deep imaging

Lecture 0: Class outline and motivation

Machine Learning and Imaging

BME 548L

Roarke Horstmeyer

Organizational stuff

- This class will be in-person
- I will hold lectures on zoom, but you must let me know by 9am the day of the class if you plan to join via zoom for a particular reason

Organizational stuff

deep imaging

- This class will be in-person
- I will hold lectures on zoom, but you must let me know by 9am the day of the class if you plan to join via zoom for a particular reason
- **Course website**: **deepimaging.github.io**
- Homeworks will be announced on **Canvas** and posted on above webpage
- Homeworks will be a mix between “hand-written” problems and code
- “Hand-written” problems can be “written up” and turned in via pdf
- Coding assignments will be shared as a Google Colab/Jupyter notebook (.ipynb)
- Code should be written up “Jupyter notebook style”, mixing code, notes, results
- Coding assignments will be submitted via a link to the Google Colab page

deep imaging

Ed Discussion Page – hopefully it works!

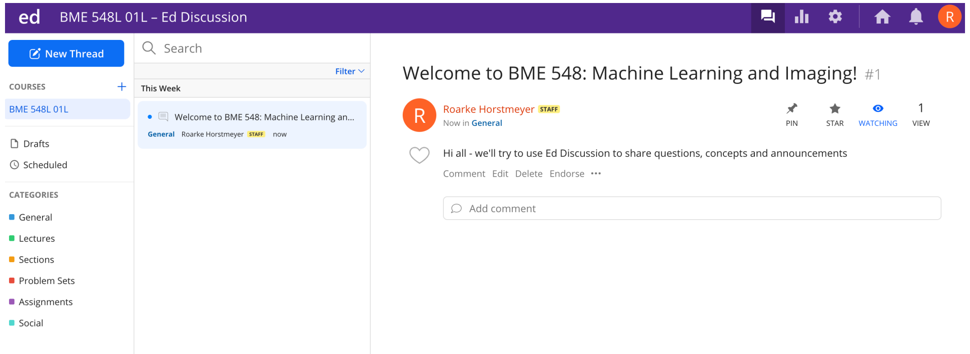

Course webpage

Machine Learning and Imaging

Syllabus Lectures Homework TA Sessions Project Info Past Projects Resources

Machine Learning and Imaging - Spring 2024

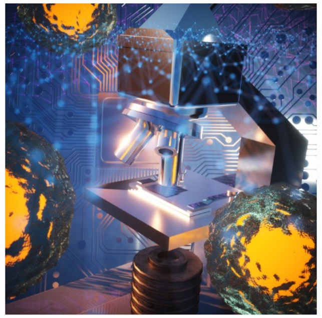

Welcome to Duke University's Machine Learning and Imaging (BME 548) class! This class aims to teach you how they to improve the performance of you deep learning algorithms, by jointly optimizing the hardware that acquired your data. It primiarly focuses on imaging data - from cameras, microscopes, MRI, CT, and ultrasound systems, for example. It begins with overview of machine learning and imaging science, and then focuses on the intersection of the two fields. This class is for you if 1) you would with imaging systems and you would like to learn more about machine learning, 2) if you are familiar with machine learning and would like to know more about how your data is gathered, 3) if you work with both imaging systems and machine learning and would like to hear a new perspective on the topic, or 4) if you work with neither imaging systems nor machine learning but have a strong mathematical background and are motivated to learn about both.

Canvas & Ed Discussion

deepimaging.github.io

Organizational stuff – Lab Sessions

- Monday and Wednesday 4:40pm-5:55pm
- Please bring your laptop! You'll need it
- Please feel free to attend whichever lab session is most convenient (there are plenty of seats)
- First lab sections will go over basics of code setup (Python, Jupyter, Google Colab, GCS)
- Then we'll get into classification & Tensorflow, homework help, etc.
- Please try out Jupyter/Google Colab on your machines before lab sections next week

Additional topics can be found in the syllabus:

https://deepimaging.github.io/syllabus/

We have a brief survey for you to please complete!

Please complete this before the next class:

https://deepimaging.github.io/assigns/

References and class material

– We will not use a textbook, but will rely on the following available resources for lecture notes:

Deep Learning, A. Goodfellow et al.: https://www.deeplearningbook.org/
Introduction to Fourier Optics, J. Goodman
Learning from Data, Y. S. Abu-Mustafa
Introduction to Linear Algebra, G. Strang

Helpful lectures:
Stanford CS231n: http://cs231n.stanford.edu/syllabus
Caltech, Learning from Data: https://work.caltech.edu/telecourse.html
Stanford CS230: http://cs230.stanford.edu/syllabus

Organizational stuff - Grading

- 5 homework assignments, maybe 2 short quizzes, 1 final project
- -20%/day for late homework
- Final project is important, mostly for you!
- Participation: come to lecture & lab & office hours if needed, self-scored
- Collaboration encouraged, but please write up your *own code* and *own solutions*

Grading:
Homework: 43%
Final project: 38%
Project proposal: 6%
Lab Workbooks: 5%
Participation: 8%

Tentative Course schedule:

Week 0 – Jan 10: Machine learning and imaging systems in a nutshell

Week 1 – Jan 17 (Holiday, no class), Jan 18: Review background mathematics

Week 2– Jan 22, Jan 24: Linear algebra, optimization and cost functions

Week 3– Jan 29, 31: From optimization to machine learning

Week 4 – Feb 5, 7: Neural networks, the chain rule and back-propagation

Week 5 – Feb 12, 14: Convolutional neural networks (CNN's)

Week 6 – Feb 19, 21: CNN's in practice

Week 7 – Feb 26, 28: Light propagation and imaging systems

Week 8 – March 4, 6: Computational models of imaging systems

Week 9 – March 11, 13: Spring Break, No Class

Week 10 – March 18, 20: Project proposals and discussions

Week 11 – March 25, 27: Designing imaging systems with CNN's

Week 12 – April 1, 3: System Design and Optimization with CNN's

Week 13 – April 8, 10: Object Detection, Autoencoders

Week 14 – April 15, 17: Reinforcement Learning, Gen. adversarial networks

Final project presentations: Monday May 2, 7pm-10pm (class finals slot), details TBD.

deep imaging

What is this class about?

Historically, hardware has been optimized to create images for <u>humans</u>

deep imaging

What is this class about?

