deep imaging

Lecture 25: Looking Ahead + Review

Machine Learning and Imaging

BME 548L
Roarke Horstmeyer

deep imaging

Announcements and schedule

- Today is the last lecture
- Homework #5 Due: Tuesday Nov 17
- November 24, 1-3 or 3-5 pm: Final projects due and presentations
    - Sign up for slot at Slack link
- Project help:
    - I will continue my office hours next week *but* they'll be 1 hour earlier
        - Wednesday 9am – 10am, Thursday 9am – 10am
    - Email me if you'd like to meet another time
    - Email TA's / reach out on Slack to have office hours/meet as well

deep imaging

Components of final project

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

40% of total grade

1. Presentation Slides – 10%
   - 7 minute presentation, 1 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

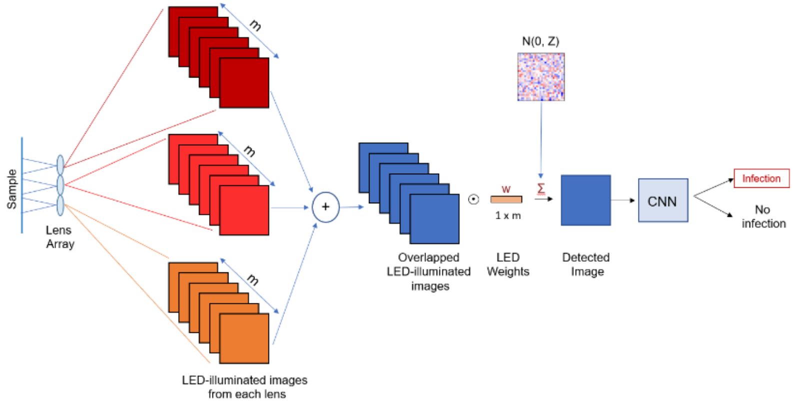

This project presents an imaging system that simultaneously captures multiple images and automatically classifies their contents to increase detection throughput. Our optical design consists of a set of multiple lenses that each image a unique field-of-view onto a single image sensor. The resulting "overlapped"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

