deep imaging

Lecture 25: Looking Ahead + Review

Machine Learning and Imaging

BME 548L

Roarke Horstmeyer

Announcements and schedule

- Thursday is the last lecture
- Homework #5 Due: Thursday April 22
- Thursday April 29, 9-noon: Final projects due and presentations
    - Sign up for slot at Slack link
    - We'll meet at the usual zoom link
- Project help:
    - I will continue my office hours this week/next week
        - Wednesday 10am – 11am, Thursday 10am – 11am
    - Part of next class I can set aside for final project questions?
    - Email me if you'd like to meet another time
    - Email TA's / reach out on Slack to have office hours/meet as well

Components of final project

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

40% of total grade

1. Presentation Slides – 10%
   - 8 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 – 5%
4. brief website template & permission to share results – 5%
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 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

[Paper PDF]

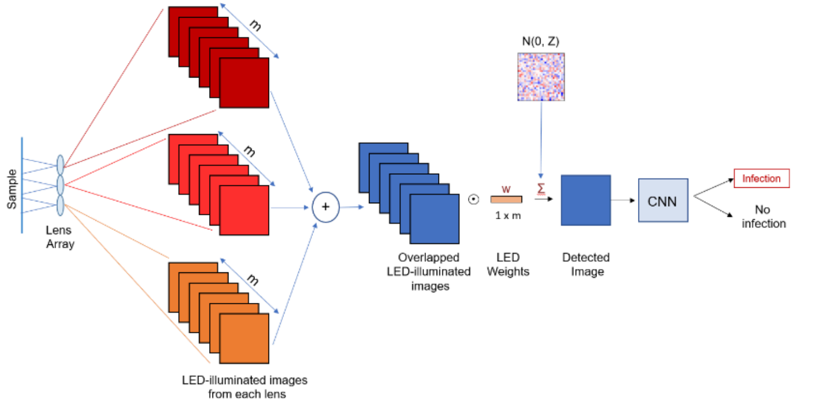

This project presents an imaging system that simultaneously captures multiple images and automatically classifies their contents to increase detection throughput. Our optical design consists of a set of multiple lenses that each image a unique field-of-view onto a single image sensor. The resulting "overlapped" 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!

Instructions:

https://assessment.trinity.duke.edu/students-course-evaluations

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

deep imaging

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

1. Proliferation of trained models, similar datasets and shared goals

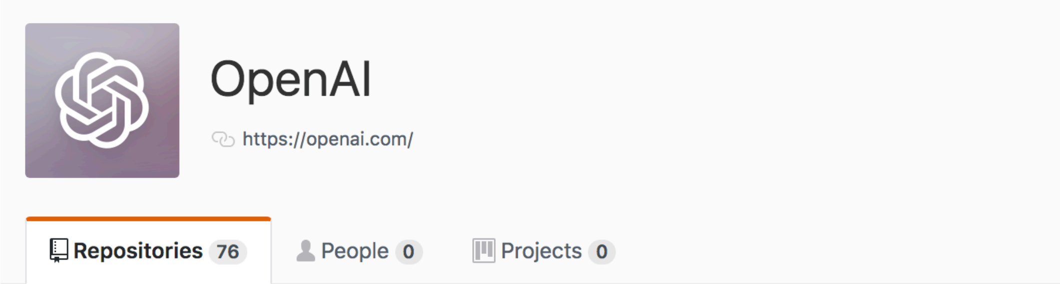

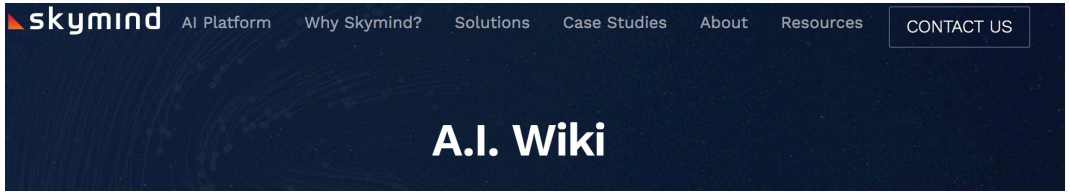

Visipedia

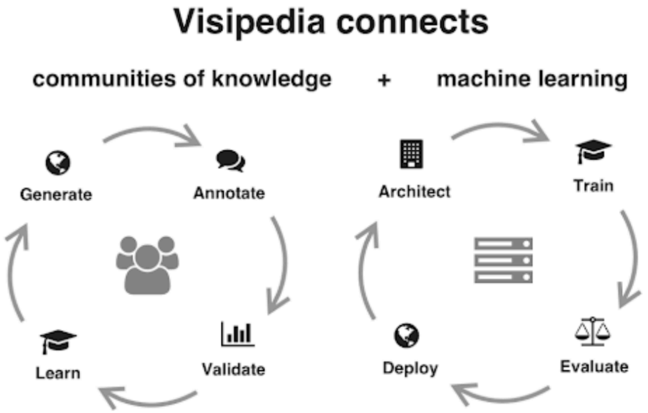

Caltech Visipedia

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

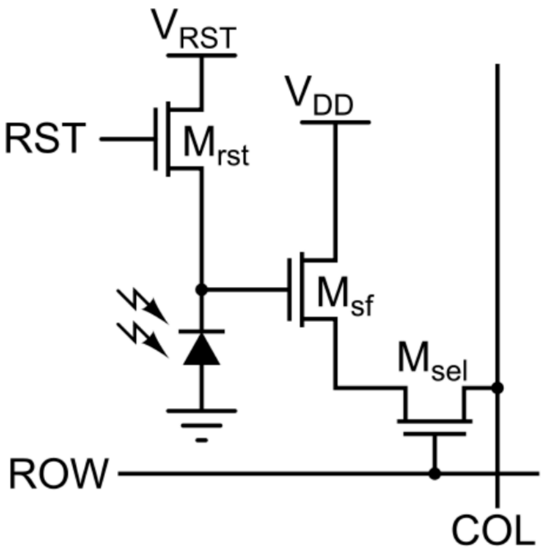

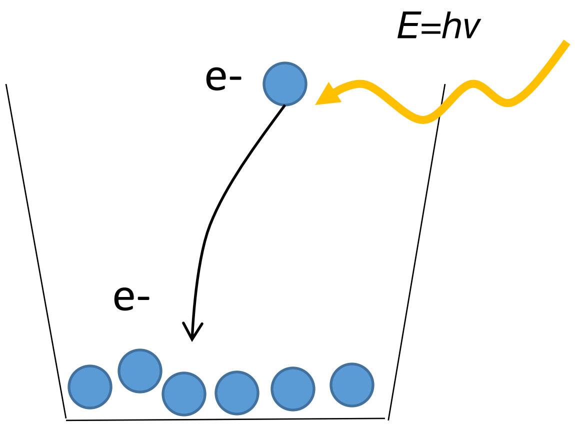

SPAD pixel: was there a photon or not?

