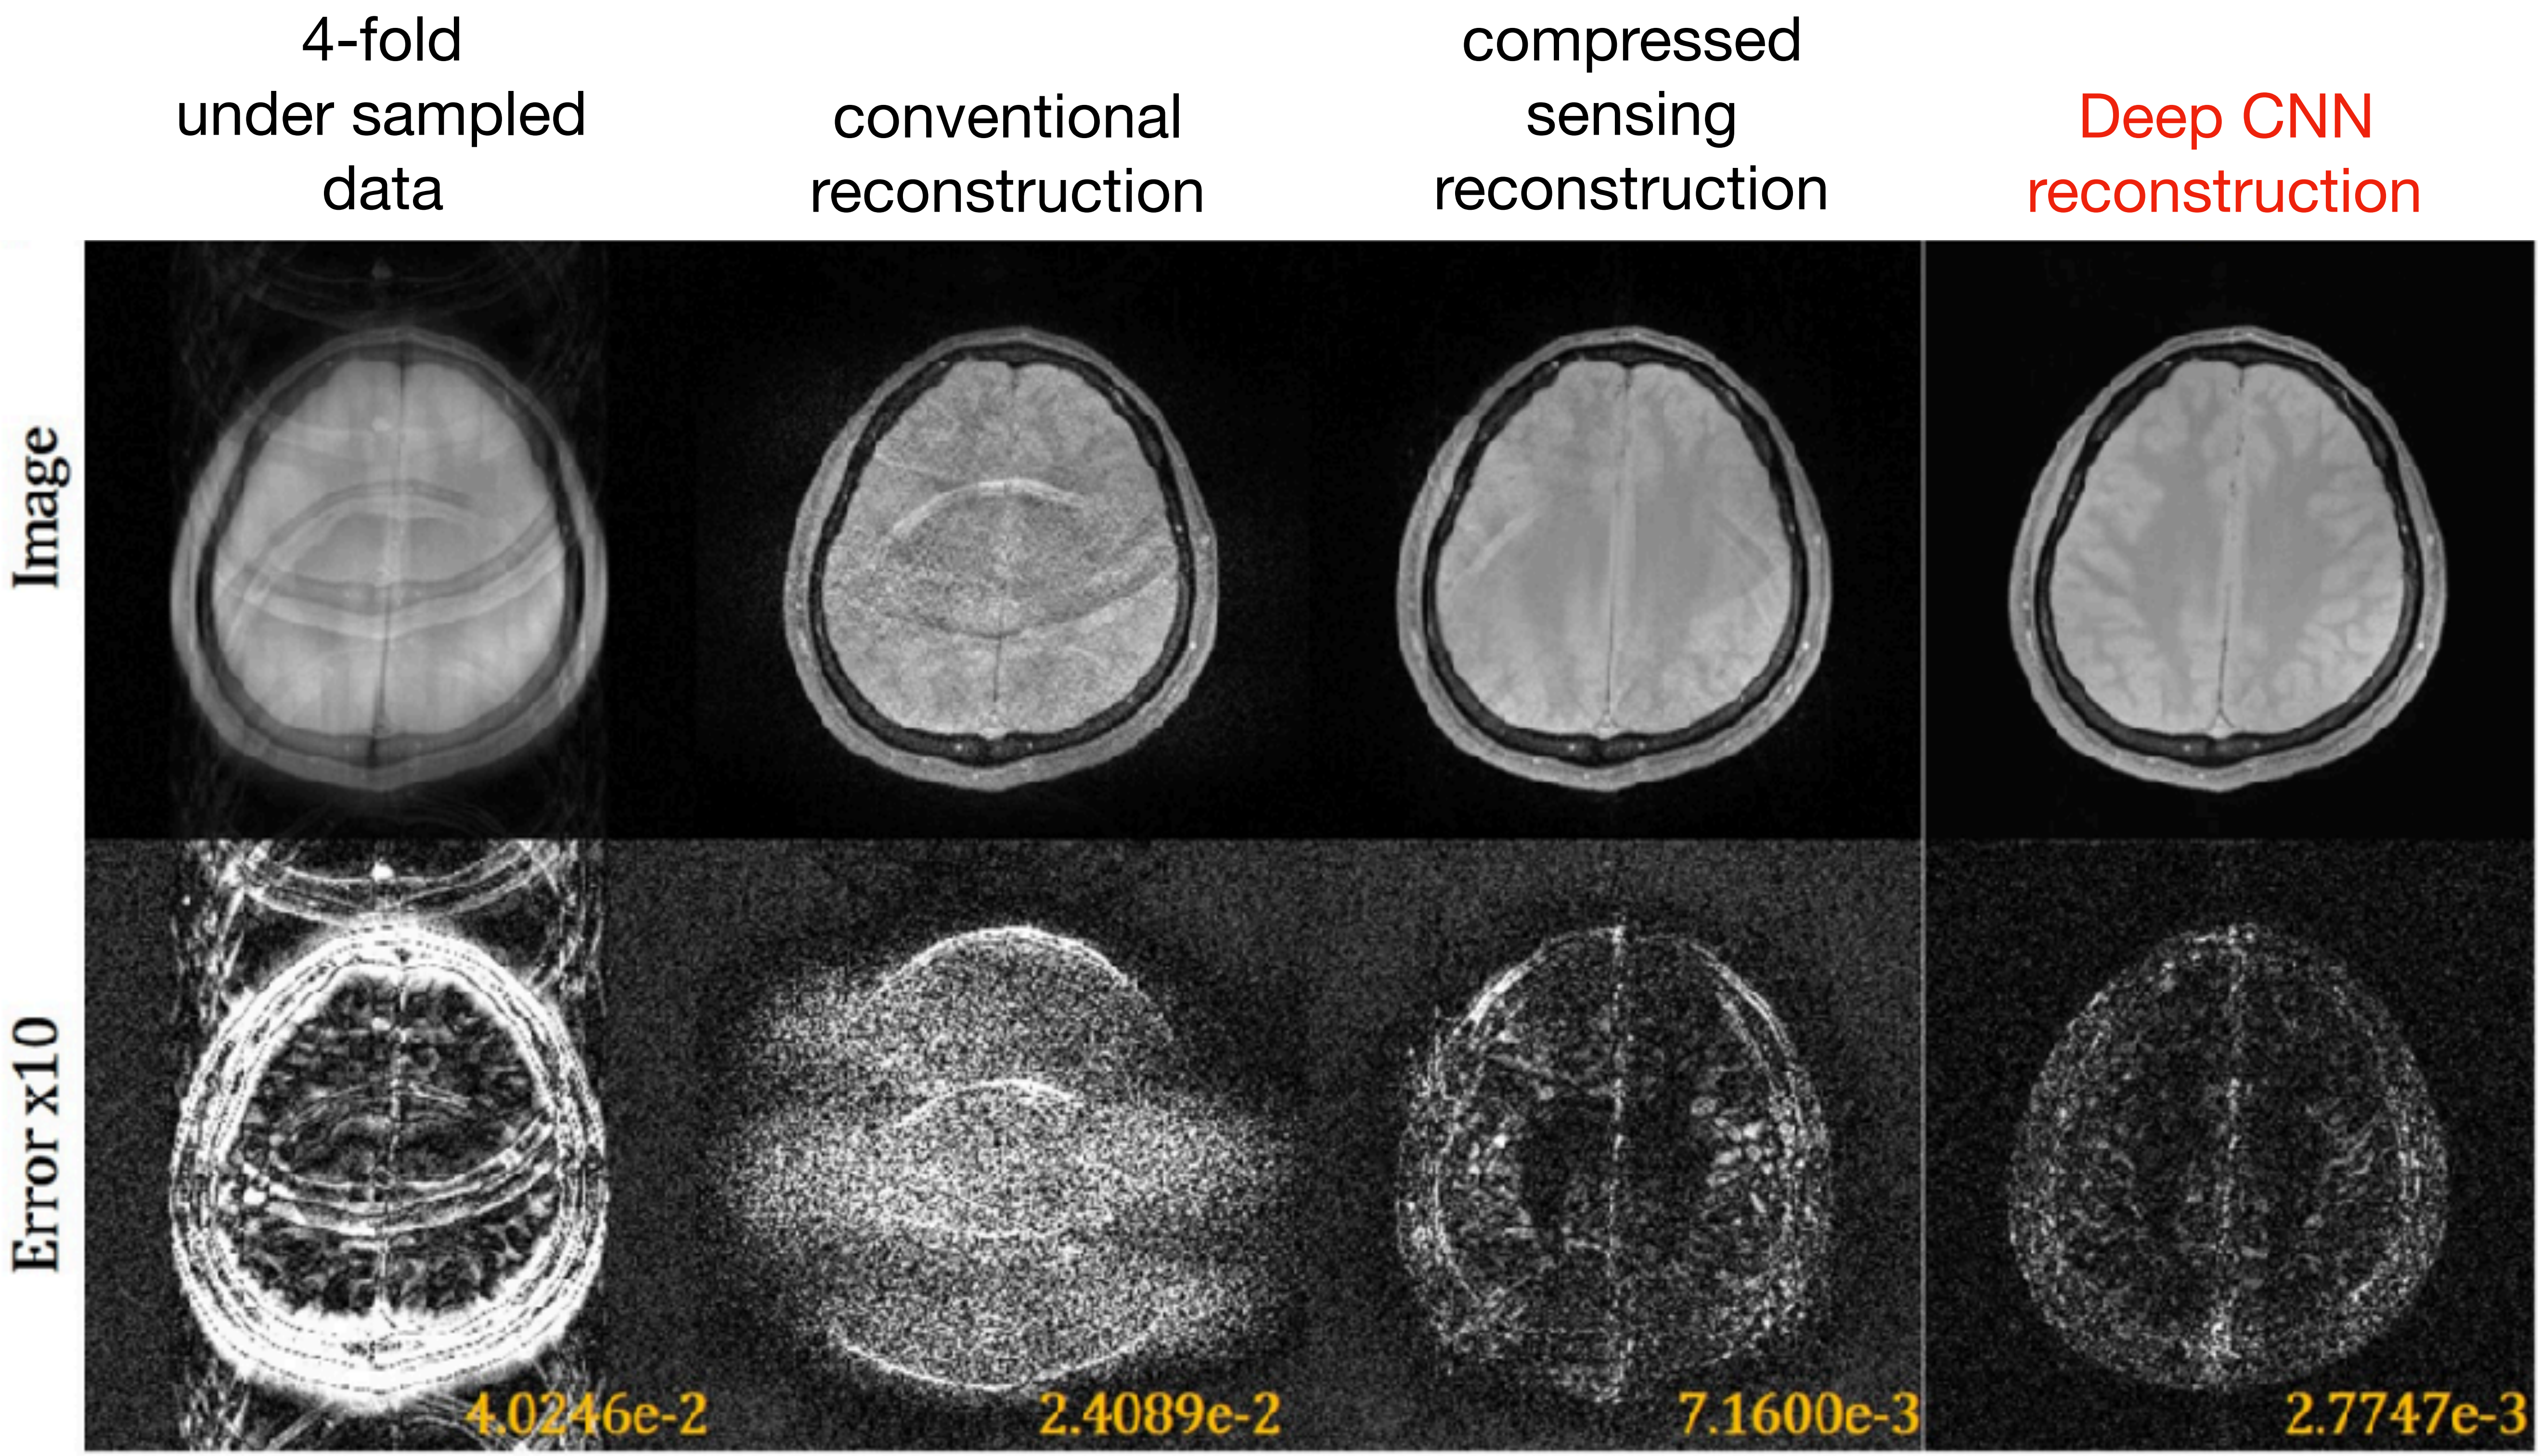

Lee, D., Yoo, J., & Ye, J. C. (2017, April). Deep residual learning for compressed sensing MRI. In 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)(pp. 15-18). IEEE.

Drawbacks to Deep CNNs

- Need to **retrain** CNN for any change in measurements
  - Undersampling rate (e.g., 2-fold, 4-fold, 10-fold, etc.)
  - Undersampling pattern (e.g., lines, spirals, radial, etc.)
  - Change in noise statistics (e.g., different scanner)

- Relatively **high sample complexity**
  - need many training images to avoid overfitting

- **Sensitive to perturbations**

On instabilities of deep learning in image reconstruction - Does AI come at a cost?

Vegard Antun[1] Francesco Renna[2] Clarice Poon[3] Ben Adcock[4] Anders C. Hansen*[5,1]

