“Deep Learning is Overhyped”…

…Is Overhyped

“Deep Learning is just the latest fad – next year it’ll be something else”

Mark I Perceptron at the Cornell Aeronautical Laboratory (1957)

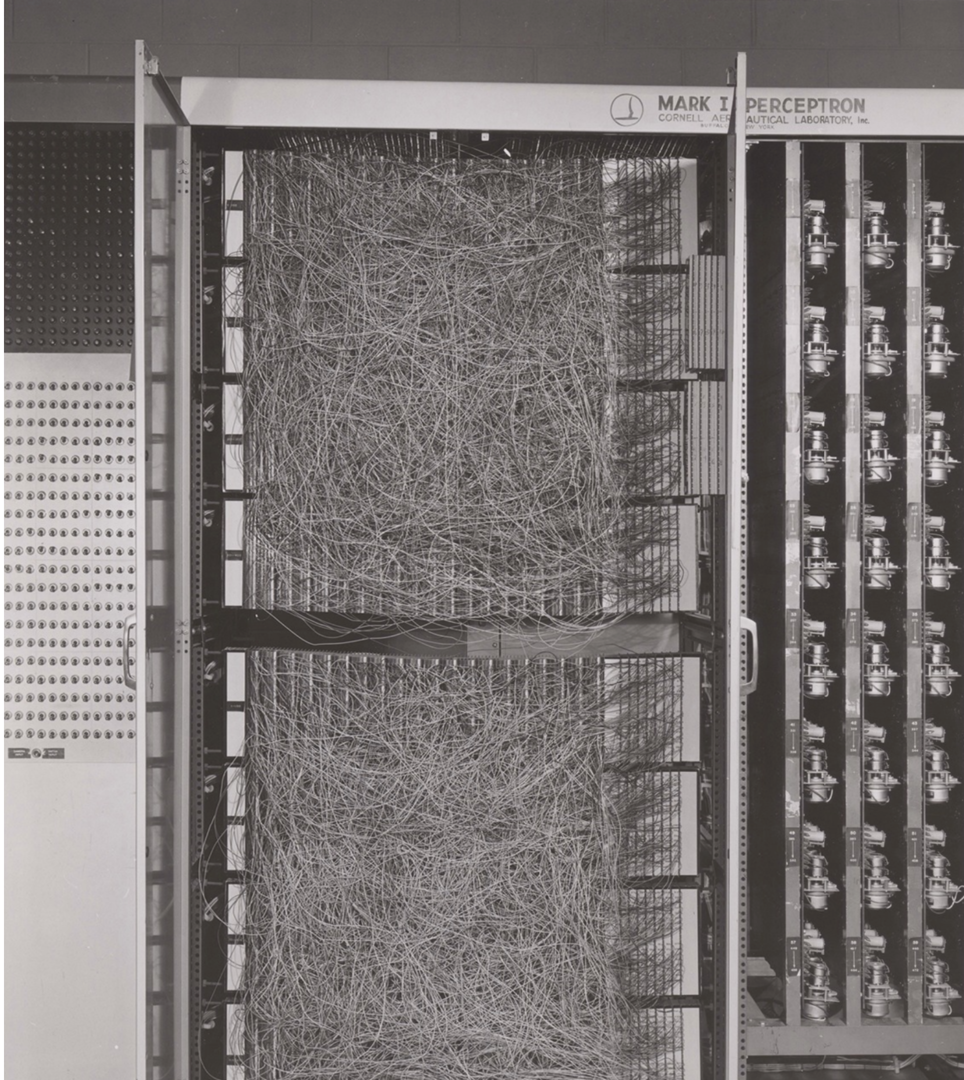

“Deep Learning is just another name for machine learning”

"Deep Learning is just another tool, like SVMs, random forests, and logistic regression"

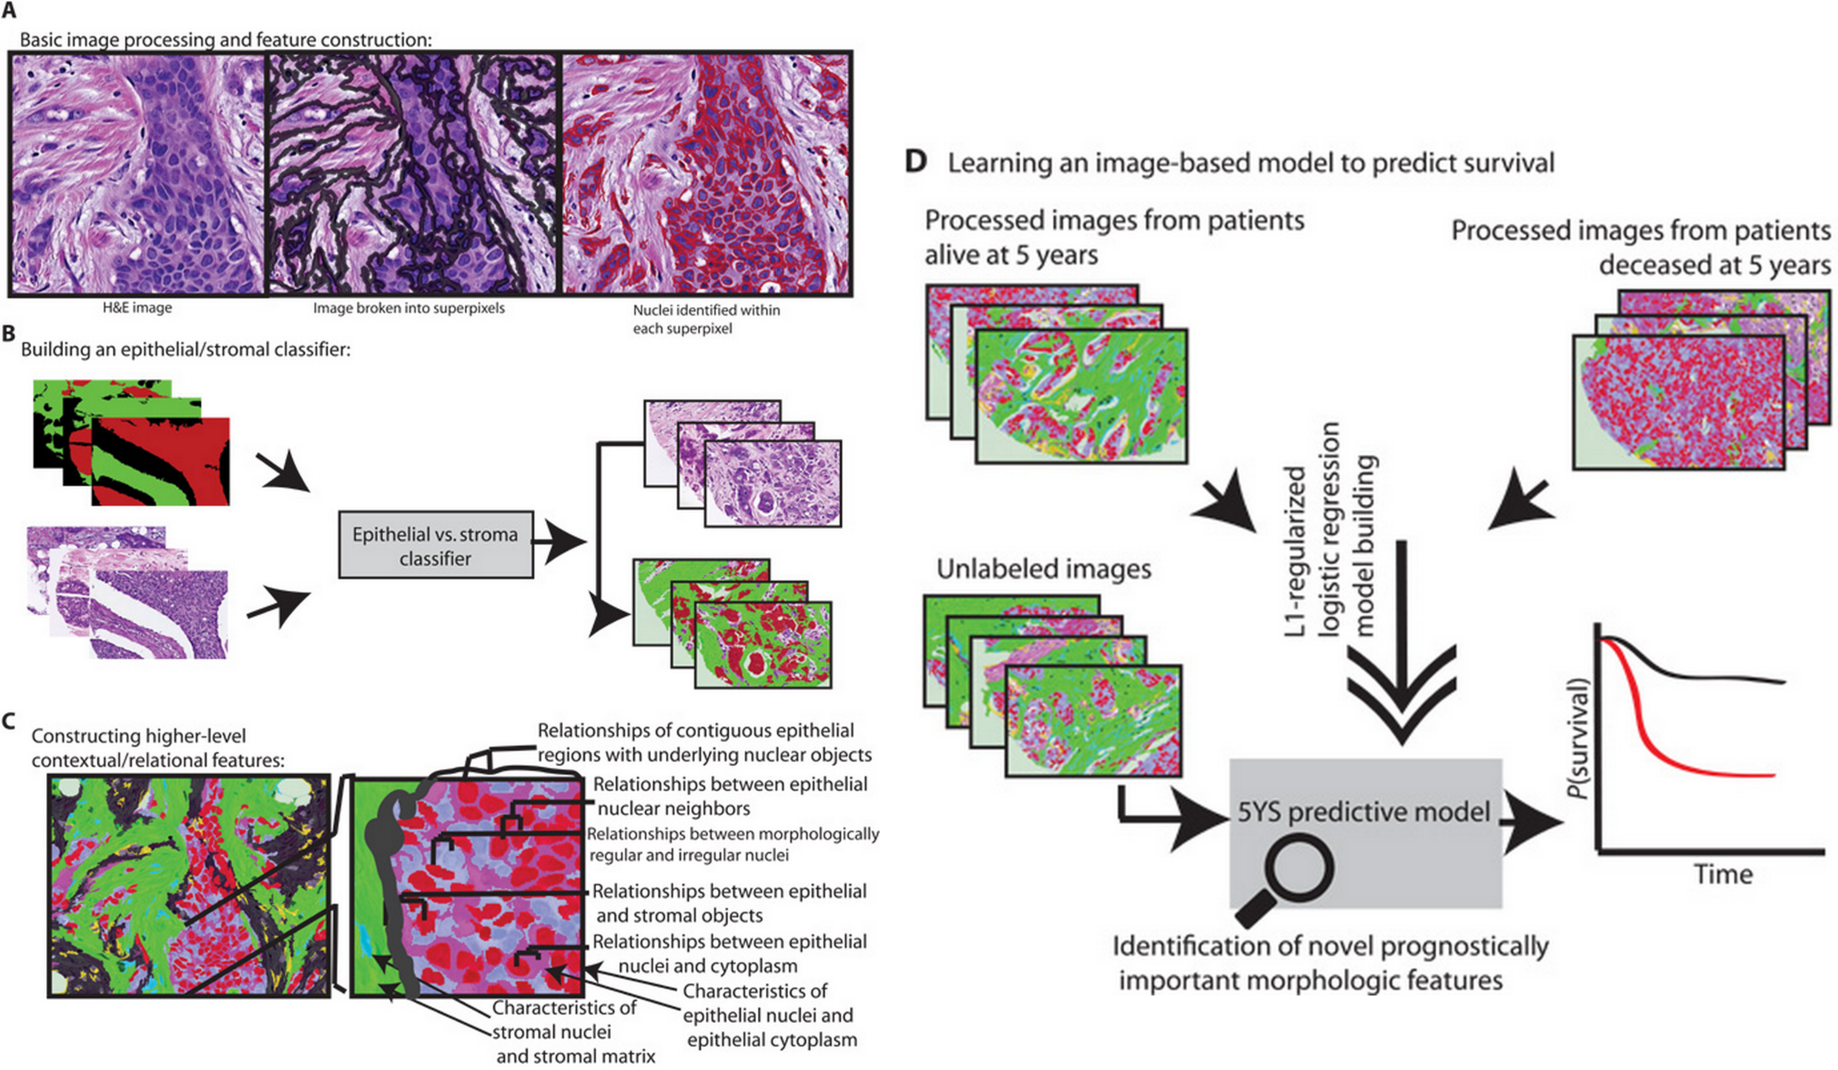
A
Basic image processing and feature construction:
H&E image
Image broken into superpixels
Nuclei identified within each superpixel
B
Building an epithelial/stromal classifier:
Epithelial vs. stroma classifier
C
Constructing higher-level contextual/relational features:
Relationships of contiguous epithelial regions with underlying nuclear objects
Relationships between epithelial nuclear neighbors
Relationships between morphologically regular and irregular nuclei
Relationships between epithelial and stromal objects
Relationships between epithelial nuclei and cytoplasm
Characteristics of epithelial nuclei and epithelial cytoplasm
Characteristics of stromal nuclei and stromal matrix
D
Learning an image-based model to predict survival
Processed images from patients alive at 5 years
Processed images from patients deceased at 5 years
L1-regularized logistic regression model building
Unlabeled images
5YS predictive model
Identification of novel prognostically important morphologic features
P(survival)
Time