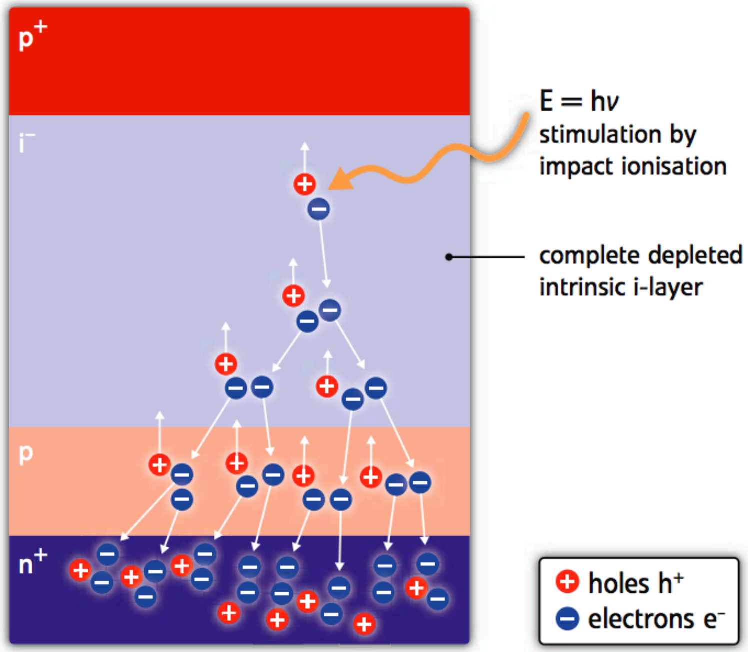

Where are things going with Machine Learning and Imaging in 10 years?

deep imaging

2. "Cameras" on many devices & new types of sensors

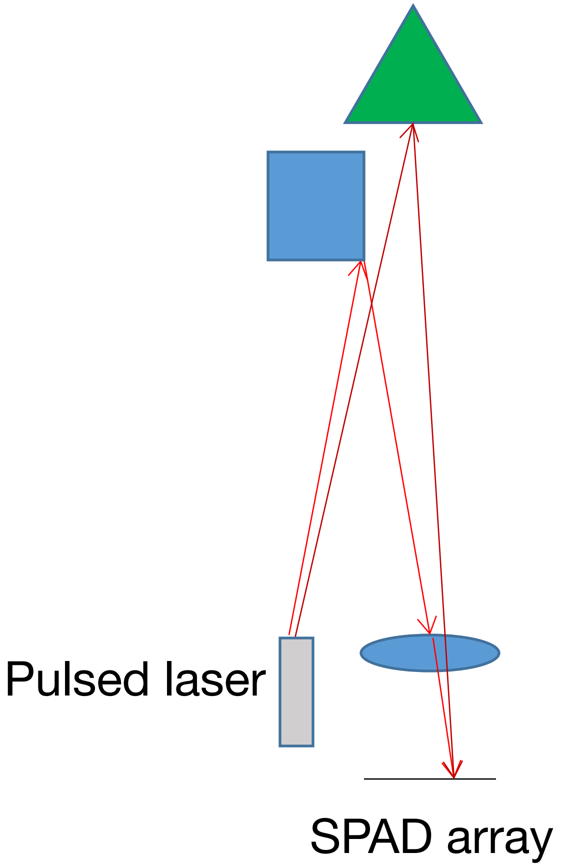

- Light travels 1 ft in 1 ns.
- SPADs can precisely photon arrival time to measure travel distance (TOF)