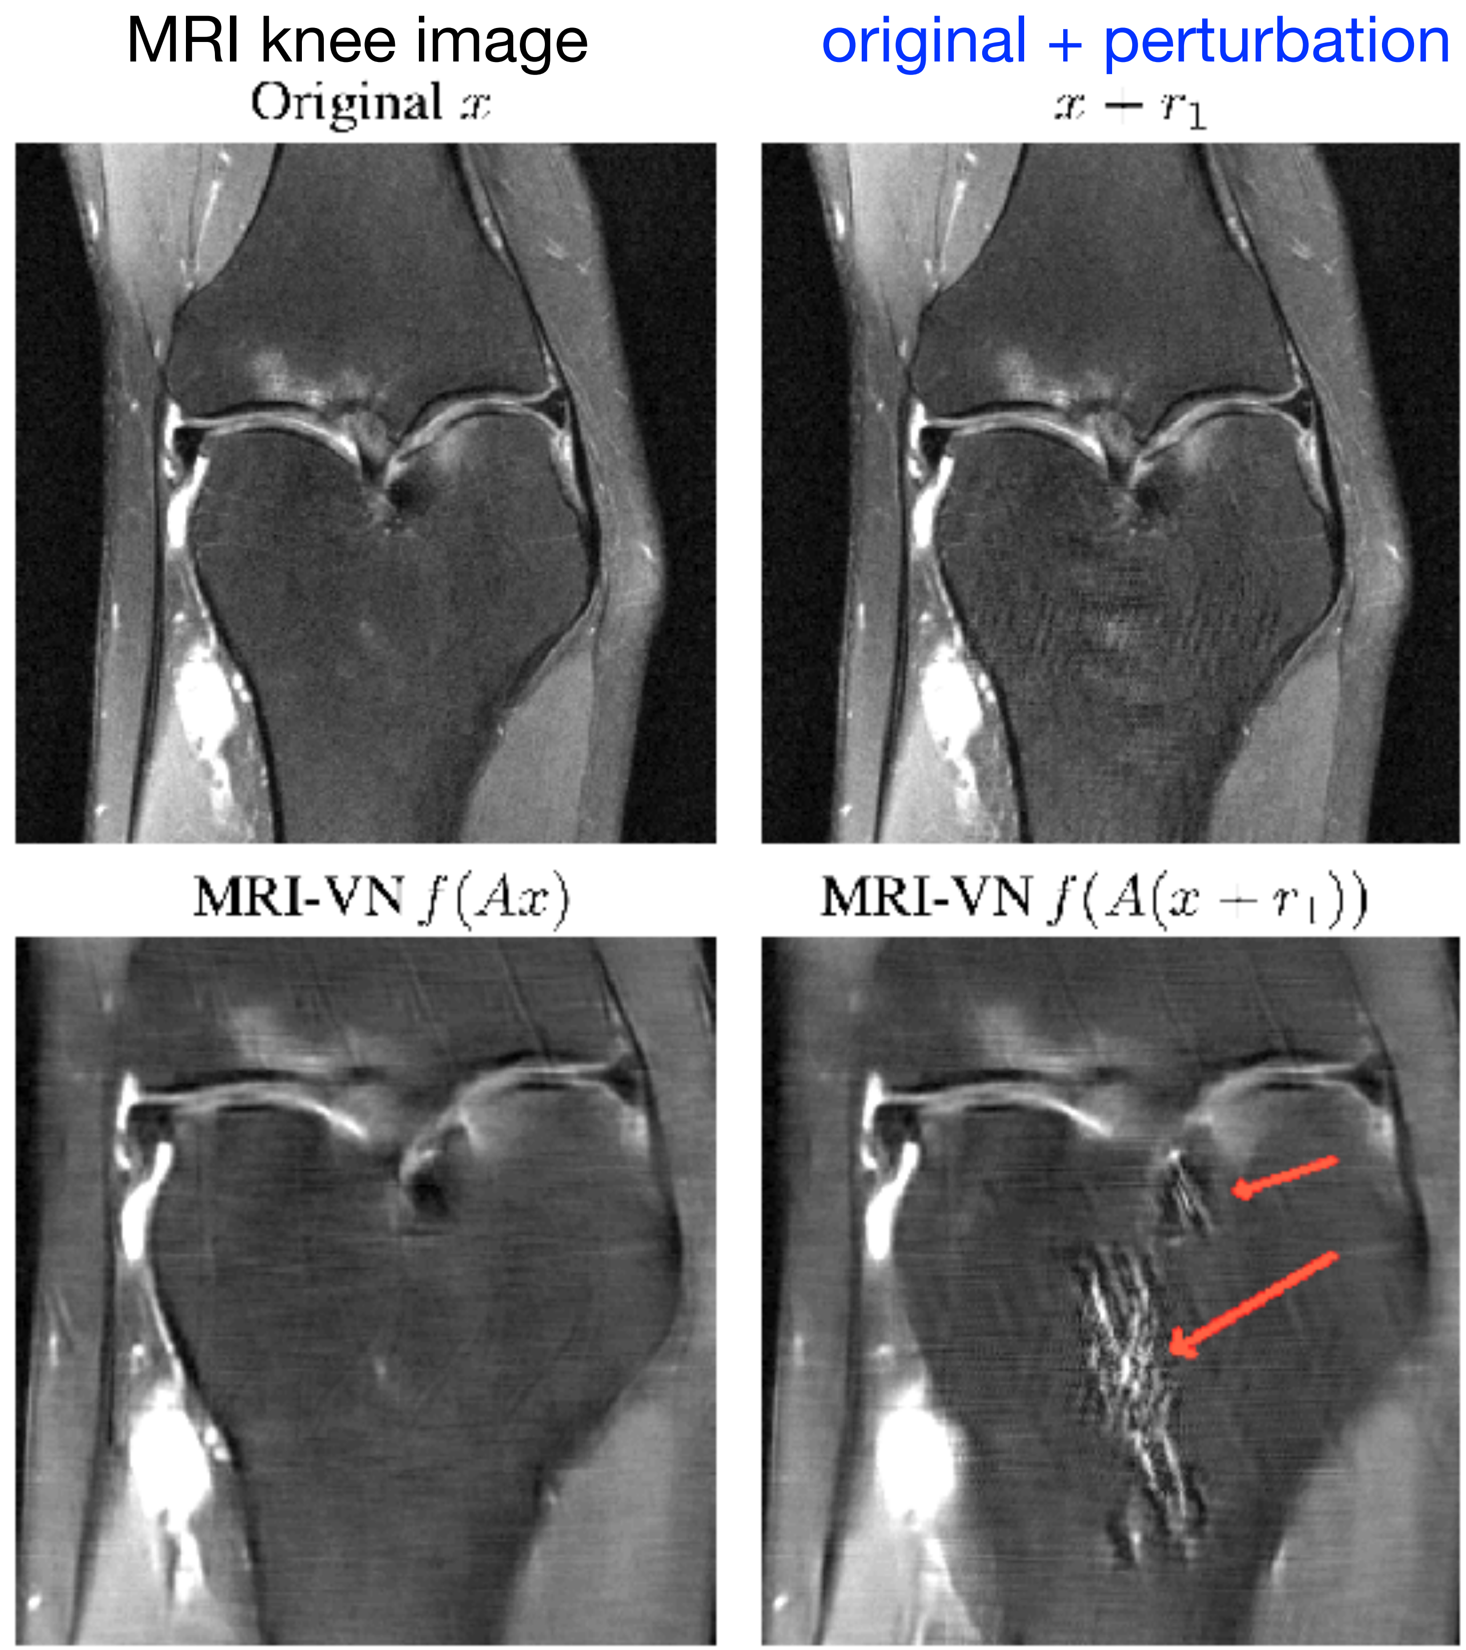

Deep CNN Reconstructions

Artifacts arise from small perturbations

Approach 2: Train a generative model (GAN)

NVIDIA "Style-based GAN" face image model

See also: https://thispersondoesnotexist.com/

Karras, Laine, & Aila 2019

https://www.youtube.com/watch?v=kSLJriaOumA

Generative adversarial networks (GANs)

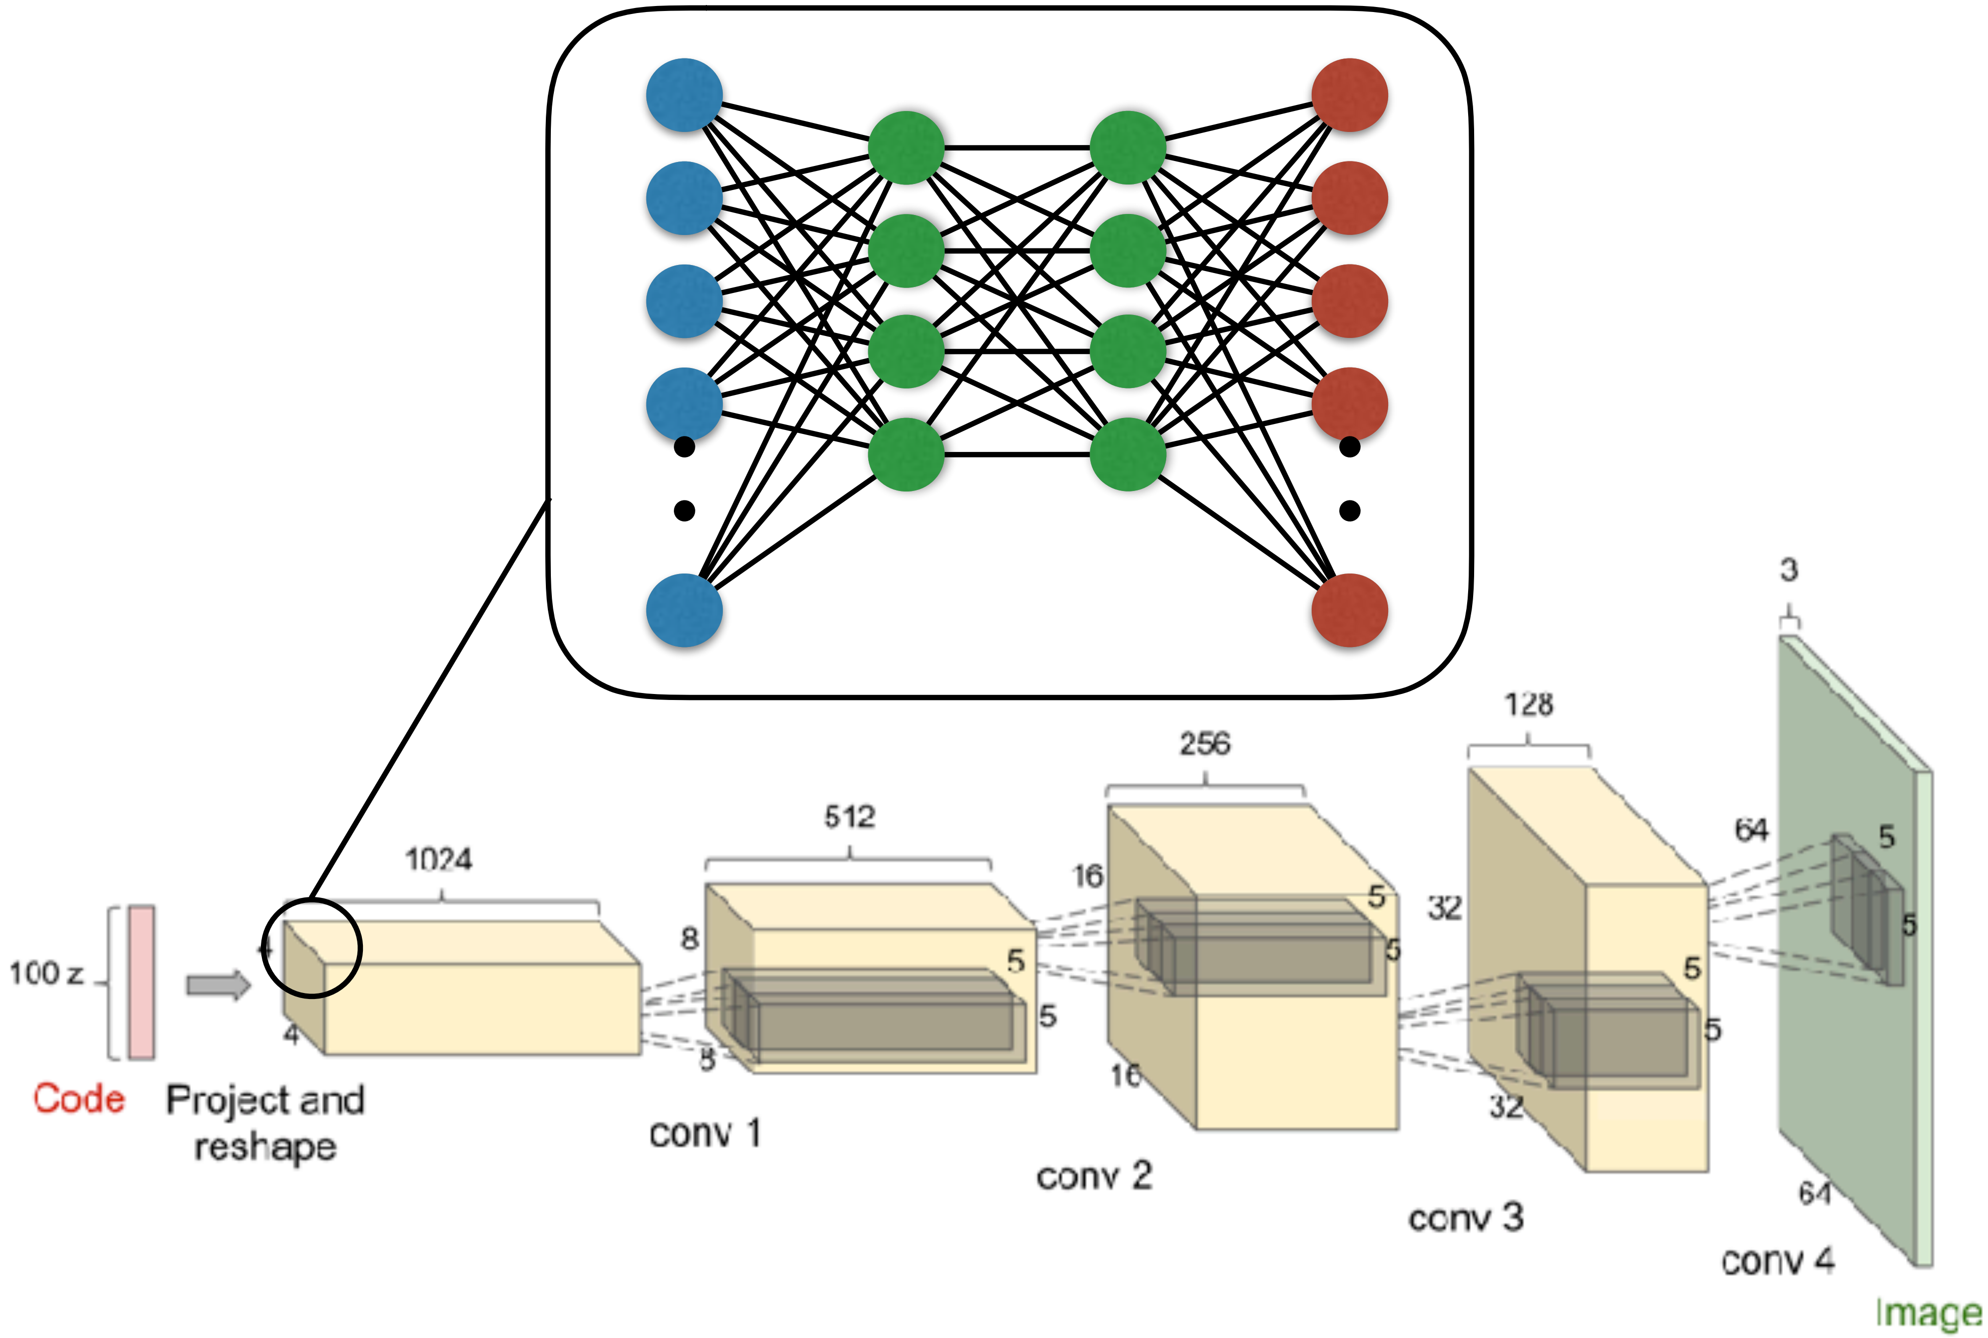

generator network — deep CNN

DCGAN - Goodfellow et al., 2016

Generative adversarial networks (GANs)