deep imaging

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 shared goals

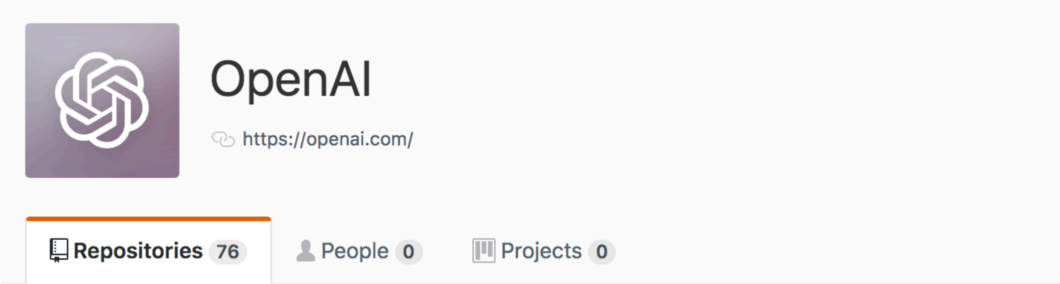

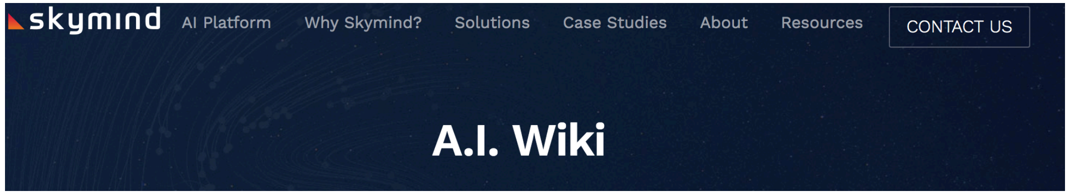

Visipedia

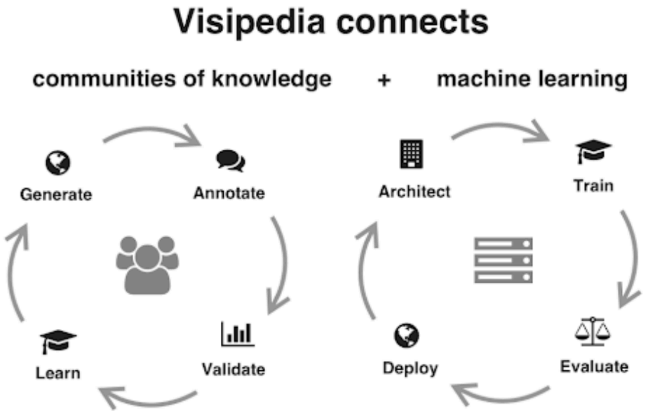

Caltech Visipedia

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

deep imaging

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

Standard CMOS pixel = bucket that collects electrons

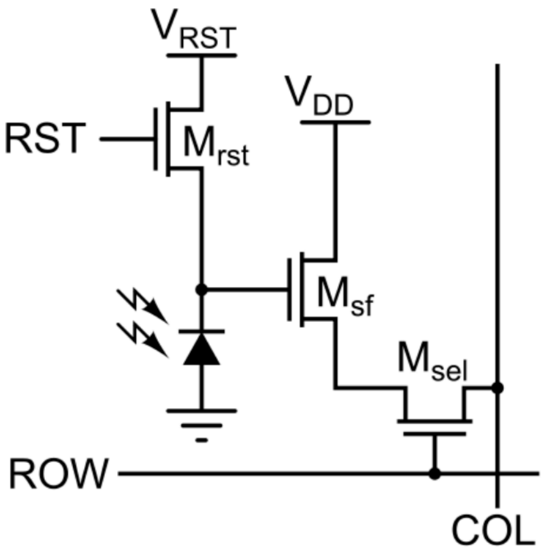

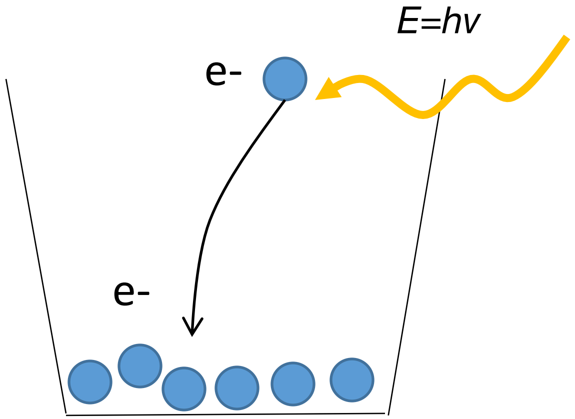

SPAD pixel: was there a photon or not?

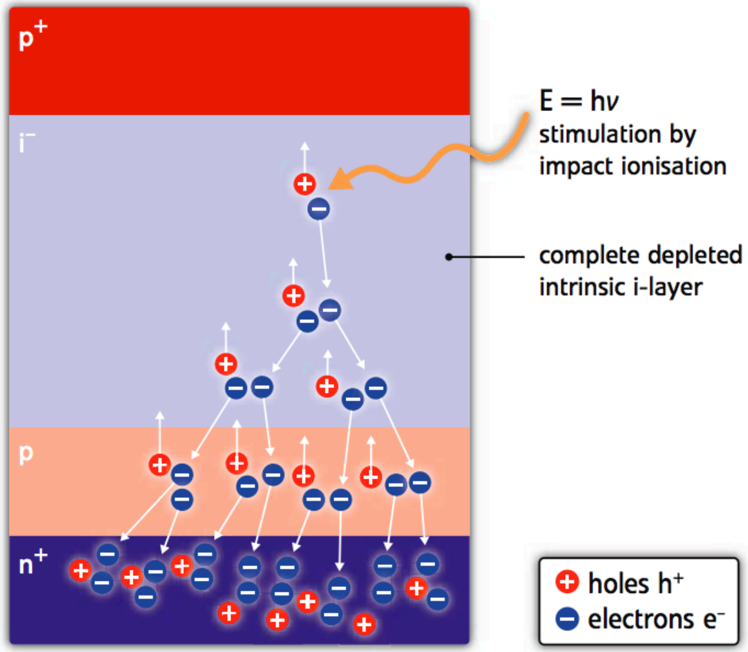

Where are things going with Machine Learning and Imaging in 10 years?

deep imaging

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

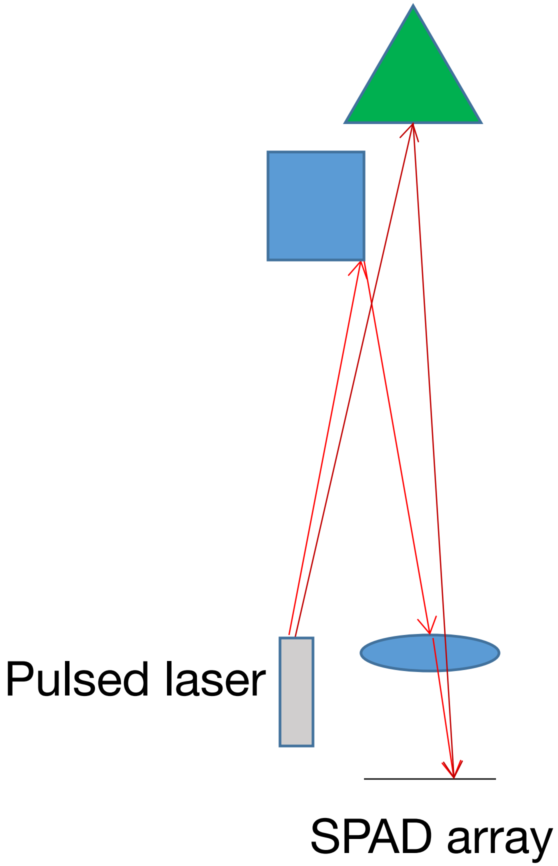

- Light travels 1 ft in 1 ns.
- SPADs can precisely photon arrival time to measure travel distance (TOF)