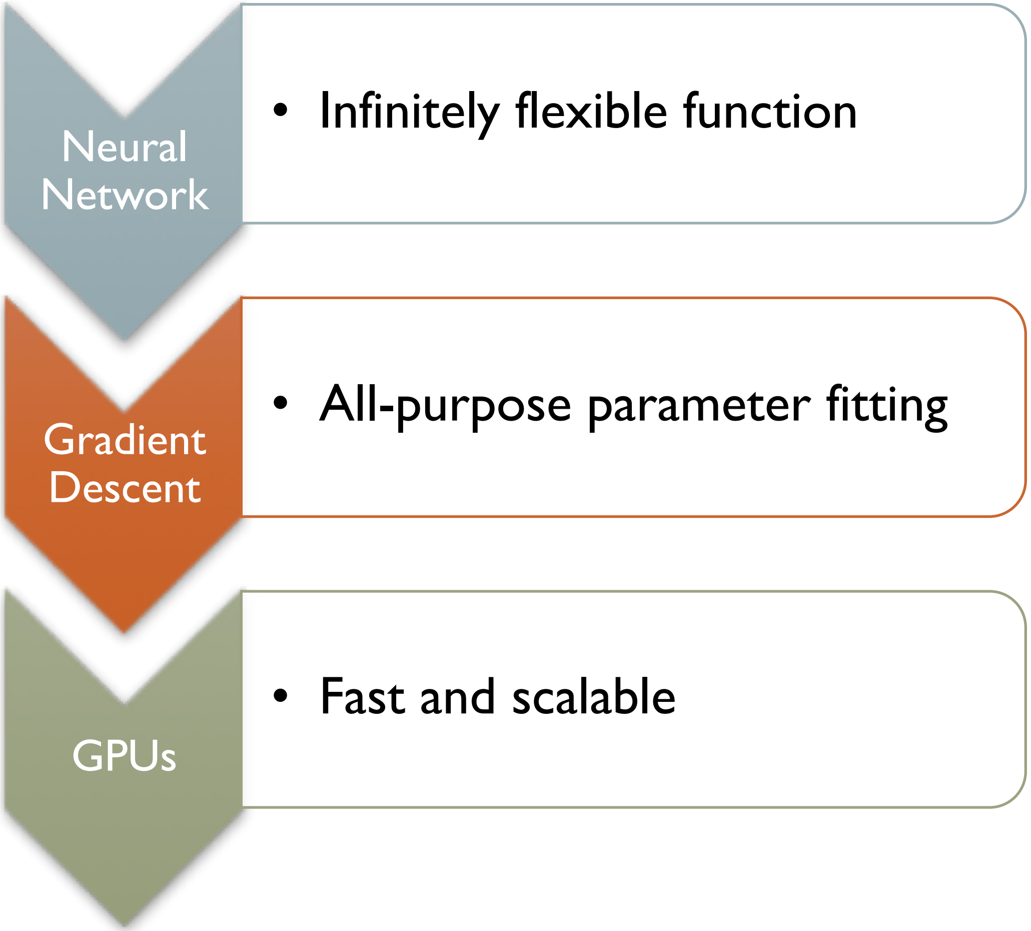

Neural Network

- Infinitely flexible function

Gradient Descent

- All-purpose parameter fitting

GPUs

- Fast and scalable

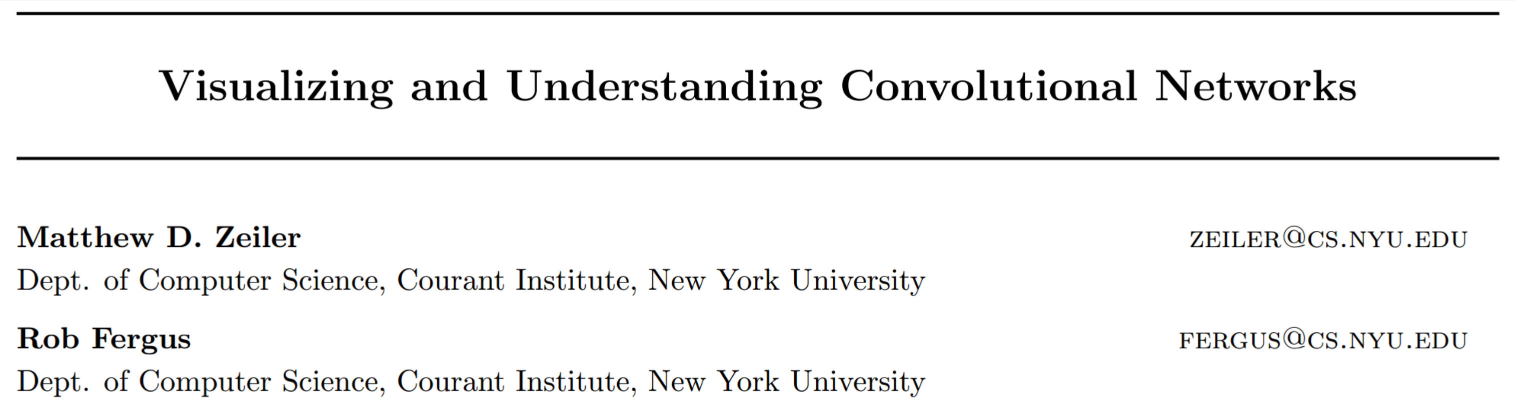

Visualizing and Understanding Convolutional Networks

Matthew D. Zeiler ZEILER@CS.NYU.EDU
Dept. of Computer Science, Courant Institute, New York University