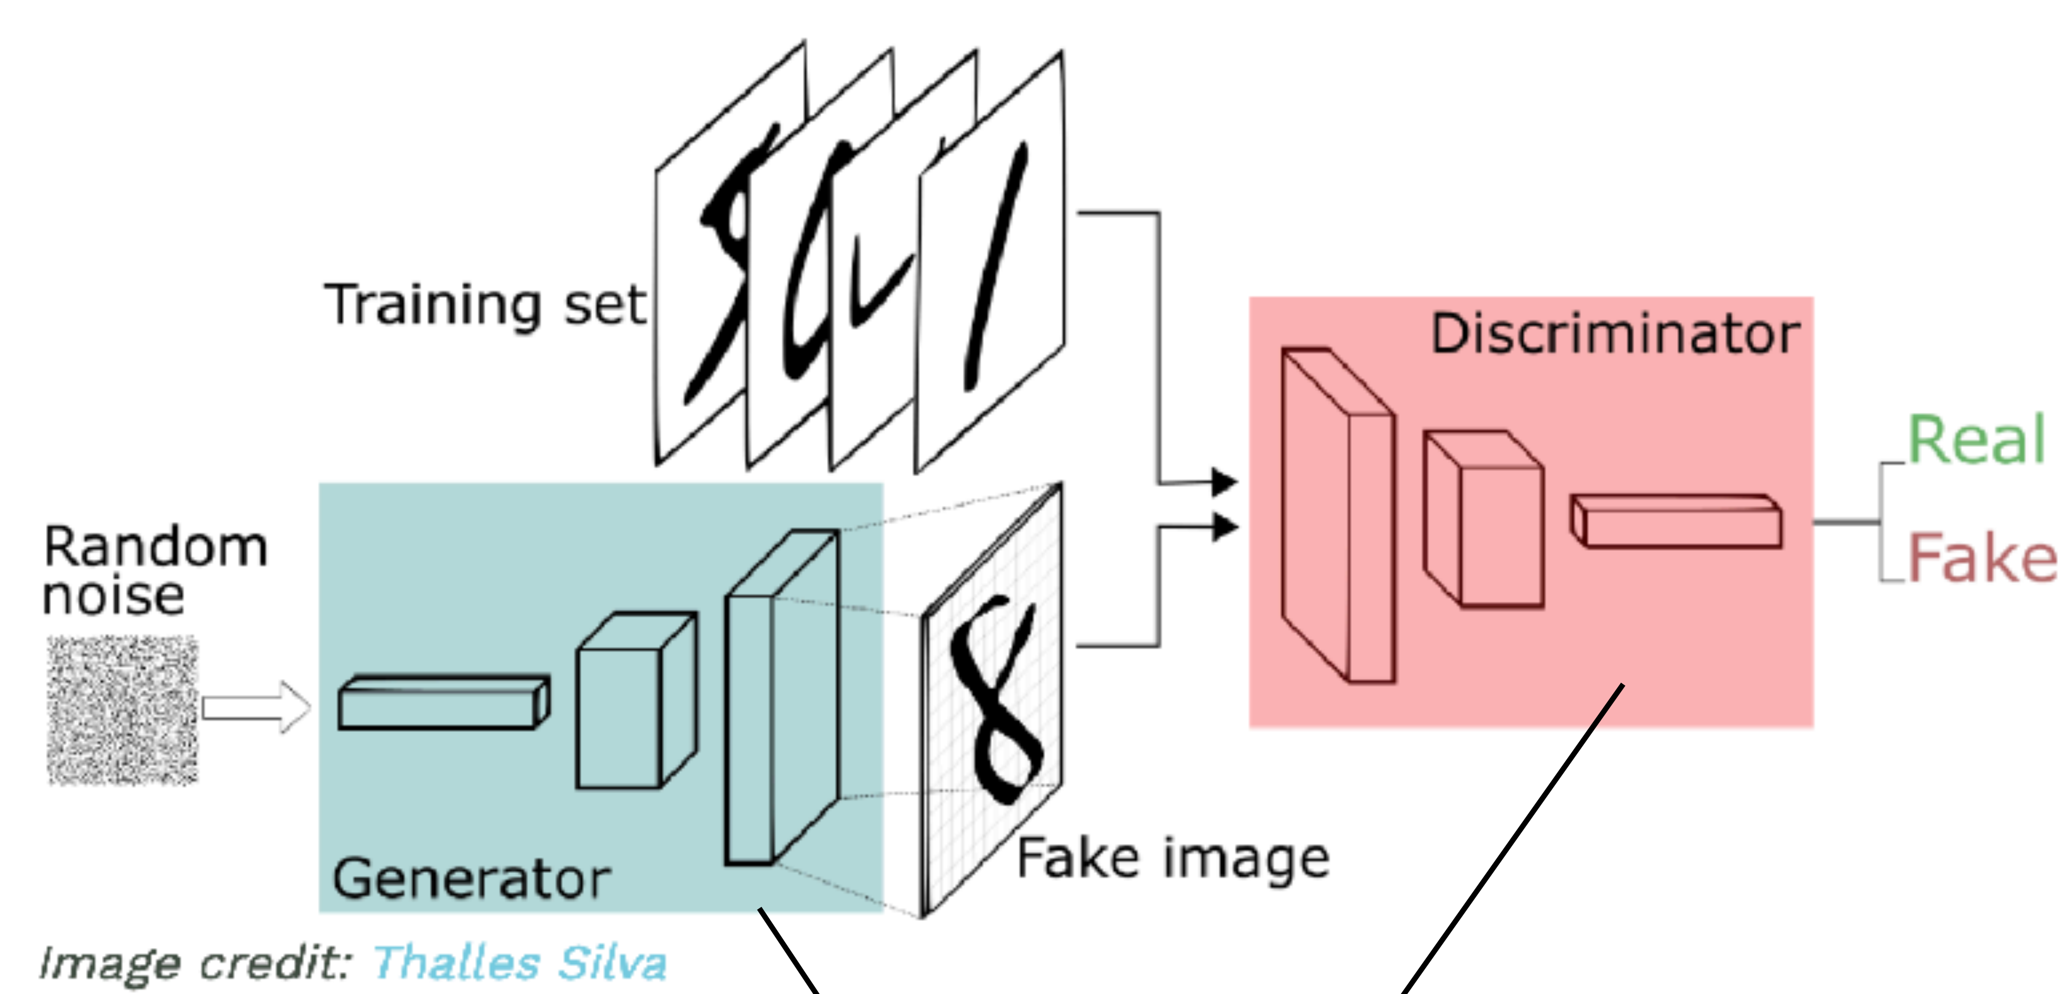

Jointly train Generator and Discriminator network

GANs for super-resolution

bicubic (21.59dB/0.6423)	deep CNN (23.53dB/0.7832)	GAN (21.15dB/0.6868)	original

Figure 2: From left to right: bicubic interpolation, deep residual network optimized for MSE, deep residual generative adversarial network optimized for a loss more sensitive to human perception, original HR image. Corresponding PSNR and SSIM are shown in brackets. [4× upscaling]

Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network

Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, Wenzhe Shi
Twitter
{cledig,ltheis,fhuszar,jcaballero,aacostadiaz,aaitken,atejani,jtotz,zehanw,wshi}@twitter.com