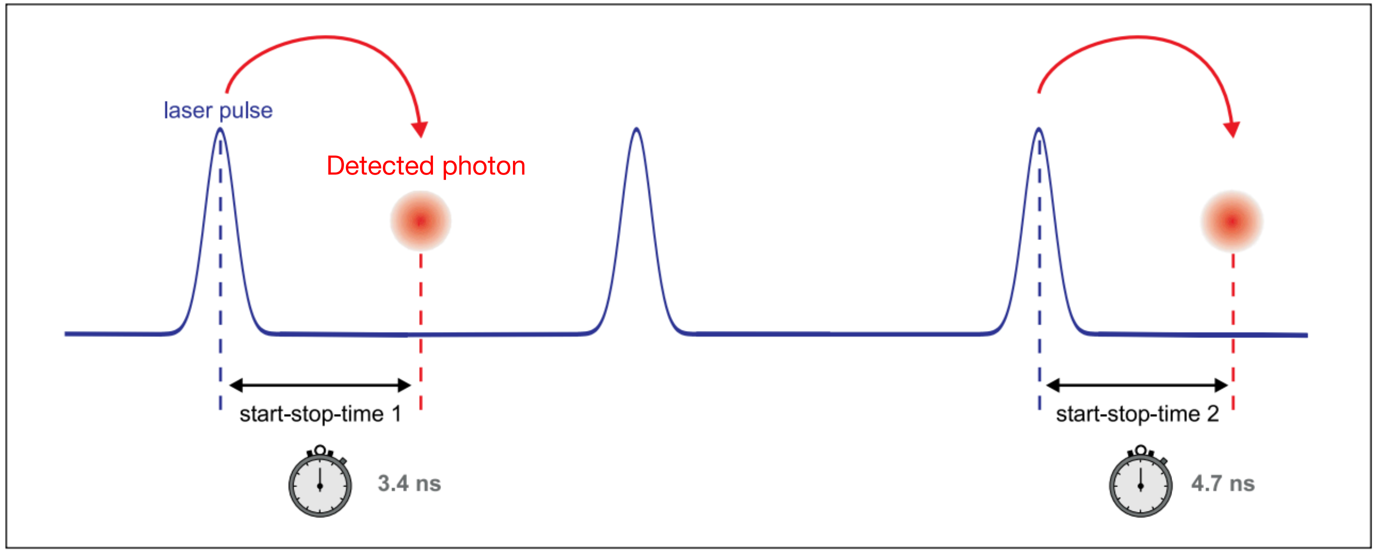

https://www.picoquant.com/images/uploads/page/files/7253/technote_tcspc.pdf

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

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

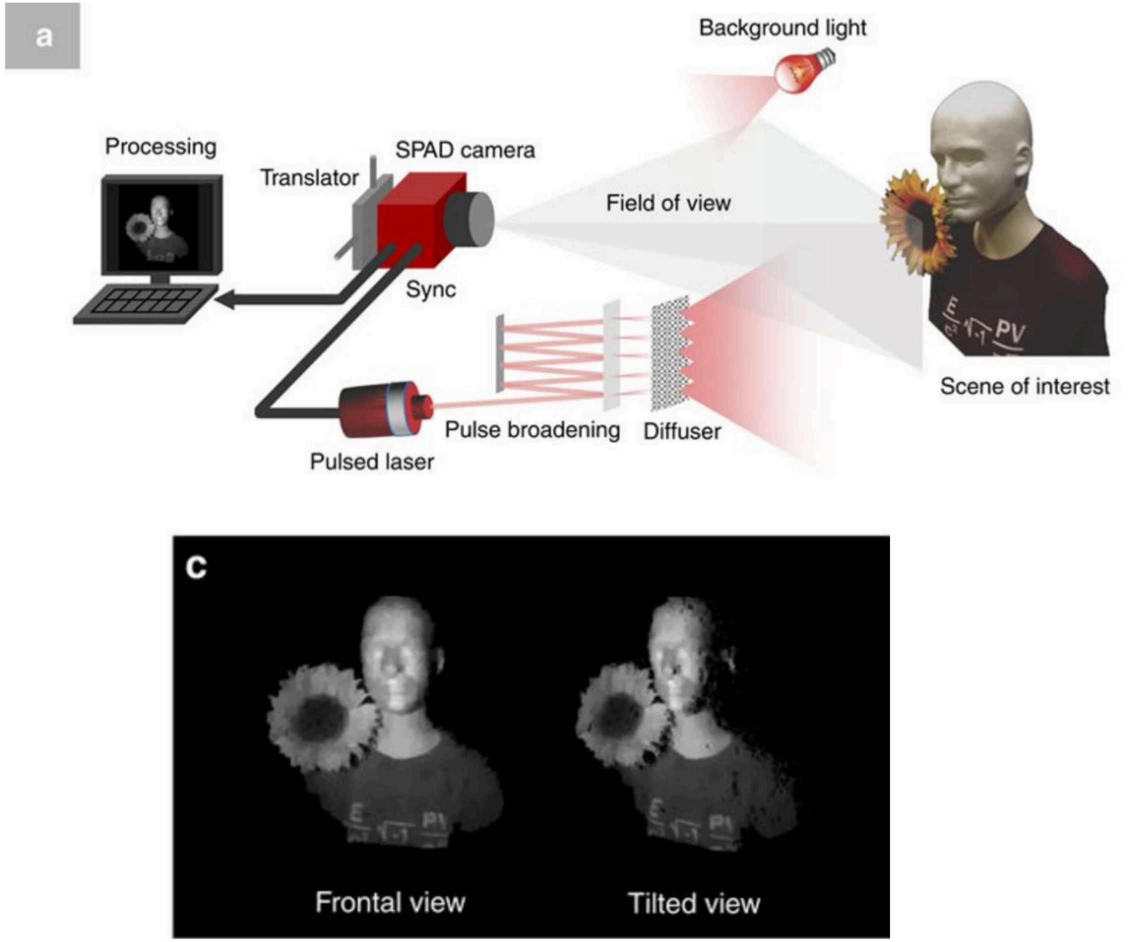

D. Shen et al, https://www.nature.com/articles/ncomms12046

