

# Lecture 24: Looking Ahead + Review

Machine Learning and Imaging

BME 548L

Roarke Horstmeyer

## Announcements and schedule

- Today is the last lecture! (no lectures next week)
- HW4 being graded, HW5 will be graded quickly once turned in
- Homework #5 Due Wed April 24 at 11:59pm
- Final code and presentation due via email: Thursday May 2, 5pm
- Final presentation time slots: Thursday May 2, 7pm - 10pm
- Sign up for slot at google sheets link will be shared via Ed Discussion
- Final presentation paper write-up, website template and permission form due: Saturday May 4 at 11:59pm
- Project help:
  - We will continue lab sessions this week/next week
    - Office hours Tuesday 11am – noon
    - Email me if you'd like to meet another time
    - Email TA's / reach out on Canvas with questions!

## **Project content details –**

- 1. Start with code-base and annotated data**
  - e.g., downloaded or from existing work
  
- 2. Make sure it works and gives reasonable results**
  - work with TA's and others to make sure code runs successfully!
  
- 3. Main project component – experiment and explore with code and data**
  - A. do not just randomly alter neural network architecture**
  
  - B. instead, explore something meaningful about how image data was acquired, properties of the image data, different ML-related goals for data**
  
  - C. One useful direction is to incorporate “physical layers” – trainable weights that optimize some aspect of image capture. But this is an effective way to explore point B above!**
  
- 4. Discuss insights gained from step 3**

## Components of final project

See <https://deepimaging.github.io/proj-info/>

38% of total grade

1. Presentation Slides – 10%
  - 7-minute presentation, 2 minute for questions
2. 4-6 page write up with at least 3 figures and 5 references – 20%
  - Introduction, related work, methods, results, discussion
3. Code used for final results in folder or .ipynb's – 4%
4. brief website template & permission to share results – 4%
5. shared annotated datasets & permissions – no grade, but would be much appreciated if using an interesting dataset

## Final project webpage

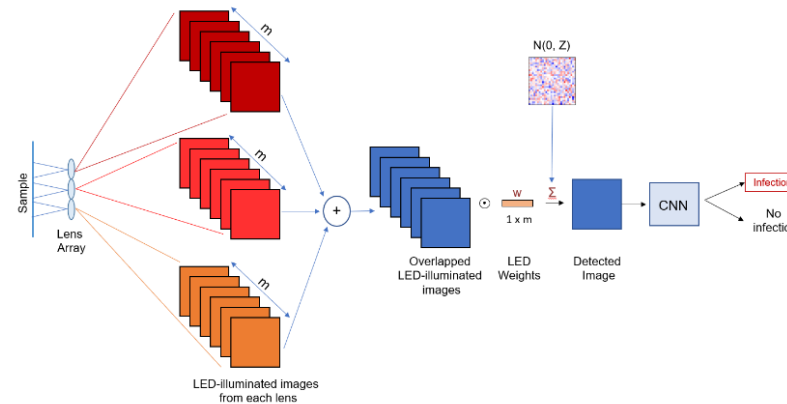
- Must be submitted
- Will share template
- Will post to [deepimaging.io](https://deepimaging.io) with permission
- Will also send permission form, which must be submitted with final project as well

## Optimizing illumination for overlapped classification

Amey A Chaware

[amey.chaware@duke.edu](mailto:amey.chaware@duke.edu)

[Paper PDF](#)



This project presents an imaging system that simultaneously captures multiple images and automatically classifies their contents to increase detection throughput. Our optical design consists of a set of multiple lenses that each image a unique field-of-view onto a single image sensor. The resulting “overlapped” image exhibits reduced contrast, but includes measurements from across a proportionally larger viewing area. We then post-process this overlapped image with a deep convolutional neural network to classify the presence or absence of certain features of interest. We examine the specific case of detecting the malaria parasite within overlapped microscope images of blood smears. We demonstrate that it is possible to overlap 7 unique images onto a common sensor while still offering accurate classification of the presence or absence of the parasite, thus offering a 7x potential speed-up for automated disease diagnosis with microscope image data. Additionally, we explore the use of supervised deep-learning network to jointly optimize the physical setup of an optical microscope to improve automatic image classification accuracy in overlapped imaging. We take advantage of the wide degree of flexibility available in choosing how a sample is illuminated in a microscope to design a specific pattern of light that leads to a better performance.

Paper:

- [Paper PDF](#)

Code and Data:

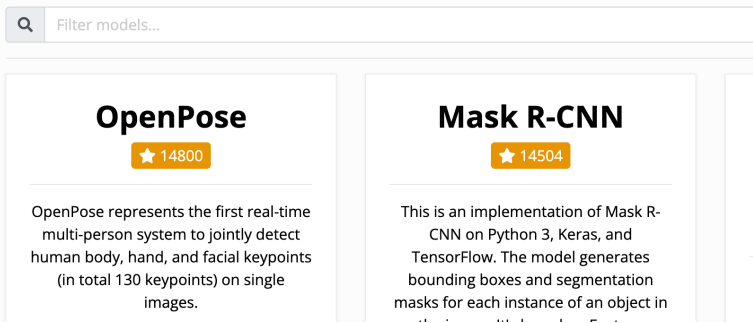
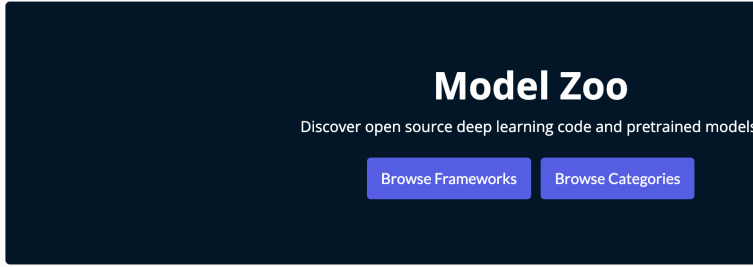

- You can provide a link to your code here: [Code](#)

**Please complete course evaluations!**

# Where are things going with Machine Learning and Imaging in 10 years?


# Where are things going with Machine Learning and Imaging in 10 years?

1. Proliferation of trained models, similar datasets, and novel programming languages



## Welcome to the OpenAI developer platform

### Start with the basics



### Build an assistant




### Explore the API



## Visipedia

### Visipedia connects





# Where are things going with Machine Learning and Imaging in 10 years?

1. Proliferation of trained models, similar datasets, and novel programming languages

Environment developed by Facebook (Meta) – arguably more popular for research use



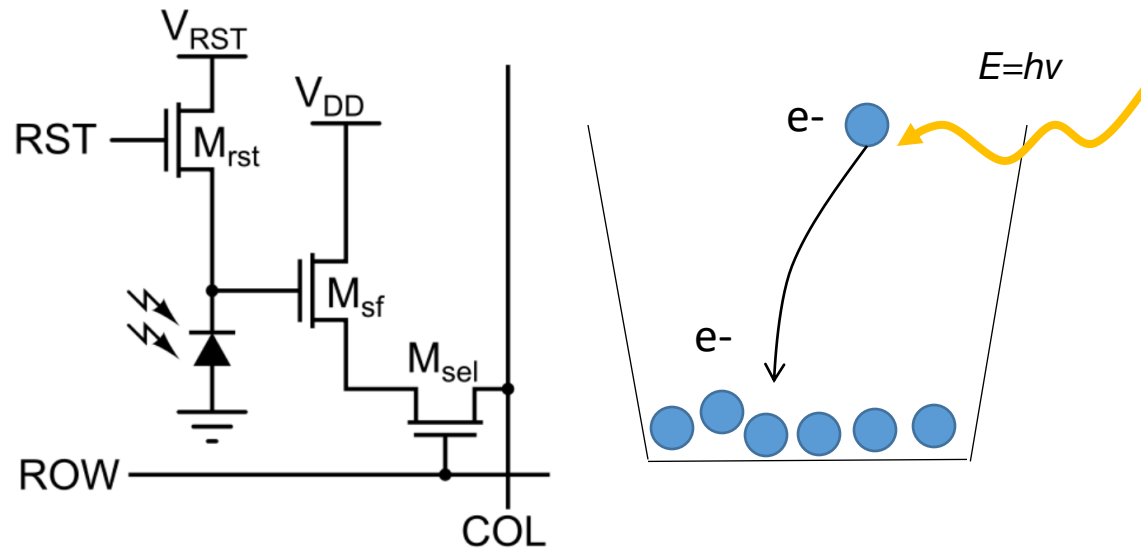
Jax: Rapidly deploy the same code on GPU's and TPU's




# Where are things going with Machine Learning and Imaging in 10 years?

## 2. “Cameras” on many devices & new types of sensors

Standard CMOS pixel = bucket that collects electrons





SPAD pixel: was there a photon or not?



# Where are things going with Machine Learning and Imaging in 10 years?

## 2. “Cameras” on many devices & new types of sensors

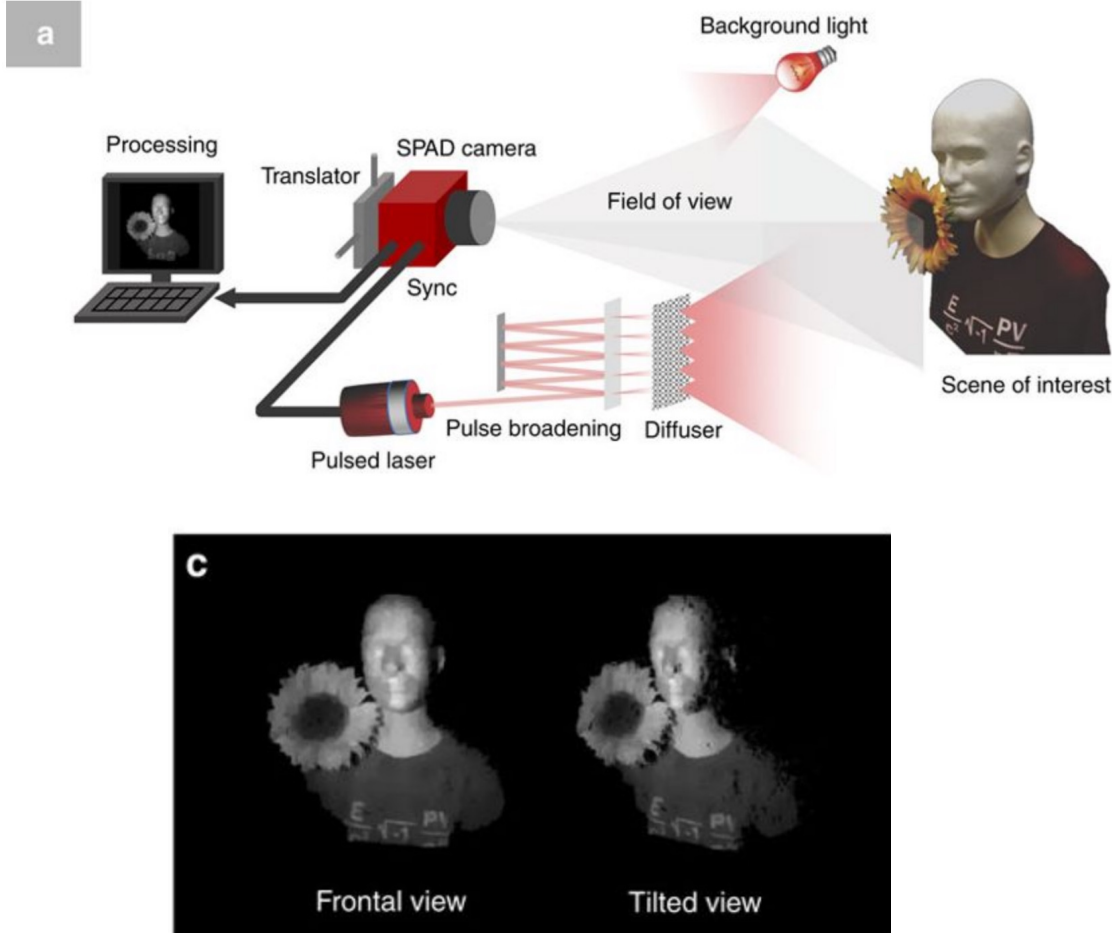
- Light travels 1 ft in 1 ns.
- SPADs can precisely photon arrival time to measure travel distance (TOF)