Real MRI images

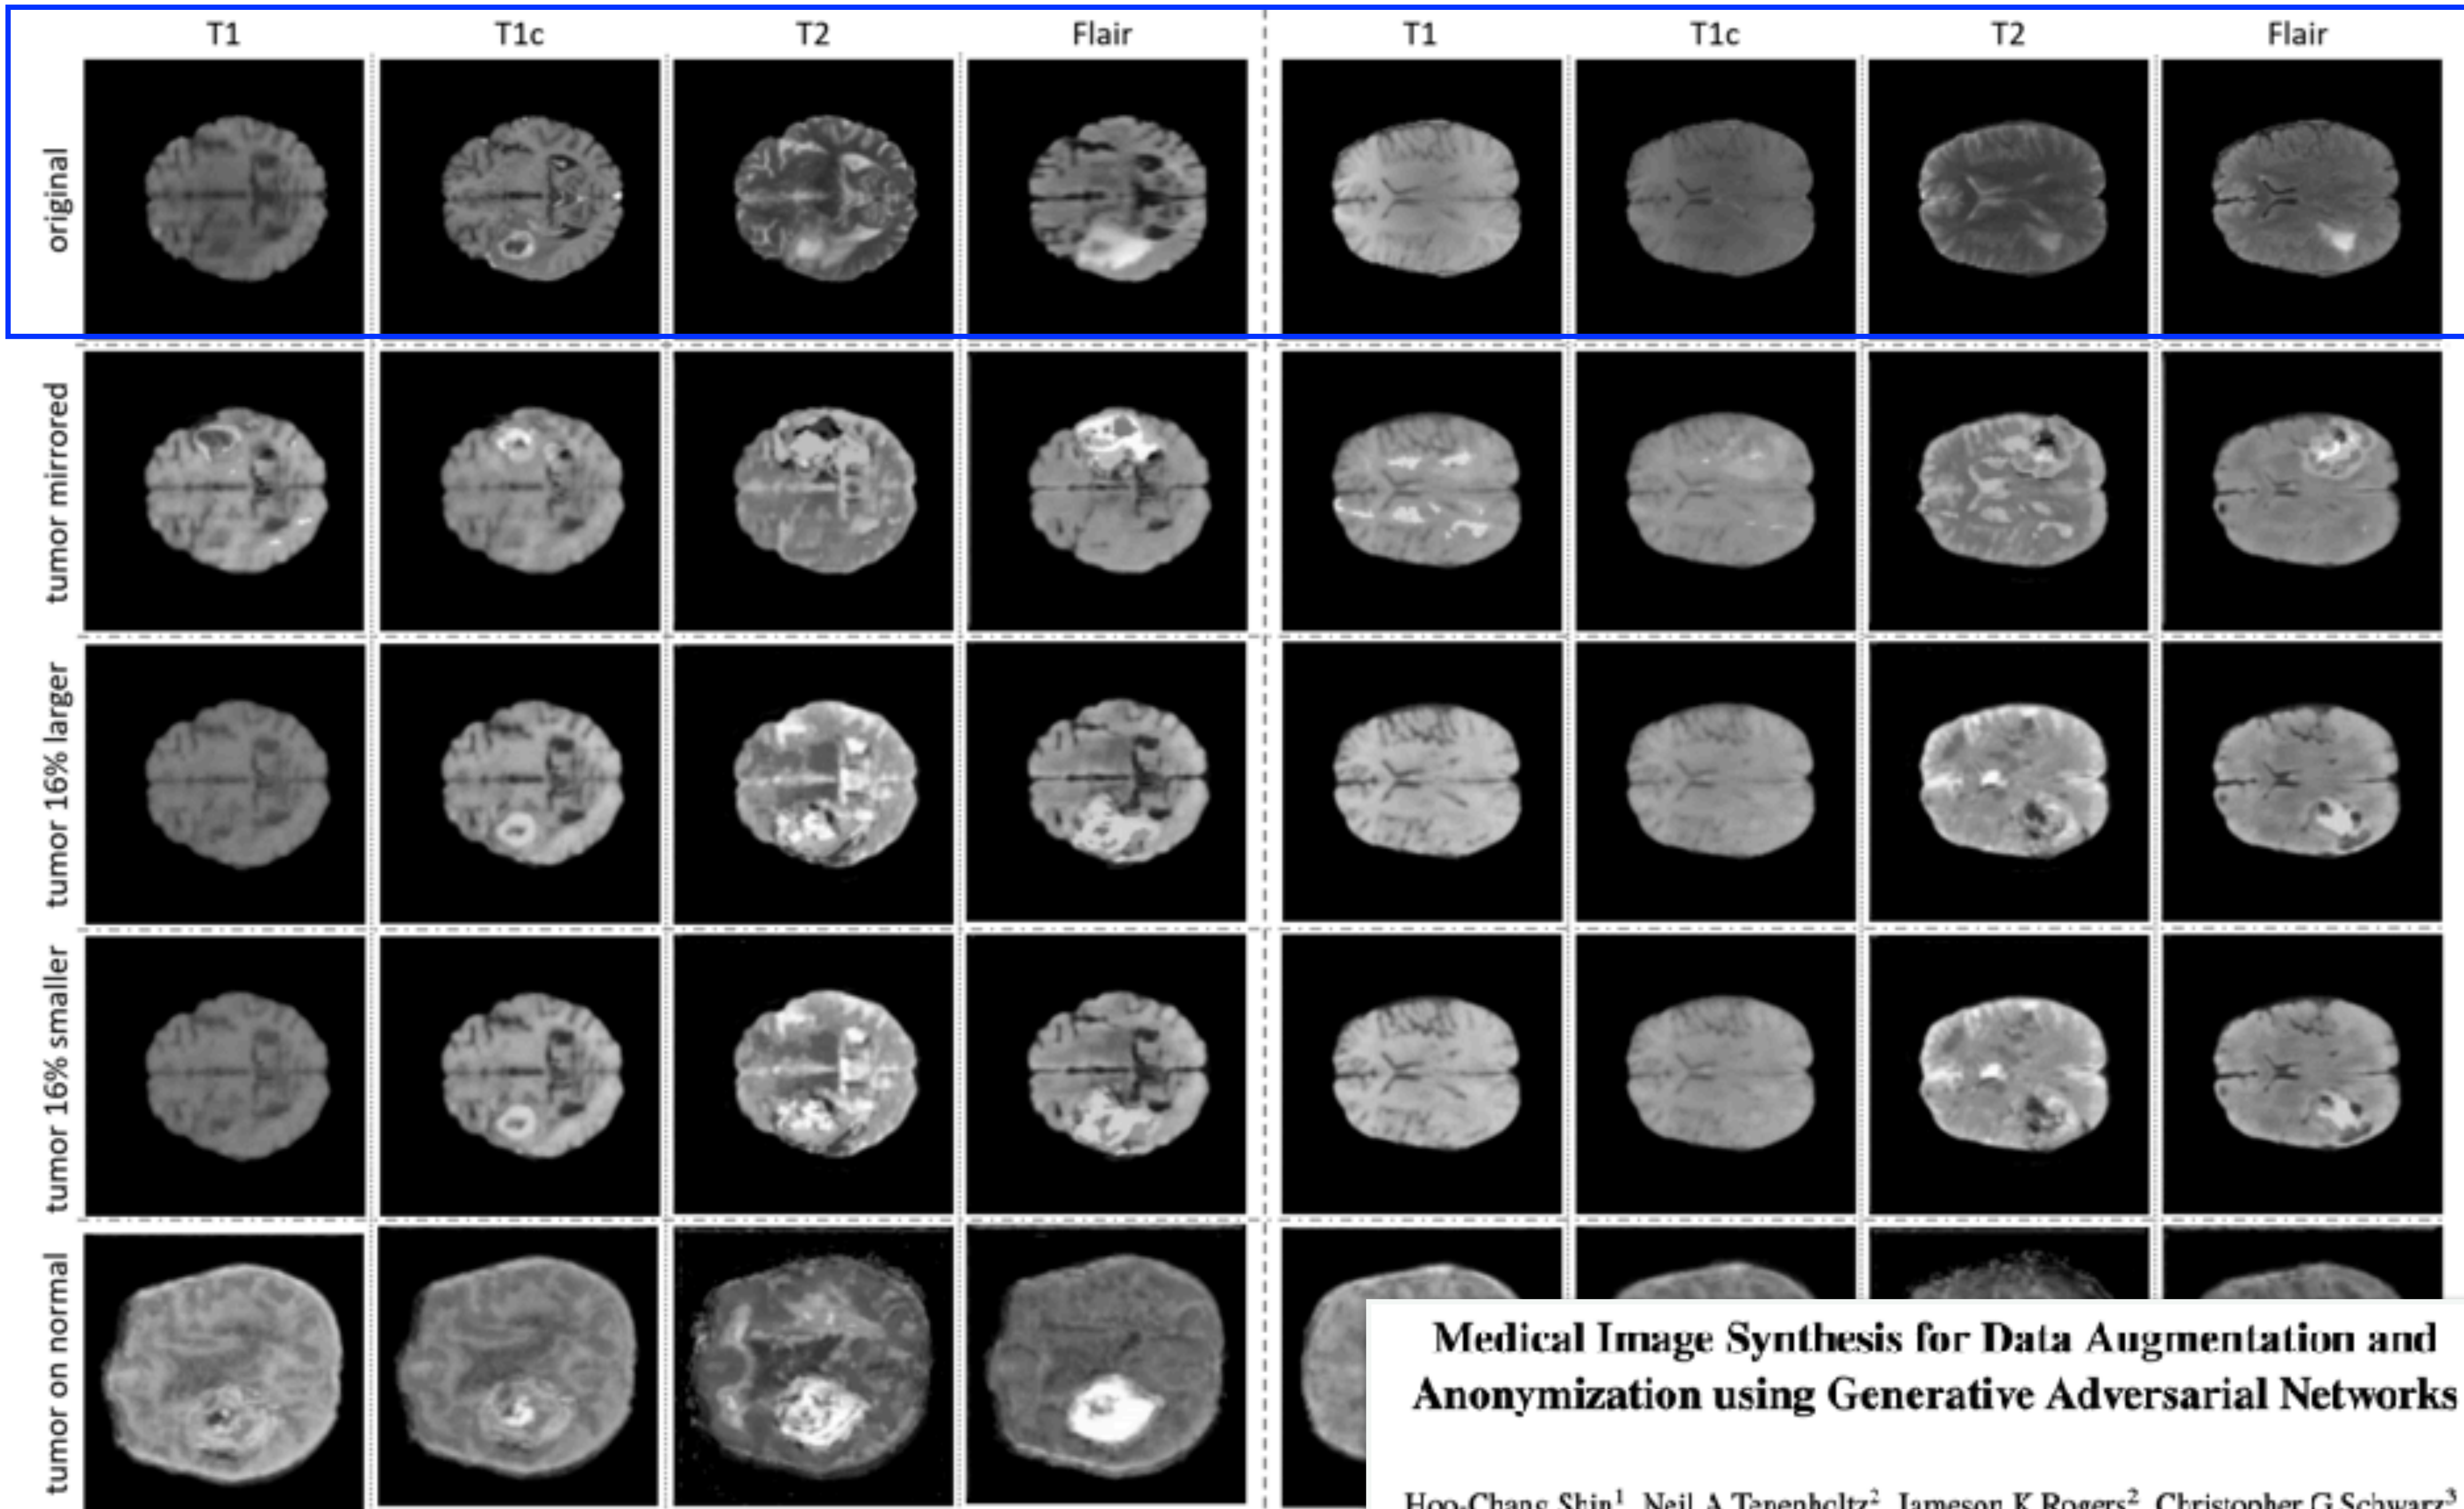

Simulated MRI Images

Medical Image Synthesis for Data Augmentation and Anonymization using Generative Adversarial Networks

Hoo-Chang Shin[1], Neil A Tenenholtz[2], Jameson K Rogers[2], Christopher G Schwarz[3], Matthew L Senjem[3], Jeffrey L Gunter[3], Katherine Andriole[2], and Mark Michalski[2]

[1] NVIDIA Corporation
[2] MGH & BWH Center for Clinical Data Science, Boston, MA, USA
[3] Mayo Clinic, Rochester, MN, USA

GANs for image reconstruction

Idea: Find image in the range of the generator that best fits the measurements

$$x_1 = G(z_1)$$

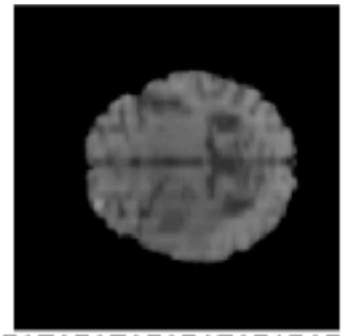

$$\|Hx_1 - y\|^2 = 100$$

$$x_2 = G(z_2)$$

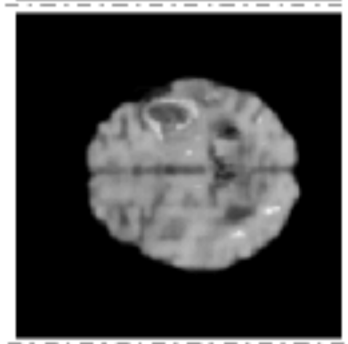

$$\|Hx_2 - y\|^2 = 1$$

$$x_3 = G(z_3)$$

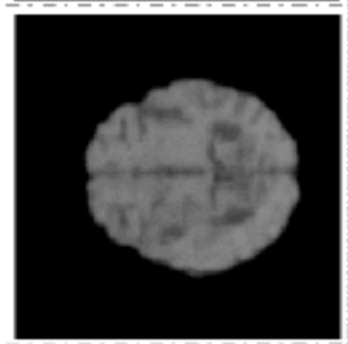

$$\|Hx_2 - y\|^2 = 10$$

Find best z by solving an optimization problem, e.g., by gradient descent

$$\underset{x \in \mathrm{range}(G)}{\text{minimize}} \|Hx - y\|^2 \longrightarrow \underset{z}{\text{minimize}} \|HG(z) - y\|^2$$

Bora, A., Jalal, A., Price, E., & Dimakis, A. G. Compressed sensing using generative models. ICML 2017

GANs for MRI Reconstruction

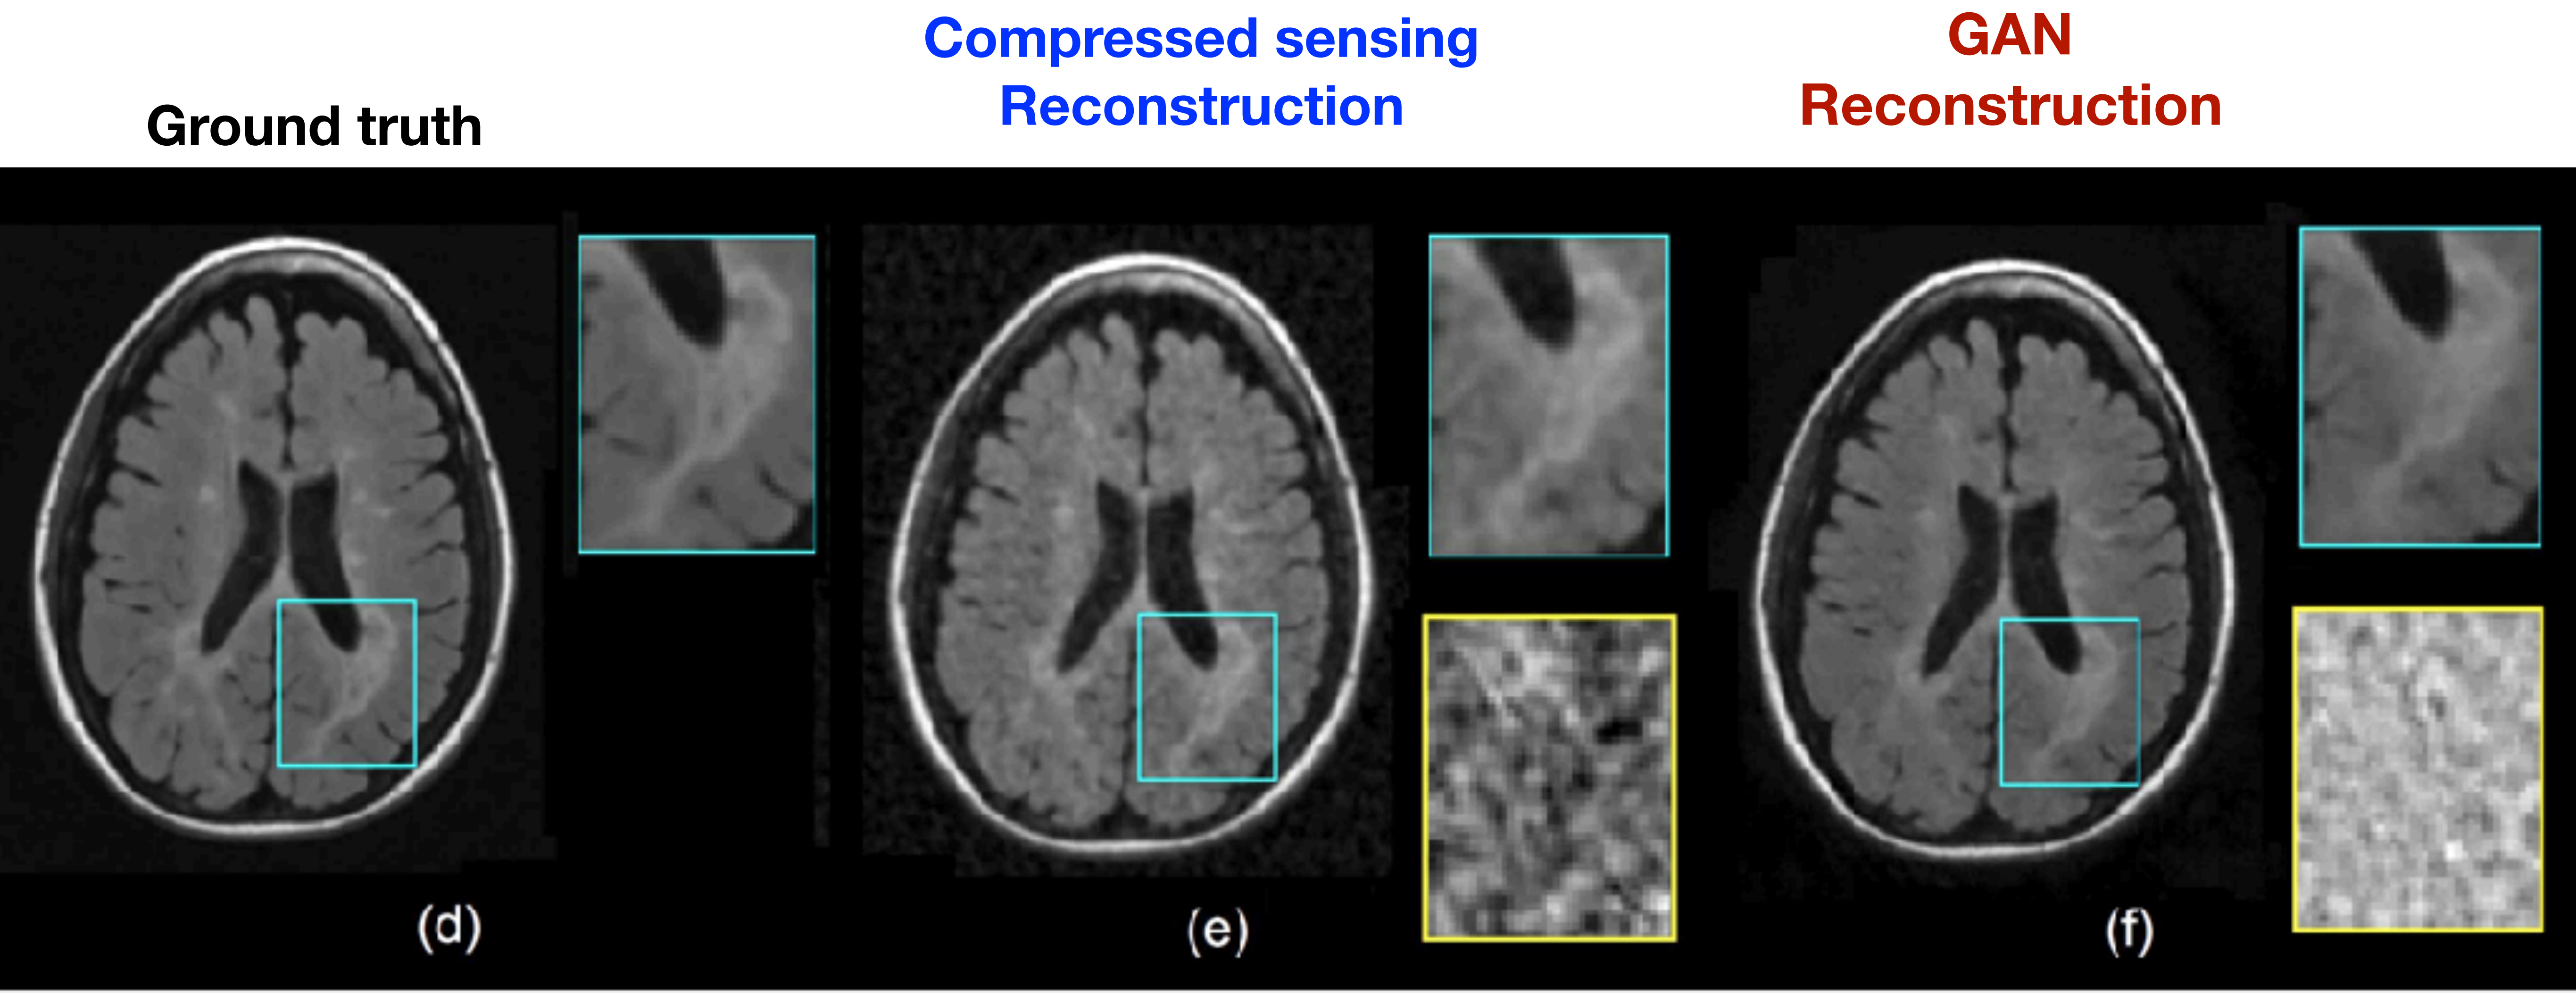

DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction'

Drawbacks to GANs

- Training an accurate GAN requires many training samples (NVIDIA faces: ~200,000 training examples)

- Reconstructed images must lie in the range of the GAN
  - If patient has abnormality not contained in training set, the abnormality may be "smoothed over" by the GAN