Figure 1: Imaging Through Thick Scattering.

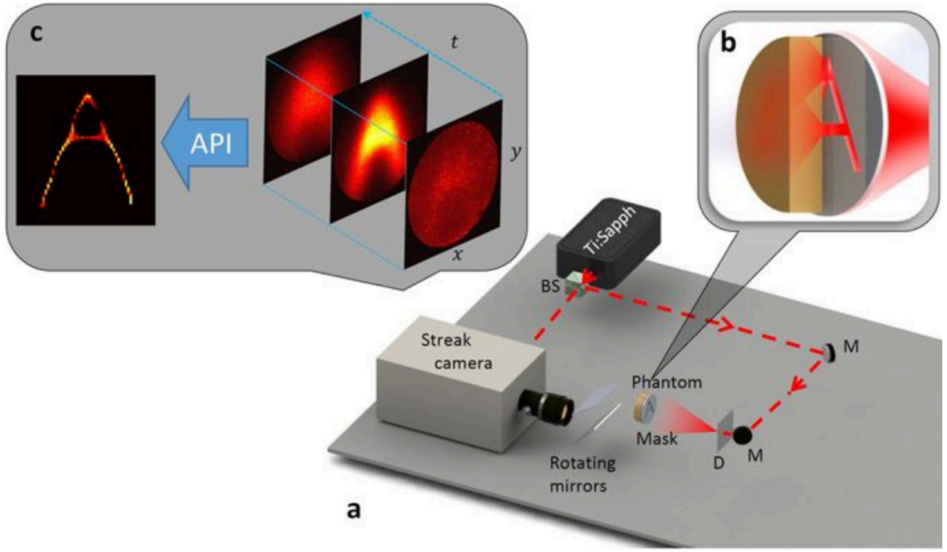

G. Satat et al, https://www.nature.com/articles/srep33946

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

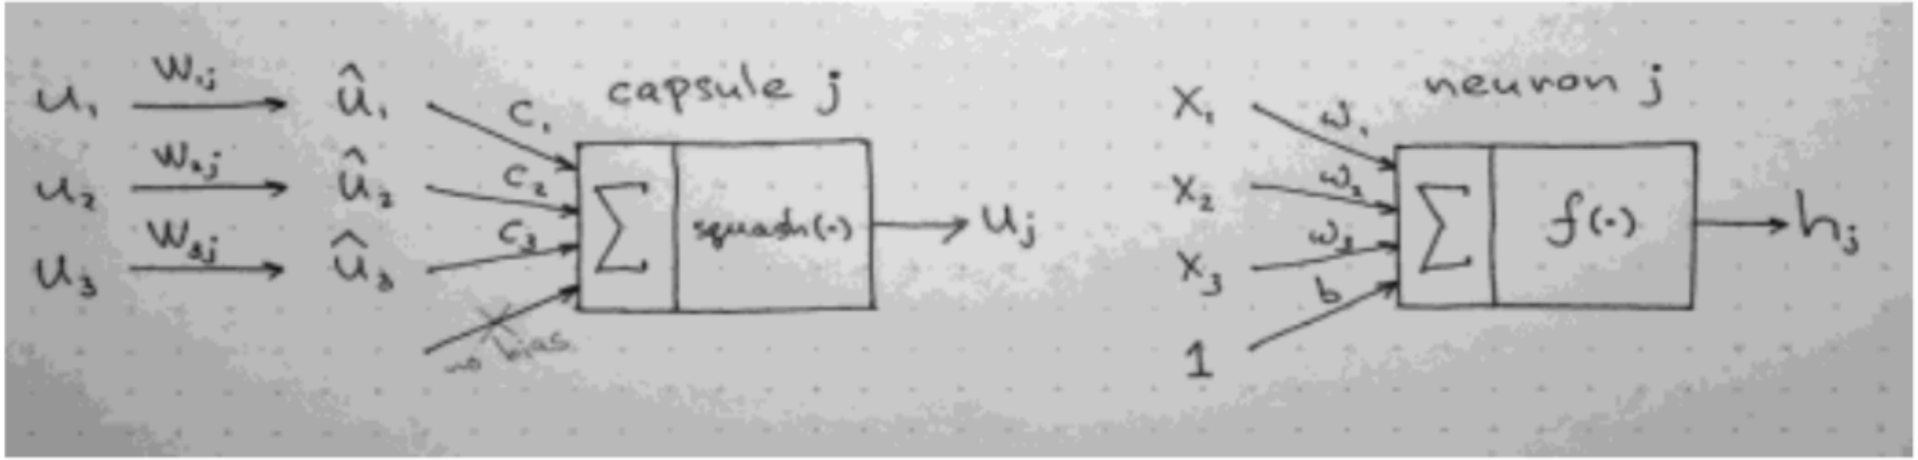

capsule j
squash(·)
no bias
neuron j
f(·)