Rob Fergus FERGUS@CS.NYU.EDU
Dept. of Computer Science, Courant Institute, New York University

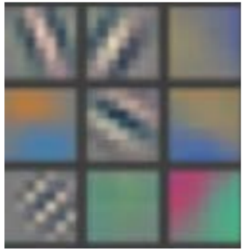

Layer 1

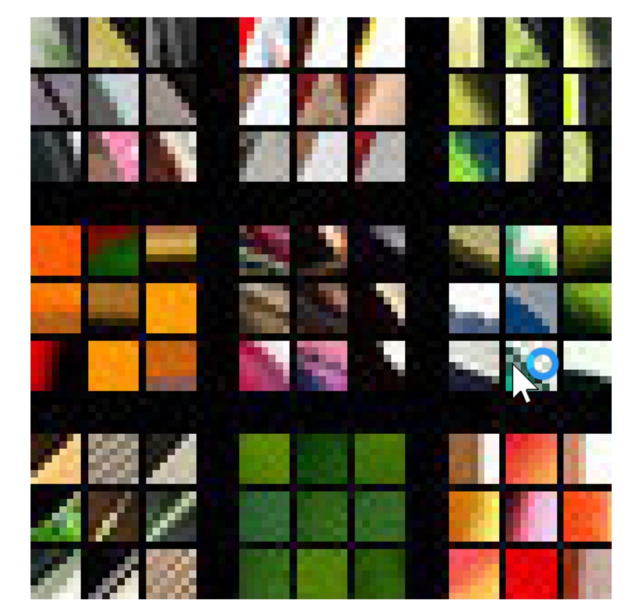