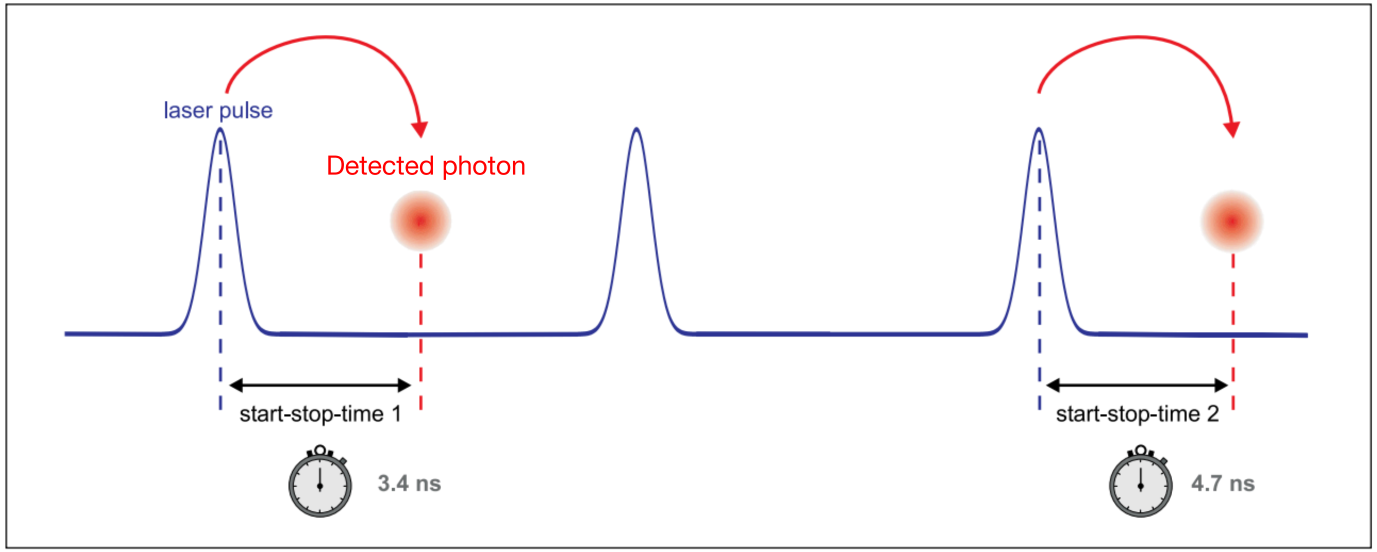

https://www.picoquant.com/images/uploads/page/files/7253/technote_tcspc.pdf

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

deep imaging

2. "Cameras" on many devices & new types of sensors

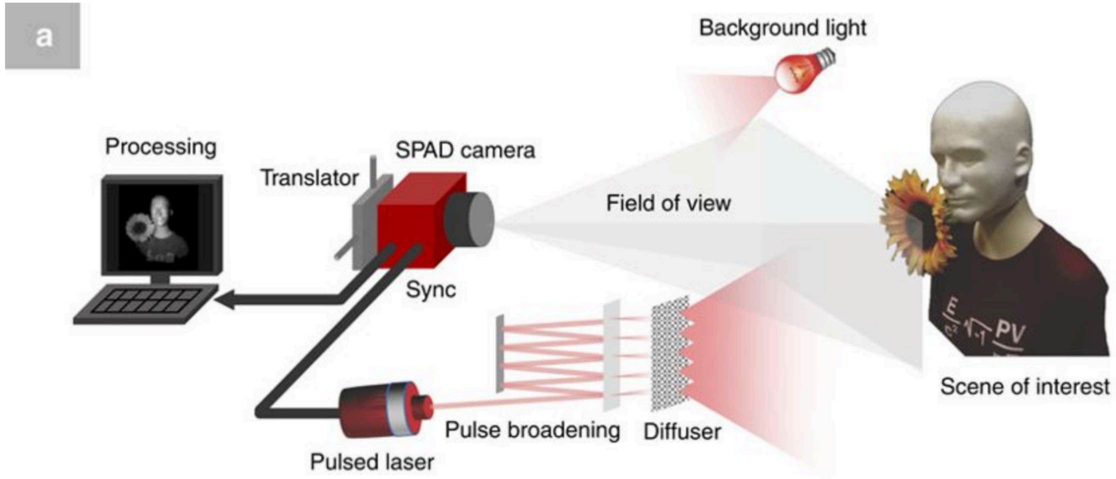

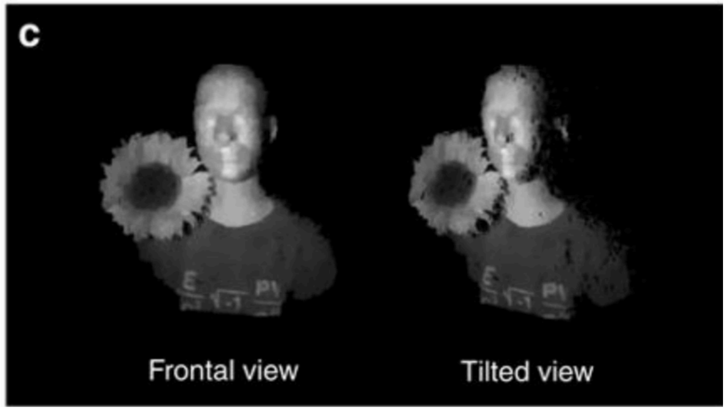

D. Shen et al, https://www.nature.com/articles/ncomms12046

Figure 1: Imaging Through Thick Scattering.

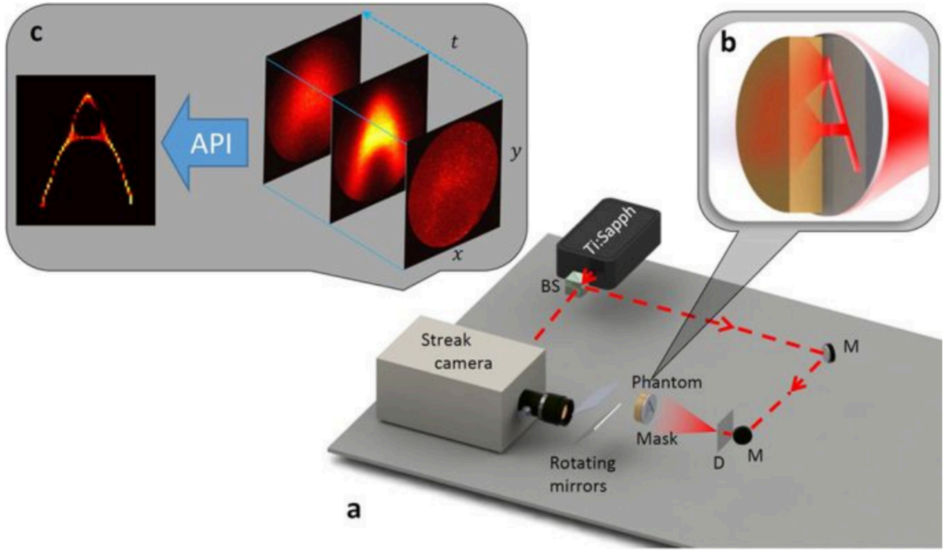

[G. Satat et al, https://www.nature.com/articles/srep33946](https://www.nature.com/articles/srep33946)

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?