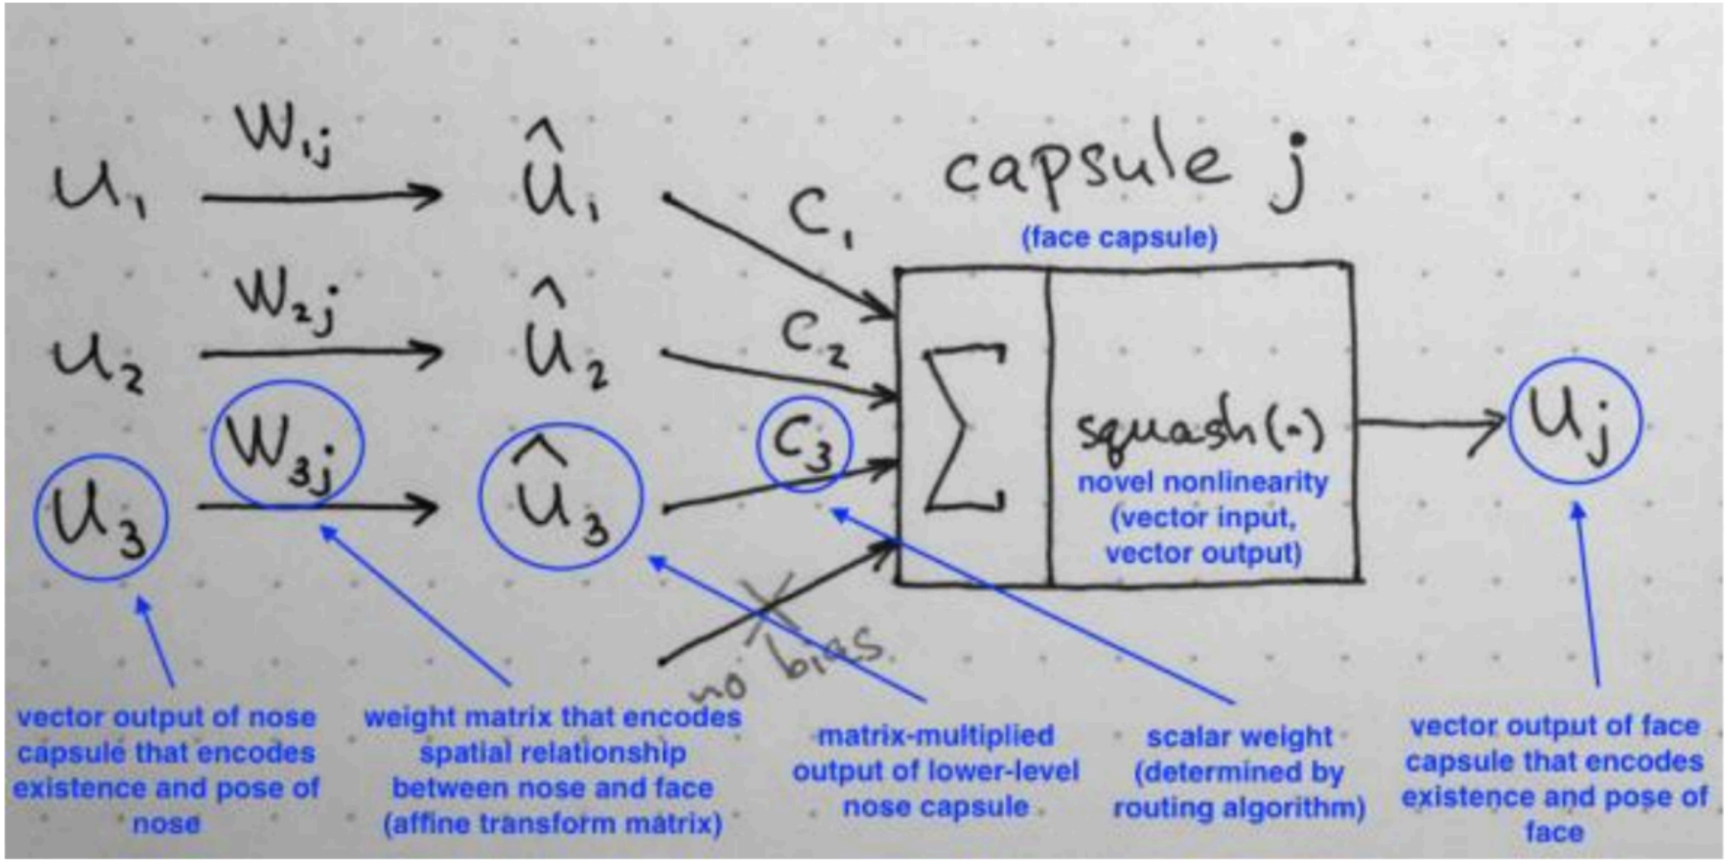

capsule j
(face capsule)
squash(·)
novel nonlinearity (vector input, vector output)
no bias
vector output of nose capsule that encodes existence and pose of nose
weight matrix that encodes spatial relationship between nose and face (affine transform matrix)
matrix-multiplied output of lower-level nose capsule
scalar weight (determined by routing algorithm)
vector output of face capsule that encodes existence and pose of face