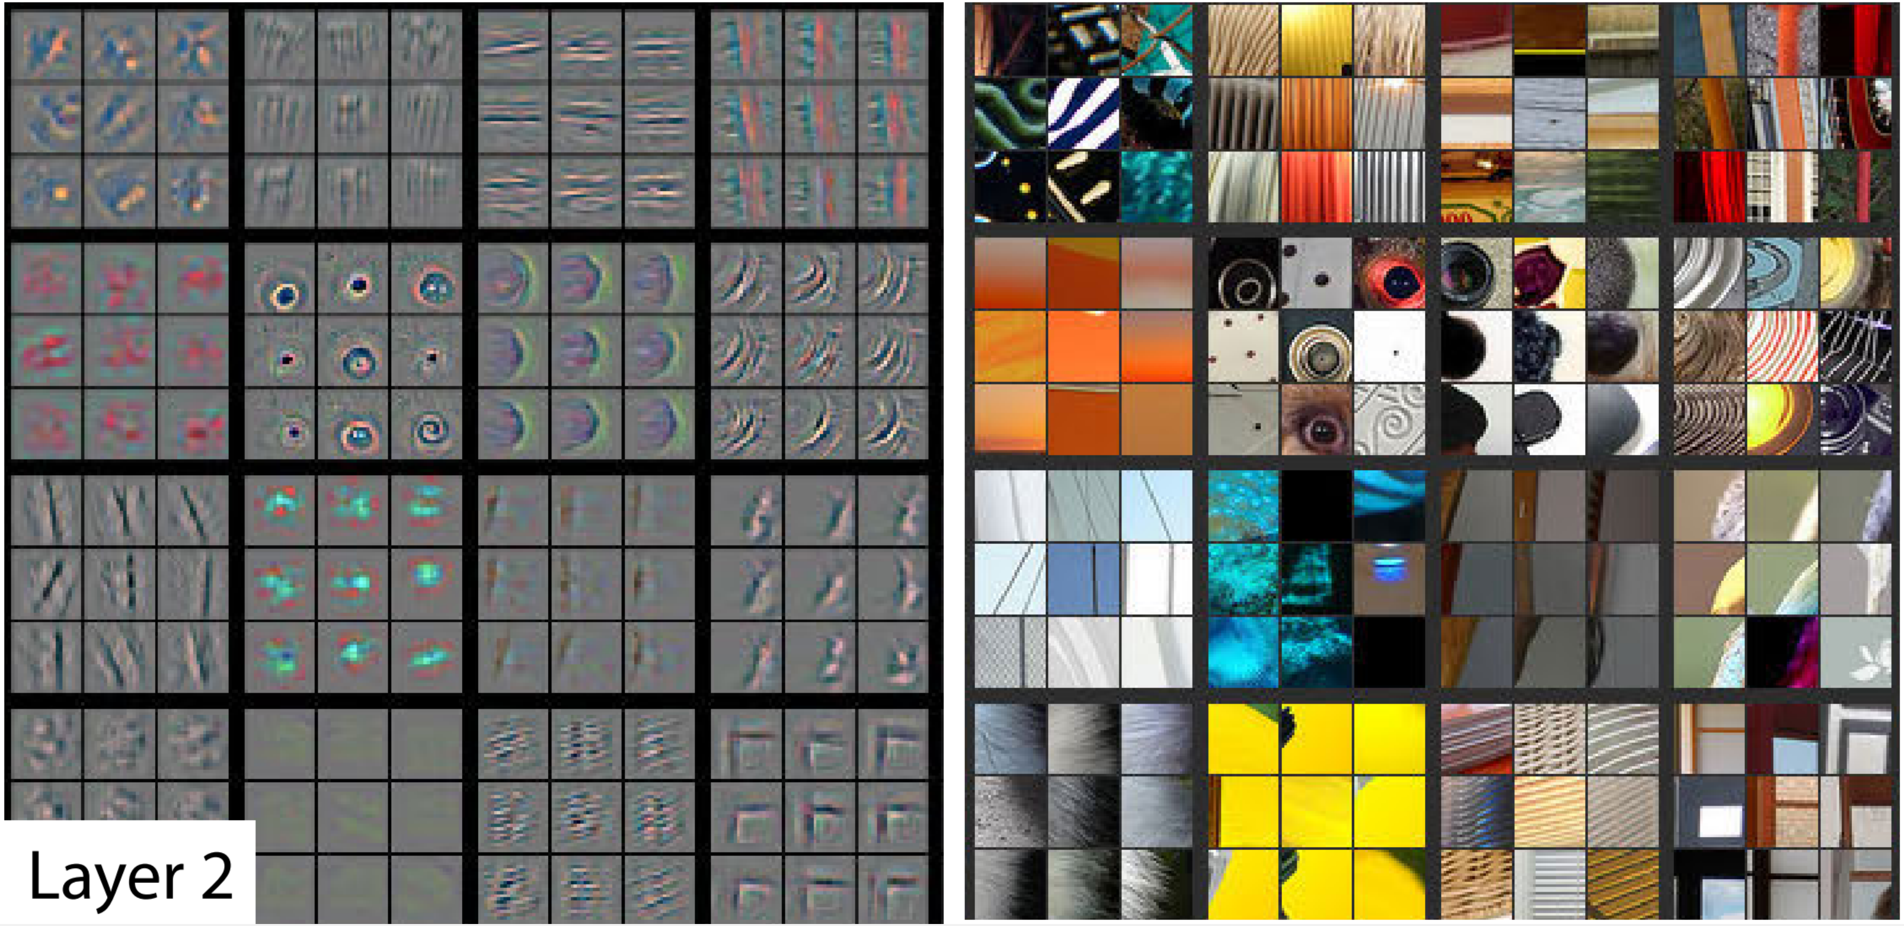

Layer 2

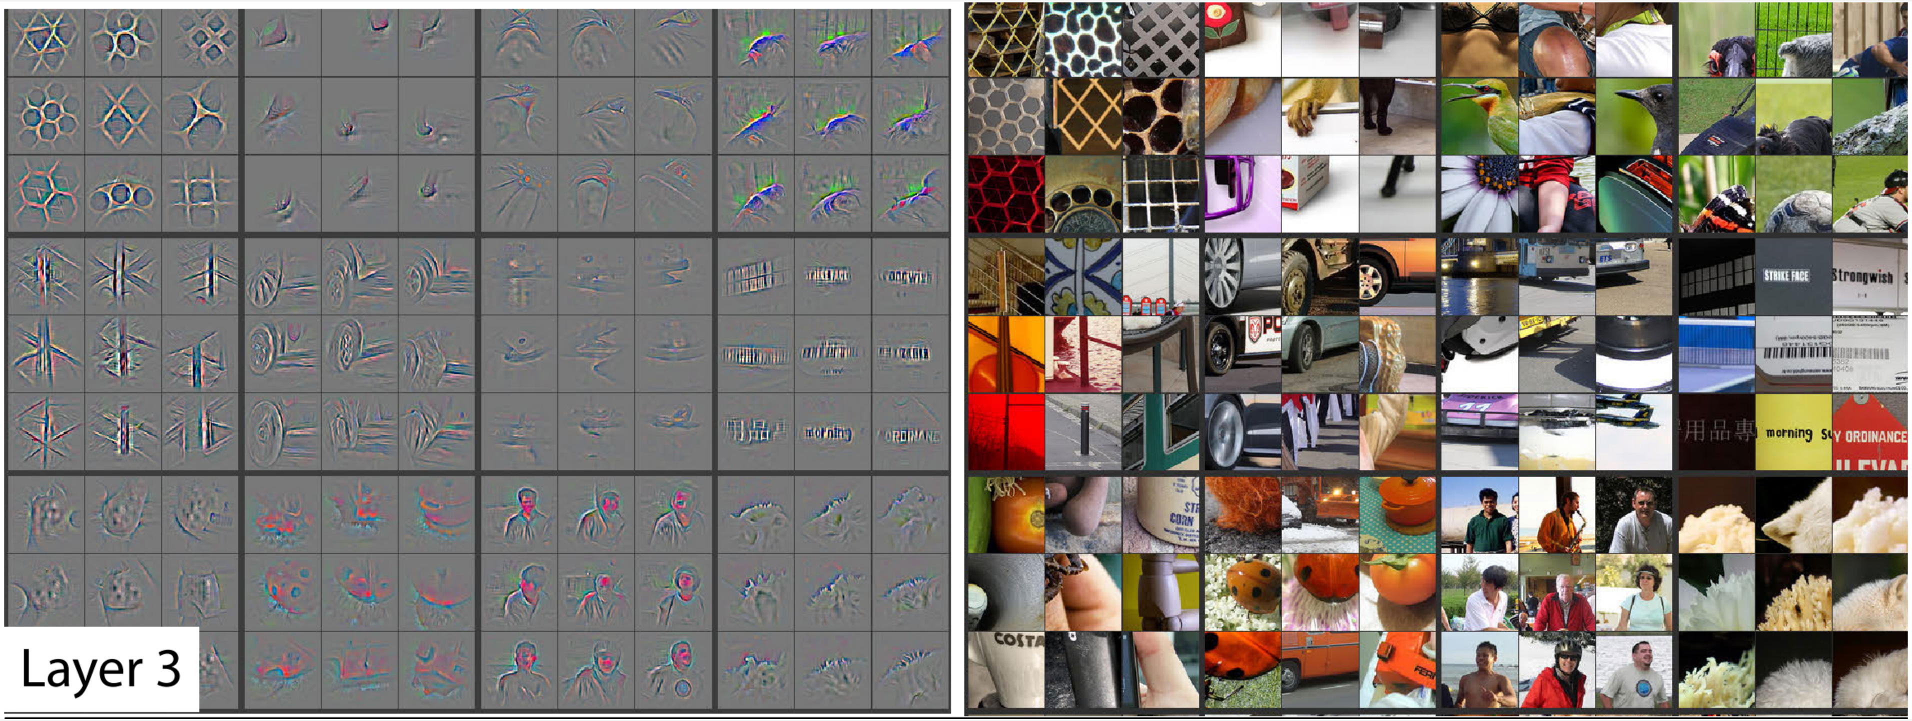

Layer 3

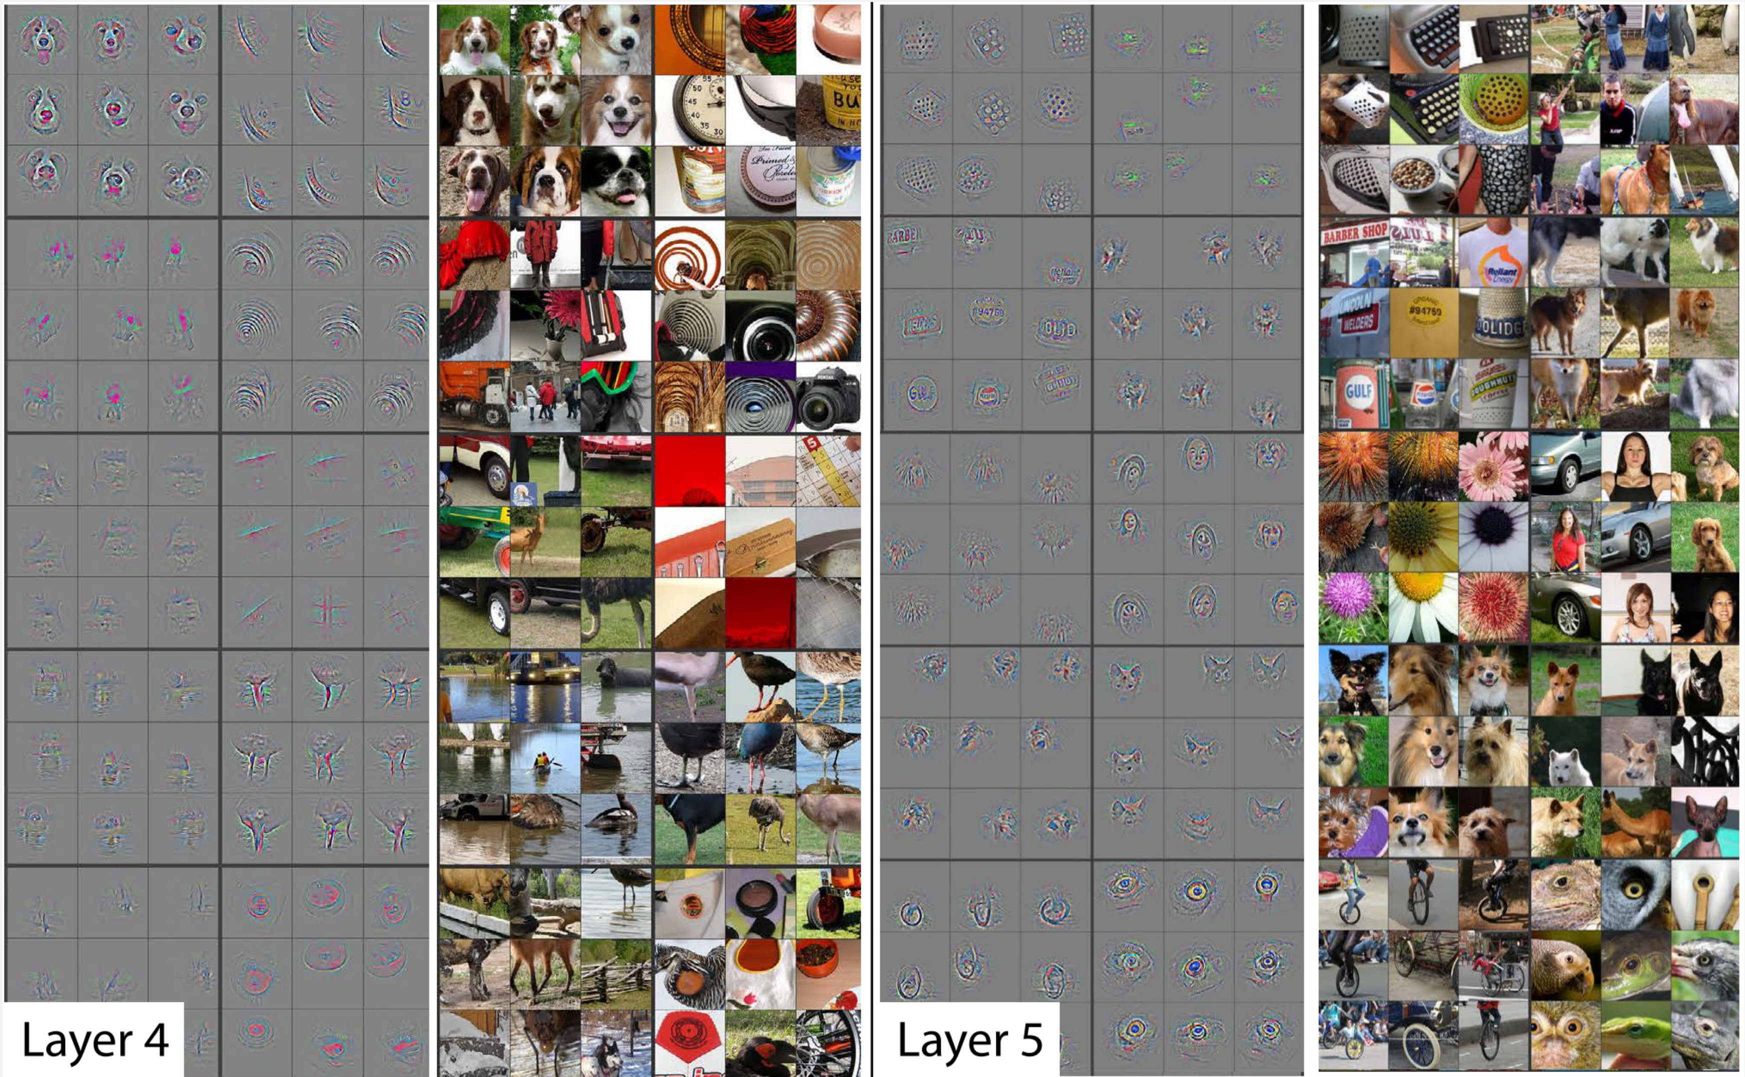
Layer 4
Layer 5

“Image recognition isn’t useful for anything much in practice”