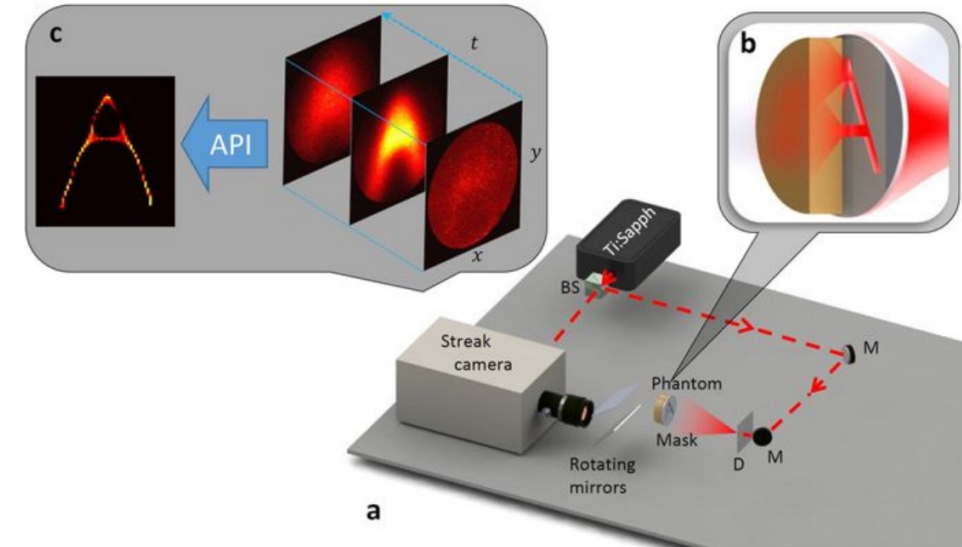
[https://www.picoquant.com/images/uploads/page/files/7253/technote\\_tcspc.pdf](https://www.picoquant.com/images/uploads/page/files/7253/technote_tcspc.pdf)

# Where are things going with Machine Learning and Imaging in 10 years?

## 2. “Cameras” on many devices & new types of sensors



**Figure 1: Imaging Through Thick Scattering.**



G. Satat et al, <https://www.nature.com/articles/srep33946>

D. Shen et al, <https://www.nature.com/articles/ncomms12046>

# LiDAR Scanner

LiDAR (Light Detection and Ranging) is used to determine distance by measuring how long it takes light to reach an object and reflect back. It is so advanced, it's being used by NASA for the next Mars landing mission. And it's now been engineered to fit in the thin and light iPad Pro.

The custom-designed LiDAR Scanner uses direct time of flight to measure reflected light from up to five meters away, both indoors and out. It works at the photon level, operates at nanosecond speeds, and opens up tremendous possibilities for augmented reality and beyond.



# Where are things going with Machine Learning and Imaging in 10 years?

3. Beyond convolutions - new constructs for deep networks

# Where are things going with Machine Learning and Imaging in 10 years?

## 3. Beyond convolutions - new constructs for deep networks

---

### Dynamic Routing Between Capsules

---

**Sara Sabour**



**Nicholas Frosst**

**Geoffrey E. Hinton**

Google Brain

Toronto

{sasabour, frosst, geoffhinton}@google.com



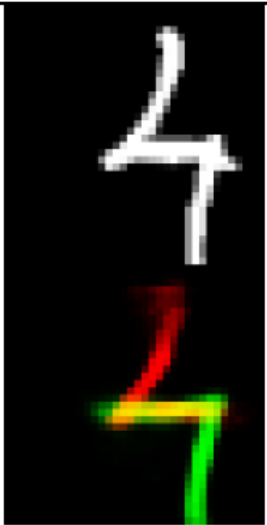

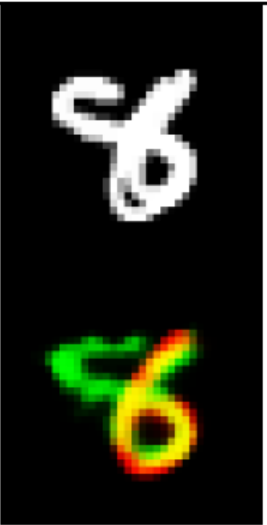
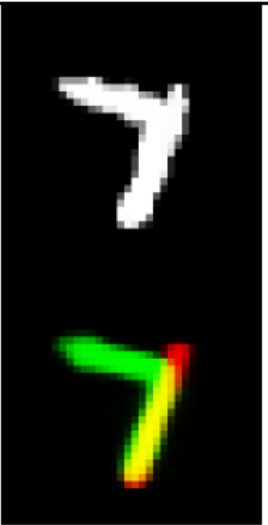
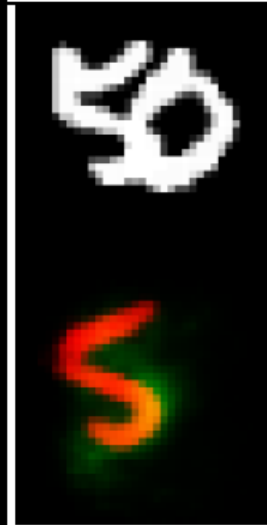

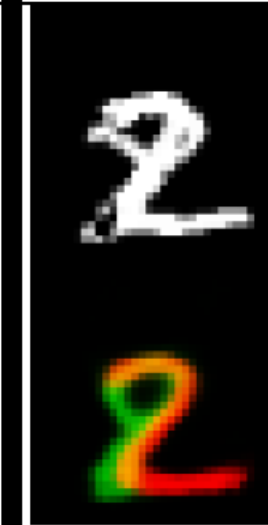

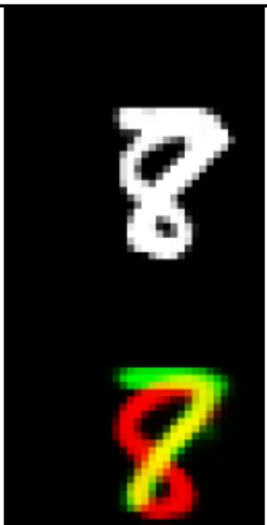
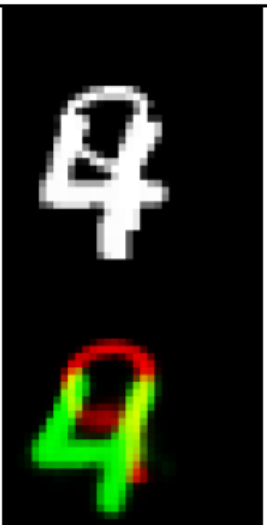
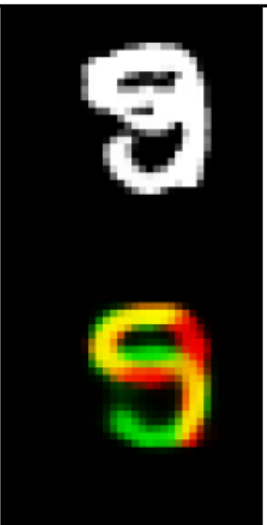
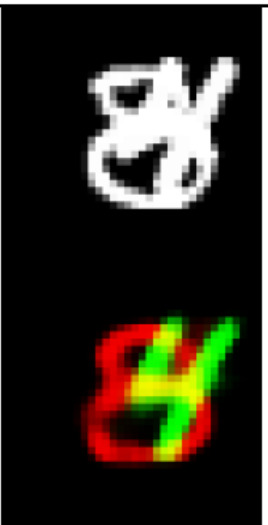
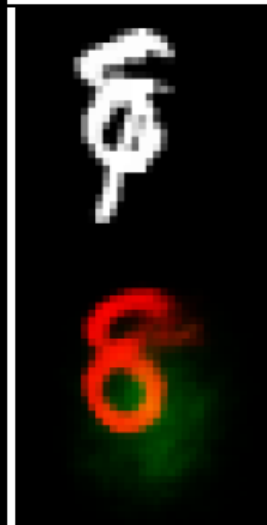
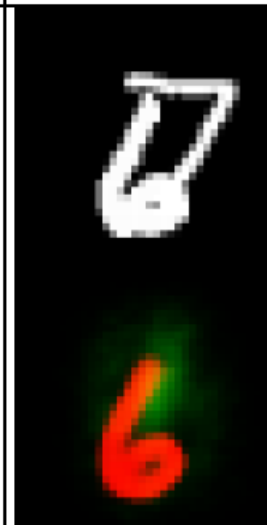
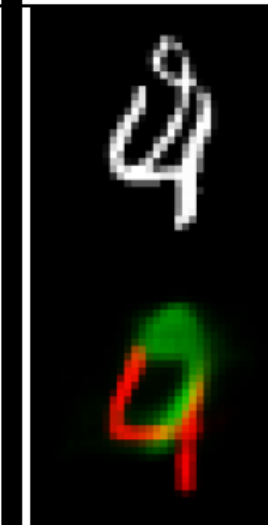



# Where are things going with Machine Learning and Imaging in 10 years?

## 3. Beyond convolutions - new constructs for deep networks

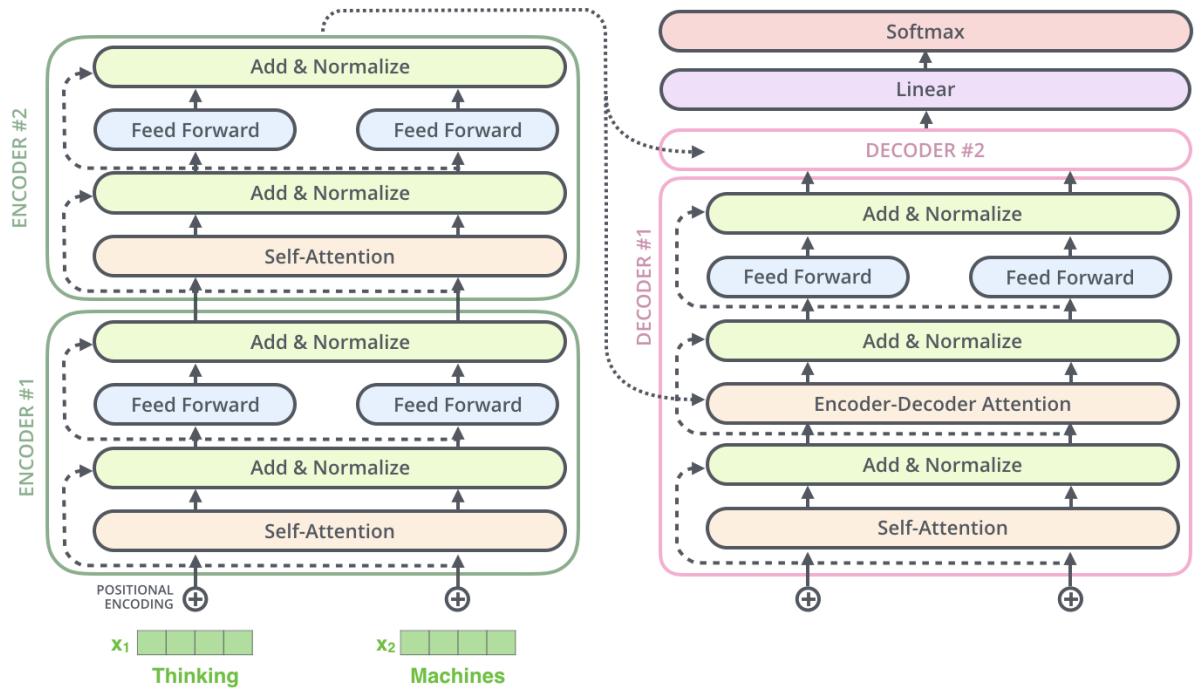
Capsule vs. Traditional Neuron			
Input from low-level capsule/neuron	vector( $\mathbf{u}_i$ )	scalar( $x_i$ )	
Operation	Affine Transform	$\hat{\mathbf{u}}_{j i} = \mathbf{W}_{ij}\mathbf{u}_i$	–
	Weighting	$\mathbf{s}_j = \sum_i c_{ij}\hat{\mathbf{u}}_{j i}$	$a_j = \sum_i w_i x_i + b$
	Sum		
	Nonlinear Activation	$\mathbf{v}_j = \frac{\ \mathbf{s}_j\ ^2}{1+\ \mathbf{s}_j\ ^2} \frac{\mathbf{s}_j}{\ \mathbf{s}_j\ }$	$h_j = f(a_j)$
Output	vector( $\mathbf{v}_j$ )	scalar( $h_j$ )	

Therefore it will not be enough to give the parts two weights if one of them does not have any other support.