deep imaging

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	$\mathbf{v}_j = \frac{\Vert\mathbf{s}_j\Vert^2}{1+\Vert\mathbf{s}_j\Vert^2}\frac{\mathbf{s}_j}{\Vert\mathbf{s}_j\Vert}$	$h_j = f(a_j)$
Output		vector$(\mathbf{v}_j)$	scalar(h_j)

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

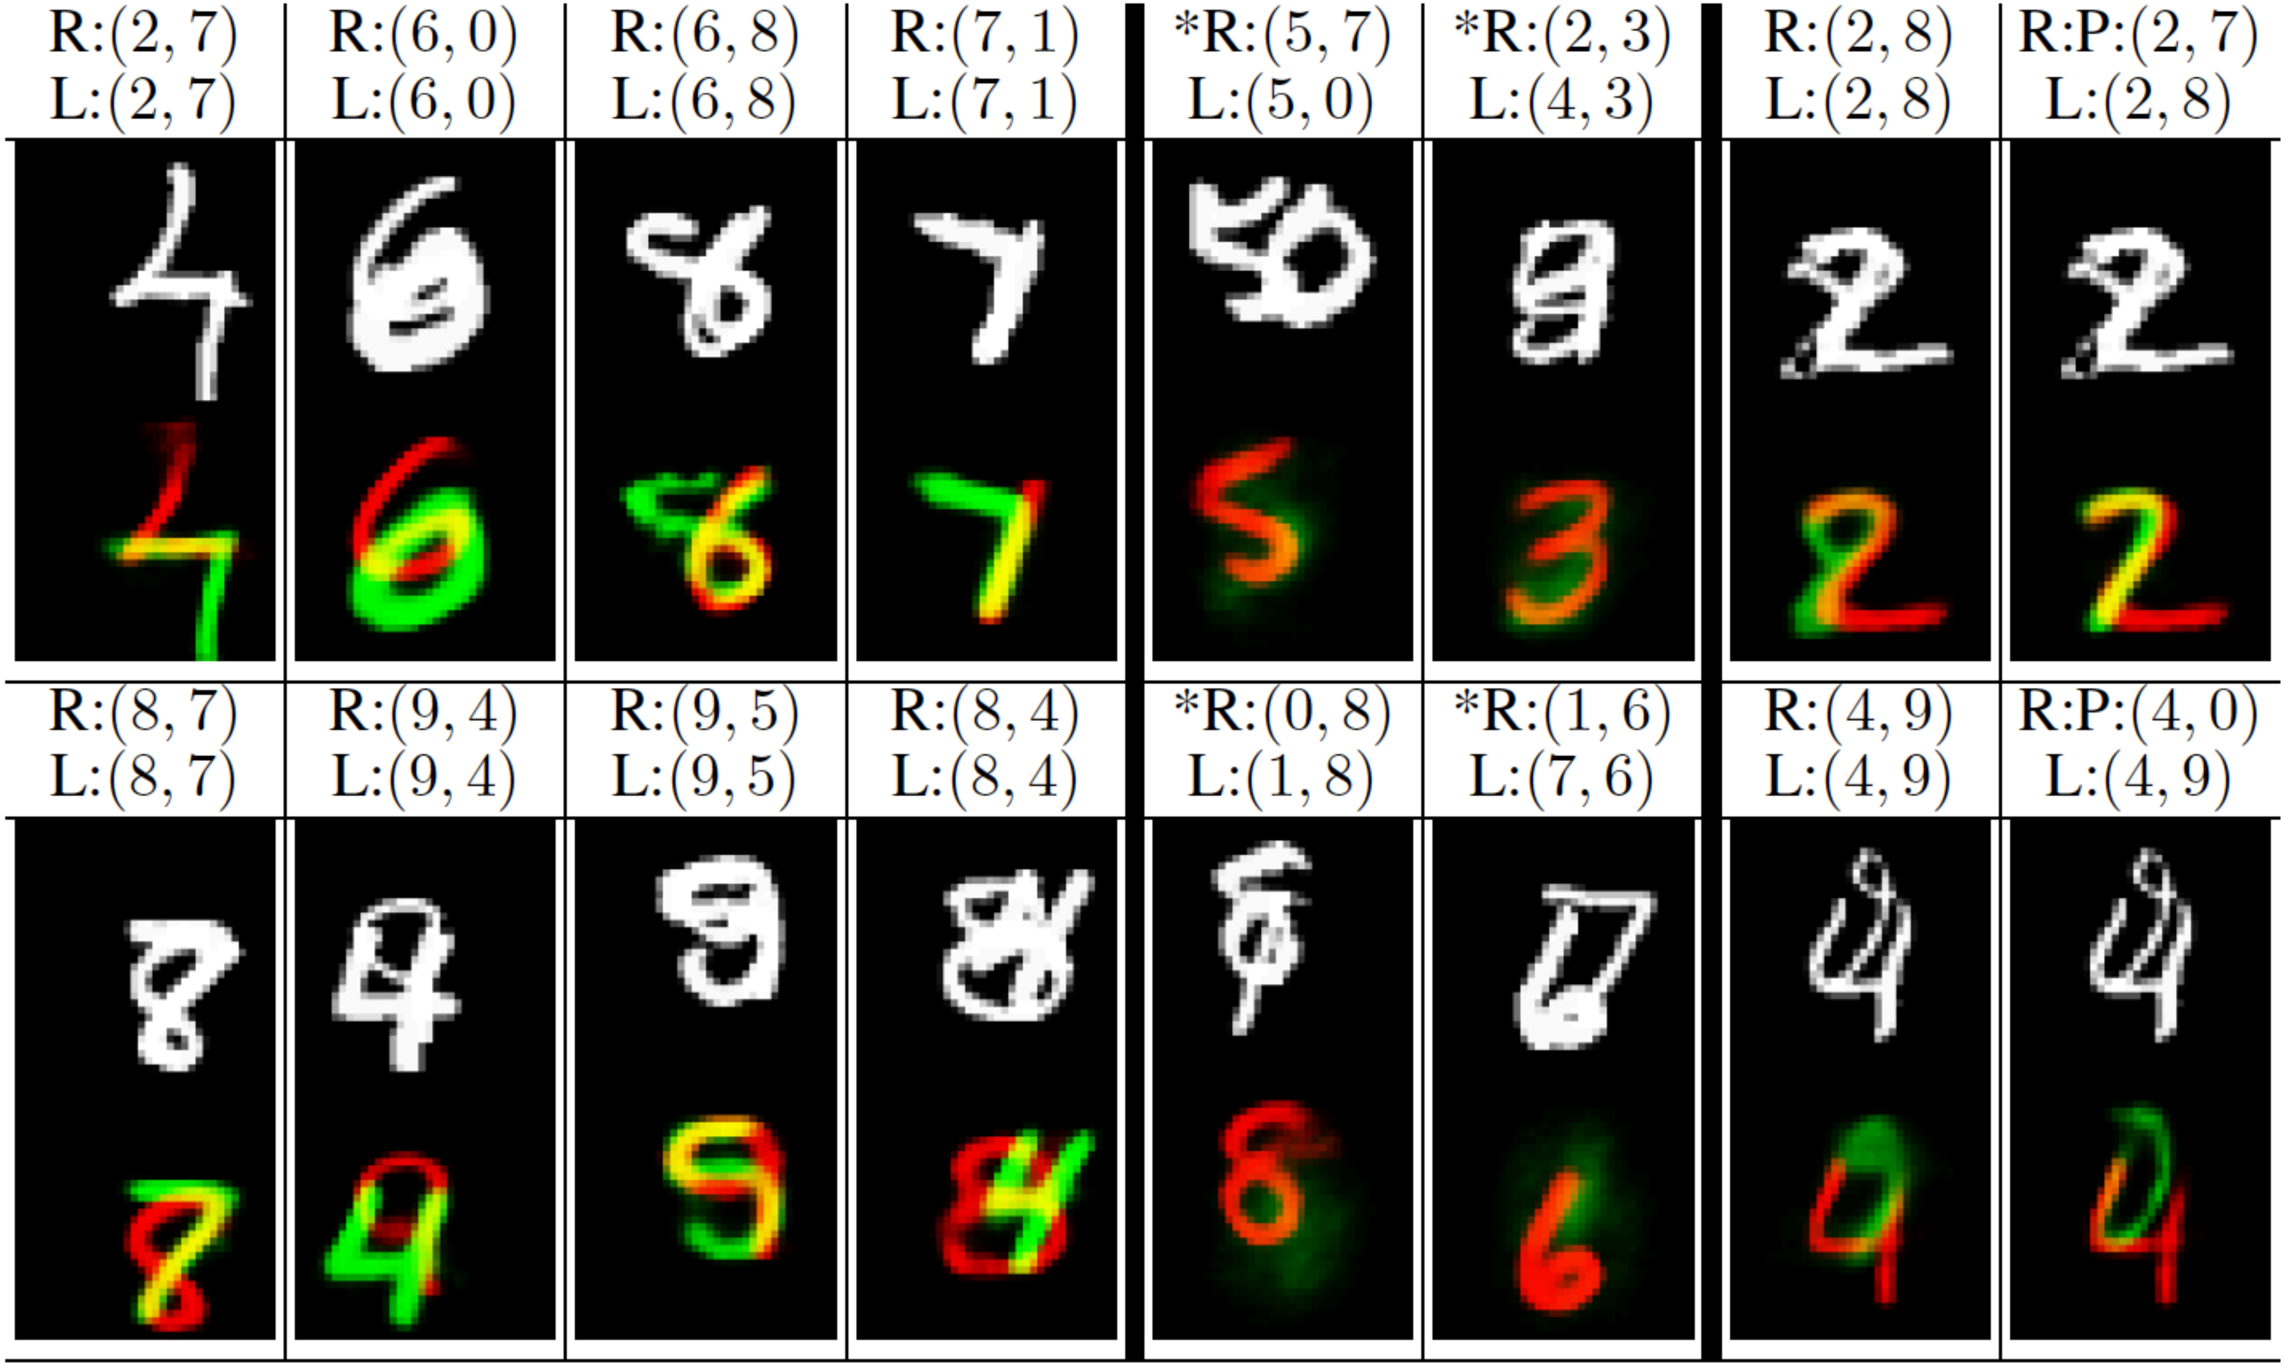

Where are things going with Machine Learning and Imaging in 10 years?

deep imaging

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

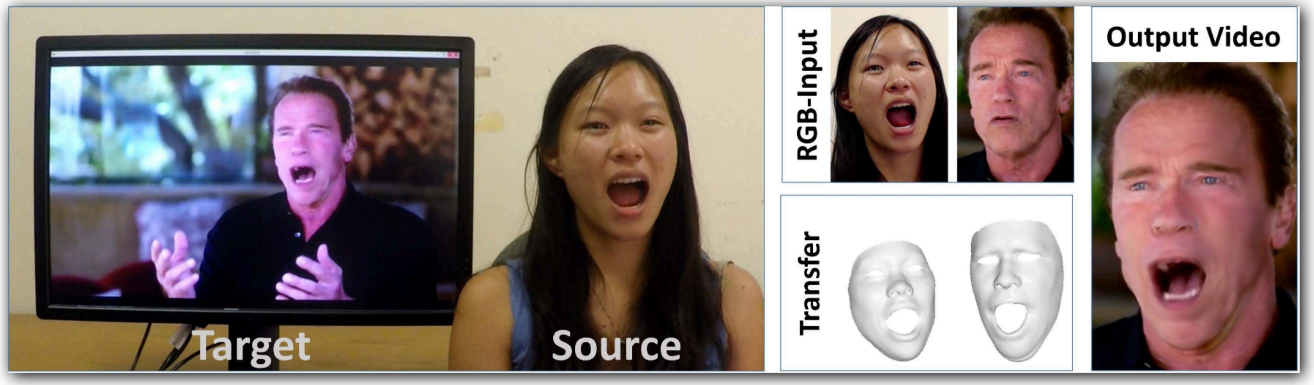

Proc. Computer Vision and Pattern Recognition (CVPR), IEEE, June 2016

What are the implications of this for medical imaging?