	False Positive Rate	False Negative Rate
Panel of 4 Human Radiologists	66.3%	7.0%
Enlitic Algorithm	47.5%	0.0%

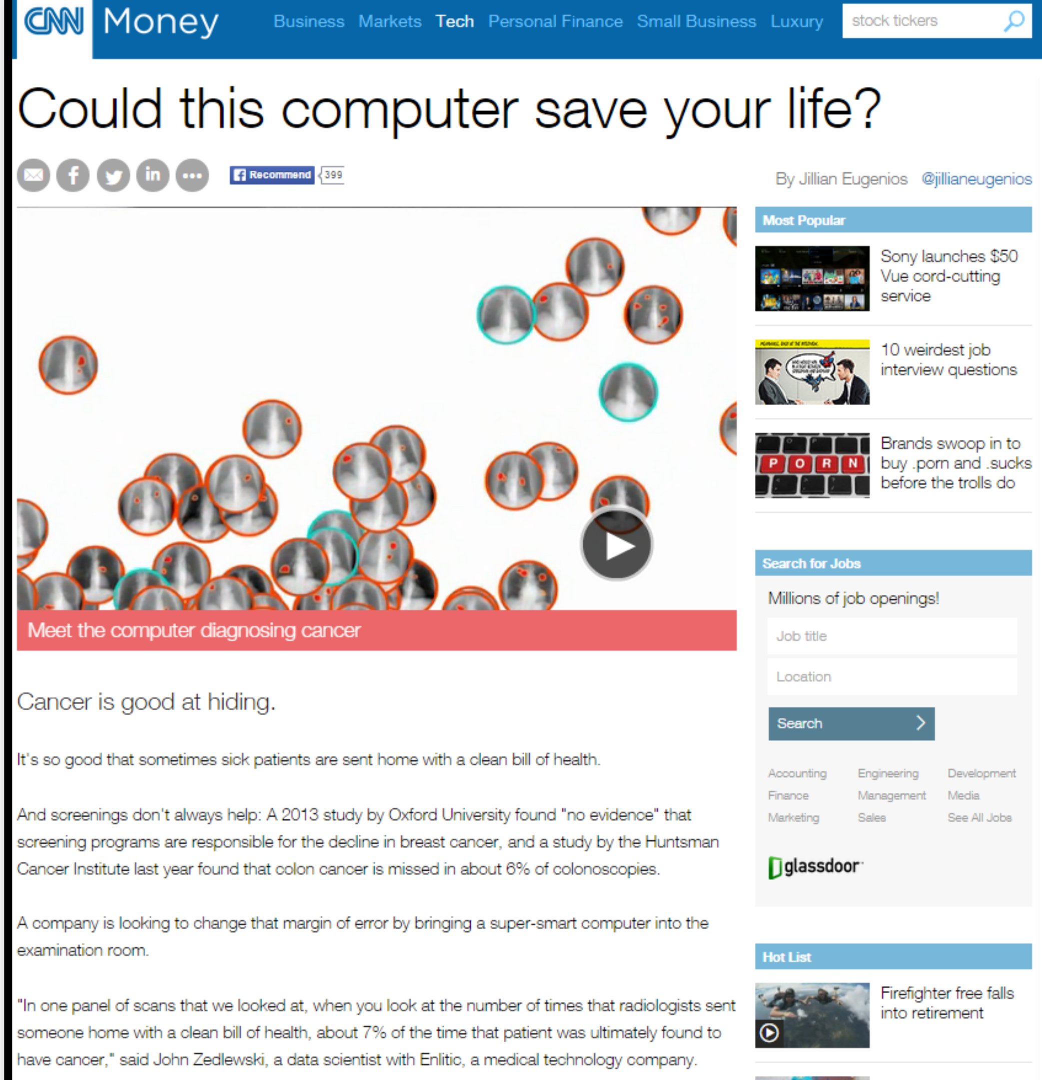

Could this computer save your life?