deep imaging

3. Beyond convolutions - new constructs for deep networks

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

deep imaging

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

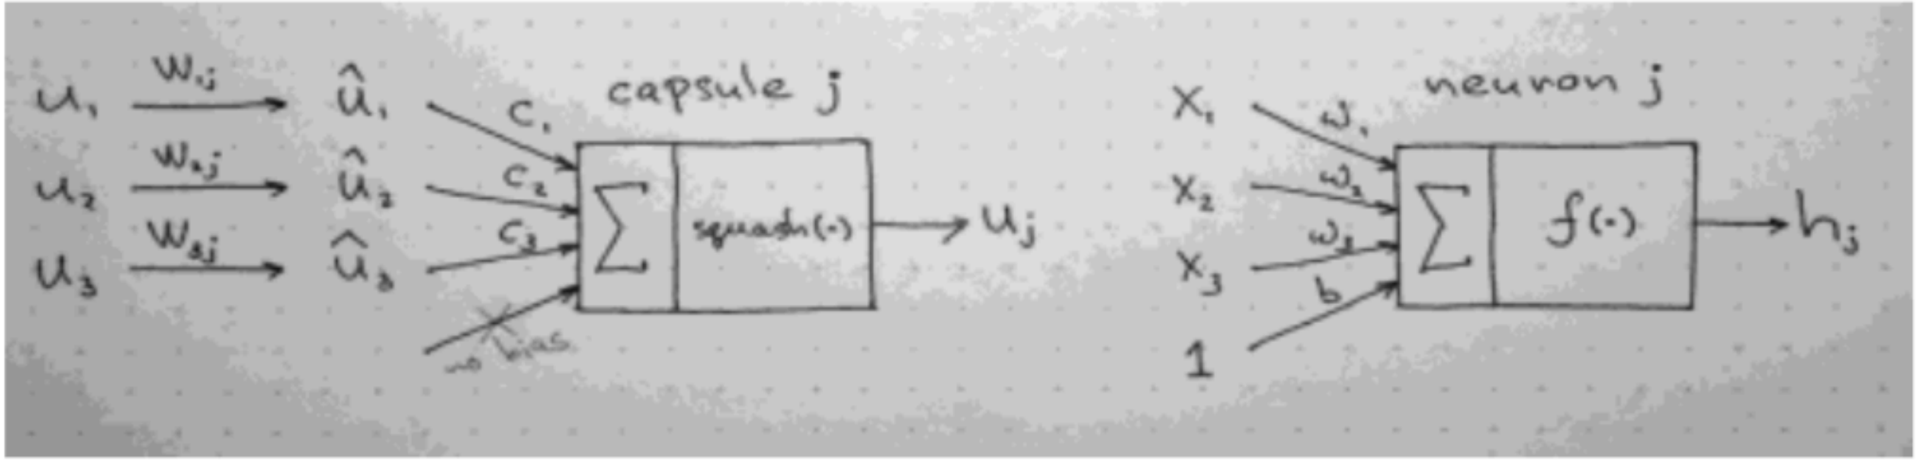

u_1 W_{1j} $\hat{u}_1$ c_1
u_2 W_{2j} $\hat{u}_2$ c_2
u_3 W_{3j} $\hat{u}_3$ c_3
capsule j
squash(·)
u_j
no bias
x_1 ω_1
x_2 ω_2
x_3 ω_3
1 b
neuron j
$f(\cdot)$
h_j

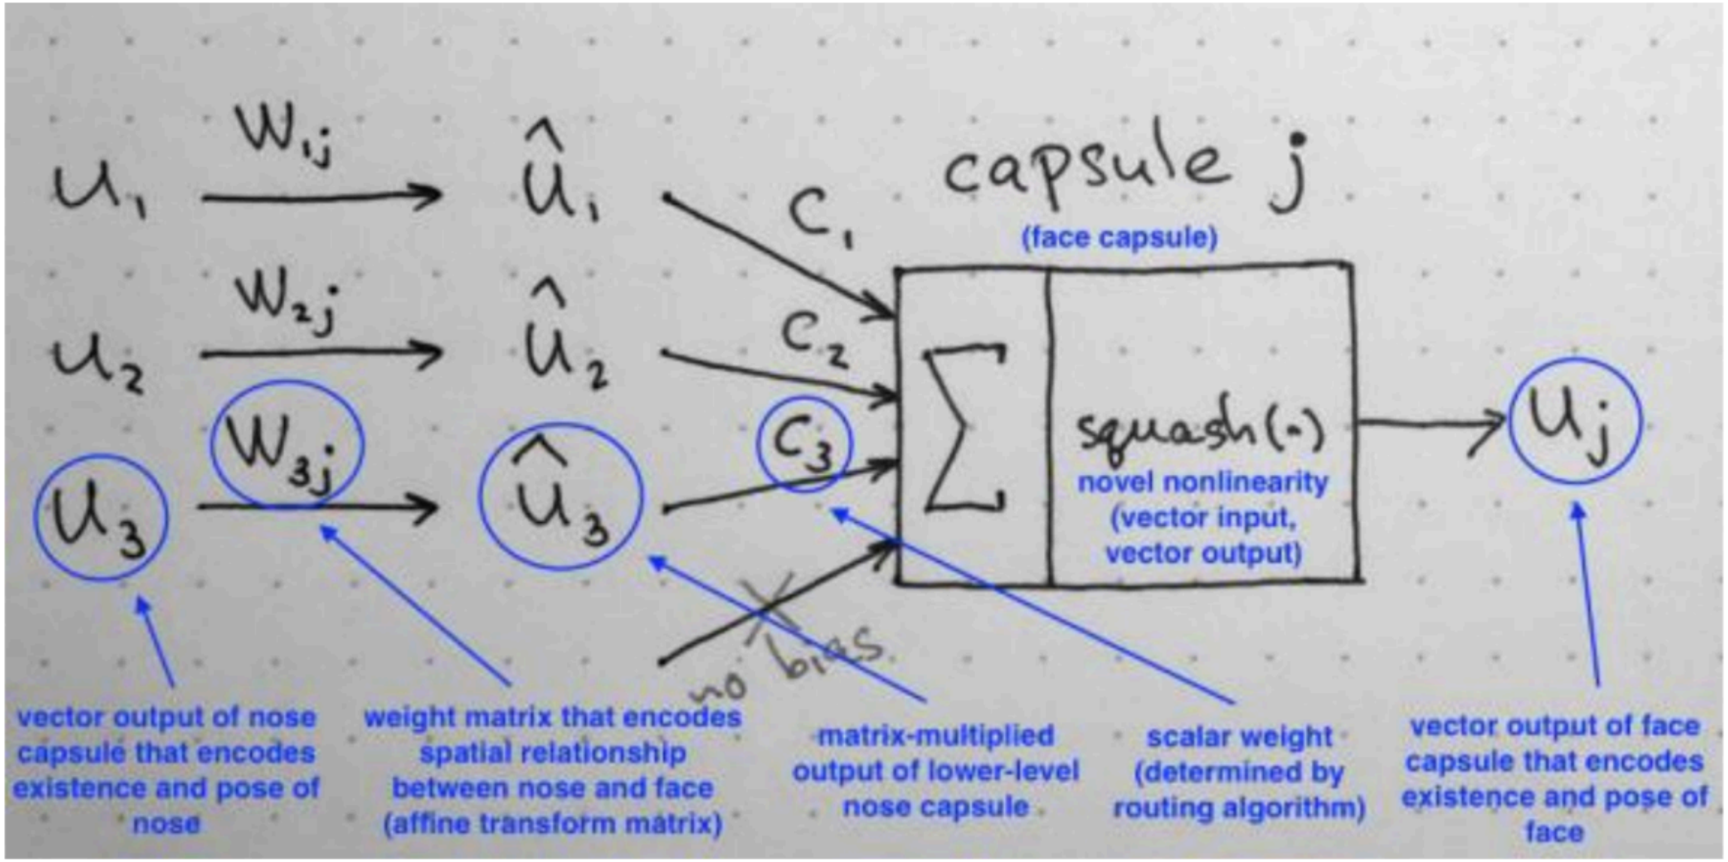

u_1 W_{1j} $\hat{u}_1$ c_1
u_2 W_{2j} $\hat{u}_2$ c_2
u_3 W_{3j} $\hat{u}_3$ c_3
capsule j
(face capsule)
squash(·)
novel nonlinearity (vector input, vector output)
u_j
no bias
vector output of nose capsule that encodes existence and pose of nose
weight matrix that encodes spatial relationship between nose and face (affine transform matrix)
matrix-multiplied output of lower-level nose capsule
scalar weight (determined by routing algorithm)
vector output of face capsule that encodes existence and pose of face

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

deep imaging

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)