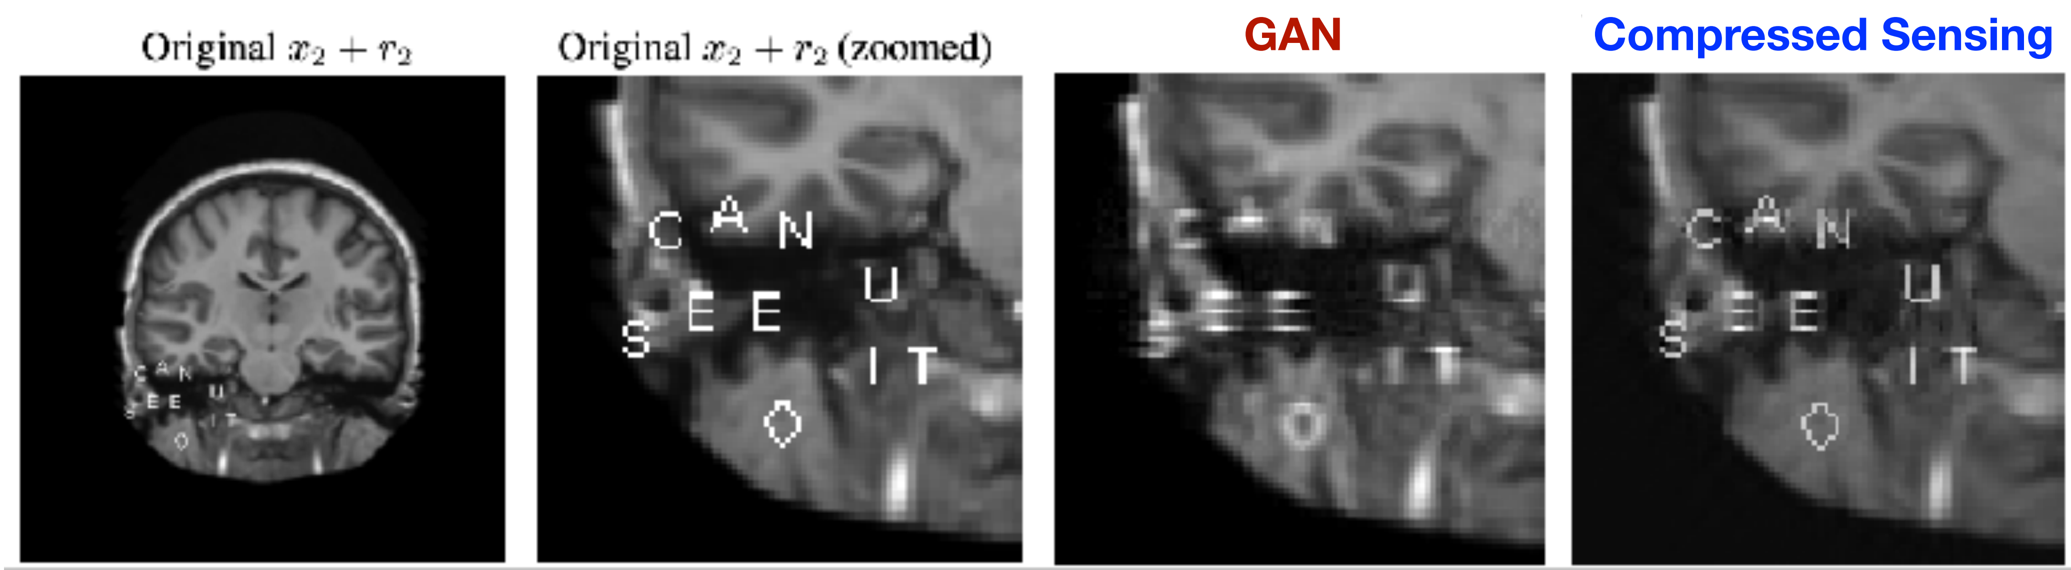

Antun, V., Renna, F., Poon, C., Adcock, B., & Hansen, A. C. (2019).
On instabilities of deep learning in image reconstruction-Does AI come at a cost?. arXiv preprint arXiv:1902.05300.

Approach 3: Unrolling algorithms

Learning to learn by gradient descent by gradient descent

Marcin Andrychowicz[1], **Misha Denil**[1], **Sergio Gómez Colmenarejo**[1], **Matthew W. Hoffman**[1], **David Pfau**[1], **Tom Schaul**[1], **Brendan Shillingford**[1,2], **Nando de Freitas**[1,2,3]

[1]Google DeepMind [2]University of Oxford [3]Canadian Institute for Advanced Research

marcin.andrychowicz@gmail.com
{mdenil,sergomez,mwhoffman,pfau,schaul}@google.com
brendan.shillingford@cs.ox.ac.uk, nandodefreitas@google.com

Learning to Optimize via Deep Learning

- **Intuition:** Compressed sensing gives good reconstructions but requires solving a computational costly optimization problem each time

- Can we learn to solve the compressed sensing optimization problem with deep learning?

compressed sensing

$$y \longrightarrow \boxed{\underset{x}{\text{minimize}}\ \|Hx - y\|^2 + r(x)} \longrightarrow \tilde{x}$$

$$y \longrightarrow \boxed{\textbf{deep neural network}} \longrightarrow \tilde{x}$$

Example: Unrolled gradient descent

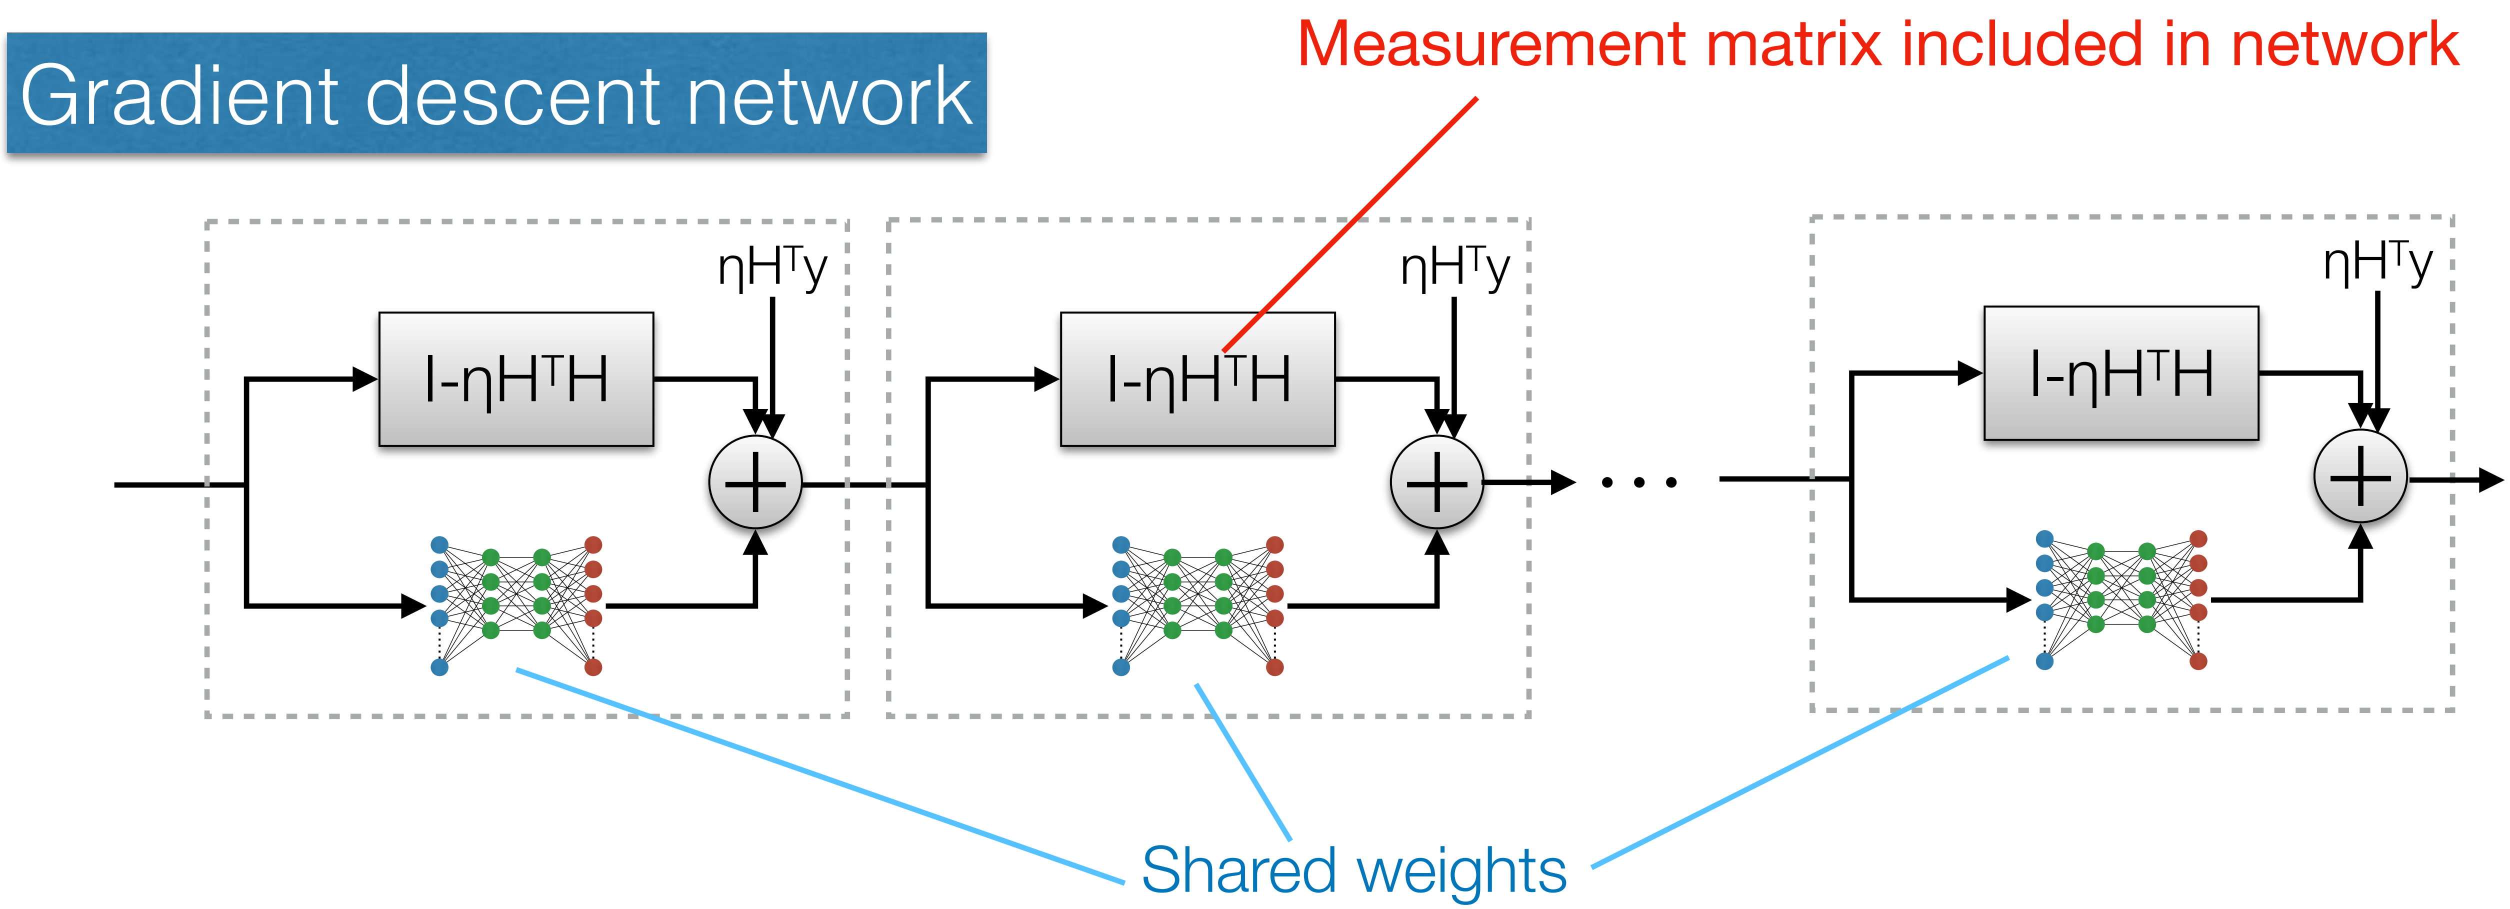

- Mimics finitely many iterations of gradient descent applied to $$\underset{x}{\text{minimize}}\ \|Hx - y\|^2 + r(x)$$
- Replace regularizer r(x) with learned neural network

Neumann networks (O., Gilton, Willett, 2019)