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

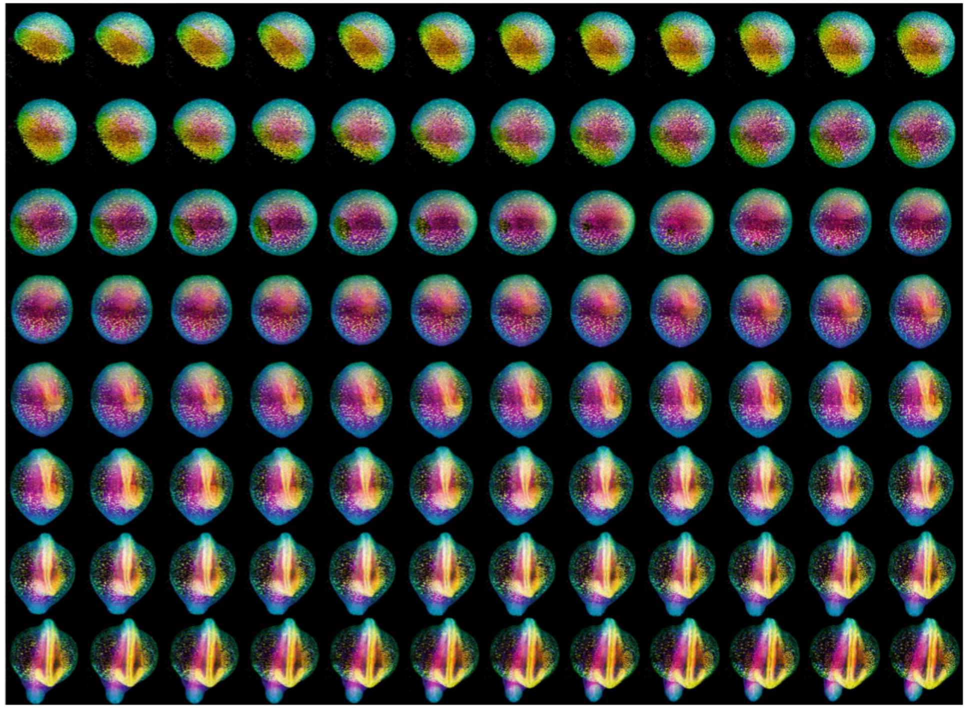

Take-aways for the future of machine learning and imaging

1. It's not going away....it works, there's a big community, and lots of $

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…