3. Therefore it will not assign one pixel to two digits if one of them does not have any other support.

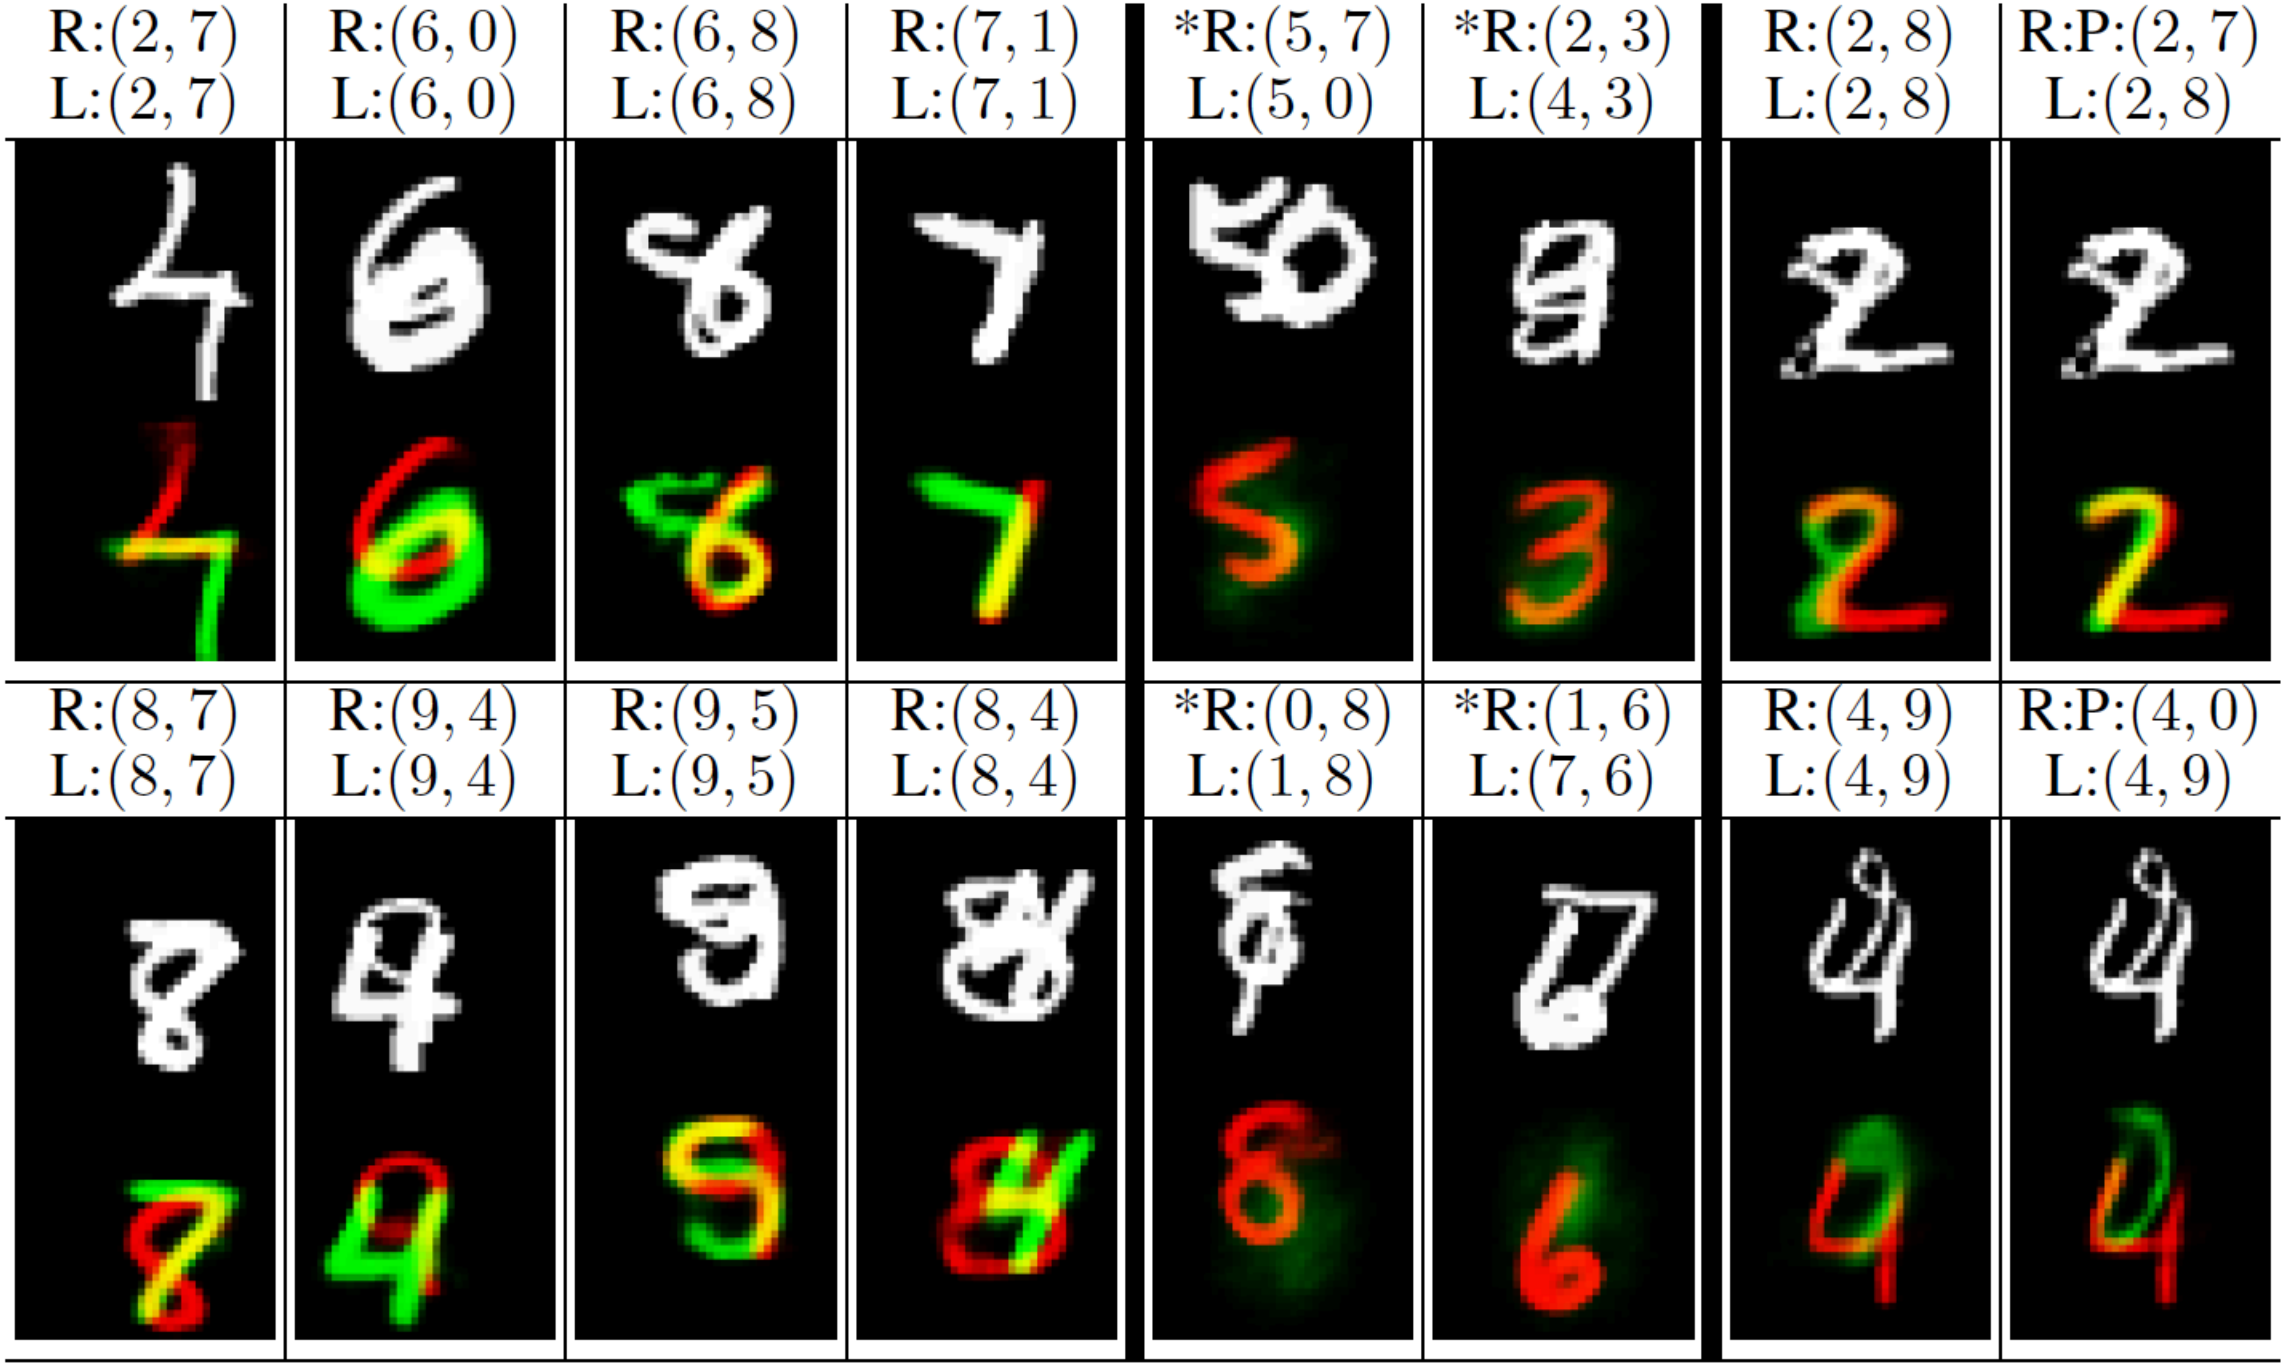

Where are things going with Machine Learning and Imaging in 10 years?

4. Generative data is getting pretty realistic...

Example of Realistic Synthetic Photographs Generated with BigGANTaken from Large Scale GAN Training for High Fidelity Natural Image Synthesis, 2018.

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