Neumann network

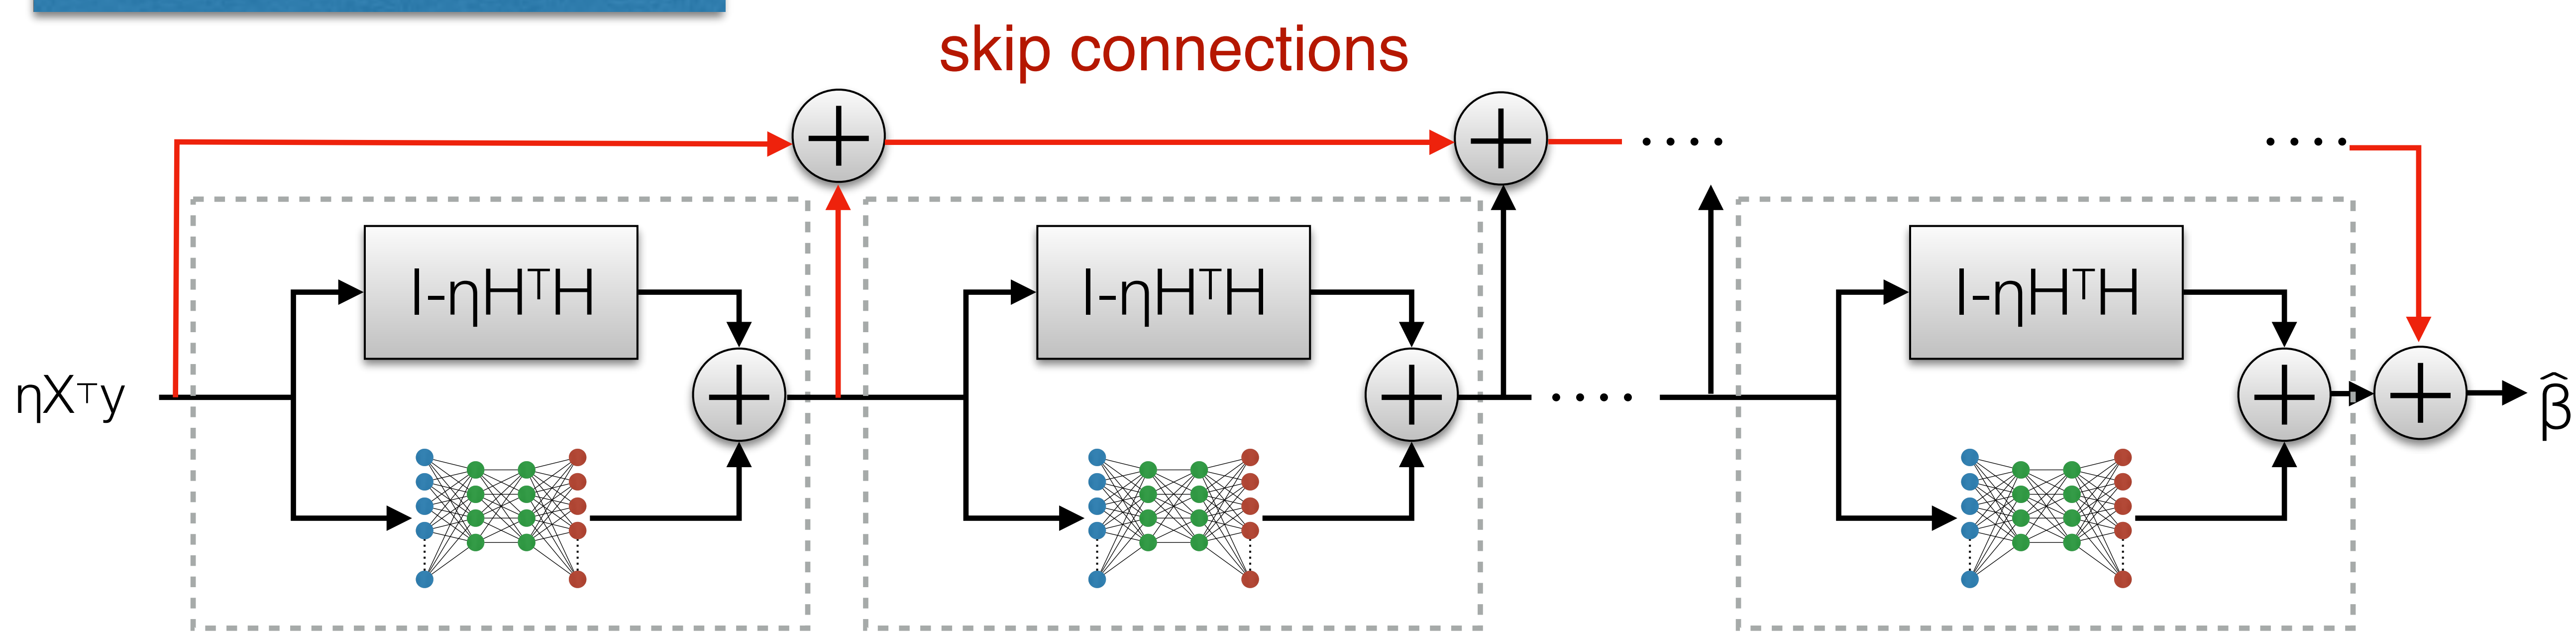

Dense Convolutional Networks (DenseNets)