By Jillian Eugenios @jillianeugenios

Cancer is good at hiding.

It's so good that sometimes sick patients are sent home with a clean bill of health.

And screenings don't always help: A 2013 study by Oxford University found "no evidence" that screening programs are responsible for the decline in breast cancer, and a study by the Huntsman Cancer Institute last year found that colon cancer is missed in about 6% of colonoscopies.

A company is looking to change that margin of error by bringing a super-smart computer into the examination room.

"In one panel of scans that we looked at, when you look at the number of times that radiologists sent someone home with a clean bill of health, about 7% of the time that patient was ultimately found to have cancer," said John Zedlewski, a data scientist with Enlitic, a medical technology company.

The South Korean professional Go player Lee Sedol reviews the match after finishing against Google's artificial-intelligence program, AlphaGo.

Lee Jin-man / AP

How Google's AlphaGo Beat a Go World Champion

Inside a man-versus-machine showdown

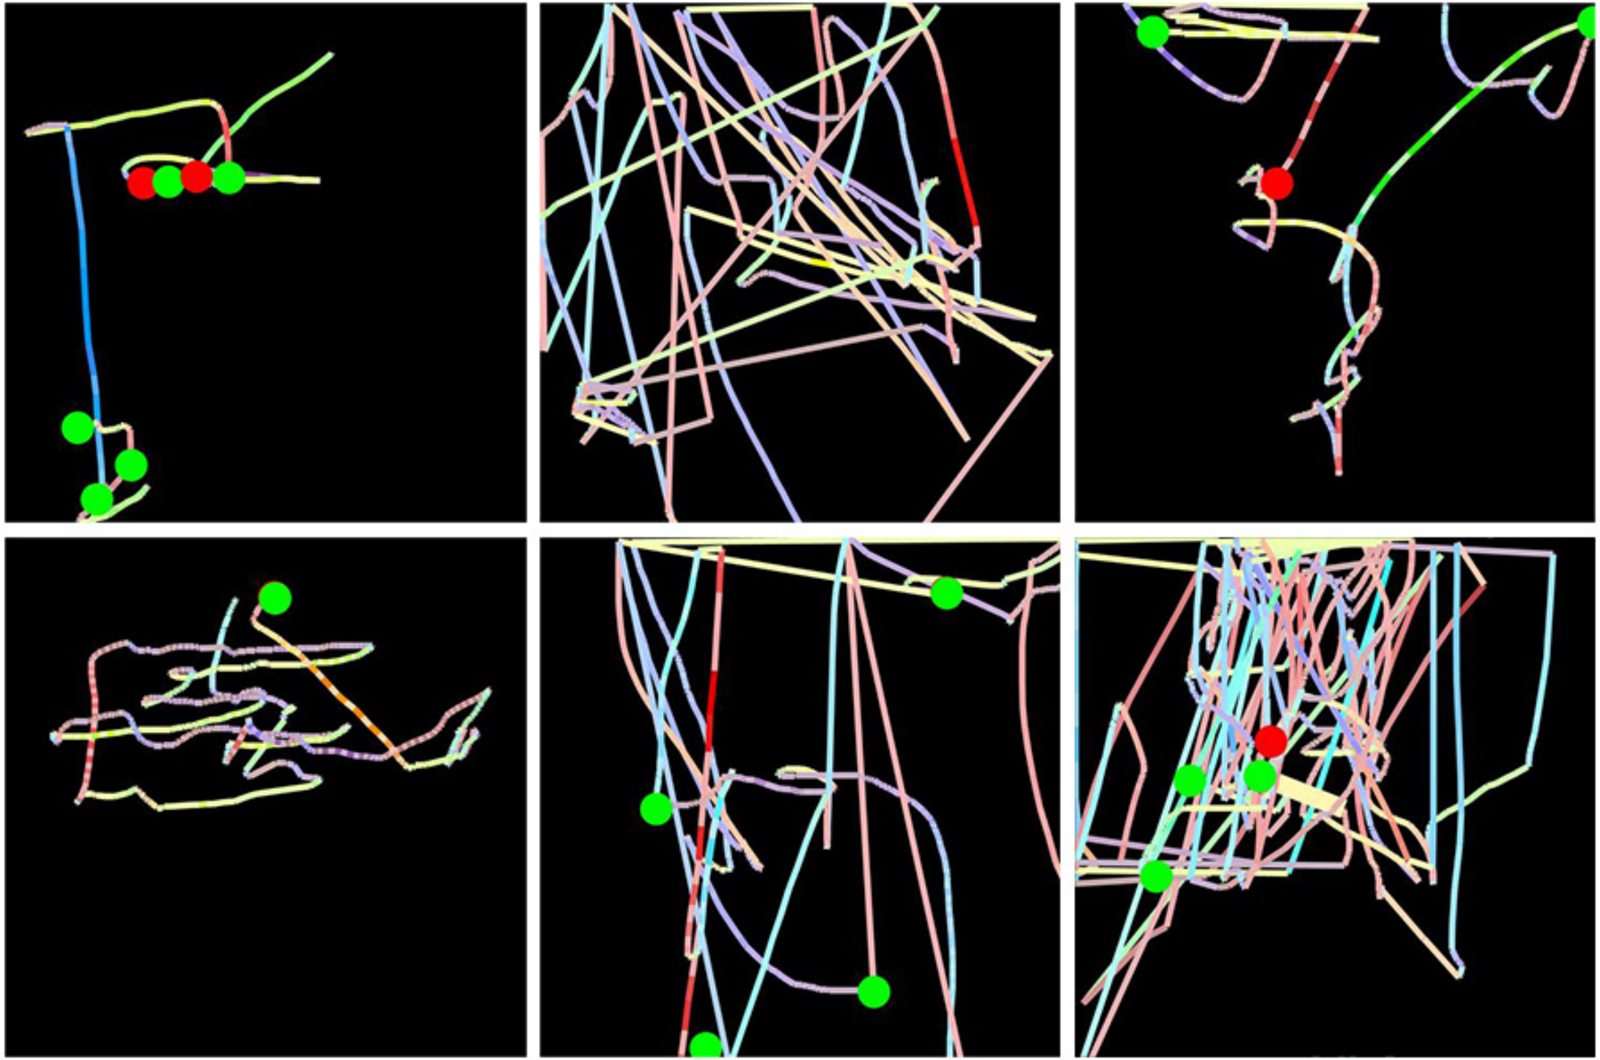

“Deep learning is only useful for image recognition”

Works now …using off-the-shelf libraries

- Computer vision

Works now …in latest research

- NLP
- Structured data
- Time series / signals

Classification & Regression

Not ready Limited success on research problems

- Reinforcement learning
- Adversarial models
- Anomaly detection

“Deep learning just isn’t working well in infosec”