Justus Thies [1] Michael Zollhöfer [2] Marc Stamminger [3] Christian Theobalt [2] Matthias Nießner [1]

[1] Technical University of Munich [2] Max Planck Institute for Informatics [3] University of Erlangen-Nuremberg

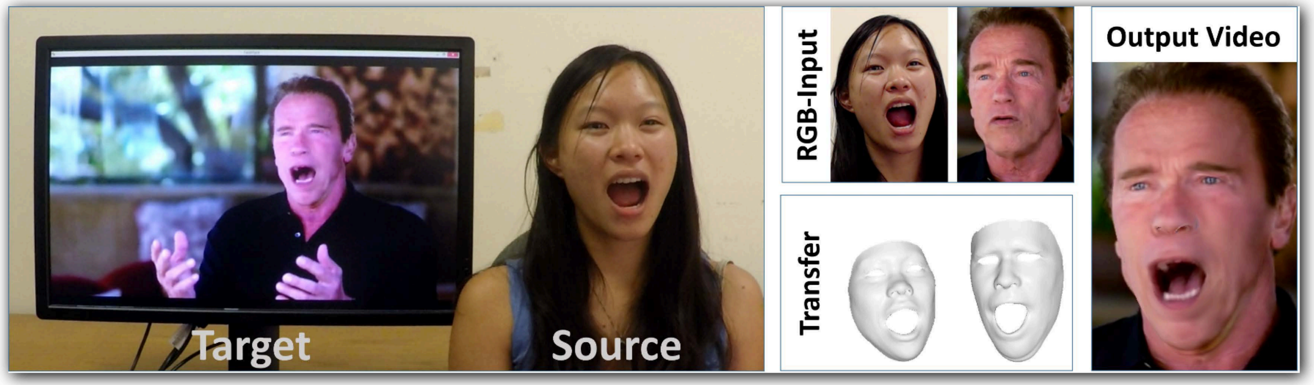

Proc. Computer Vision and Pattern Recognition (CVPR), IEEE, June 2016

What are the implications of this for medical imaging?

deep imaging

II.A - Motivation

Probability distributions:

Samples from the "real data distribution"

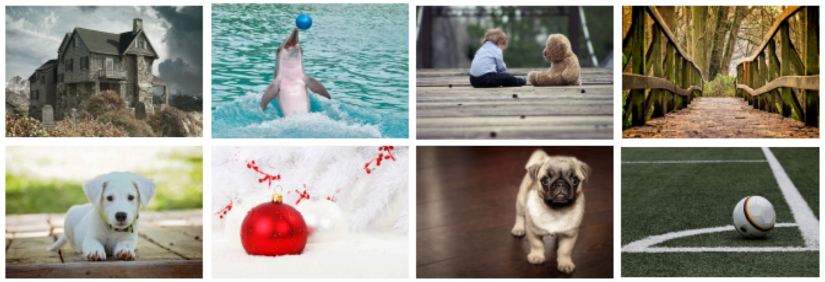

Samples from the "generated distribution"