Over the last 10 years, computers have become “really good” at automatically processing image data

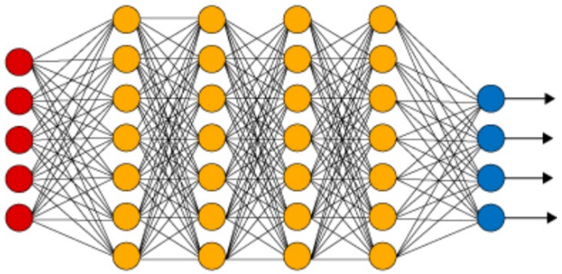

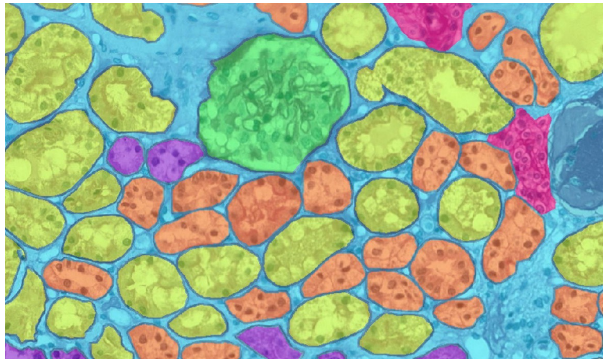

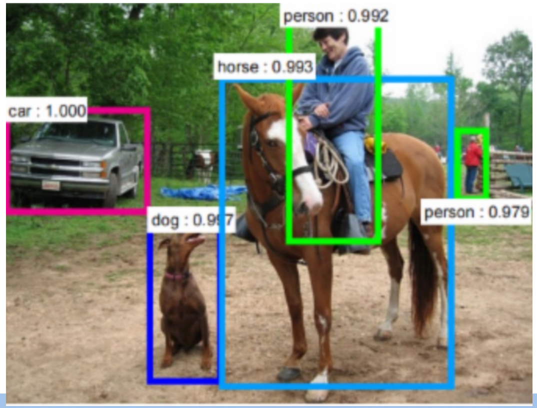

Nearly all new advances enabled by deep neural networks

Convolutional neural networks for automated annotation of cellular cryo-electron tomograms

Muyuan Chen[1,2], Wei Dai[2,4], Stella Y Sun[2], Darius Jonasch[2], Cynthia Y He[3], Michael F Schmid[2], Wah Chiu[2] & Steven J Ludtke[2]

deep imaging

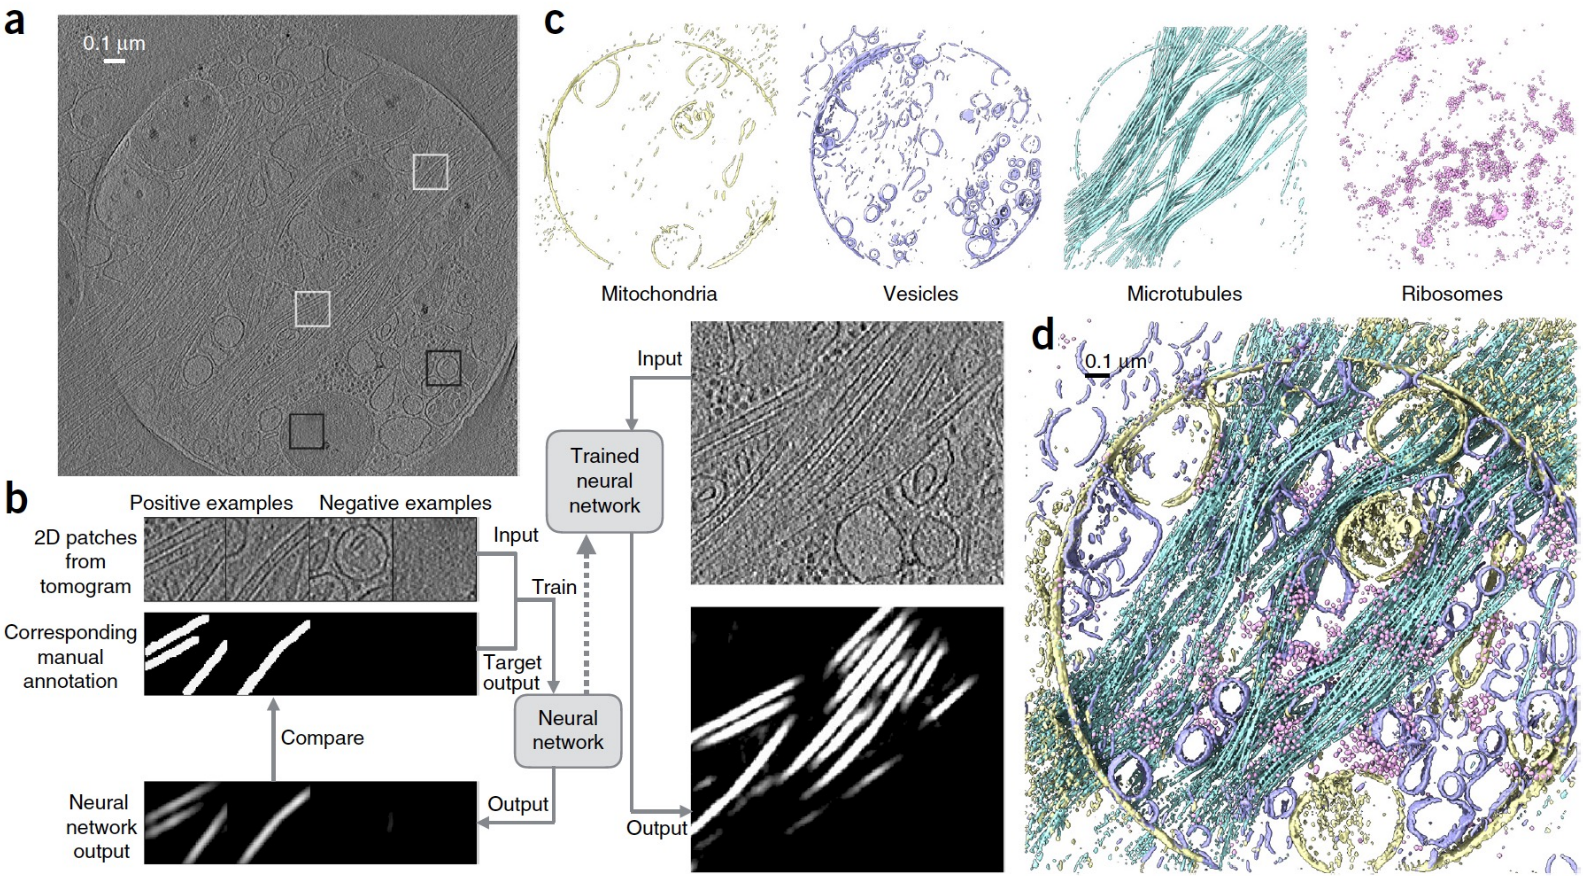

Automatic Neuron Detection in Calcium Imaging Data Using Convolutional Networks

Noah J. Apthorpe[1*] Alexander J. Riordan[2*] Rob E. Aguilar[1] Jan Homann[2]
Yi Gu[2] David W. Tank[2] H. Sebastian Seung[12]
[1]Computer Science Department [2]Princeton Neuroscience Institute
Princeton University