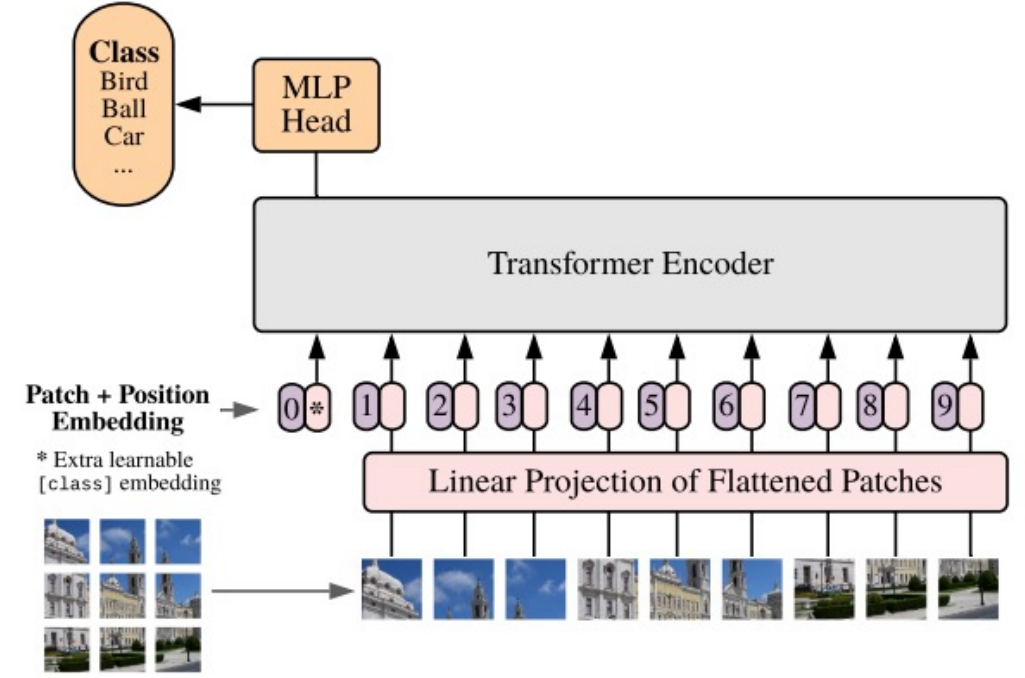
R:(2, 7) L:(2, 7)	R:(6, 0) L:(6, 0)	R:(6, 8) L:(6, 8)	R:(7, 1) L:(7, 1)	*R:(5, 7) L:(5, 0)	*R:(2, 3) L:(4, 3)	R:(2, 8) L:(2, 8)	R:P:(2, 7) L:(2, 8)
							
R:(8, 7) L:(8, 7)	R:(9, 4) L:(9, 4)	R:(9, 5) L:(9, 5)	R:(8, 4) L:(8, 4)	*R:(0, 8) L:(1, 8)	*R:(1, 6) L:(7, 6)	R:(4, 9) L:(4, 9)	R:P:(4, 0) L:(4, 9)
							

## Example 2: Transformers for image analysis

Transformers for text analysis/generation




### Vision Transformer (ViT)




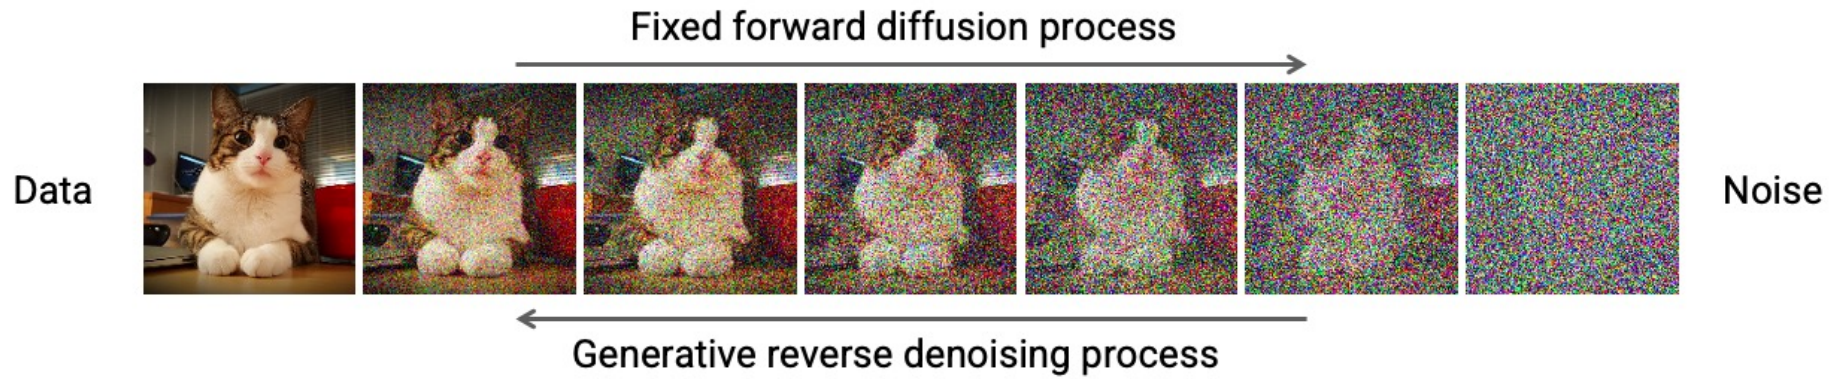
A. Dosovitsky et al, "AN IMAGE IS WORTH 16X16 WORDS: TRANSFORMERS FOR IMAGE RECOGNITION AT SCALE"


# A visual introduction to transformers

- <https://www.youtube.com/watch?v=wjZofJX0v4M>



## Example 2: Diffusion models for image generation and analysis

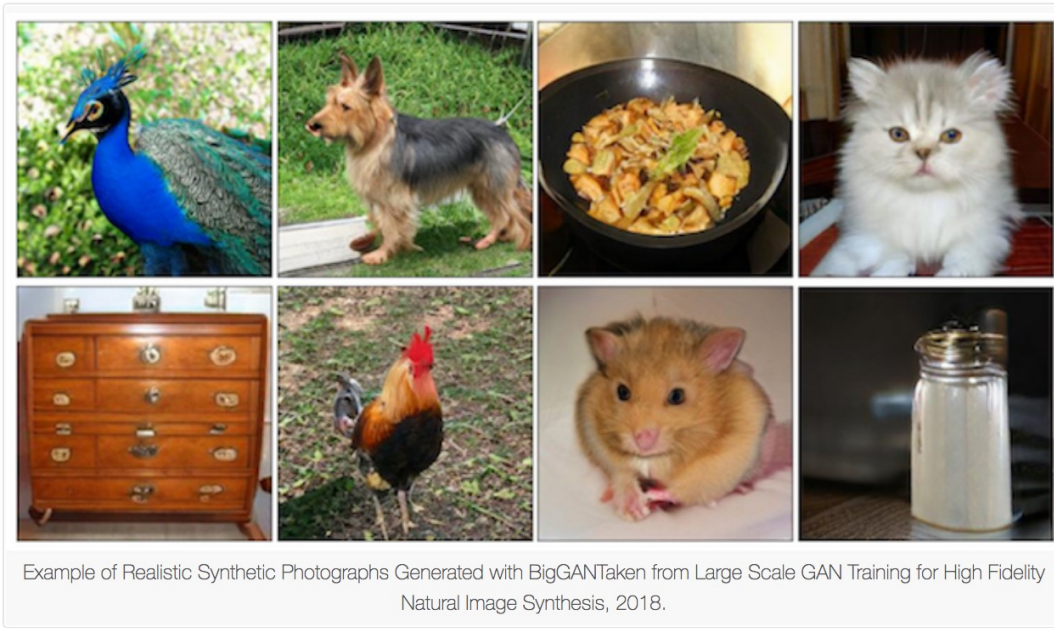




<https://research.nvidia.com/labs/toronto-ai/VideoLDM/>

# Where are things going with Machine Learning and Imaging in 10 years?

## 4. Generative data is getting pretty realistic...




### Face2Face: Real-time Face Capture and Reenactment of RGB Videos

[Justus Thies](#)<sup>1</sup>   [Michael Zollhöfer](#)<sup>2</sup>   [Marc Stamminger](#)<sup>3</sup>   [Christian Theobalt](#)<sup>2</sup>   [Matthias Nießner](#)<sup>1</sup>

<sup>1</sup> Technical University of Munich

<sup>2</sup> Max Planck Institute for Informatics

<sup>3</sup> University of Erlangen-Nuremberg



Proc. Computer Vision and Pattern Recognition (CVPR), IEEE, June 2016

## What are the implications of this for medical imaging?

## 5. Models will no longer be created in an ad-hoc manner



# Model Search

Model search (MS) is a framework that implements AutoML algorithms for model architecture search at scale. It aims to help researchers speed up their exploration process for finding the right model architecture for their classification problems (i.e., DNNs with different types of layers).

Input





## 6. Joint optimization of hardware and software is proliferating

Check for updates

**technology feature**


# Smart solutions for automated imaging

Algorithms trained to interpret microscope data can greatly extend the information that can be derived from the resulting images, or even optimize how imaging experiments are conducted.

Michael Eisenstein

While buzzing about in search of food, a fruit fly encounters a deadly wasp. Fortunately, its brain reacts to the threat by initiating a cascade of responses across a network of neurons that help it to flee. Philipp Keller's group at the Howard Hughes Medical Institute's Janelia Research Campus has developed a variety of sophisticated strategies for deconvolving the circuitry underlying this and other complex functions of the *Drosophila* nervous system, using a combination of optogenetic manipulation and cutting-edge light-sheet microscopy to simulate various stimuli in living tissue and analyze the response. But perhaps the most remarkable aspect of this project is the extent to which the instruments themselves are running the show. "The microscope can basically do these experiments completely on its own," says Keller.

This work is a particularly advanced example of an emerging field of computer-assisted imaging known as 'smart microscopy'. In these configurations, the



## Take-aways for the future of machine learning and imaging


1. It's not going away....it works, there's a big community, and lots of funding
2. Hardware *and* software are rapidly evolving
3. CNN's work very well, but they are not the final solution...Transformer models and Diffusion networks are using CNN's as building blocks
4. There is currently a lack of safeguards and not enough consideration for how to ensure processed results are accurate, secure and trustworthy
5. Merger of hardware and software for key applications is inevitable...

## Ethical questions surrounding deep convolutional networks


1. What are your expectations for an image reconstruction algorithm used in a clinical setting?
2. What types of “guarantees” should we be able to make, if any, to a patient?
3. How should we guide future development of ML software to meet any guarantees?
4. How should we guide future development of ML-designed hardware to meet any guarantees?
5. Thoughts towards a system of checks and balances?