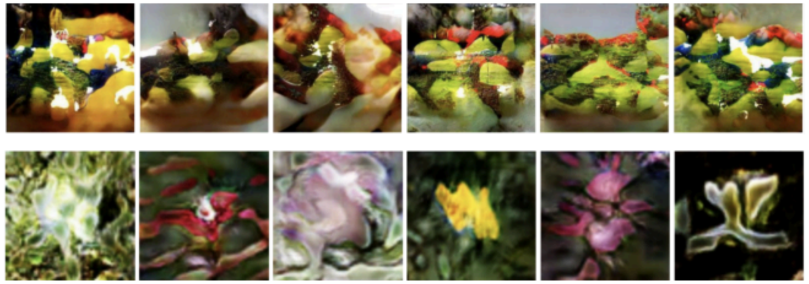

"real data distribution"

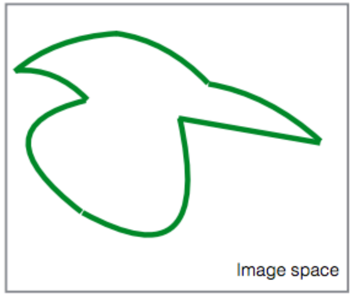

"generated distribution"

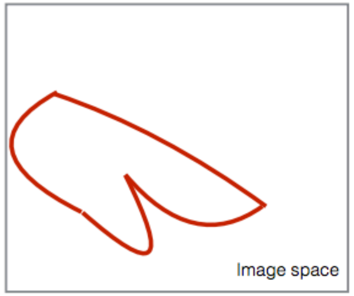

Goal

Matching distributions

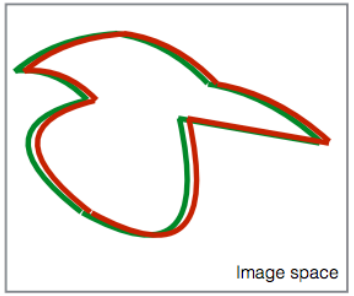

[Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, Dimitris Metaxas (2017): StackGAN: Text to Photo-realistic Image Synthesis with Stacked Generative Adversarial Networks]

Kian Katanforoosh, Andrew Ng, Younes Bensouda Mourri

II.B - G/D Game

100-d
random code

$$z = \begin{pmatrix} 0.47 \\ \vdots \\ 0.19 \end{pmatrix}$$

$\Rightarrow$

Generator "G"
(Neural Network)

$\Rightarrow$

(64,64,3)
generated image

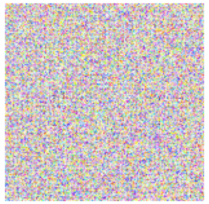

$\neq$



How can we train G to generate images from the true data distributions?

[Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, Dimitris Metaxas (2017): StackGAN: Text to Photo-realistic Image Synthesis with Stacked Generative Adversarial Networks]

Kian Katanforoosh, Andrew Ng, Younes Bensouda Mourri

Stanford CS230, Lecture 3

II.B - G/D Game

100-d
random code

$$z = \begin{pmatrix} 0.47 \\ \vdots \\ 0.19 \end{pmatrix}$$

**Generator "G"
(Neural Network)**

(64,64,3)
generated image

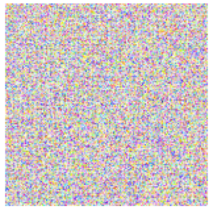

Run Adam simultaneously on two minibatches
(true data / generated data)

x

**Discriminator "D"
(Neural Network)**

Gradients

Binary classification

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

**Real images
(database)**

Probability distributions

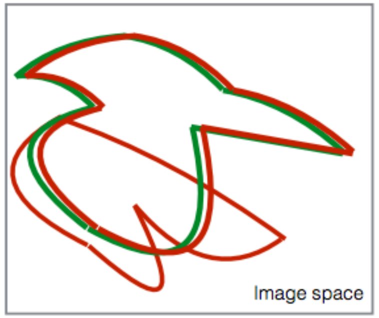

Kian Katanforoosh, Andrew Ng, Younes Bensouda Mourri

II.B - G/D Game

100-d
random code

$$z = \begin{pmatrix} 0.47 \\ \vdots \\ 0.19 \end{pmatrix}$$

Generator "G"
(Neural Network)

(64,64,3)
generated image

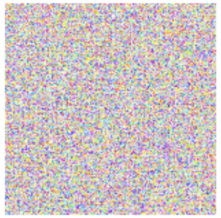

Real images
(database)

x

Discriminator "D"
(Neural Network)

Gradients

Binary classification

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

End goal: G is outputting images that are indistinguishable from real images for D

Probability distribution

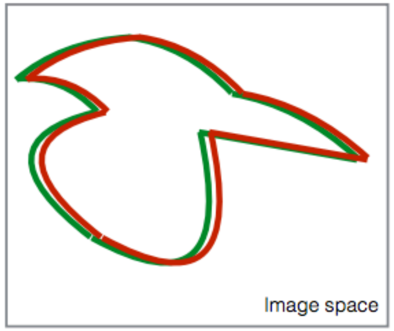

Kian Katanforoosh, Andrew Ng, Younes Bensouda Mourri

Stanford CS230, Lecture 3

II.B - G/D Game

naging

Training procedure, we want to minimize:

Labels: $\begin{cases} y_{real} \text{ is always } 1 \\ y_{gen} \text{ is always } 0 \end{cases}$

- The cost of the discriminator

$$J^{(D)} = \underbrace{-\frac{1}{m_{real}}\sum_{i=1}^{m_{real}} y_{real}^{(i)}.\log(D(x^{(i)}))}_{\substack{\text{cross-entropy 1:} \\ \text{“D should correctly label real data as 1”}}} \underbrace{-\frac{1}{m_{gen}}\sum_{i=1}^{m_{gen}} (1-y_{gen}^{(i)}).\log(1-D(G(z^{(i)})))}_{\substack{\text{cross-entropy 2:} \\ \text{“D should correctly label generated data as 0”}}}$$

- The cost of the generator

$$J^{(G)} = -J^{(D)} = \frac{1}{m_{gen}}\sum_{i=1}^{m_{gen}} \log(1-D(G(z^{(i)})))$$

“G should try to fool D: by minimizing the opposite of what D is trying to minimize”

“Maximize probability that the discriminator is wrong and labels the fake example as a real example”