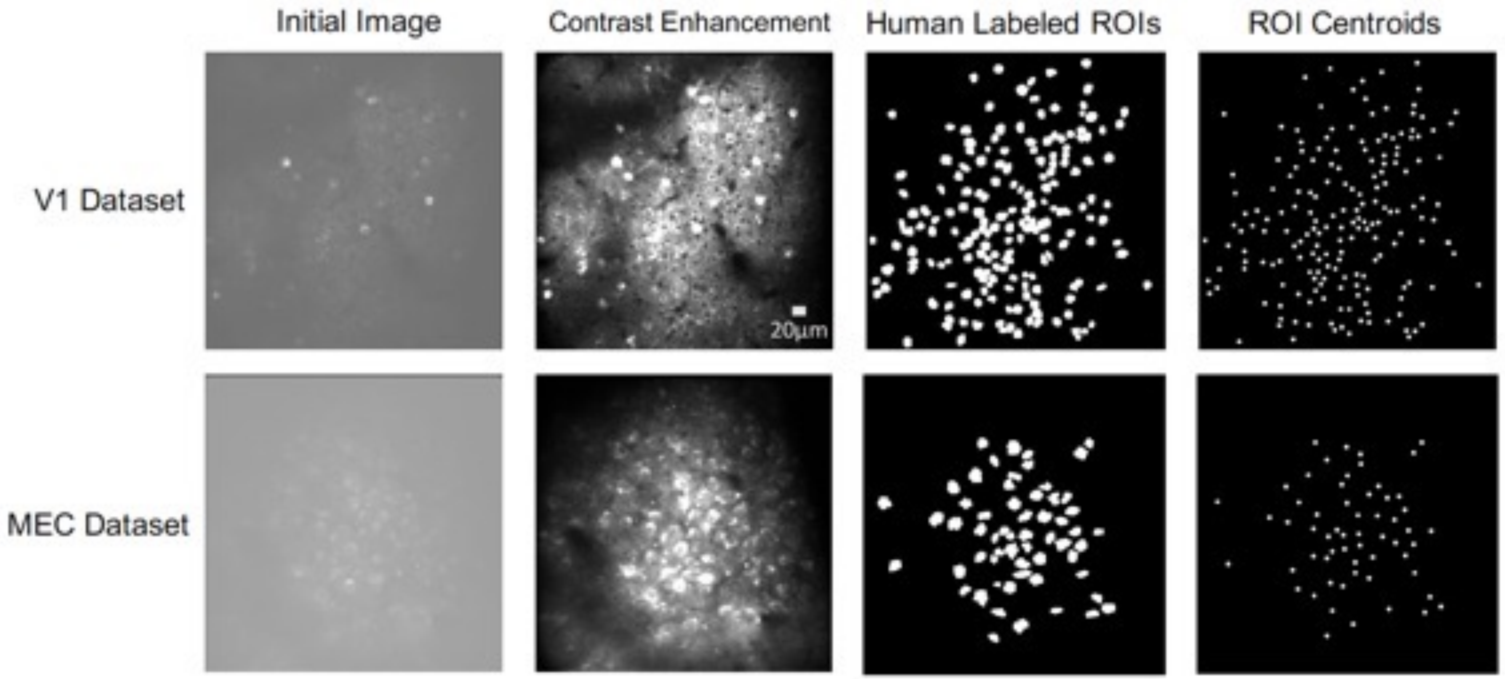

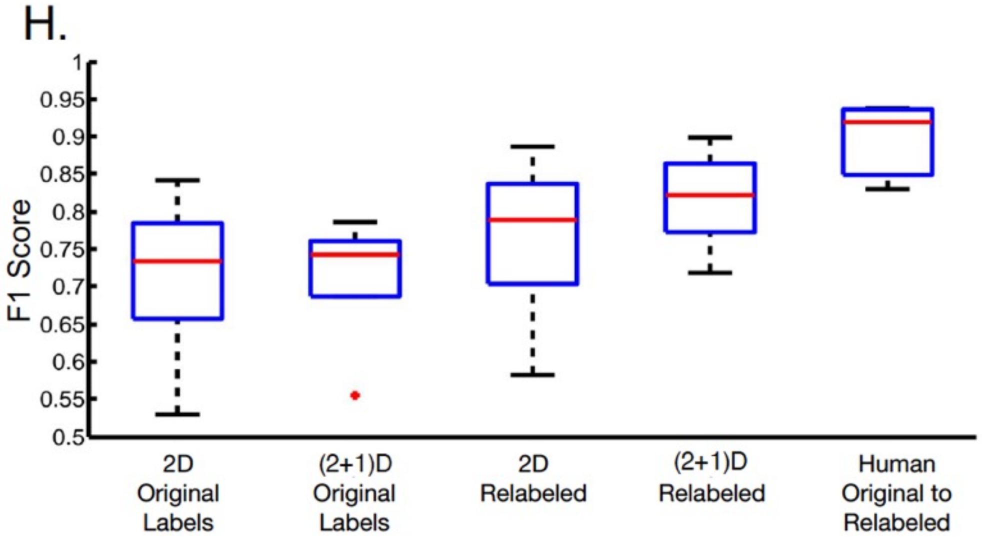

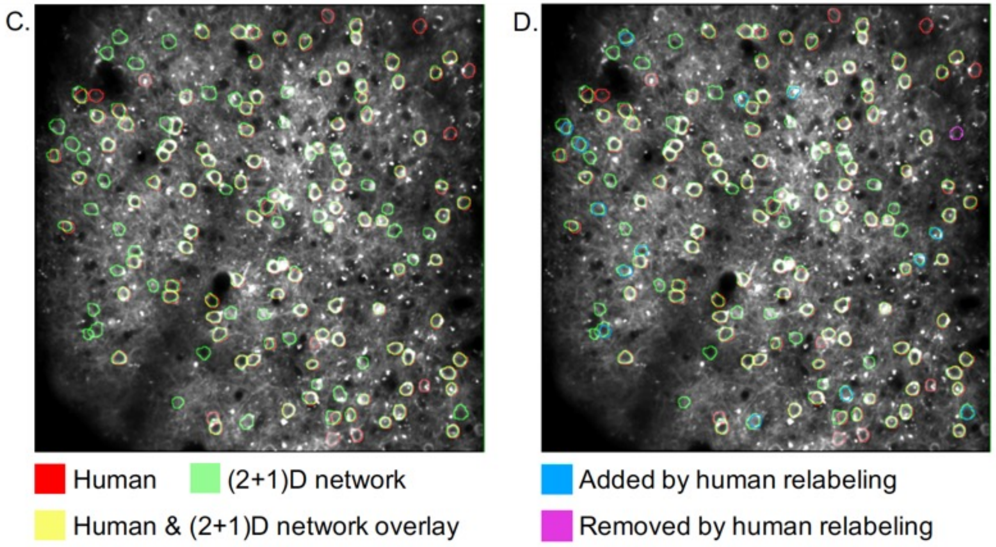

Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva[1]*, Brett Kuprel[1]*, Roberto A. Novoa[2,3], Justin Ko[2], Susan M. Swetter[2,4], Helen M. Blau[5] & Sebastian Thrun[6]