# What was this class about?

Where your final project is aimed



*Computer-centered hardware + software*



Human-centered hardware design

Computer-centered software design

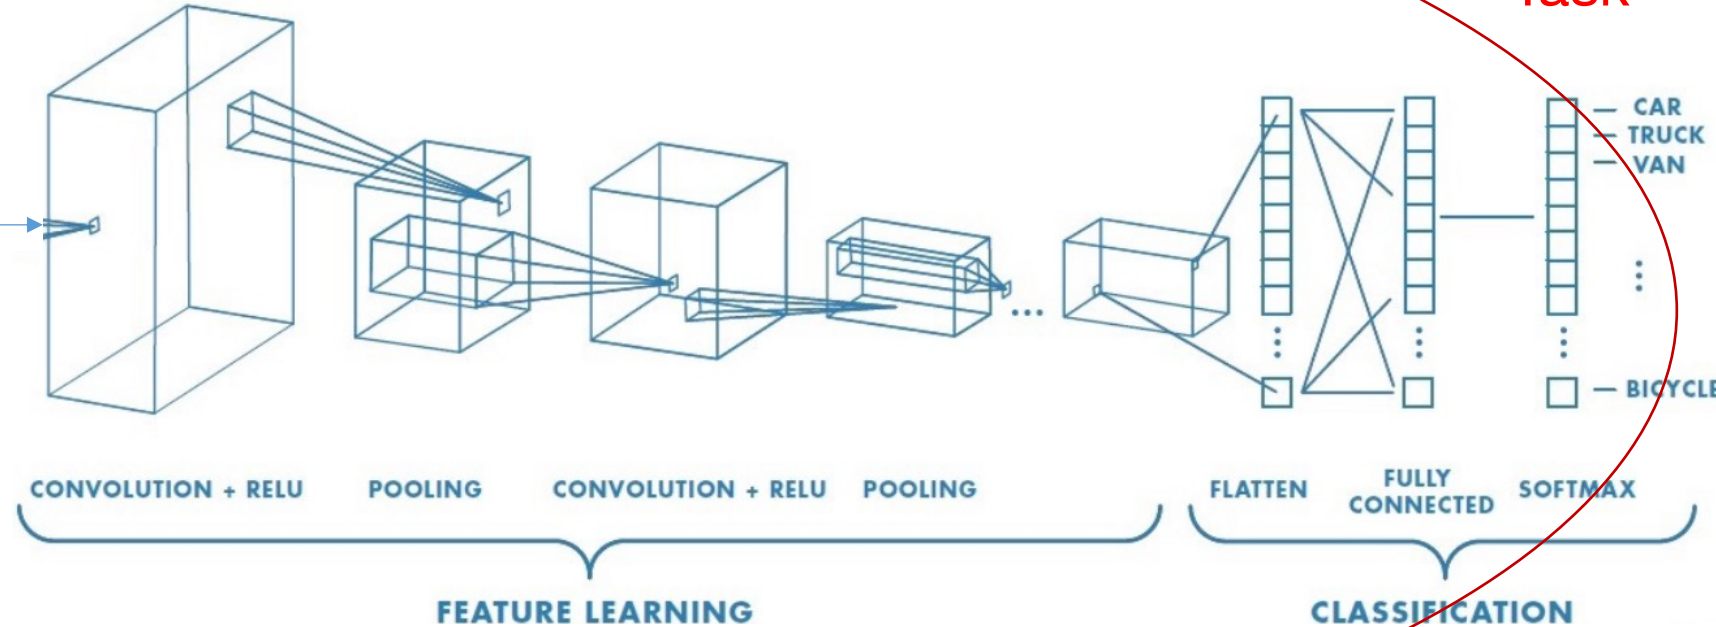
# Bringing together physical and digital image representations

Physical world  $I_0$




Hardware

Image  $I_1$




Final project: try to optimize all of this together!