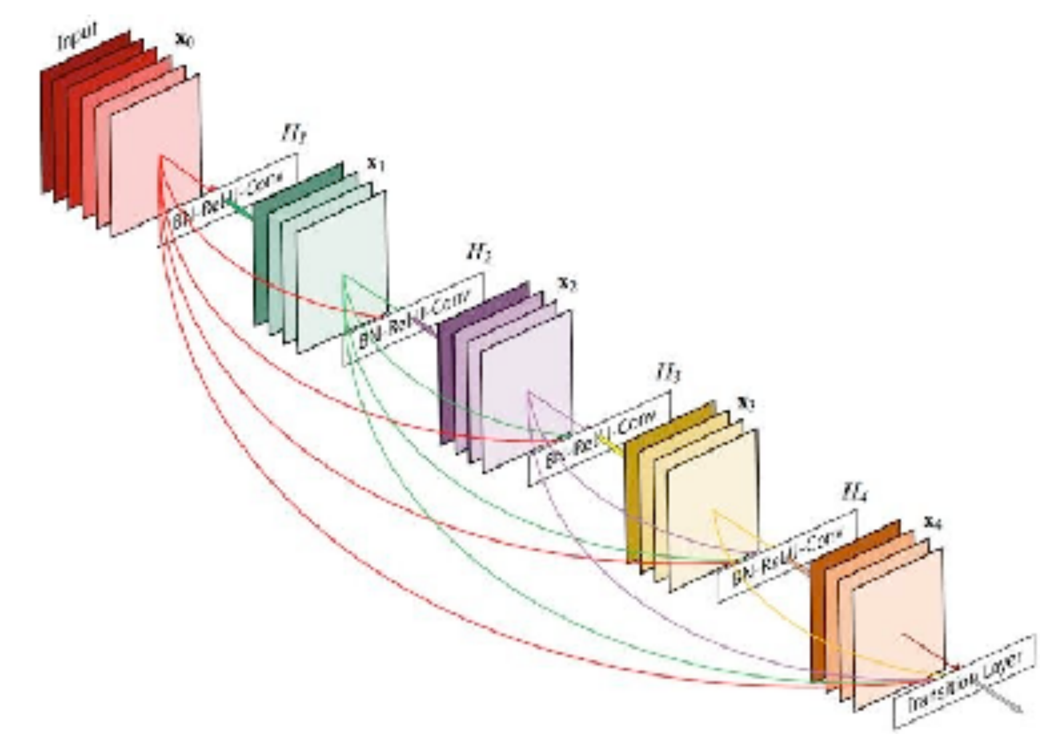

Sample Complexity - Deblurring task

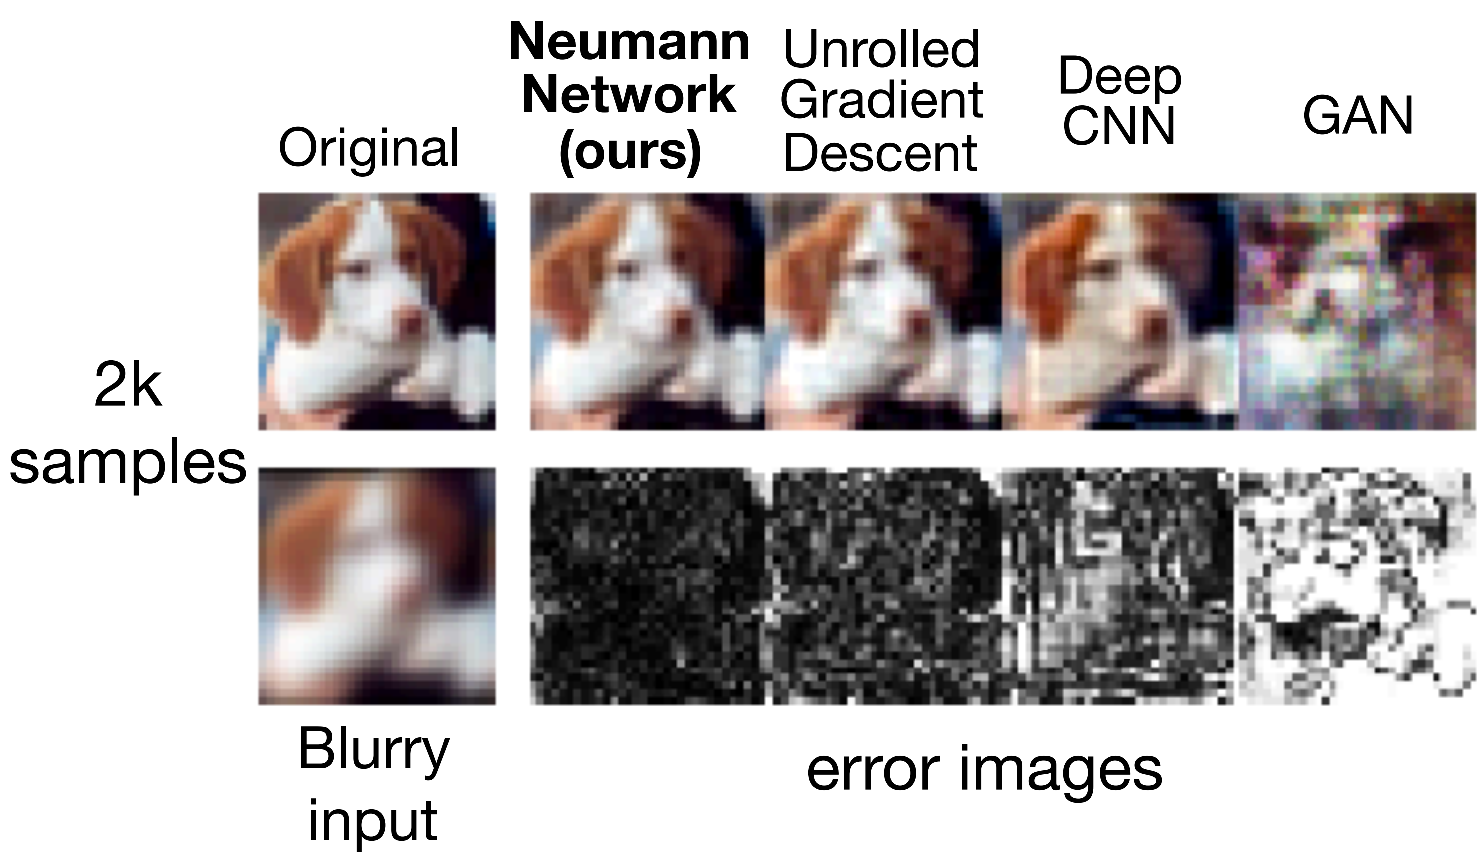

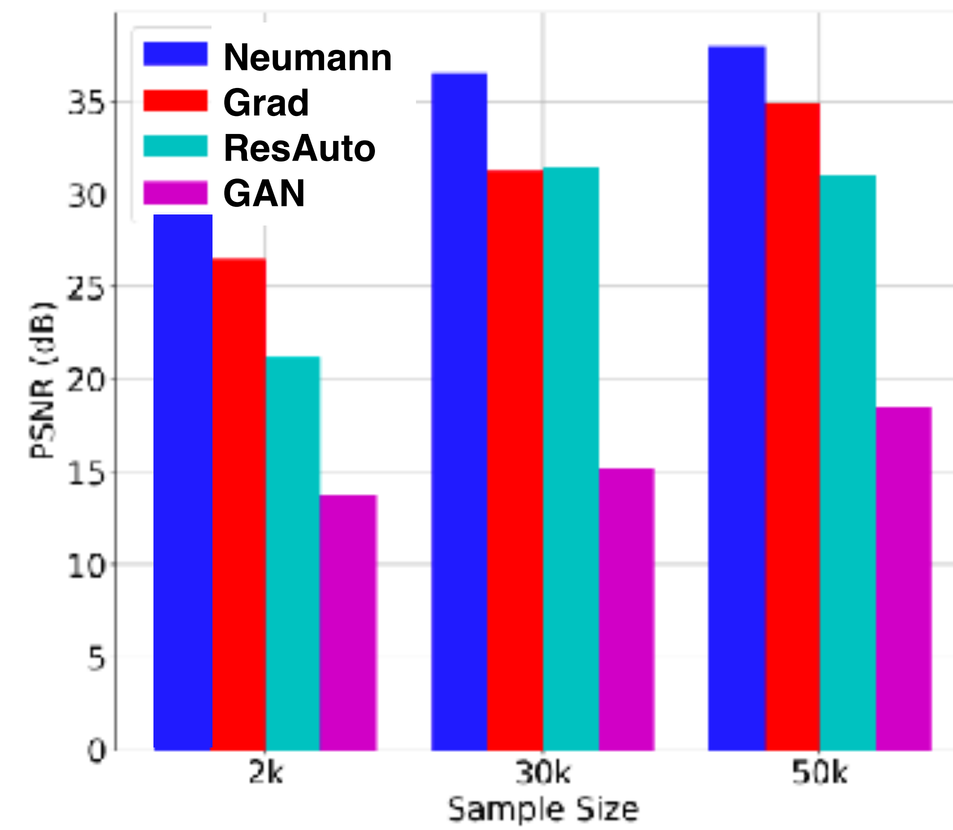

Neumann Networks for Inverse Problems in Imaging

Davis Gilton, Greg Ongie, Rebecca Willett*

January 15, 2019

MRI Reconstruction Results

Original	Neumann Network	Compressed Sensing	CNN

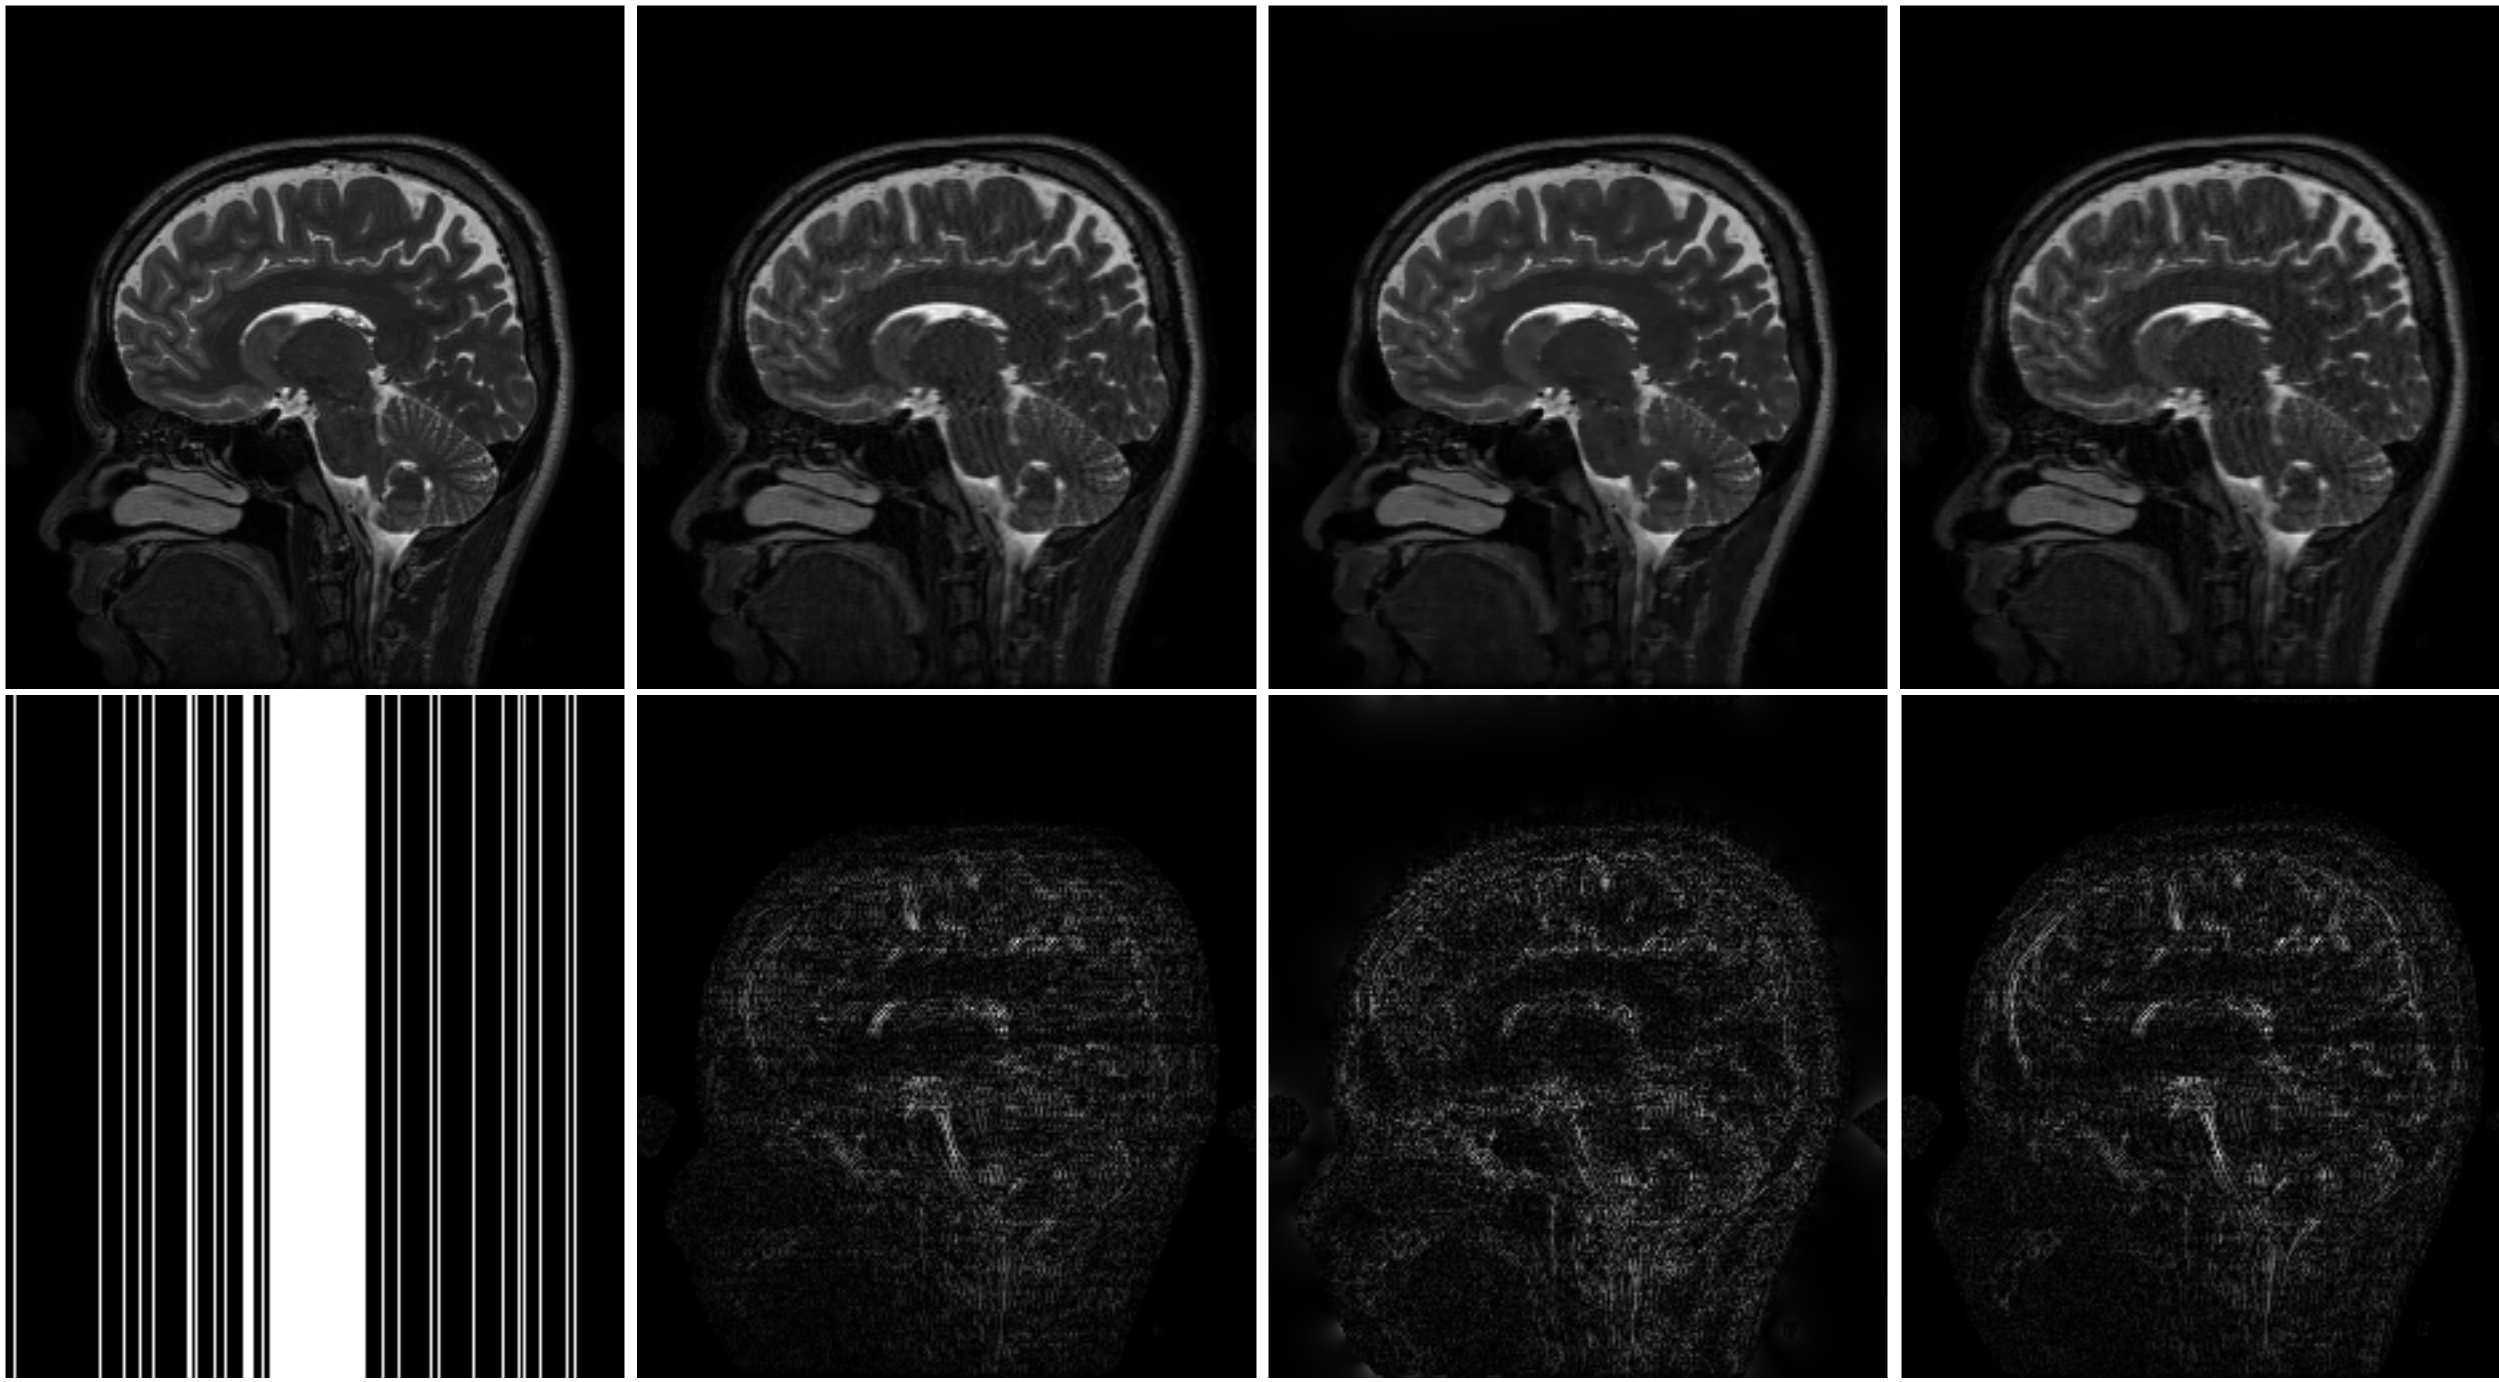

k-space Sampling Mask	PSNR: 34.95 dB Time: 16.3 sec	PSNR: 32.29 dB Time: 349.2 sec	PSNR: 32.39 dB Time: 1.6 sec

Recap and Outlook

Challenges in deep learning for biomedical imaging

- **Challenge 1: Limited Training Data**
  - Unrolling — incorporate forward model into network
- **Challenge 2: Complex Input Formats**
  - Use “approximate inverse” as input to network, rather than raw measurements
- **Challenge 3: Beyond Classification**
  - Adapt CNN’s to perform image restoration tasks
  - Use GAN’s or to model image distribution

Going Forward: Uncertainty Quantification

How do we know we are not hallucinating features in the reconstruction?

Can we learn a full posterior? p(x|y)

conditional mean

point-wise standard deviation

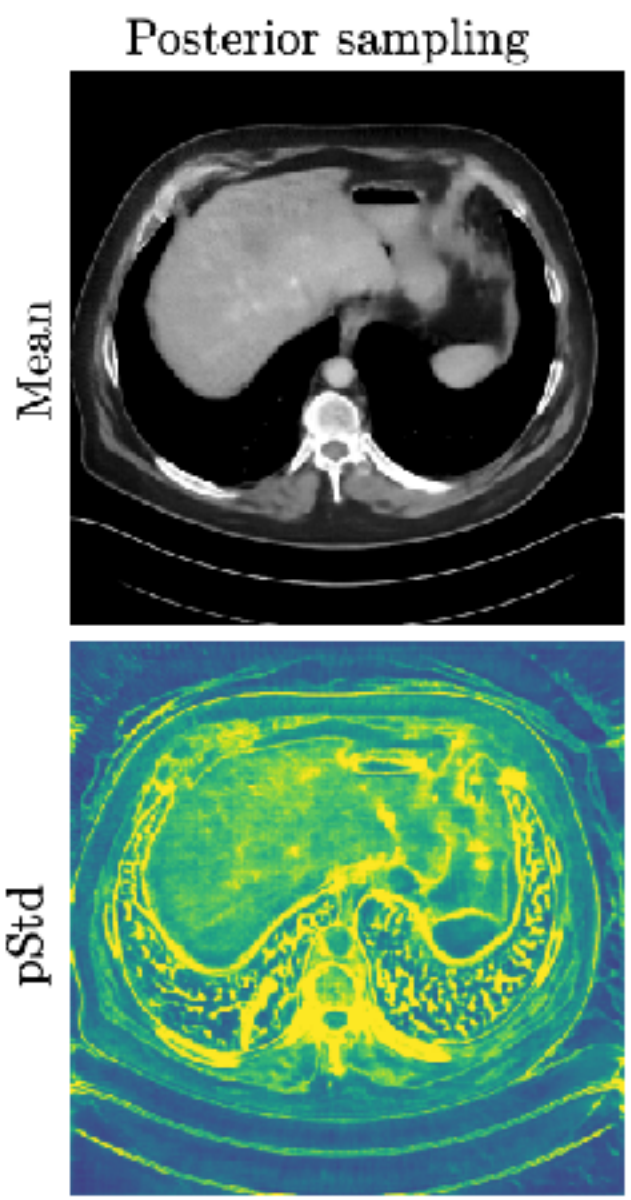

Hypothesis testing

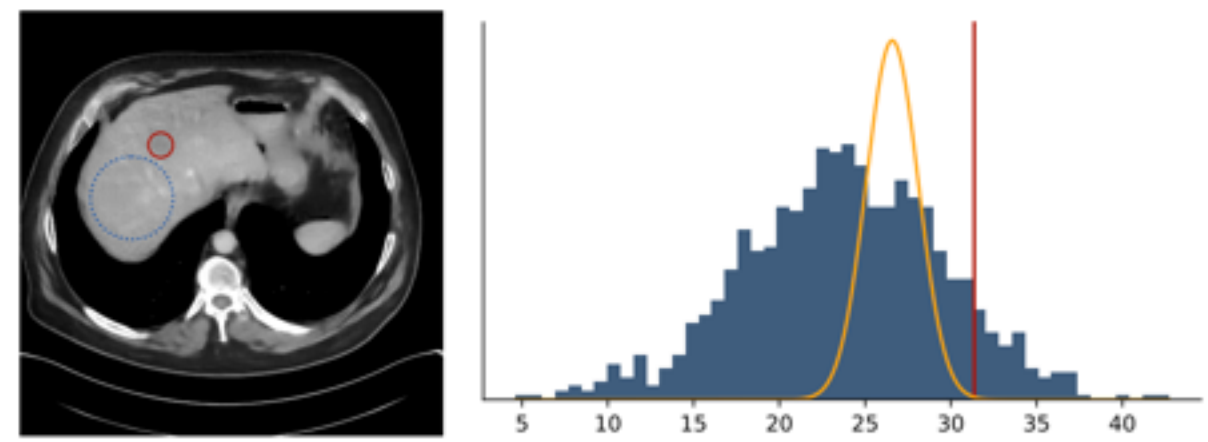

Fig. 4: The suspected tumor (red) and the reference region (blue) shown in the sample posterior mean image. Right plot shows average contrast differences between the tumor and reference region. The histogram is computed by posterior sampling applied to test data (fig. 2), the yellow curve is from direct estimation, and the true value is the red threshold.