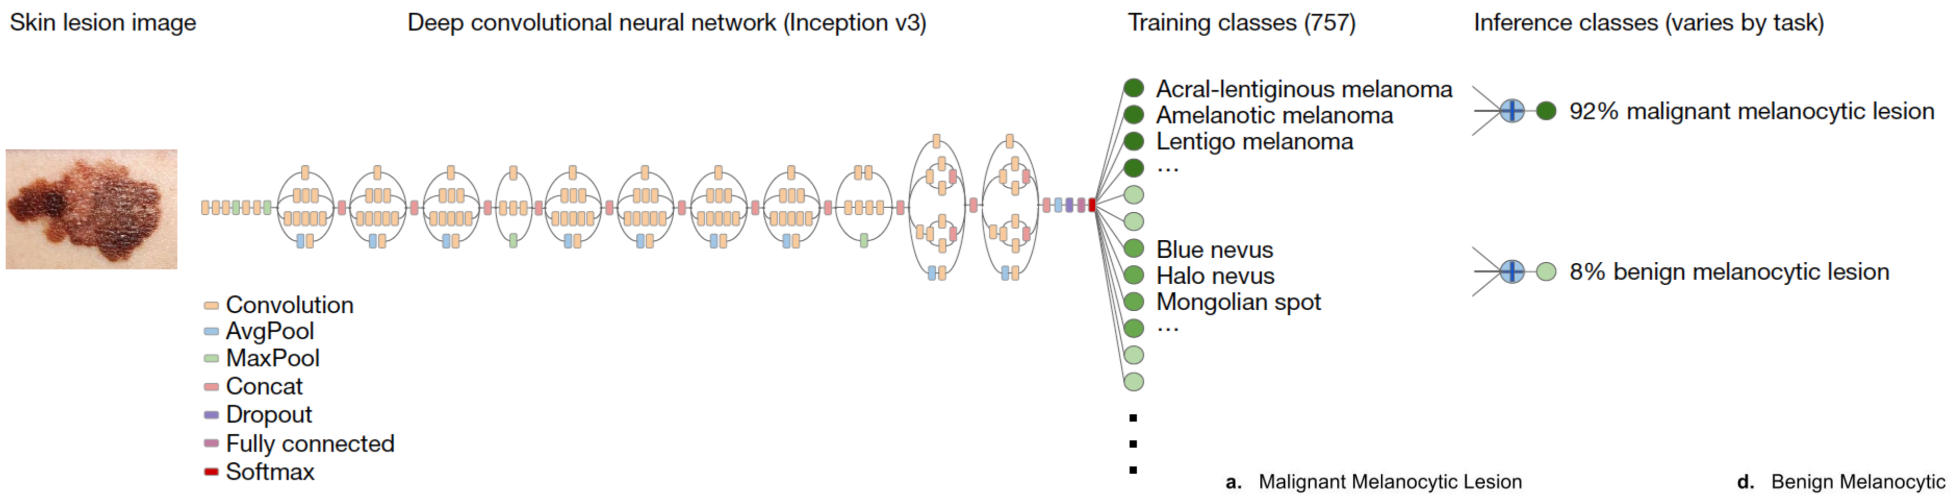

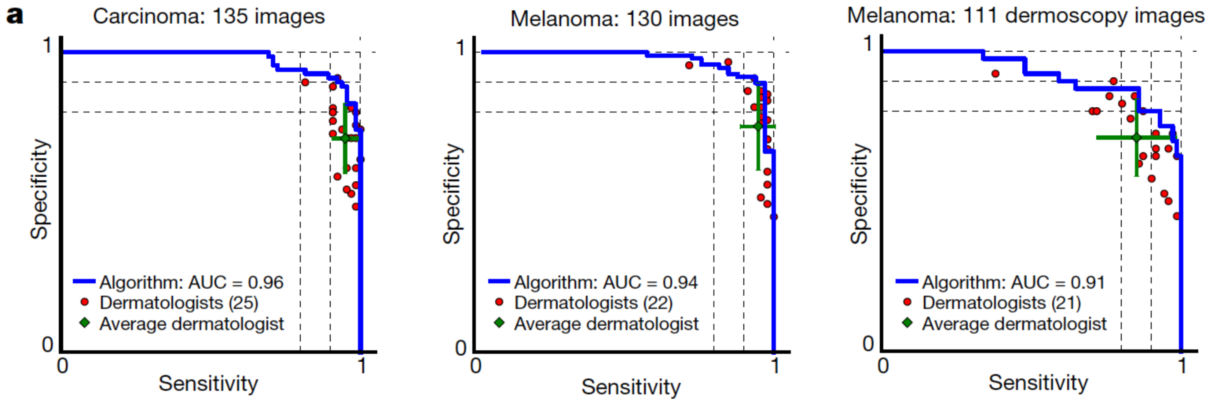

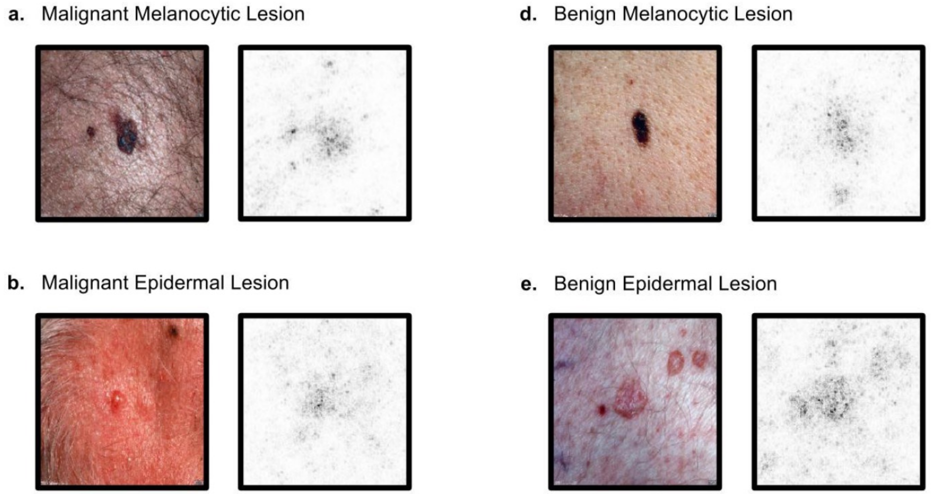

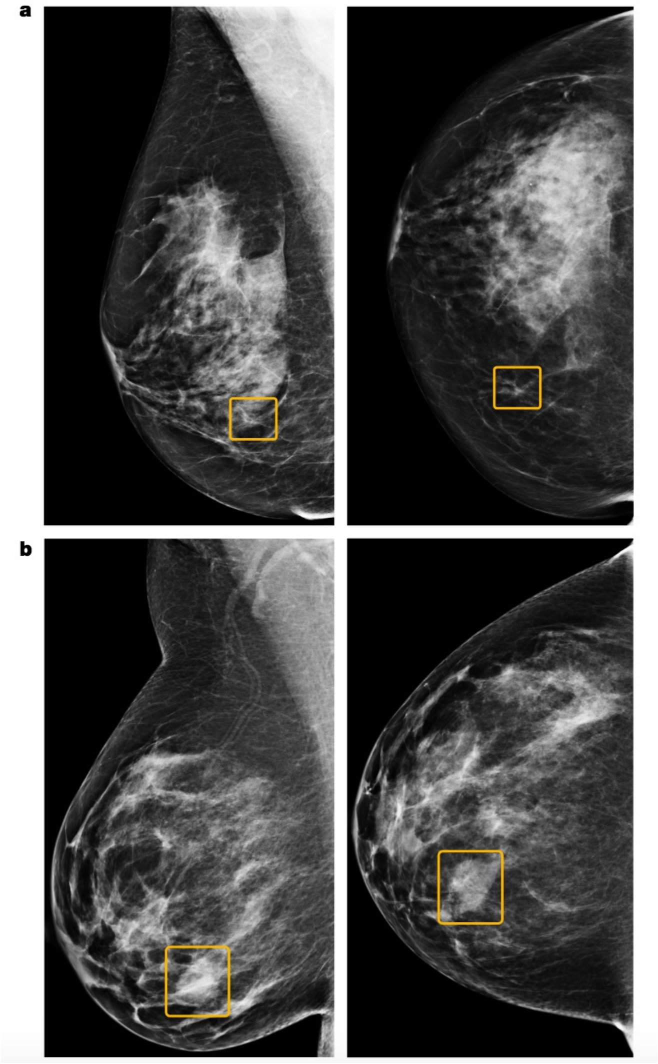

Article

International evaluation of an AI system for breast cancer screening

https://doi.org/10.1038/s41586-019-1799-6

Received: 27 July 2019

Accepted: 5 November 2019

Published online: 1 January 2020

Scott Mayer McKinney[1,14*], Marcin Sieniek[1,14], Varun Godbole[1,14], Jonathan Godwin[2,14], Natasha Antropova[2], Hutan Ashrafian[3,4], Trevor Back[2], Mary Chesus[2], Greg C. Corrado[1], Ara Darzi[3,4,5], Mozziyar Etemadi[6], Florencia Garcia-Vicente[6], Fiona J. Gilbert[7], Mark Halling-Brown[8], Demis Hassabis[2], Sunny Jansen[9], Alan Karthikesalingam[10], Christopher J. Kelly[10], Dominic King[10], Joseph R. Ledsam[2], David Melnick[6], Hormuz Mostofi[1], Lily Peng[1], Joshua Jay Reicher[11], Bernardino Romera-Paredes[2], Richard Sidebottom[12,13], Mustafa Suleyman[2], Daniel Tse[1*], Kenneth C. Young[8], Jeffrey De Fauw[2,15] & Shravya Shetty[1,15*]

MENU

nature medicine

Letter | Published: 19 May 2020

Artificial intelligence–enabled rapid diagnosis of patients with COVID-19

Xueyan Mei, Hao-Chih Lee, [...] Yang Yang

Fig. 3: Examples of chest CT images of patients with COVID-19 and visualization of features correlated to COVID-19 positivity.