# ML+Imaging pipeline introduction



Continuous complex fields


Measurement device



Black box transformations

- Convolution
- Fourier Transform


Digitization



Sampling Theorem

Discrete math & Linear algebra

Machine Learning



Optimization

Linear classification



Logistic classifier

Neural networks

Convolutional NN's

# Physical models for light propagation to sensor

- Interpretation #1: Radiation (*Incoherent*)
- Model: Rays





- Real, non-negative

$$I_s = H B S_0$$

- Sample absorption **S**
- Illumination brightness **B**
- Blur in **H**

- Interpretation #2: Electromagnetic wave (*Coherent*)
- Model: Waves



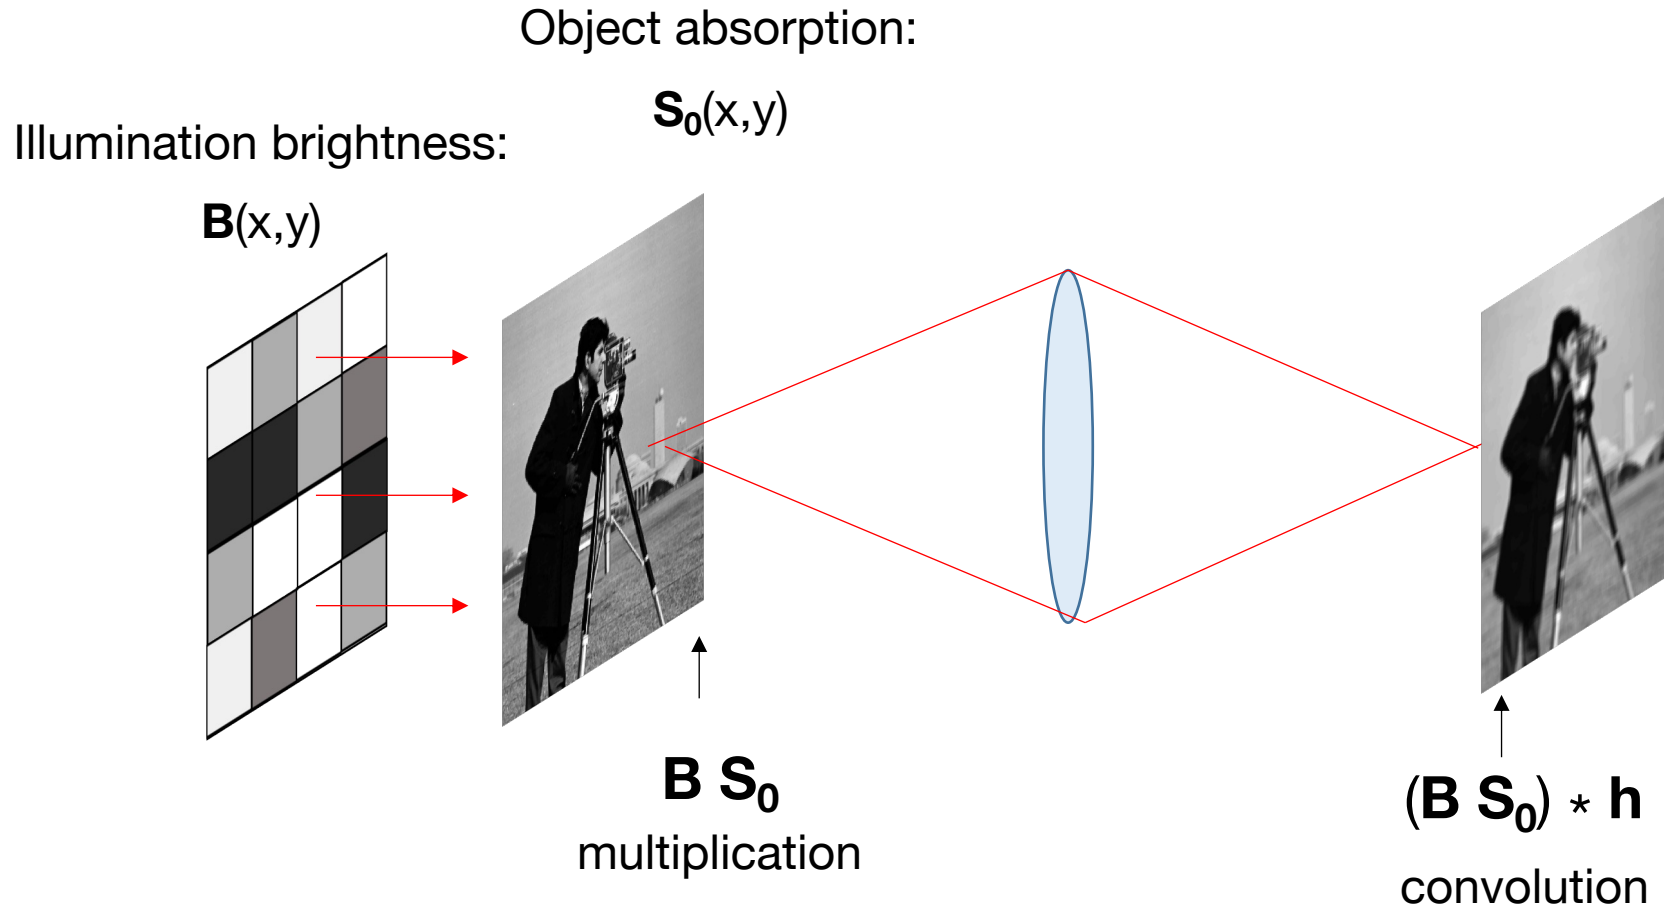
- Complex field

$$I_c = |H C S_c|^2$$

- Sample abs./phase **S**
- Illumination wave **B**
- Blur in **H**


# Mathematical model of for incoherent image formation

- All quantities are real, and non-negative






# Model of image formation for wave optics (coherent light):



# Machine Learning



Optimization


Linear classification

Logistic classifier

Neural networks

Convolutional NN's


## Digitization



Sampling Theorem

Discrete math & Linear algebra

## Measurement device




Black box transformations

- Convolution
- Fourier Transform

$$I_{Inc} = H B S_0$$

$$I_{Coh} = |H C S_c|^2$$

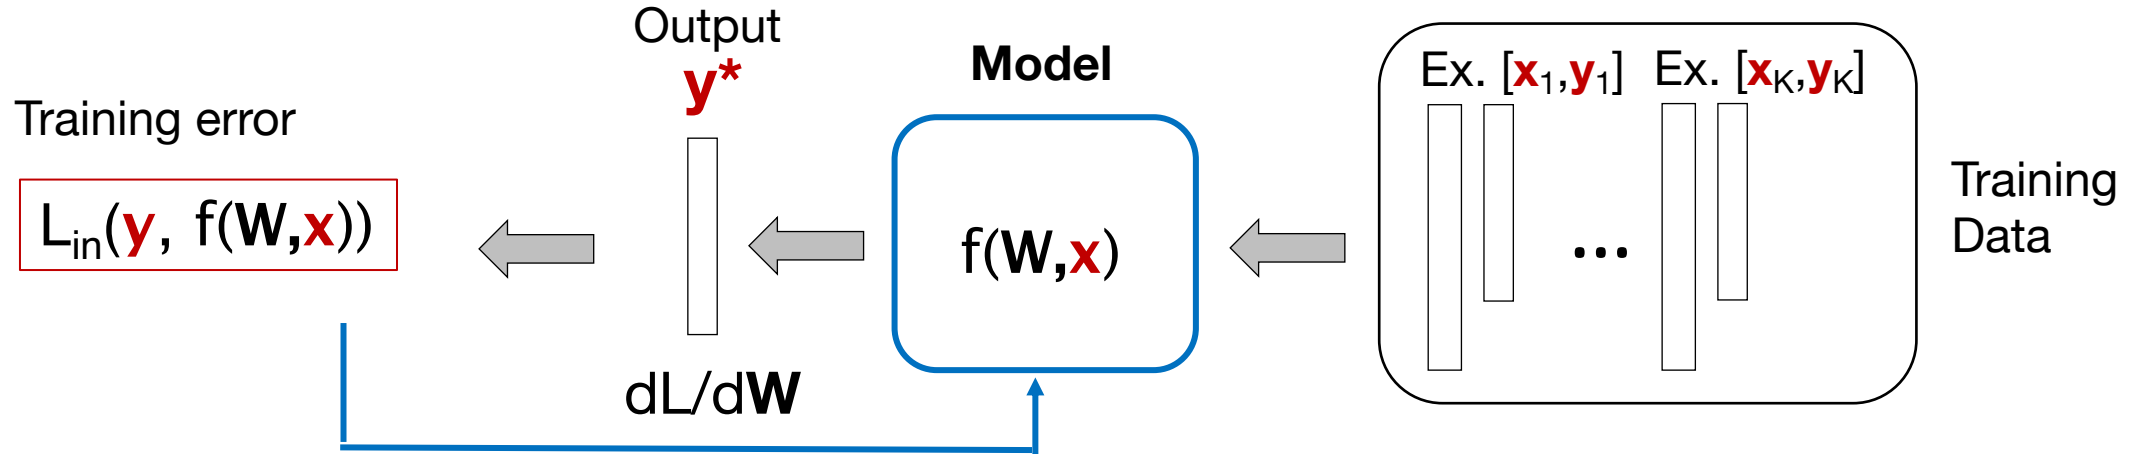
## Real World



Continuous complex fields

# Summary of machine learning pipeline:

## 1. Network Training

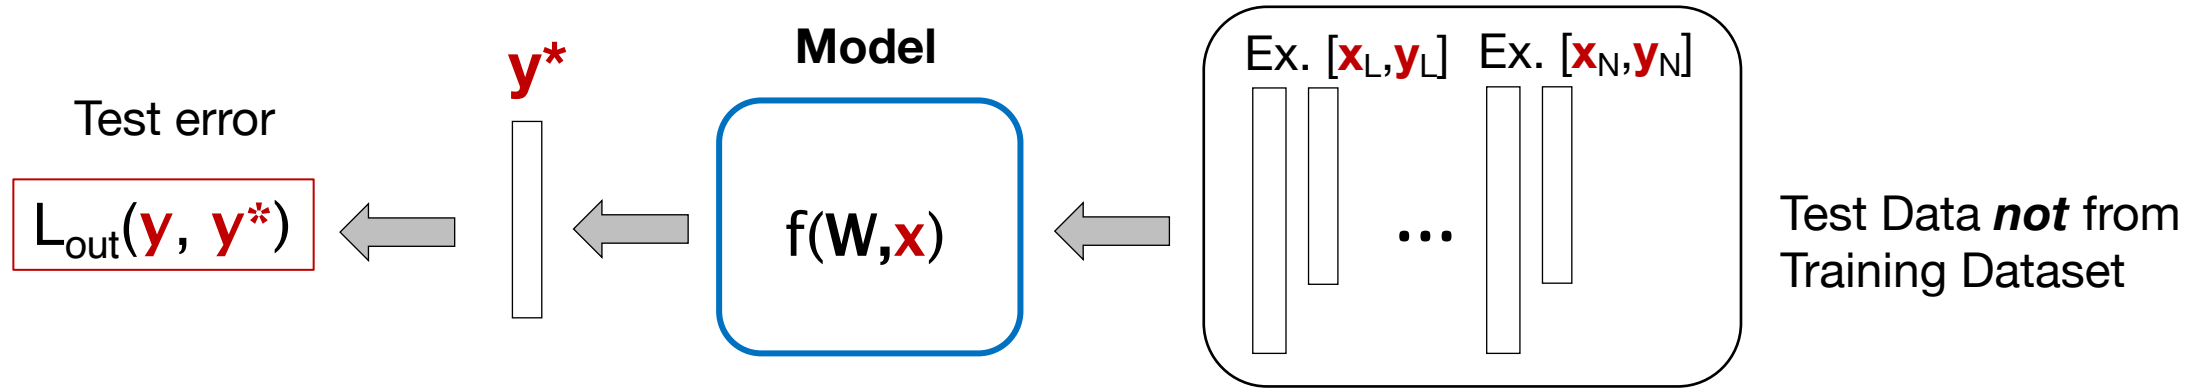


What we need for network training:

1. Labeled examples
2. A model and loss function
3. A way to minimize the loss function  $L$

# Summary of machine learning pipeline:

## 2. Network Testing



What we need for network testing:


- 4. **Unique** labeled test data
- 5. **Evaluation** of model error

# Let's start with a simpler approach: linear regression

General linear model:


$$L = \frac{1}{N} \sum_{i=1}^N L_i(f(x_i, \mathbf{W}), y_i)$$

$$L = \frac{1}{N} \sum_{i=1}^N L_i(\mathbf{W}\mathbf{x}_i, y_i)$$



Assume 1 class = 1 linear fit

$$L = \frac{1}{N} \sum_{i=1}^N L_i(w^T x_i - y_i)$$



Use MSE error model

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$

Where labels determined by thresholding


$$f(\mathbf{x}_i) = y_i^* = \text{sgn}(\mathbf{w}^T \mathbf{x}_i) \quad \text{sgn}(x) := \begin{cases} -1 & \text{if } x < 0, \\ 0 & \text{if } x = 0, \\ 1 & \text{if } x > 0. \end{cases}$$

# Let's start with a simpler approach: linear regression

General linear model:


$$L = \frac{1}{N} \sum_{i=1}^N L_i(f(x_i, \mathbf{W}), y_i)$$

$$L = \frac{1}{N} \sum_{i=1}^N L_i(\mathbf{W}\mathbf{x}_i, y_i)$$



Assume 1 class = 1 linear fit

$$L = \frac{1}{N} \sum_{i=1}^N L_i(w^T x_i - y_i)$$




Use MSE error model

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$

Where labels determined by thresholding

$$f(\mathbf{x}_i) = y_i^* = \text{sgn}(\mathbf{w}^T \mathbf{x}_i) \quad \text{sgn}(x) := \begin{cases} -1 & \text{if } x < 0, \\ 0 & \text{if } x = 0, \\ 1 & \text{if } x > 0. \end{cases}$$

# Why does linear regression with $\text{sgn}()$ achieve classification?




Without  $\text{sgn}()$ : regression for best fit

$$f(\mathbf{x}_i) = \mathbf{w}^T \mathbf{x}_i$$

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$

- If  $y_i$  can be anything, minimizing  $L$  makes  $\mathbf{w}$  the plane of best fit

# Why does linear regression with $\text{sgn}()$ achieve classification?



Without  $\text{sgn}()$ : regression for best fit


$$f(\mathbf{x}_i) = \mathbf{w}^T \mathbf{x}_i$$

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$

- $y_i$  can only be -1 or +1, which defines its class



# Why does linear regression with $\text{sgn}()$ achieve classification?




Without  $\text{sgn}()$ : regression for best fit

$$f(\mathbf{x}_i) = \mathbf{w}^T \mathbf{x}_i$$

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$

- $y_i$  can only be -1 or +1, which defines its class
- Can still find plane of best fit

# Why does linear regression with $\text{sgn}()$ achieve classification?




With  $\text{sgn}()$  operation:

$$f(\mathbf{x}_i) = y_i^* = \text{sgn}(\mathbf{w}^T \mathbf{x}_i)$$

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$

- Anything point to one side of  $y=0$  intersection is class +1, anything on the other side of intersection is class -1

# Why does linear regression with $\text{sgn}()$ achieve classification?



With  $\text{sgn}()$  operation:

$$f(\mathbf{x}_i) = y_i^* = \text{sgn}(\mathbf{w}^T \mathbf{x}_i)$$

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$


- $y$  axis isn't really needed now & can view this decision boundary in 2D

# Why does linear regression with $\text{sgn}()$ achieve classification?


With  $\text{sgn}()$  operation:

$$f(\mathbf{x}_i) = y_i^* = \text{sgn}(\mathbf{w}^T \mathbf{x}_i)$$

$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$



Sign operation takes linear regression and makes it a classification operation!



## How to minimize $L_{\text{in}}$

For logistic regression,

$$L_{\text{in}}(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N \ln \left( 1 + e^{-y_n \mathbf{w}^T \mathbf{x}_n} \right)$$

← **iterative** solution

Compare to linear regression:

$$L_{\text{in}}(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N (\mathbf{w}^T \mathbf{x}_n - y_n)^2$$

← **closed-form** solution

# In the rest of this class: solve via gradient descent



With a matrix, compute this for each entry:

$$\frac{dL(W_i)}{dW_i} = \lim_{h \rightarrow 0} \frac{L(W_i + h) - L(W_i)}{h}$$

Example:

$$W = [1,2;3,4]$$
$$L(W, x, y) = 12.79$$

$$W_{1+h} = [1.001,2;3,4]$$
$$L(W_{1+h}, x, y) = 12.8$$

$$dL(W_1)/dW_1 = 12.8 - 12.79 / 0.001$$

$$dL(W_1)/dW_1 = 10$$



- Repeat for all entries of  $\mathbf{W}$ ,  $dL/d\mathbf{W}$  will have  $N \times M$  entries for  $N \times M$  matrix

## Steepest descent and the best step size $\epsilon$

1. Evaluate function  $f(\mathbf{x}^{(0)})$  at an initial guess point,  $\mathbf{x}^{(0)}$
2. Compute gradient  $\mathbf{g}^{(0)} = \nabla_{\mathbf{x}} f(\mathbf{x}^{(0)})$
3. Next point  $\mathbf{x}^{(1)} = \mathbf{x}^{(0)} - \epsilon^{(0)} \mathbf{g}^{(0)}$
4. Repeat –  $\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - \epsilon^{(n)} \mathbf{g}^{(n)}$ , until  $|\mathbf{x}^{(n+1)} - \mathbf{x}^{(n)}| < \text{threshold } t$



$$L = \frac{1}{N} \sum_{i=1}^N (w^T x_i - y_i)^2$$
$$\nabla L(w) = \frac{2}{N} X^T (Xw - y) = 0$$

```
while previous_step_size > precision and iters < max_iters:  
    prev_x = cur_x  
    cur_x -= gamma * df(prev_x)  
    previous_step_size = abs(cur_x - prev_x)  
    iters+=1
```





$$f = W_1x$$

Learned  $f$ : not flexible



$$f = W_2\max(W_1x, 0)$$

Learned  $f$ : a bit flexible



$$f = W_3\max(0, W_2\max(W_1x, 0))$$

Learned  $f$ : more flexible

Does it generalize???

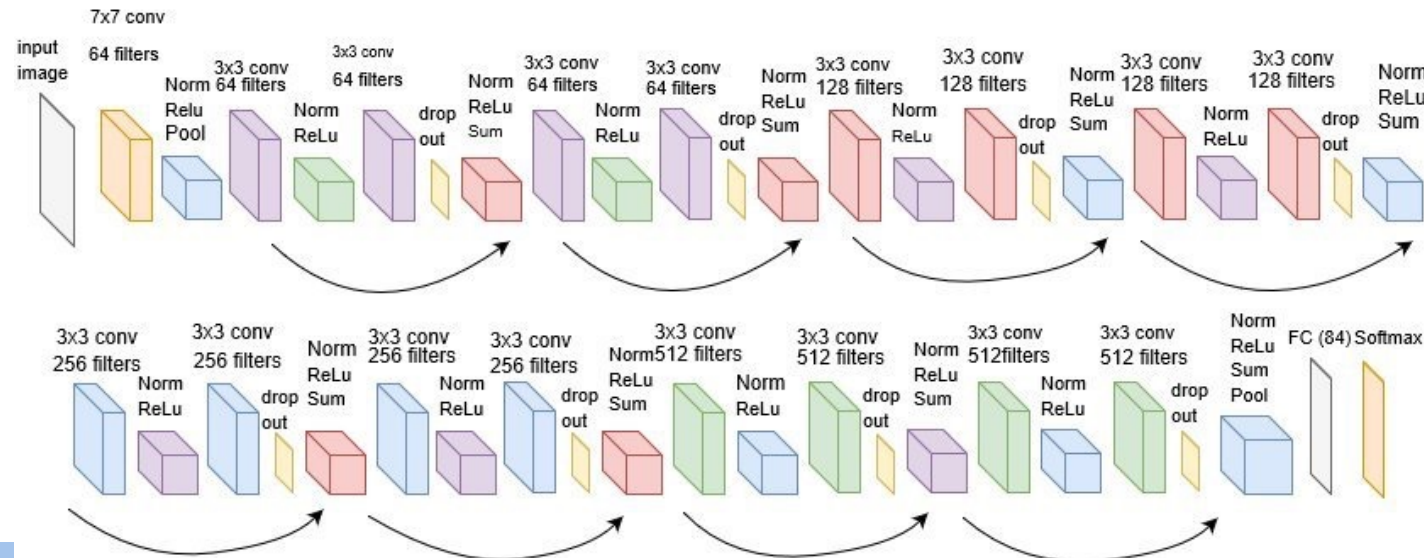

↓  
We can keep adding these "layers"...