Deep Bayesian Inversion

Computational uncertainty quantification for large scale inverse problems

Jonas Adler
Department of Mathematics
KTH - Royal institute of Technology
jonasadl@kth.se
Research and Physics, Elekta

Ozan Öktem
Department of Mathematics
KTH - Royal institute of Technology
ozan@kth.se

Access to Datasets

- Further advances will require standardized training and test sets
- Facebook/NYU FastMRI: 900 3-D knee MRI images

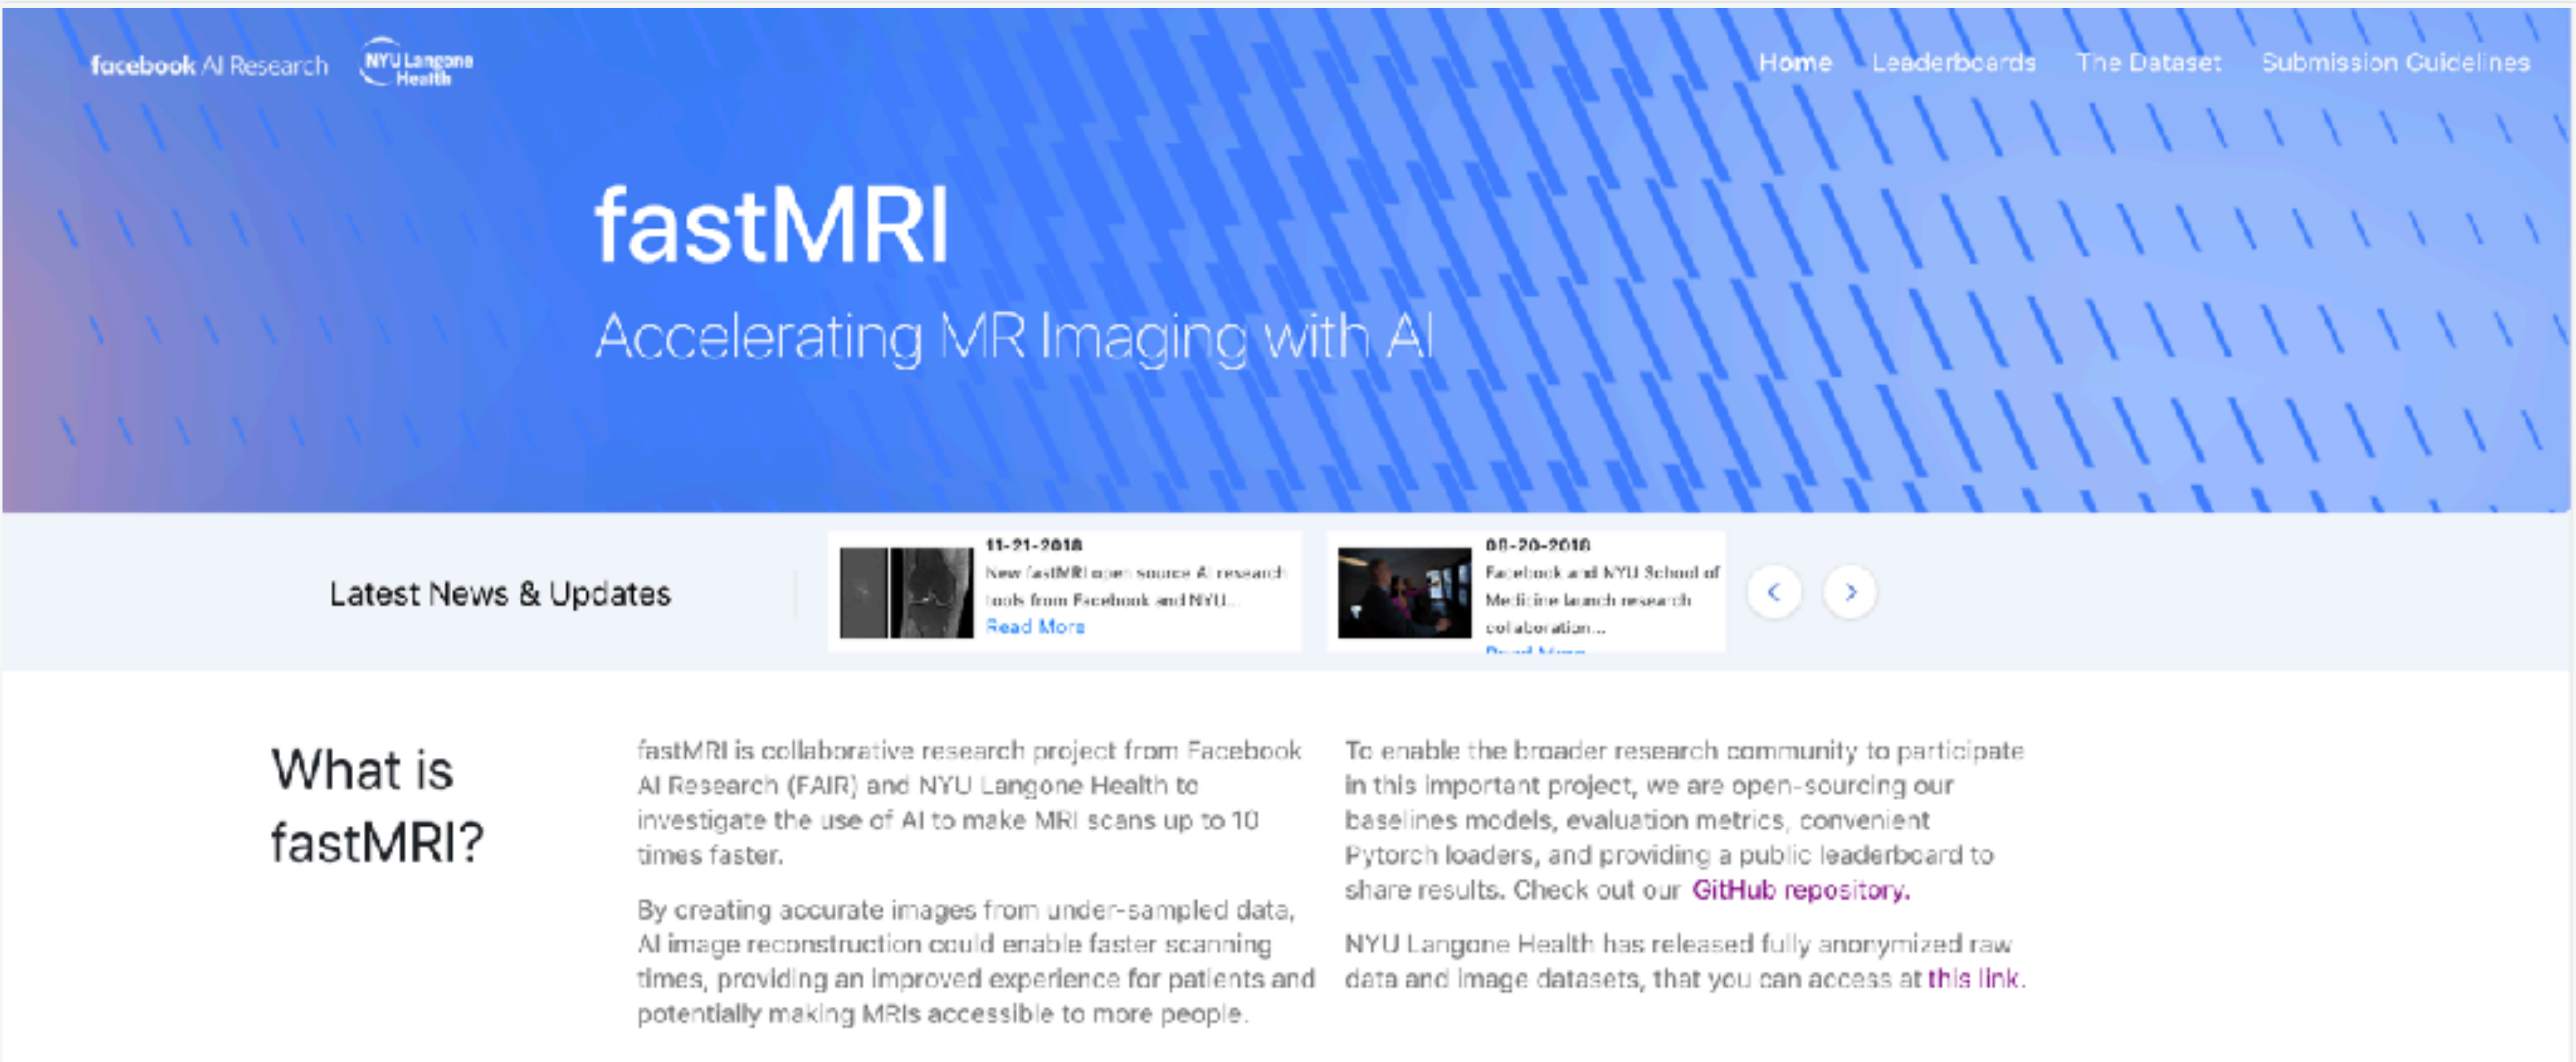

http://fastmri.org/

Is reconstruction even necessary?

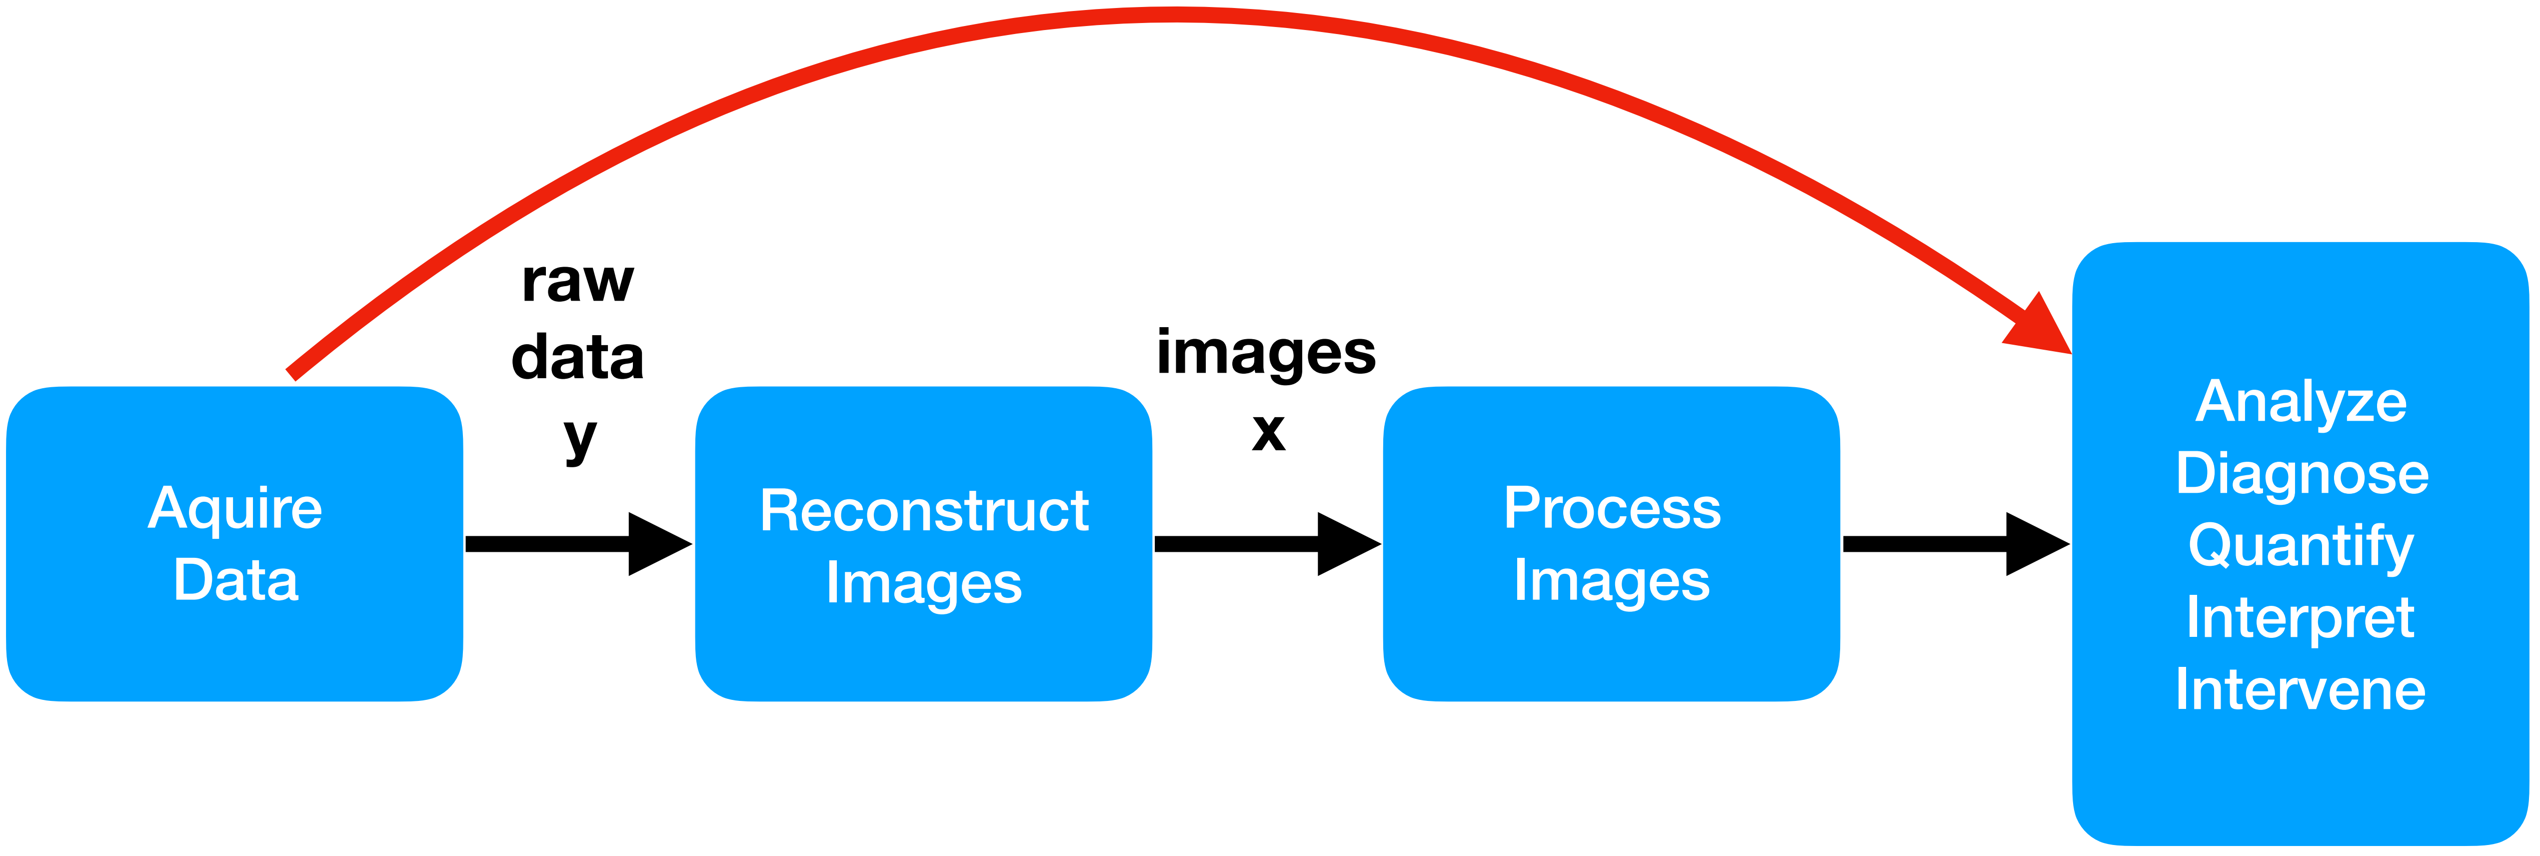

Thanks!

Additional reading:

Litjens, G., et al. (2017). A survey on deep learning in medical image analysis. Medical image analysis, 42, 60-88.

Lundervold, A. S., & Lundervold, A. (2019). An overview of deep learning in medical imaging focusing on MRI. Zeitschrift für Medizinische Physik, 29(2), 102-127.

Kaggle data science bowl 2017: Lung nodule classification
https://www.kaggle.com/c/data-science-bowl-2017/overview

McCann, M. T., Jin, K. H., & Unser, M. (2017). Convolutional neural networks for inverse problems in imaging: A review. IEEE Signal Processing Magazine, 34(6), 85-95.

email: gongie@uchicago.edu
web: gregongie.github.io