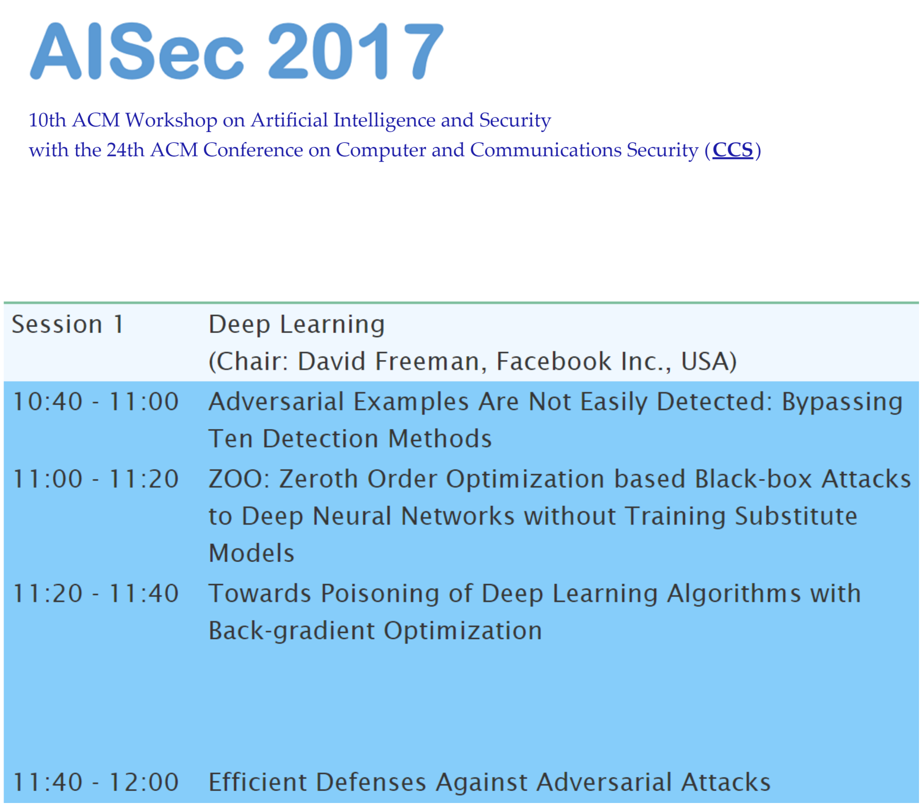

AISec 2017

10th ACM Workshop on Artificial Intelligence and Security
with the 24th ACM Conference on Computer and Communications Security (CCS)

Session 1	Deep Learning (Chair: David Freeman, Facebook Inc., USA)
10:40 - 11:00	Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods
11:00 - 11:20	ZOO: Zeroth Order Optimization based Black-box Attacks to Deep Neural Networks without Training Substitute Models
11:20 - 11:40	Towards Poisoning of Deep Learning Algorithms with Back-gradient Optimization
11:40 - 12:00	Efficient Defenses Against Adversarial Attacks

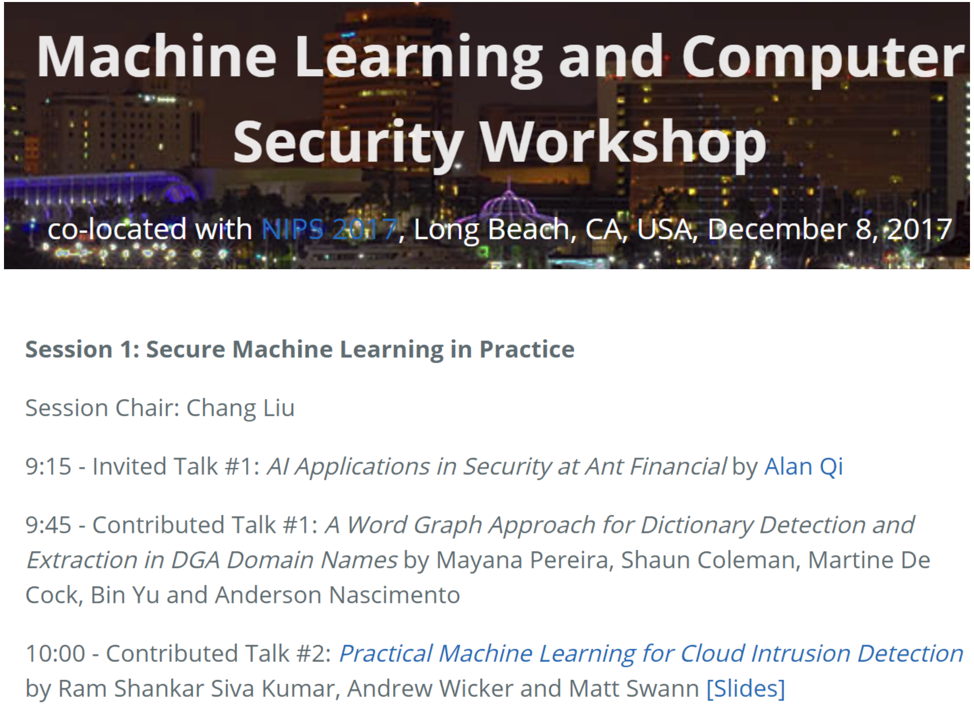

Machine Learning and Computer Security Workshop

co-located with NIPS 2017, Long Beach, CA, USA, December 8, 2017

Session 1: Secure Machine Learning in Practice

Session Chair: Chang Liu

9:15 - Invited Talk #1: *AI Applications in Security at Ant Financial* by Alan Qi

9:45 - Contributed Talk #1: *A Word Graph Approach for Dictionary Detection and Extraction in DGA Domain Names* by Mayana Pereira, Shaun Coleman, Martine De Cock, Bin Yu and Anderson Nascimento

10:00 - Contributed Talk #2: *Practical Machine Learning for Cloud Intrusion Detection* by Ram Shankar Siva Kumar, Andrew Wicker and Matt Swann [Slides]

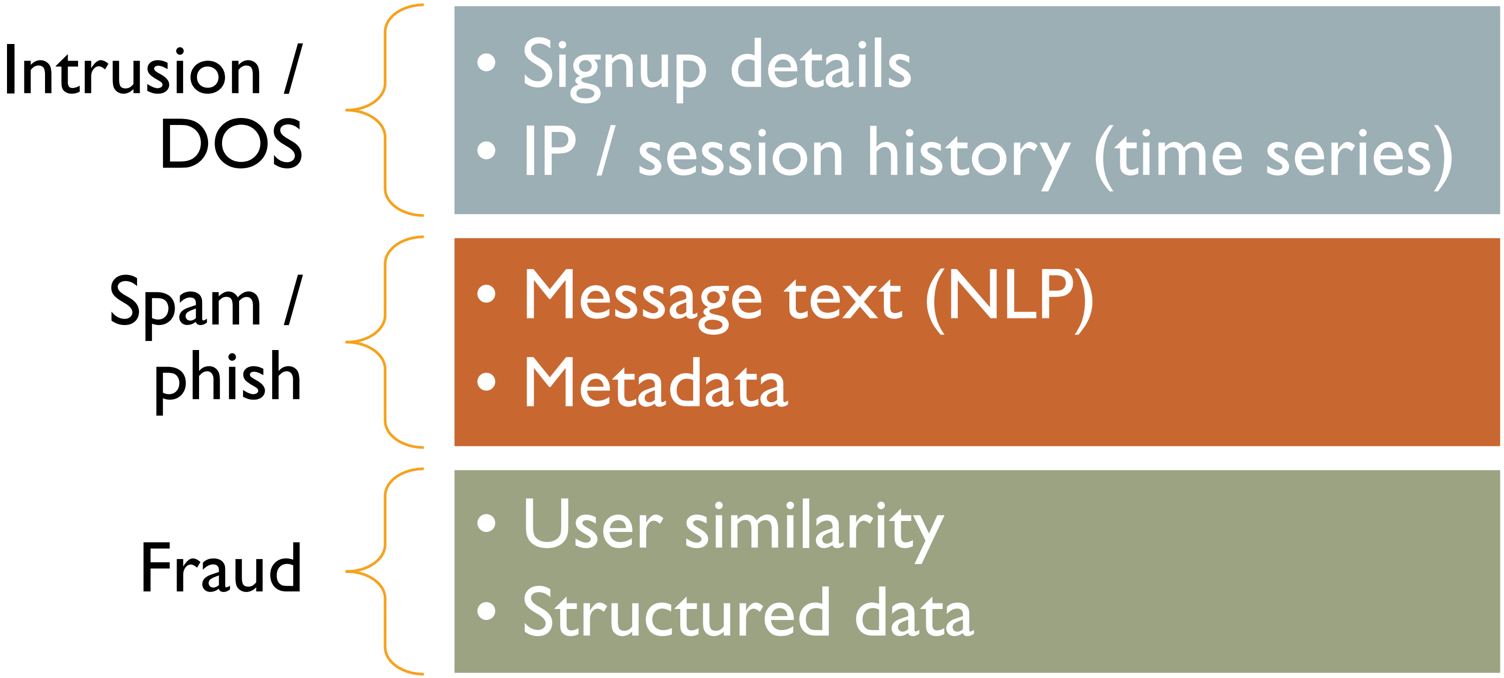

Intrusion / DOS

- Signup details
- IP / session history (time series)