## Before CNN's – understand two competing goals in machine learning

1. Can we make sure the in-sample error  $L_{in}(y, f(x,W))$  is small enough?
  - Appropriate cost function
  - “complex enough” model
  
2. Can we make sure that  $L_{out}(y, f(x,W))$  is close enough to  $L_{in}(y, f(x,W))$ ?
  - Probabilistic analysis says yes!
  - $|L_{in} - L_{out}|$  bounded from above
  - Bound grows with model capacity (bad)
  - Bound shrinks with # of training examples (good)

# Gets us to Convolutional Neural Networks



# Why should we use these convolutions?



$x = \text{cat image}$


Image interpretation



This type of matrix can dramatically reduce the number of weights that are used while still allowing *local* regions to mix:

Full matrix:  $O(n^2)$

**Banded matrix:  $k \cdot O(n)$**



Mix all the pixels in the red box, with associated weights, to form this entry of  $S$

# Important components of a CNN

## CNN Architecture

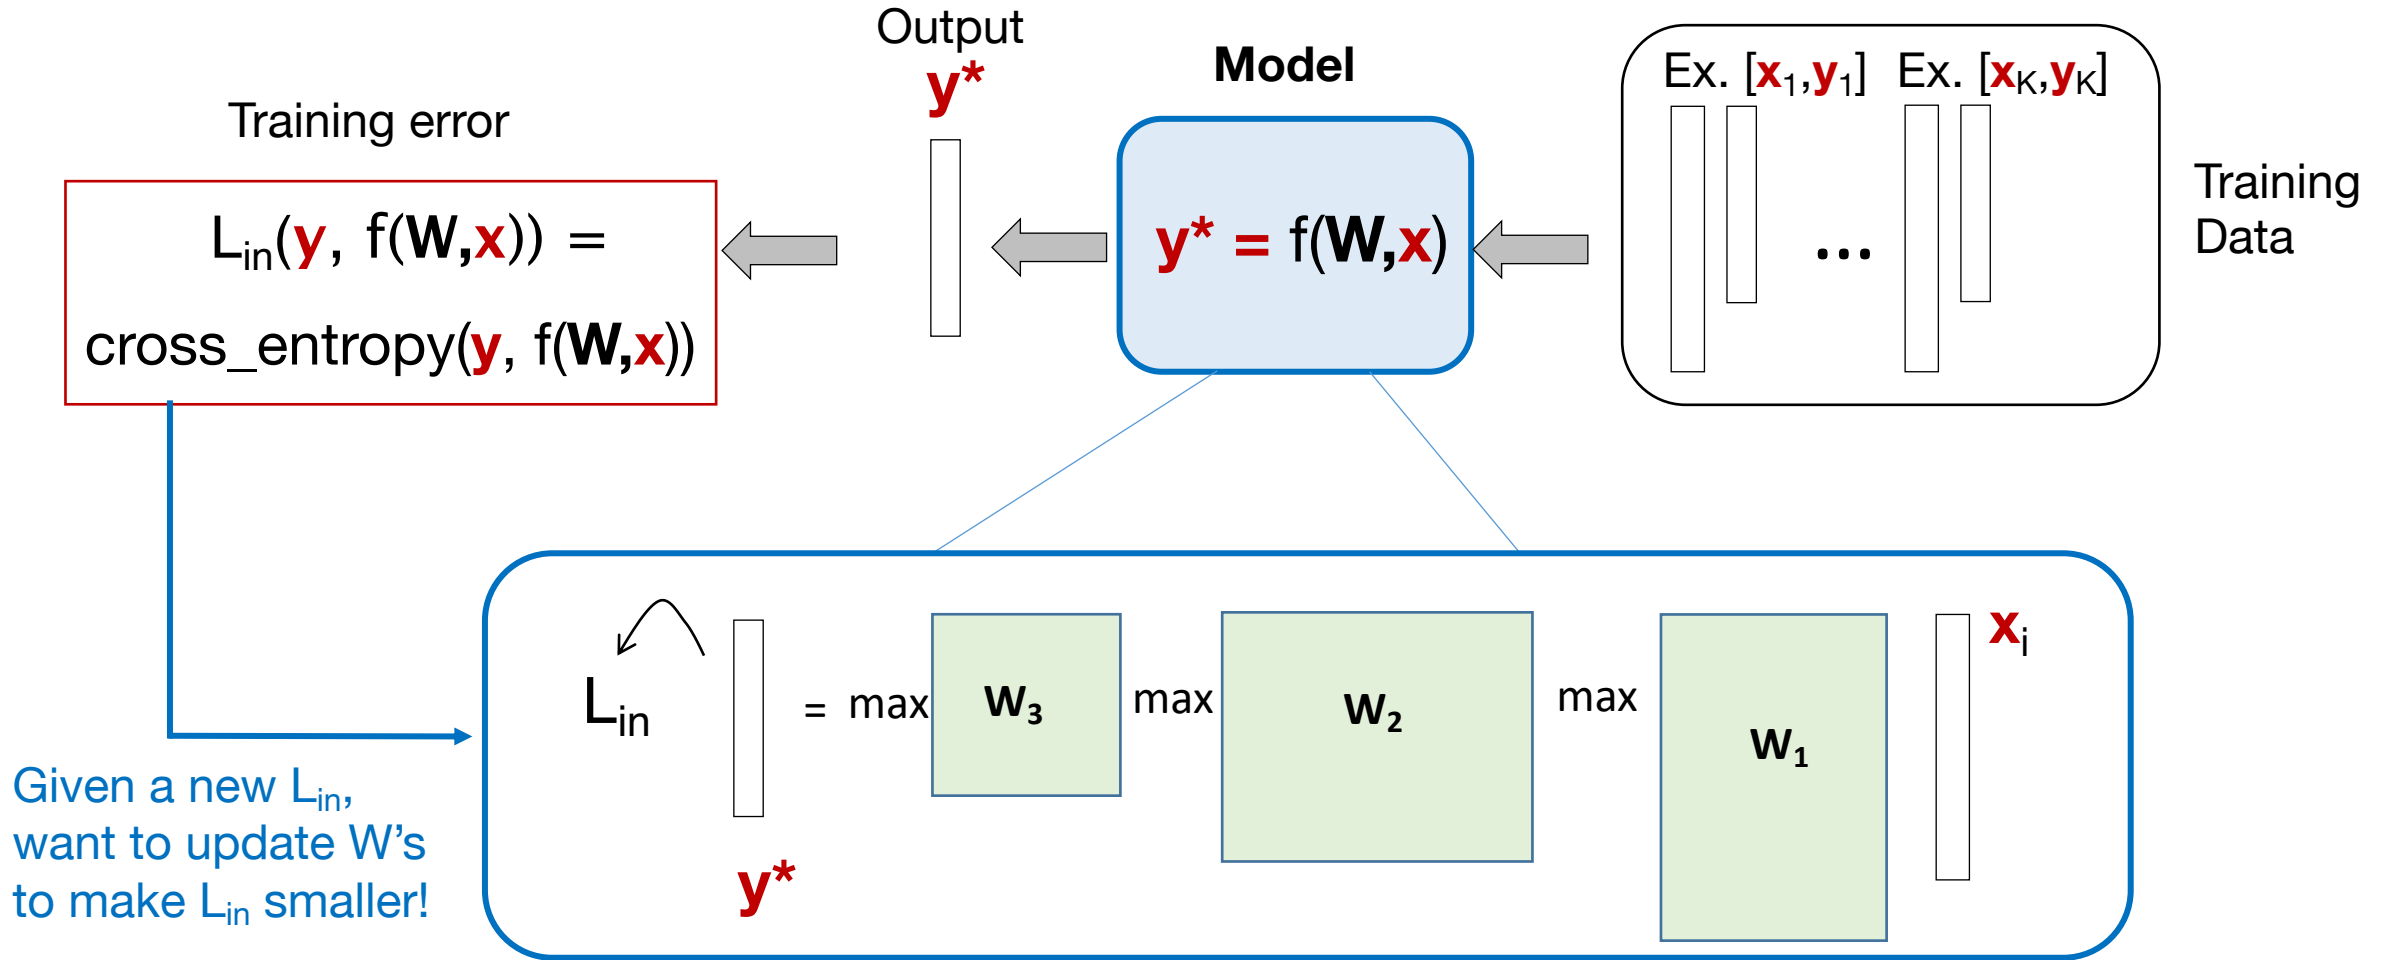
- CONV size, stride, pad, depth
- ReLU & other nonlinearities
- POOL methods
- # of layers, dimensions per layer
- Fully connected layers

## Loss function & optimization

- Type of loss function
- Regularization
- Gradient descent method
- SGD batch and step size


**Other specifics:** Pre-processing, initialization, dropout, batch normalization, augmentation

# Our very basic convolutional neural network



Backwards pass uses new  $L_{in}$  to update  $\mathbf{W}$ 's – **backpropagation!**

# Automatic differentiation on computational graphs



To both determine  $f$  and find  $df/dx_i$  :

- Create graph of local operations
- Compute analytic (symbolic) gradient at each node (unit) in graph
- Use inter-relationships to establish final desired gradient,  $df/dx_1$ 
  - Forward differentiation
  - Backwards differentiation = Backpropagation

# Let's go through an example:

$$L = \| \mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 \mathbf{X}) \|_2^2$$

(2-layer network with MSE where we neglect labels  $\mathbf{y}$  for now)


$$dL/d\mathbf{W}_1 = ? \quad dL/d\mathbf{W}_2 = ?$$

$$z_1 = XW_1$$

$$h_1 = \text{ReLU}(z_1)$$

$$\hat{y} = h_1W_2$$

$$L = \|\hat{y}\|_2^2$$



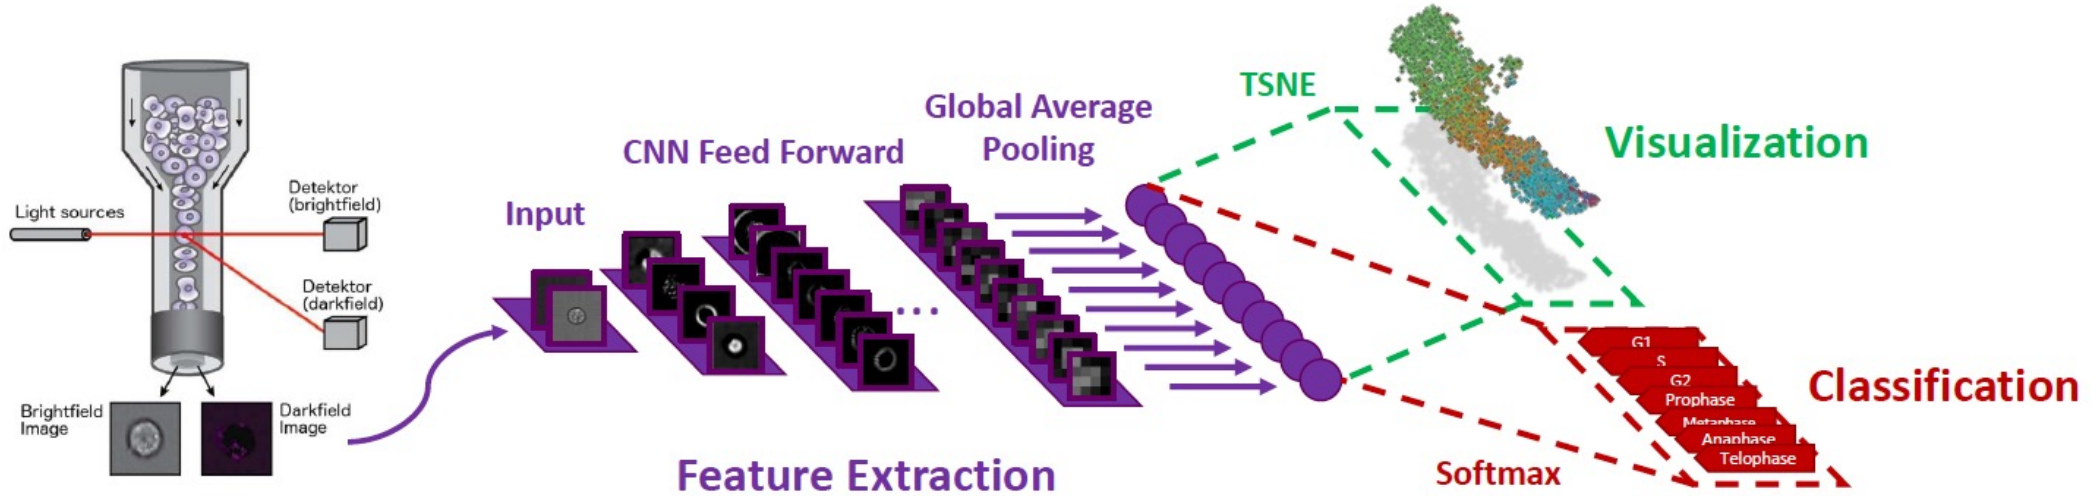


$$\frac{\partial L}{\partial W_2} = \frac{\partial \hat{y}}{\partial W_2} \frac{\partial L}{\partial \hat{y}} = 2h_1^T \hat{y}$$

$$\frac{\partial \hat{y}}{\partial W_2} = h_1^T \quad \frac{\partial L}{\partial \hat{y}} = 2\hat{y}$$

# CNNs for classification


P. Eulenberg et al., "Reconstructing cell cycle and disease progression using deep learning"





# Other Computer Vision Tasks


## Semantic Segmentation



GRASS, CAT,  
TREE, SKY

No objects, just pixels


## Classification + Localization



CAT

Single Object


## Object Detection



DOG, DOG, CAT

Multiple Object

## Instance Segmentation



DOG, DOG, CAT

## Super-resolution




Figure 1: Super-resolved image (left) is almost indistinguishable from original (right). [4x upscaling]

This image is CC0 public domain

## Post-processing of your results: a few options at different stages


Options to examine your test data after processing:

- ROC curve, Precision-Recall
- Confusion matrix
- Sliding window visualization
- Layer visualizations
- Saliency maps etc.
- tSNE visualization

# Instead, *compress x-y dimensions of input image*

- Compress spatial features into learned filters
- Then, decompress learned filters back into same spatial dimensions
- **Can be an autoencoder**
- Analogous to image compression
- A very powerful idea...