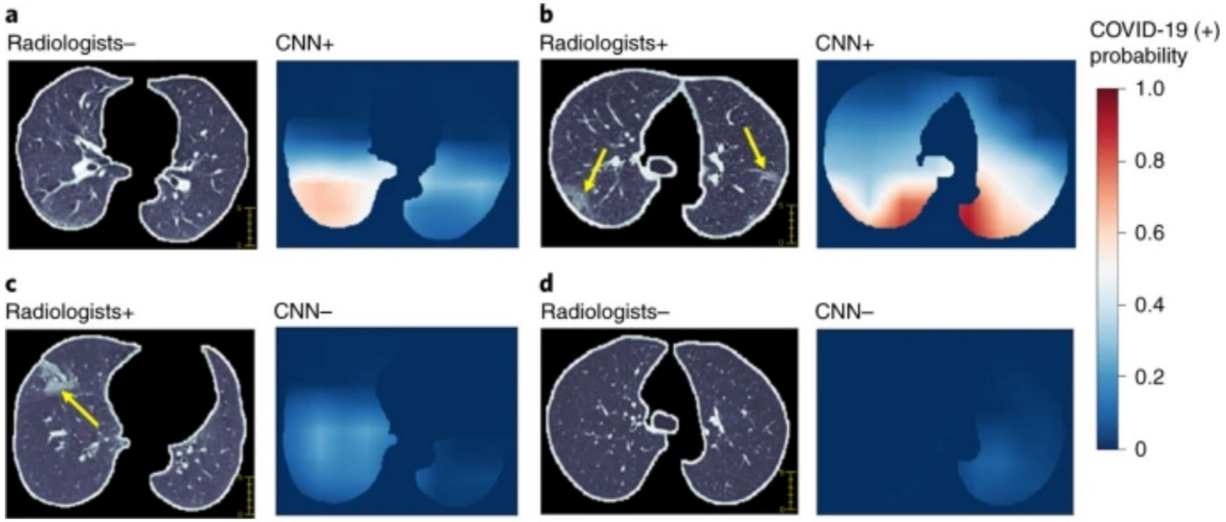

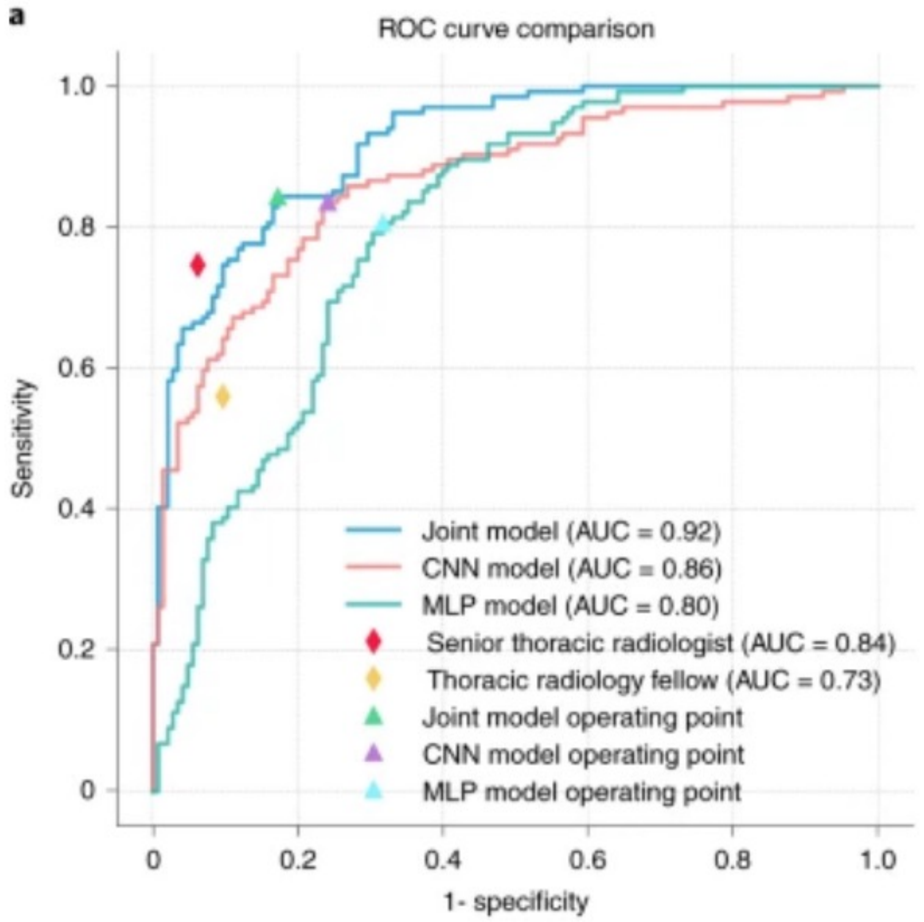

Computational Optics Lab – Duke University

deep imaging

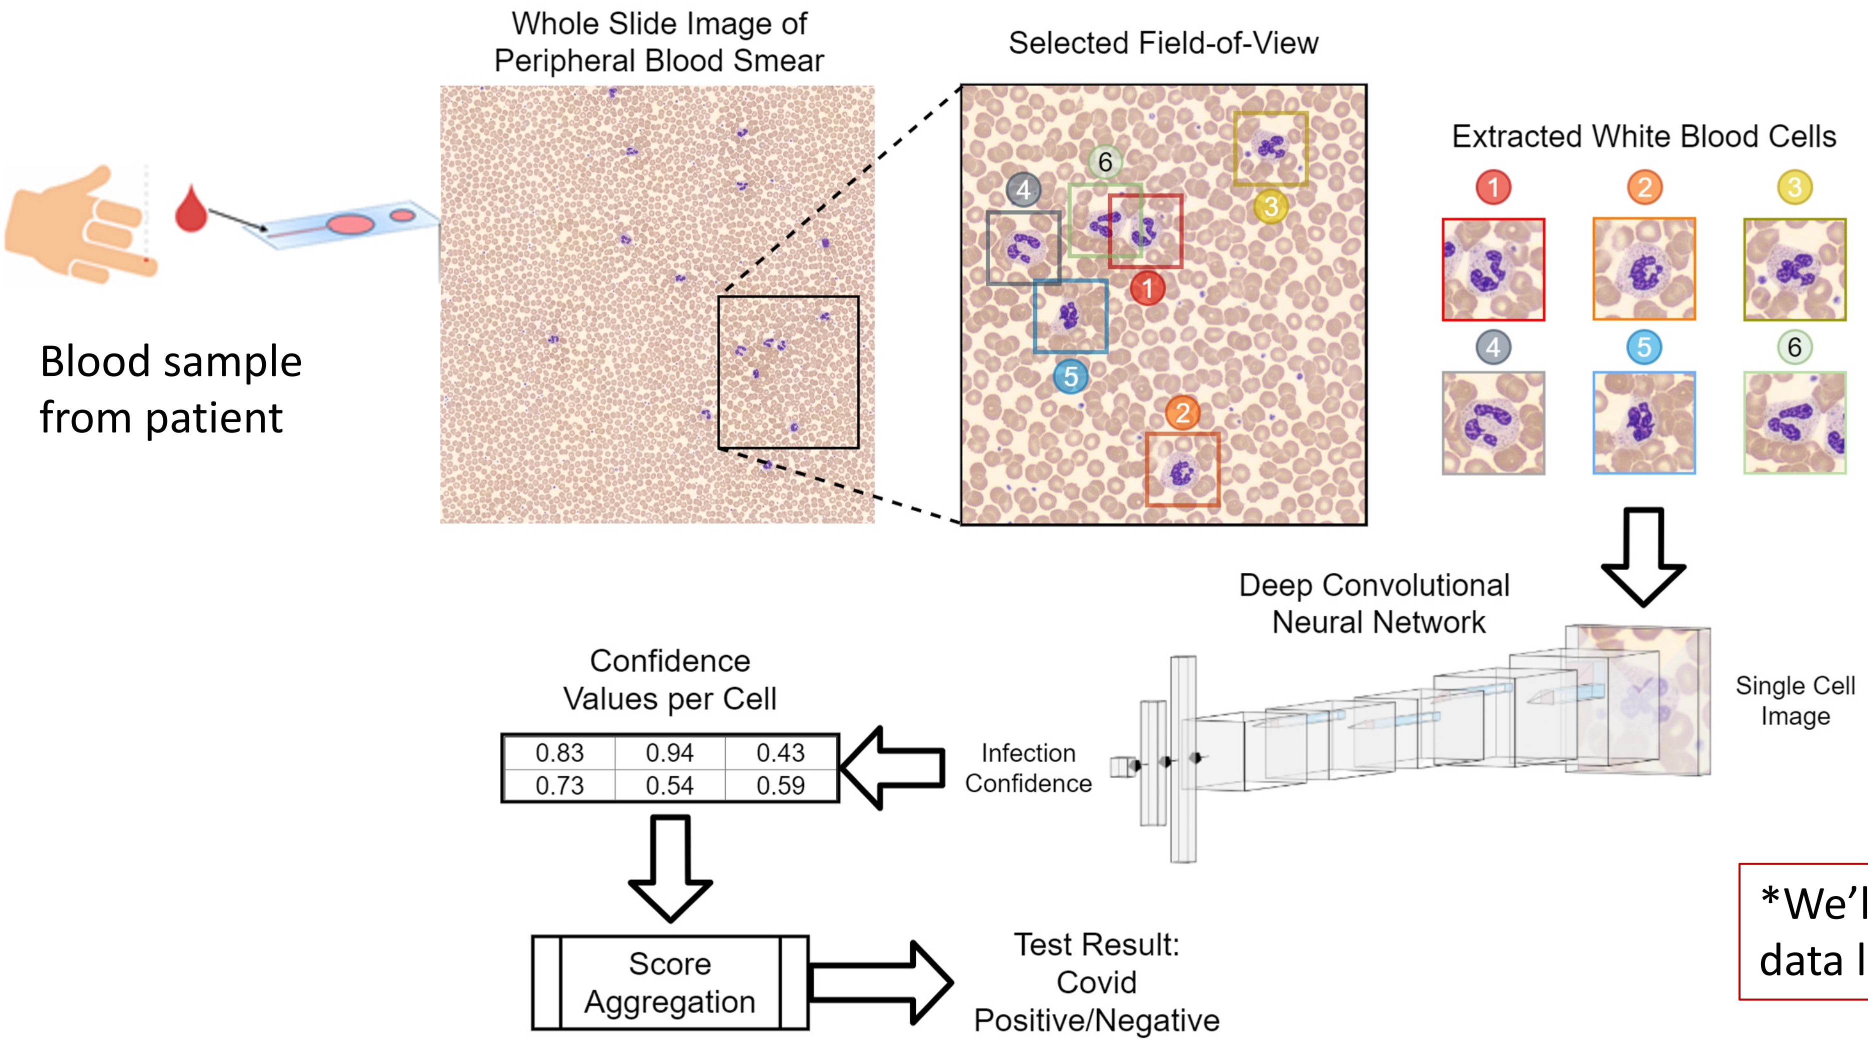

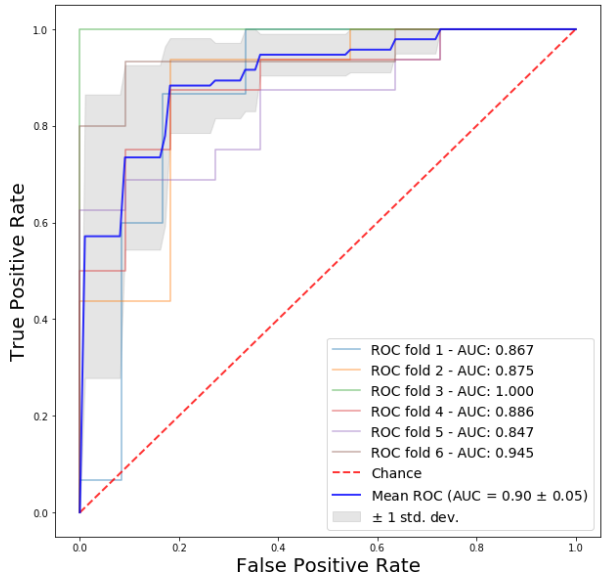

*We'll be working a bit with data like this during the class!

C. L. Cooke et al, "A multiple instance learning approach for detecting COVID-19 in peripheral blood smears," PLOS Digital Health (2022)

What is this class about?

deep imaging

We will figure out what is going to be right here

Computer-centered hardware + software

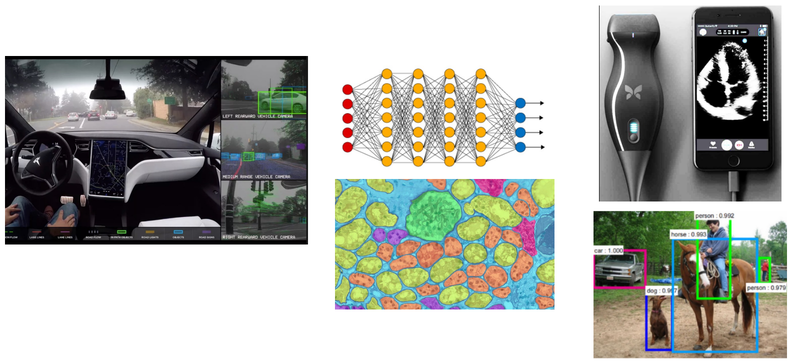

Human-centered hardware design

Computer-centered software design

Some examples of computer-centered hardware?