Spam / phish

- Message text (NLP)
- Metadata

Fraud

- User similarity
- Structured data

“I can’t use deep learning because…”

Black box	• Interpretable ML • Visualize gradients and activations
Needs too much data	• Transfer learning • Share pre-trained nets
Needs ML PhD	• No longer true • fastai & keras libs, MOOCs, etc
Only for vision	• No longer true • SoTA for speech, structured data, time series…
Needs lots of GPUs	• Was never true • …except for some research projects
“Not really AI”	• Who cares? • Do you really want to build a brain?

“I don’t know how to get started…”

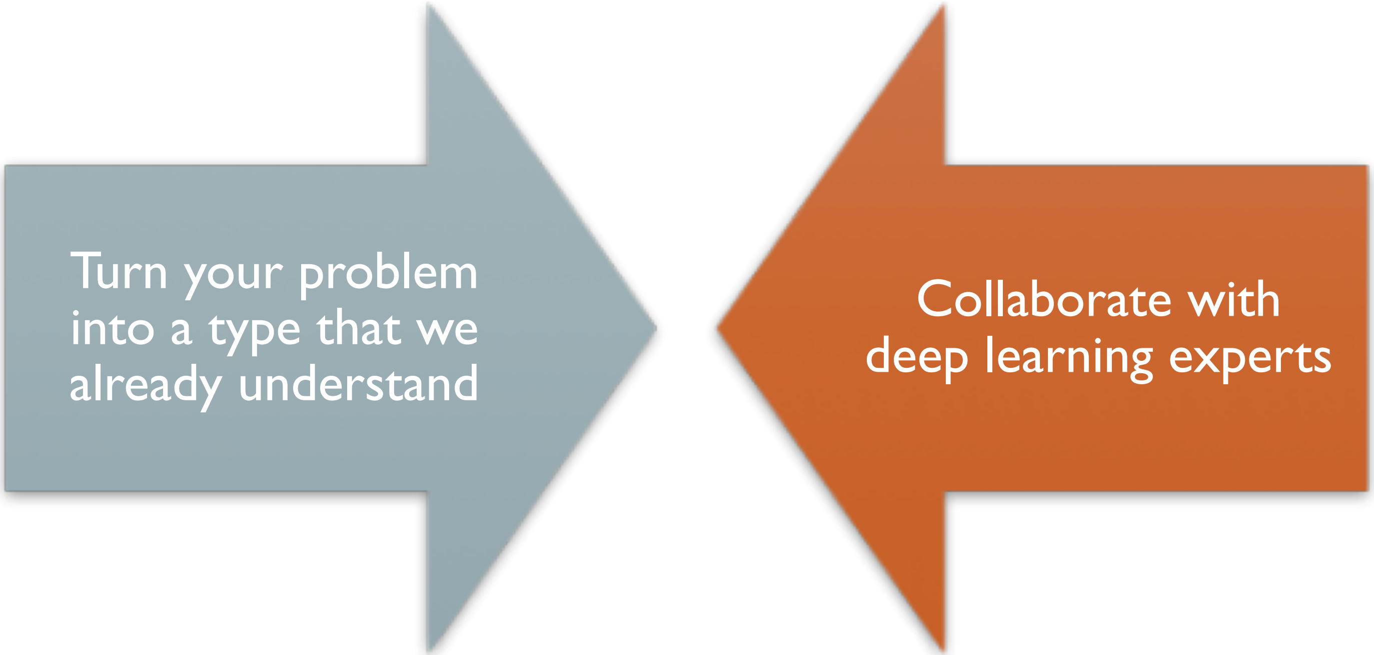

Turn your problem into a type that we already understand

Collaborate with deep learning experts

‹ ▦ ›

Overview:

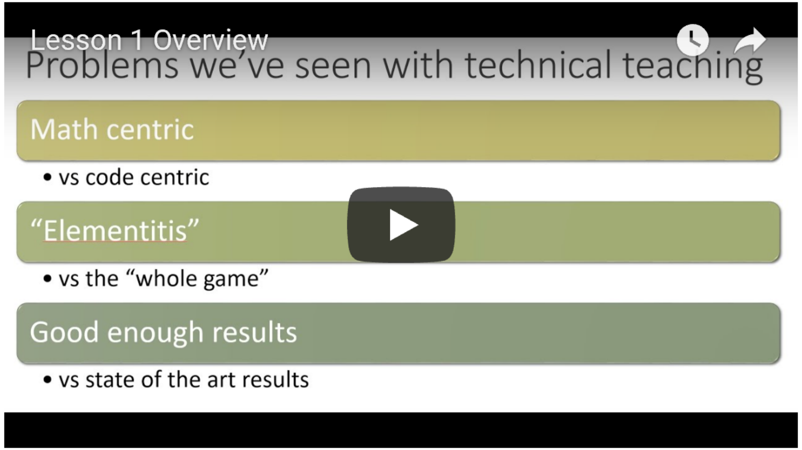

Lesson 1 Overview

Problems we've seen with technical teaching

Math centric
- vs code centric

"Elementitis"
- vs the "whole game"

Good enough results
- vs state of the art results

Lesson: Timeline / Wiki / Notes / Forum / Youtube

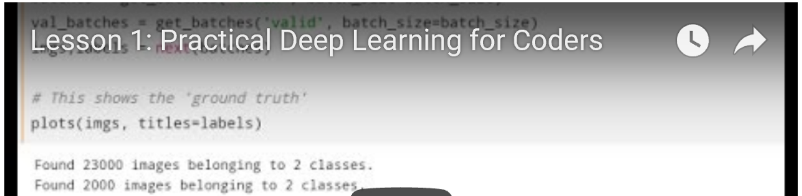

Lesson 1: Practical Deep Learning for Coders

```
val_batches = get_batches('valid', batch_size=batch_size)

# This shows the 'ground truth'
plots(imgs, titles=labels)

Found 23000 images belonging to 2 classes.
Found 2000 images belonging to 2 classes.
```

1—RECOGNIZING CATS AND DOGS

Important note: All files in the course are now available from files.fast.ai, rather than platform.ai, as shown in the videos. We have attempted to update all mentions of platform.ai to files.fast.ai on the wiki, forums, etc, but youtube does not allow us to change the videos themselves.

Welcome to the first full lesson of Practical Deep Learning For Coders! Before you start this lesson, be sure to have completed setup of your deep learning server. See the AWS Lesson to learn how to do this, if you haven't already.

Each lesson page includes links to course notes, forum discussion, and (most importantly) a wiki page. Nearly all the participants in the original in-person course said that they found these resources very important for successfully completing the course. So be sure to make the most of them! And be sure to carefully read the Getting Started page to find out how this course is designed and how to get the most out of it. (Also, apologies that the questions from the audience are hard to hear - we get a special audience mic from lesson 3 onwards which resolves that problem.)

SYNOPSIS

The 30 minute overview video introduces you to the course and explains how to get the most out of each lesson. We also pass on some tips from previous