## U-Net Architecture



## U-Net: Convolutional Networks for Biomedical Image Segmentation

Olaf Ronneberger, Philipp Fischer, and Thomas Brox

Computer Science Department and BIOS Centre for Biological Signalling Studies,  
University of Freiburg, Germany  
ronneber@informatik.uni-freiburg.de,  
WWW home page: <http://lmb.informatik.uni-freiburg.de/>


# Bringing together physical and digital image representations

Physical Layers

Digital Layers

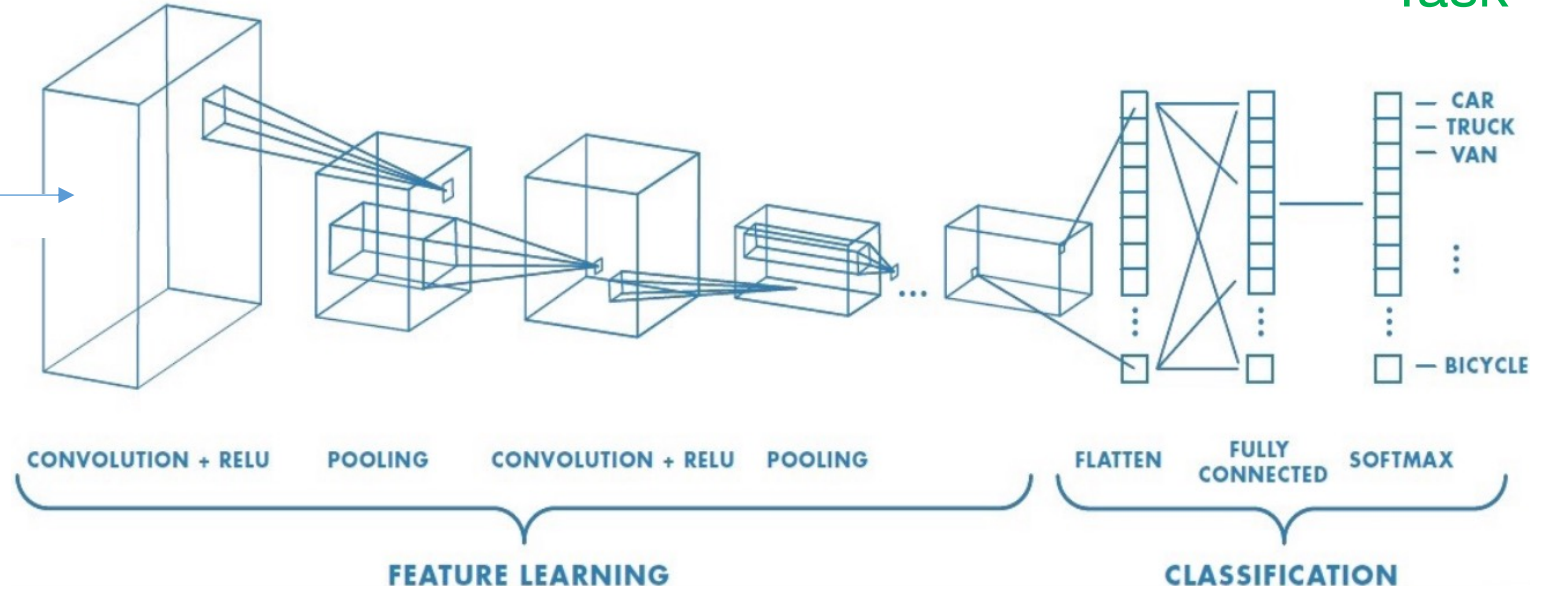
Task

Physical Function  $I_0$



Digitized  $I_s$

$$I_s = f[I_0]$$




Digital layers

Physical layers

$$\text{Task} = \mathbf{W}_n \dots \text{ReLU}[\mathbf{W}_1 \text{ReLU}[\mathbf{W}_0 f[I_0]] \dots]$$

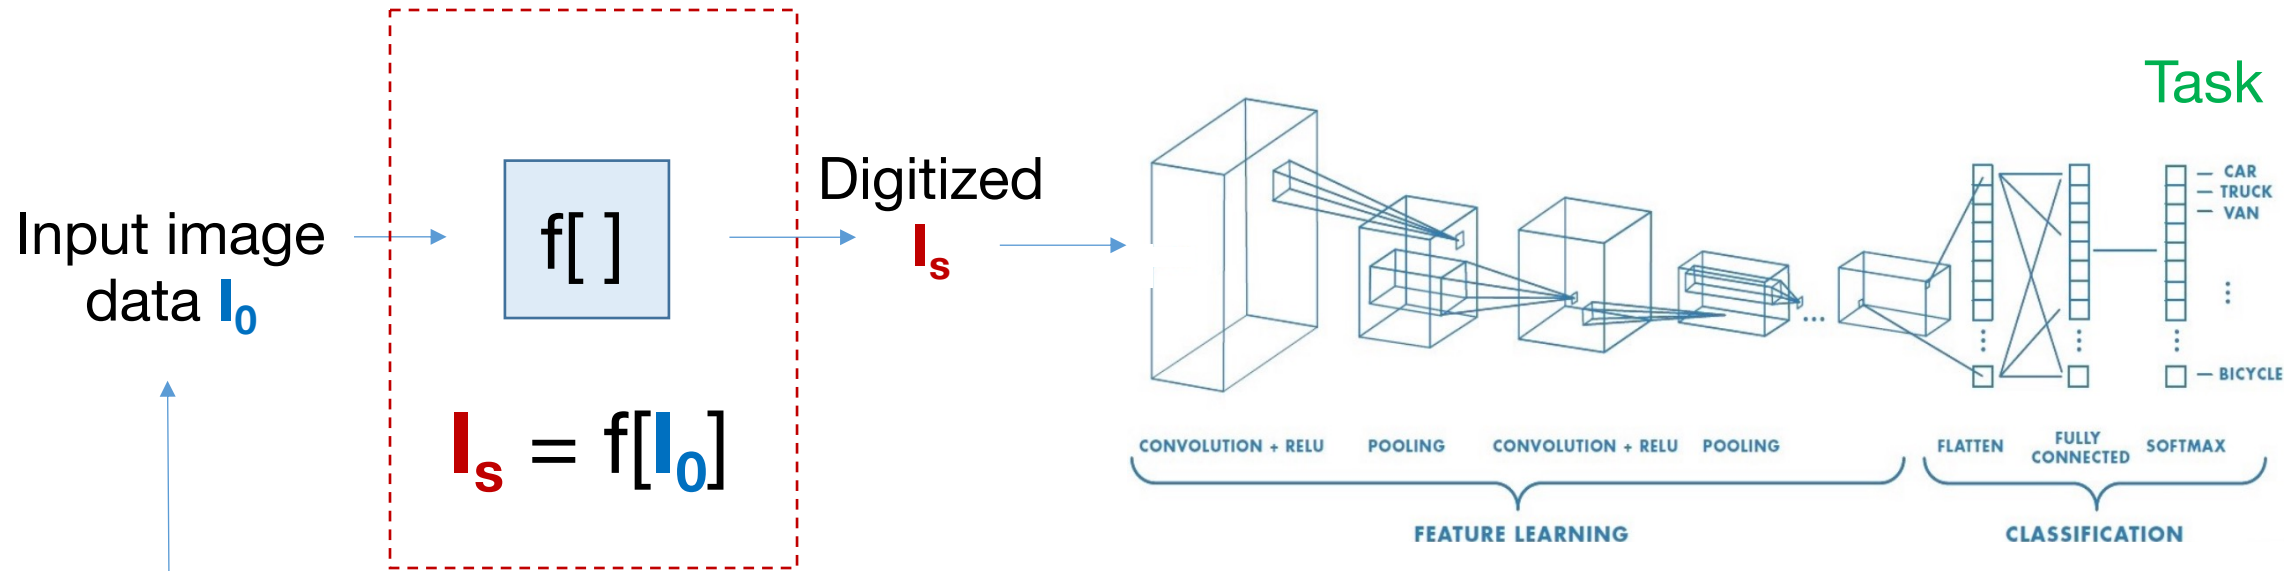
# What physical parameters effect image formation?

- **Illumination**
  - Spatial pattern
  - Angle of incidence
  - Color, polarization
- **Lens and optics**
  - Position/orientation
  - Shape
  - Focus
  - Transparency
- **Detector**
  - Pixel size
  - Pixel shape & fill factor
  - Color filters
  - Other filters
- **Digitization**
  - E to P curves
  - Digitization schemes/thresholds
  - Data transmission, multiplexing
- Physical object



## Physical Layers

## Digital Layers

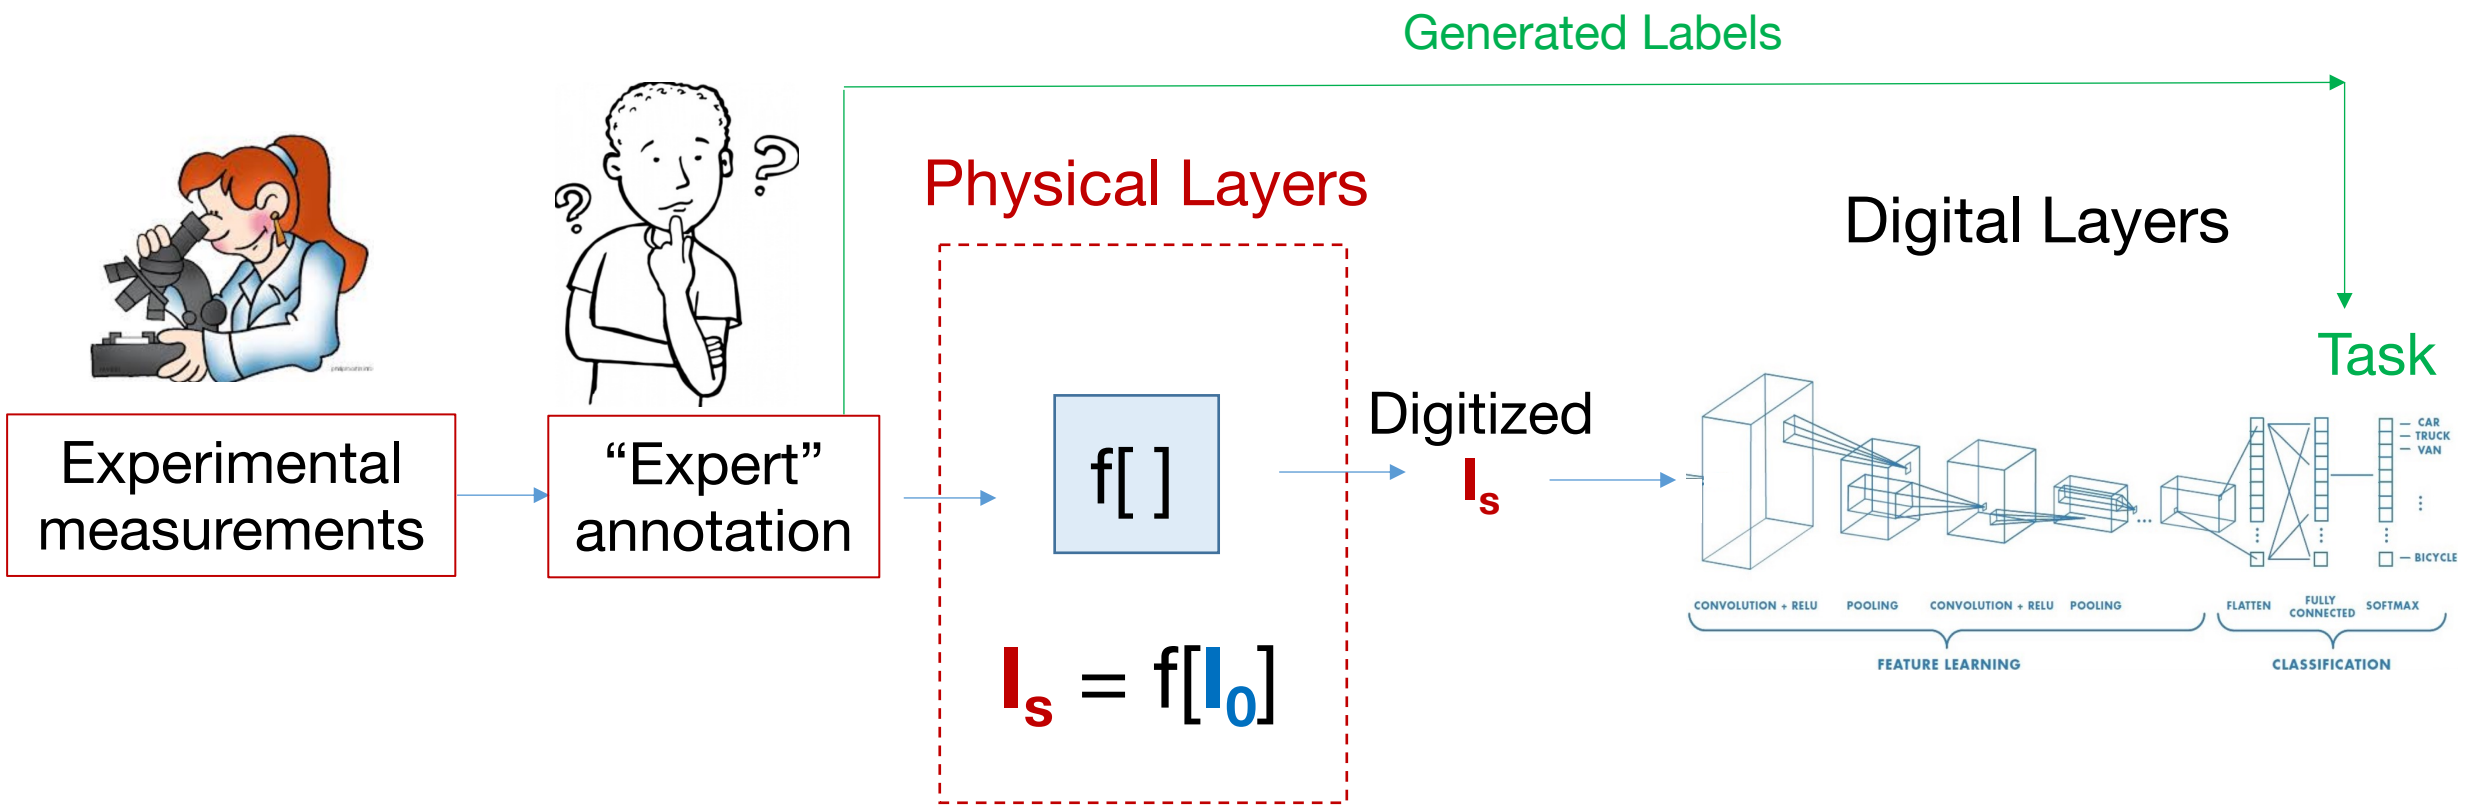


Q: Where and how should I implement my physical layer?

A: It depends on your data and implementation

- Situation #1: Fully simulated physical layers
- Situation #2: Experimentally-driven physical layers


# Situation #2: Fully simulated physical layers




Pro’s of experimental measurements: Don’t need to worry about making your simulations match the setup! (HUGE)

Con’s of experimental measurements: You’ll need to label them, limited access to desired sample information, often need to exploit some fundamental physical property

# Many-to-one recurrent neural network



Backpropagate to minimize  $dL/dW$ ,  $dL/dU$



State  $h_t$



Output  $o_t$

Simple network structure:

$$h_t = \text{ReLU}[\mathbf{W}h_{t-1} + \mathbf{U}x_t + b]$$

Label  $y$   $\longrightarrow$  Loss function  $L(o_t, y)$



# Taxonomy of Generative Models

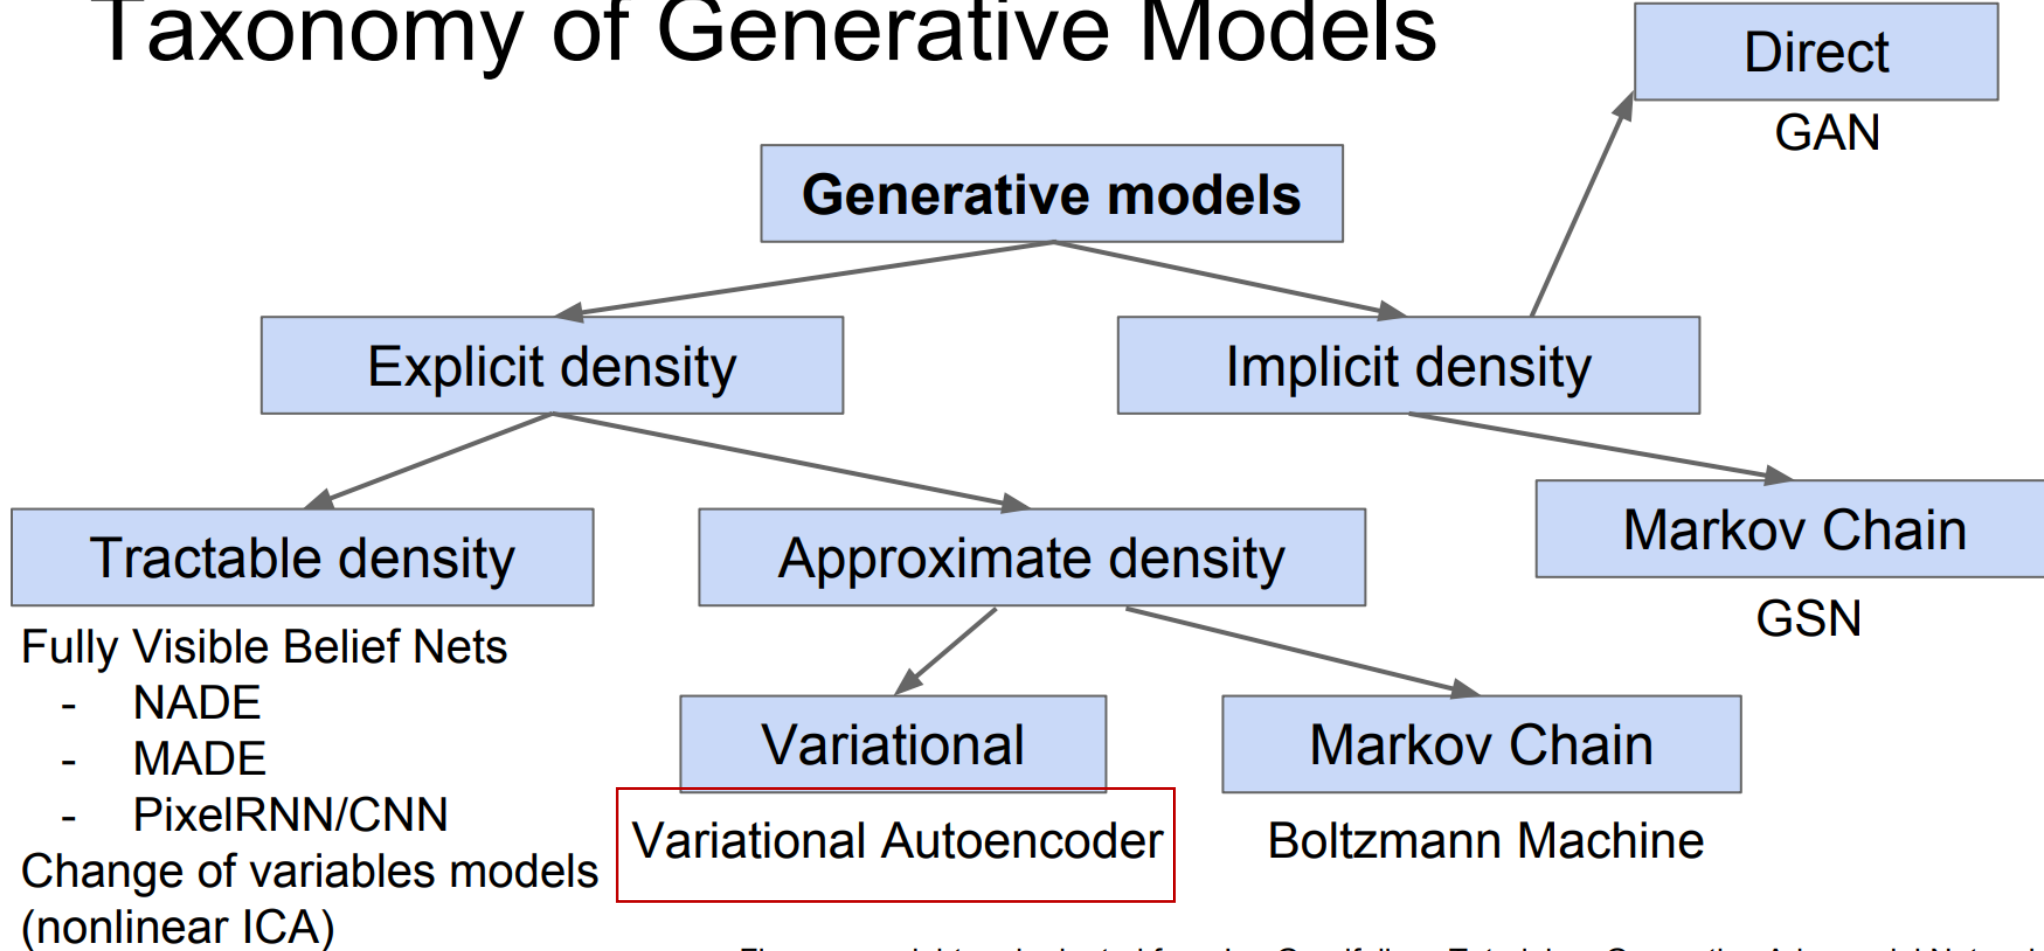



Figure copyright and adapted from Ian Goodfellow, Tutorial on Generative Adversarial Networks, 2017.

# Example: Variational Autoencoder (VAE)

Force this vector to follow a Gaussian PDF




Minimize (KL) distance between latent vector and Gaussian normal

Generative Example (once trained):


- Encode image with glasses, obtain latent vector PDF  $P_g$
- Encode image without glasses, obtain PDF  $P_{ng}$
- Compute  $\mathbf{diff} = P_g - P_{ng}$
- Encode new image to obtain  $P_{new}$ , add in  $\mathbf{diff}$
- Decode  $P_{new} + \mathbf{diff}$  to get guy with glasses!

- With Gaussian PDF, can start to add/subtract latent vector in a normalized vector space




Adding new features to samples

Glasses



Exploring a specific variation of input data[1]

# II.B - G/D Game



Run Adam simultaneously on two minibatches (true data / generated data)

Gradients

$$\begin{cases} y = 0 & \text{if } x = G(z) \\ y = 1 & \text{otherwise} \end{cases}$$



Thanks for a great semester everyone!