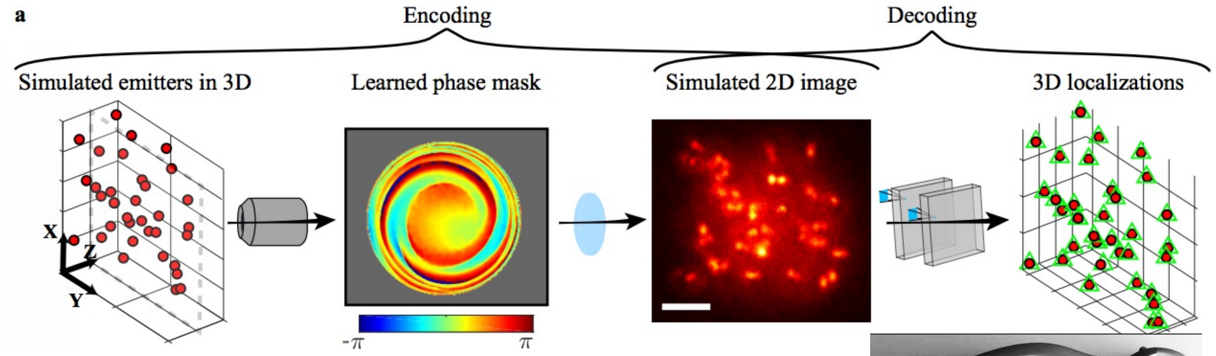

DeepStorm3D

Butterfly

Low-dose CT – Chen et al.

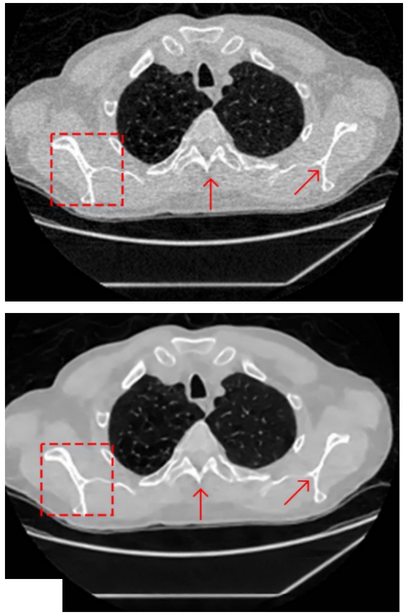

FastMRI.org

Under-sampled MRI challenge

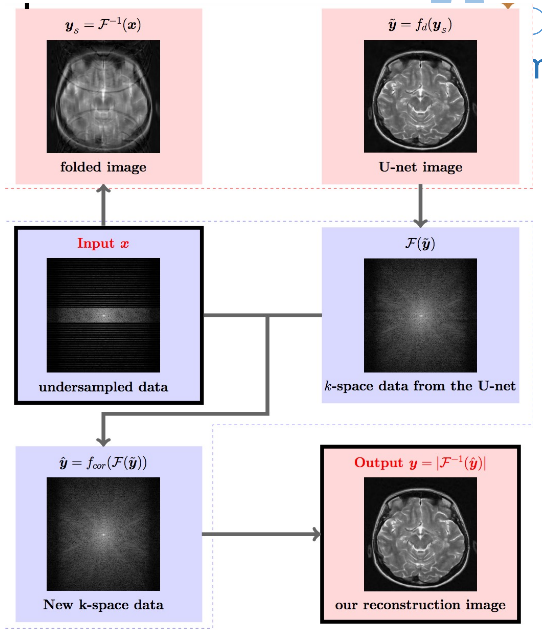

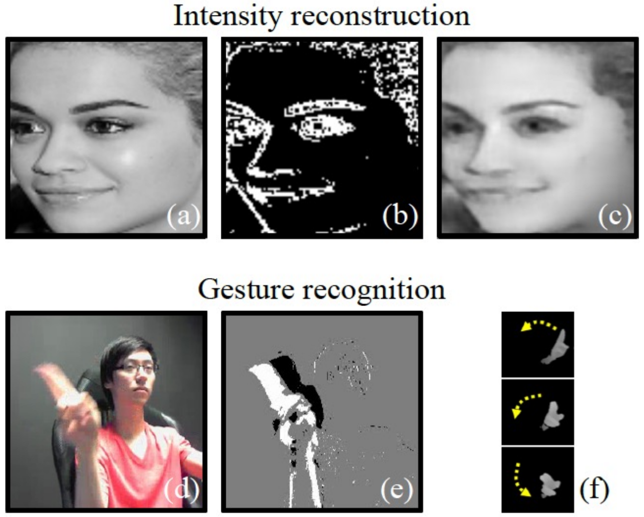

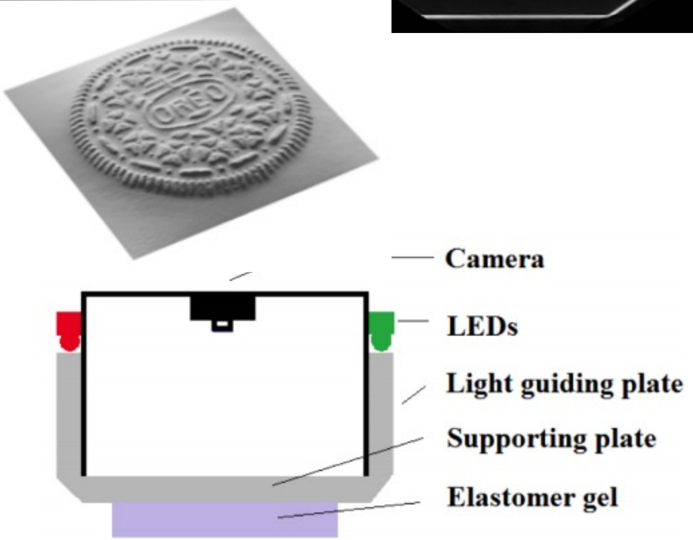

Nvidia

https://youtu.be/BIW_jq3dOEE

deep imaging

General layout for this class:

- Here is what I plan to cover:
    - Fundamental concepts behind machine learning
    - Current methods in machine learning for image analysis
        - Deep neural neworks & CNN’s
        - Classification, segmentation, translation, super-resolution
    - How to model *simplified* imaging systems (cameras, microscopes, ultrasound, CT, etc.)
    - How to optimize imaging system hardware with and for machine learning

What is this class *not* about?

deep imaging

“I want to get this score from 0.982 up to 0.999”

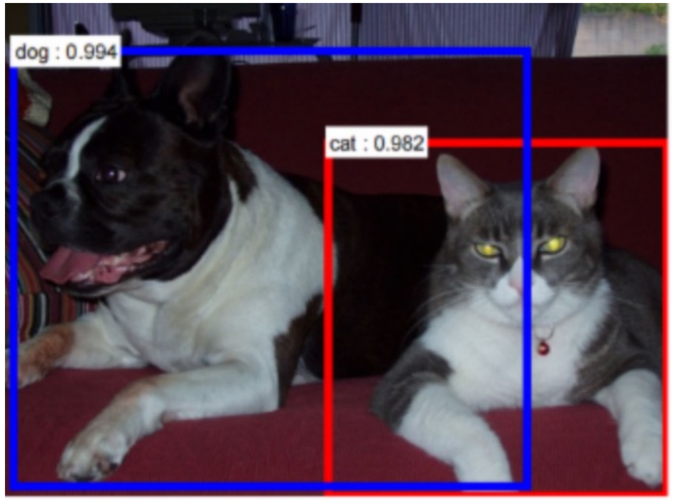

“Can we create a comprehensive mathematical framework to understand how deep learning algorithms work?”

“I’d like an in-depth explanation of how my (fill-in-the-blank) imaging system works”

“I want to really understand this one operation here”

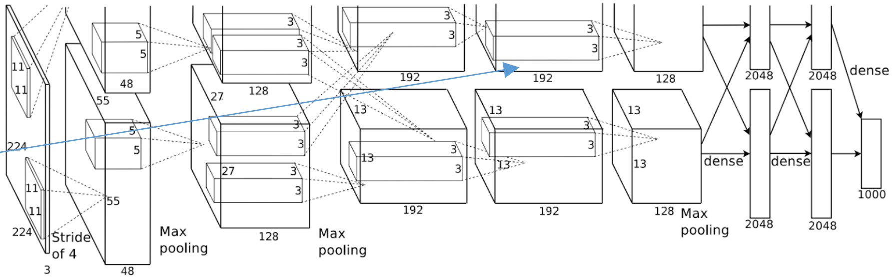

“I want to program something to make cool pictures like this”

What do I need to know about beforehand to succeed in this class?

deep imaging

- Python coding experience ideally, but MATLAB programming experience will probably be ok
- Linear algebra & Calculus
  - Vector/matrix operations
  - Matrix inversion, pseudo-inverse
  - Gradients, partial derivatives
- Signal processing
  - Complex-valued signals
  - Fourier transforms
  - Convolutions
- Optimization
  - Differences between Linear, convex, non-linear optimization
  - Gradient descent

This class is interdisciplinary (by design)

deep imaging

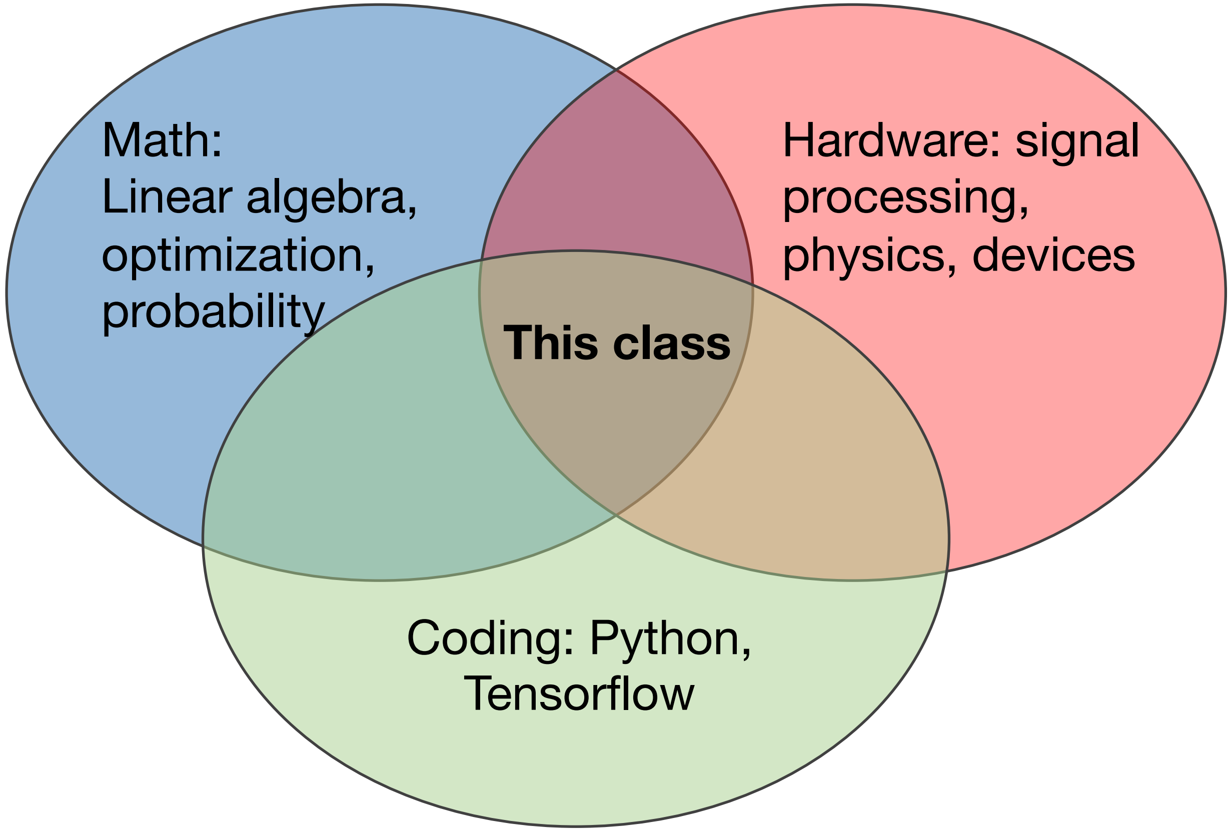

- We will move pretty fast
- We will jump between subjects I assume most have had "some" exposure to
- End goal: Meaningful project with Tensorflow
- This is the fifth time teaching this, so we have some proof it works…

What should I expect to gain from this class?

- Comfort with general mathematical principles behind machine learning
- Comfort with how to simply model generalized imaging systems (math and simulation)
- Ability to program in Python and Tensorflow
- Hands on experience with current “deep” ML algorithms (convolutional neural networks, GANs maybe)
- An ability to reason thought architecture choices for deep CNNs
- Coding experience with adding imaging system hardware optimization into a CNN or other NN structure

deep imaging

Example Final Projects from Machine Learning and Imaging

Finding Ultrasound Sub-apertures for Liver Vessel Segmentation
Single-Pixel, Single-Frequency Hand Gesture Recognition with a Dynamic Metasurfaces
Going Deeper: Depth Image Classification via simulated SPAD array images
Trained Blur Kernel for histology slide segmentation using a Deep Neural Network
Classification of Tuberculosis Bacilli With and Without Staining
A deep learning approach to improving ultrasonic plane wave imaging
Automated Image Focus Detecting Algorithm for Low-Cost Handheld Microscope
Optimal shift-variant point-spread function for improved classification
Deep Learning for Motion Tracking on the Micron Scale with Ultrasound
Sensor Multiplexing and Reconstruction for Color Images
Noise Reduction in Optical Coherence Tomography using a Deep Image Prior
Optimization of illumination for Unet-Base Cervix Segmentation
HDR image reconstruction with filters over pixels – What is the optimal design?
Detection of Lesions in Variably Noisy Ultrasound Images Using Machine Learning
Methods for Segmentation of Fine Structure in Rodent Histological Specimens
Direct reconstruction network for photoacoustic imaging with fewer measurements
Machine Learning for Ultrasound Lesion Mapping with Apodization Optimization
Resolution versus Precision in X-ray detection of Pneumonia
Optimizing illumination for overlapped image classification