II. D. In terms of code

imaging

```python
# Build and compile the discriminator
self.discriminator = self.build_discriminator()
self.discriminator.compile(loss='binary_crossentropy',
    optimizer=optimizer,
    metrics=['accuracy'])

# Build the generator
self.generator = self.build_generator()

# The generator takes noise as input and generates imgs
z = Input(shape=(self.latent_dim,))
img = self.generator(z)

# For the combined model we will only train the generator
self.discriminator.trainable = False

# The discriminator takes generated images as input and determines validity
validity = self.discriminator(img)

# The combined model  (stacked generator and discriminator)
# Trains the generator to fool the discriminator
self.combined = Model(z, validity)
self.combined.compile(loss='binary_crossentropy', optimizer=optimizer)
```

```python
def build_discriminator(self):

    model = Sequential()

    model.add(Flatten(input_shape=self.img_shape))
    model.add(Dense(512))
    model.add(LeakyReLU(alpha=0.2))
    model.add(Dense(256))
    model.add(LeakyReLU(alpha=0.2))
    model.add(Dense(1, activation='sigmoid'))
    model.summary()

    img = Input(shape=self.img_shape)
    validity = model(img)

    return Model(img, validity)
```

Erik Linder-Norén (Github): eriklindernoren/Keras-GAN: link]

II.E - Nice results

Operation on codes

Code 1

$$\begin{pmatrix} 0.12 \\ \vdots \\ 0.92 \end{pmatrix} \Rightarrow$$ **Generator "G" (Neural Network)** $\Rightarrow$

(64,64,3)
generated image

Code 2

$$\begin{pmatrix} 0.47 \\ \vdots \\ 0.19 \end{pmatrix} \Rightarrow$$ **Generator "G" (Neural Network)** $\Rightarrow$

(64,64,3)
generated image

Code 3

$$\begin{pmatrix} 0.42 \\ \vdots \\ 0.07 \end{pmatrix} \Rightarrow$$ **Generator "G" (Neural Network)** $\Rightarrow$

(64,64,3)
generated image

Code 1 Code 2 Code 3

$$\begin{pmatrix} 0.12 \\ \vdots \\ 0.92 \end{pmatrix} - \begin{pmatrix} 0.47 \\ \vdots \\ 0.19 \end{pmatrix} + \begin{pmatrix} 0.42 \\ \vdots \\ 0.07 \end{pmatrix} \Rightarrow$$ **Generator "G" (Neural Network)** $\Rightarrow$

Man with glasses - man + woman = woman with glasses

[Radford et al. (2015): UNSUPERVISED REPRESENTATION LEARNING WITH DEEP CONVOLUTIONAL GENERATIVE ADVERSARIAL NETWORKS] https://colab.research.google.com/github/tensorflow/gan/blob/master/tensorflow_gan/examples/colab_notebooks/tfgan_tutorial.ipynb?utm_source=ss-gan&utm_campaign=colab-external&utm_medium=referral&utm_content=tfgan-intro

5. Joint optimization of hardware and software is proliferating

technology feature

Check for updates

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

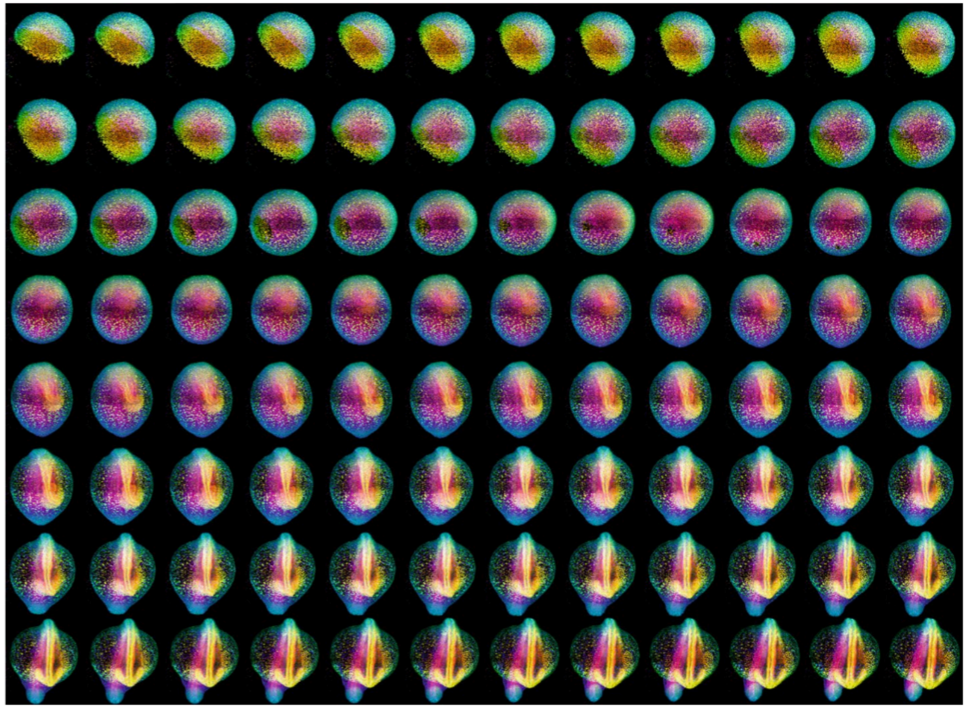

deep imaging

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…

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?

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?

The Machine Learning in Imaging Ethics Questionnaire

Situation 1: In 5 years, you walk into a clinic because you have a spot on your skin that you are concerned about. The clinician is too busy, so you step over to a terminal with a standard microscope and it images your arm. It says you are fine.

Are you comfortable with leaving the office?

Yes:

No:

Why or why not? What might change how you feel?

The Machine Learning in Imaging Ethics Questionnaire

Situation 2: The same thing happens. But this time, the machine reports that it is 99% confident in its diagnosis, given previous examples of skin marks that have been verified by doctors as benign. It also gives you the opportunity to take a look at some of these previous example images it is basing its decision on. You notice that they don't look 100% like the mark on your arm, as is expected, but they look pretty similar.

Are you comfortable with leaving the office?

Yes:

No:

Why or why not? What might change how you feel?

The Machine Learning in Imaging Ethics Questionnaire

Situation 3: In 5 years, the same thing happens. But this time, a doctor comes up after the machine makes its suggested diagnosis. He takes a very cursory look (10 seconds) and then confirms the machine's opinion.

Are you now comfortable with leaving the office?

Yes:

No:

deep imaging

The Machine Learning in Imaging Ethics Questionnaire

Situation 4: In 10 years, you go up to a modified microscope, “the Tissue Scanner 3000”, that has a number of fancy lenses and lights. As a machine learning expert by now, you’re aware that this microscope is optimized for looking at skin lesions. It performs a scan with a particular lighting configuration and reports a score of 98% confident that the lesion is benign, allowing you to look through other examples. It asks If you’d like another scan for additional confidence or a different outcome, at which point the illumination changes and it does some more scanning and reports a 99% confidence level. You can continue with another scan, but…

Are you now comfortable with leaving the office?

Yes:

No: