SAND2019-2292PE

Machine Learning Applications at Sandia National Laboratories

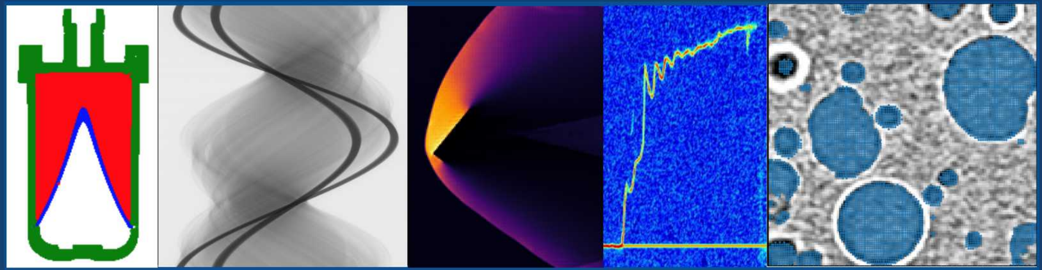

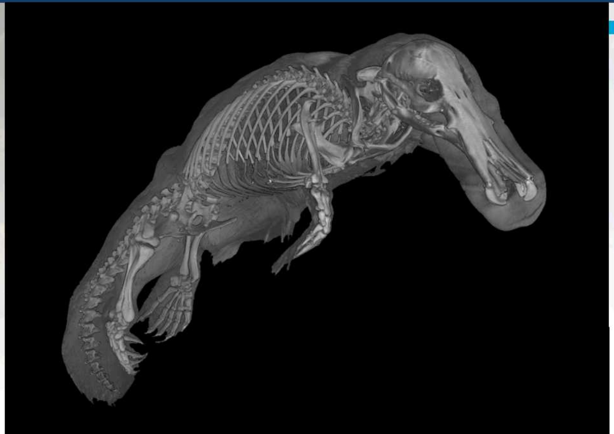

Cari Martinez

Sandia Interdisciplinary Machine Learning Research (SIMLR) Team

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Multi-mission Laboratory

INTRODUCTION TO SANDIA NATIONAL LABS

About Sandia National Laboratories

Federally Funded Research and Development Center

- Nuclear Weapons
- Defense Systems & Assessments
- Energy
- Global Security

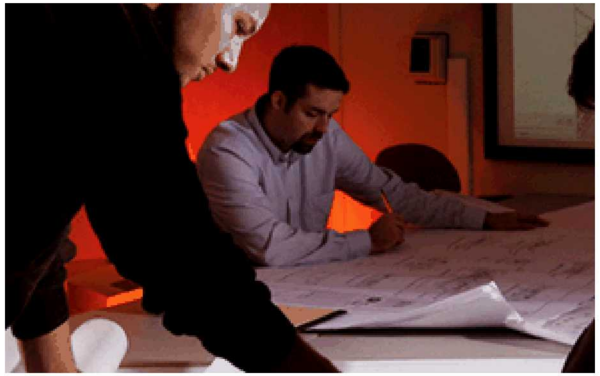

Two types of Machine Learning research at Sandia

SIMLR
BUILD TRAIN INFER

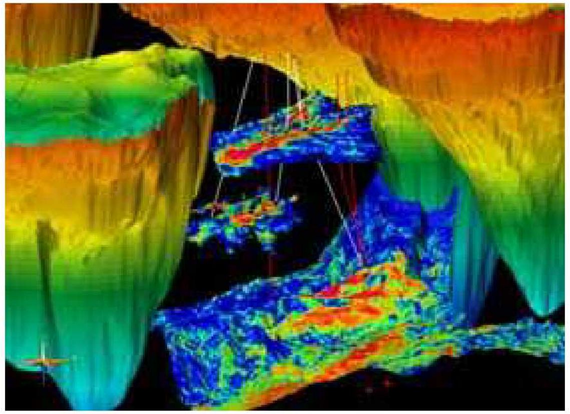

Applied ML Research

- Collaborate with other engineers
  - Mechanical Engineering
  - Chemical Engineering
- Understand their problem space
- Apply latest machine learning research from academia and industry

"Pure" ML Research:

- Develop novel ML algorithms
- Understand theoretical foundations
- Improve U.S. capability in ML

The progress of science and technology…

- **1st Paradigm** – Empirical science [1000's of years ago]
  - Experimentation and observation
- **2nd Paradigm** – Theoretical science [100's of years ago]
  - Laws and equations
- **3rd Paradigm** – Computational science [10's of years ago]
  - Numerical implementation of laws and equations
- **4th Paradigm** – Big data driven science [today]
  - Real-time assessments based on machine learning of paradigms 1-3

https://en.wikipedia.org/wiki/Antikythera_mechanism

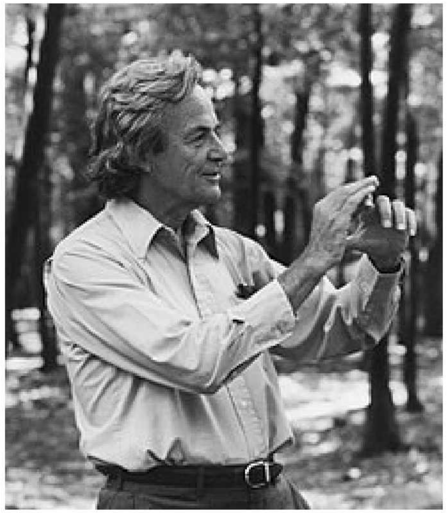

https://en.wikipedia.org/wiki/Richard_Feynman

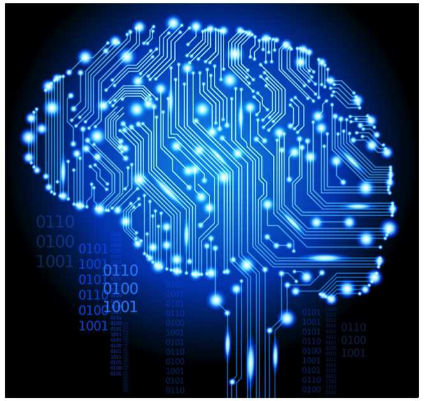

https://towardsdatascience.com/deep-learning-tips-and-tricks-1ef708ec5f53

http://www.dtic.mil/dtic/tr/fulltext/u2/1010508.pdf
https://www.slideshare.net/Delft_Software_Days/dsd-int-2014-data-science-symposium-4th-paradigm-a-research-perspective-prof-arnold-bregt-wageningen-ur-amsterdam-institute-for-advanced-metropolitan-solutions

Deep Learning is an evolution of Artificial Neural Networks

- 1940s-60s – Cybernetics
  - 1943 – Math model of a neuron
  - 1958 – Perceptron learning
- 1980s-90s – Connectionism
  - 1986 – Backpropagation learning
  - 1989 – Convolutional Neural Networks
- Limitations
  - Networks are too hard to train
  - Not enough computing power
  - Not enough memory for large datasets
  - Exploding/vanishing gradients
- 2006 – Deep learning

https://www.deeplearningbook.org/

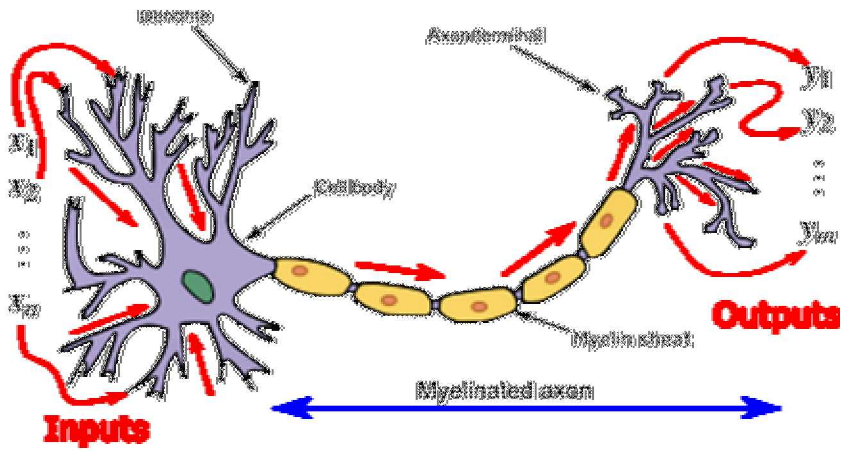

Attribution: Egm4313.s12 at English Wikipedia under license https://creativecommons.org/licenses/by-sa/3.0/deed.en

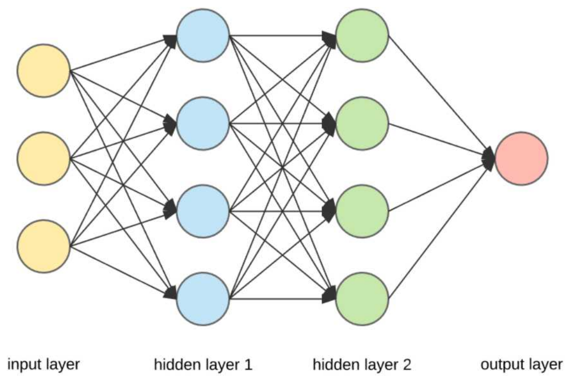

https://medium.com/@ksusorokina/image-classification-with-convolutional-neural-networks-496815db12a8

Why Machine Learning?

- Some problems are difficult to solve with a directly-coded algorithm
  - Do not generalize well
  - Write rules-based programs for each specific task
- Machine learning (ML) is good at:
  - Recognizing patterns
  - Anomaly detection
  - Prediction

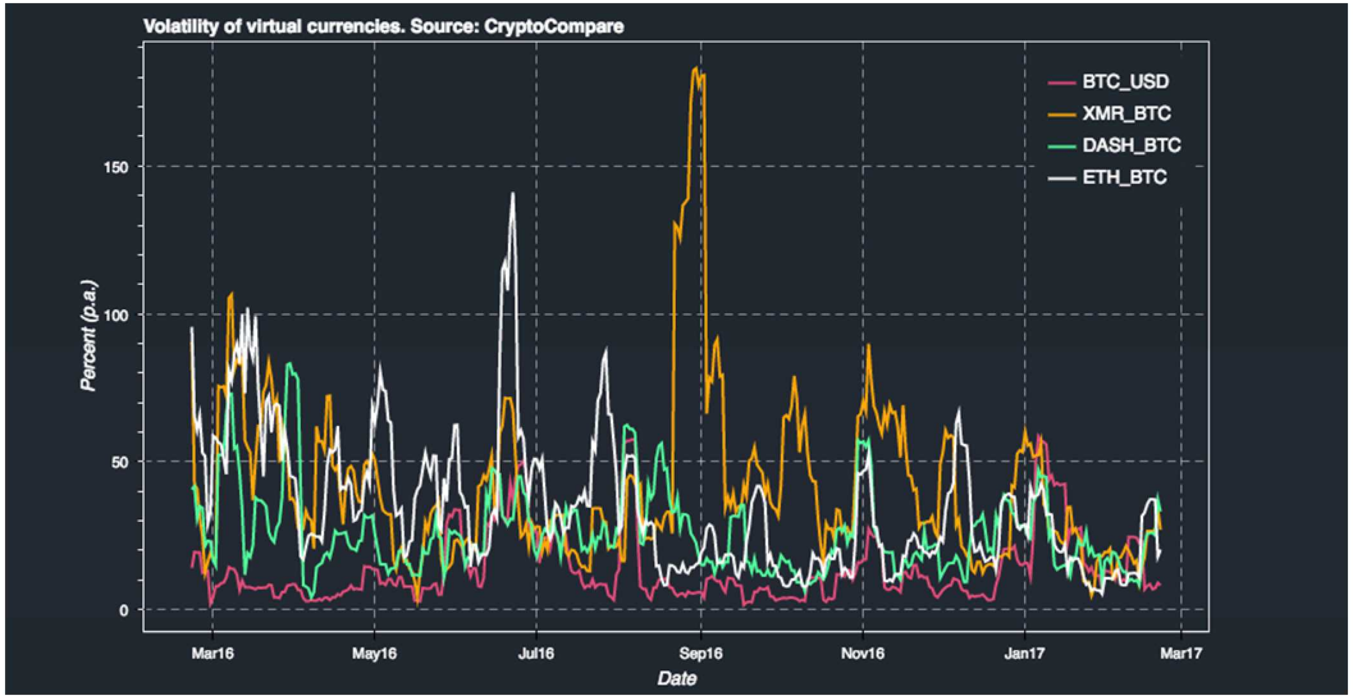

https://www.cryptocompare.com/coins/guides/5-things-you-need-to-know-about-bitcoin-volatility/

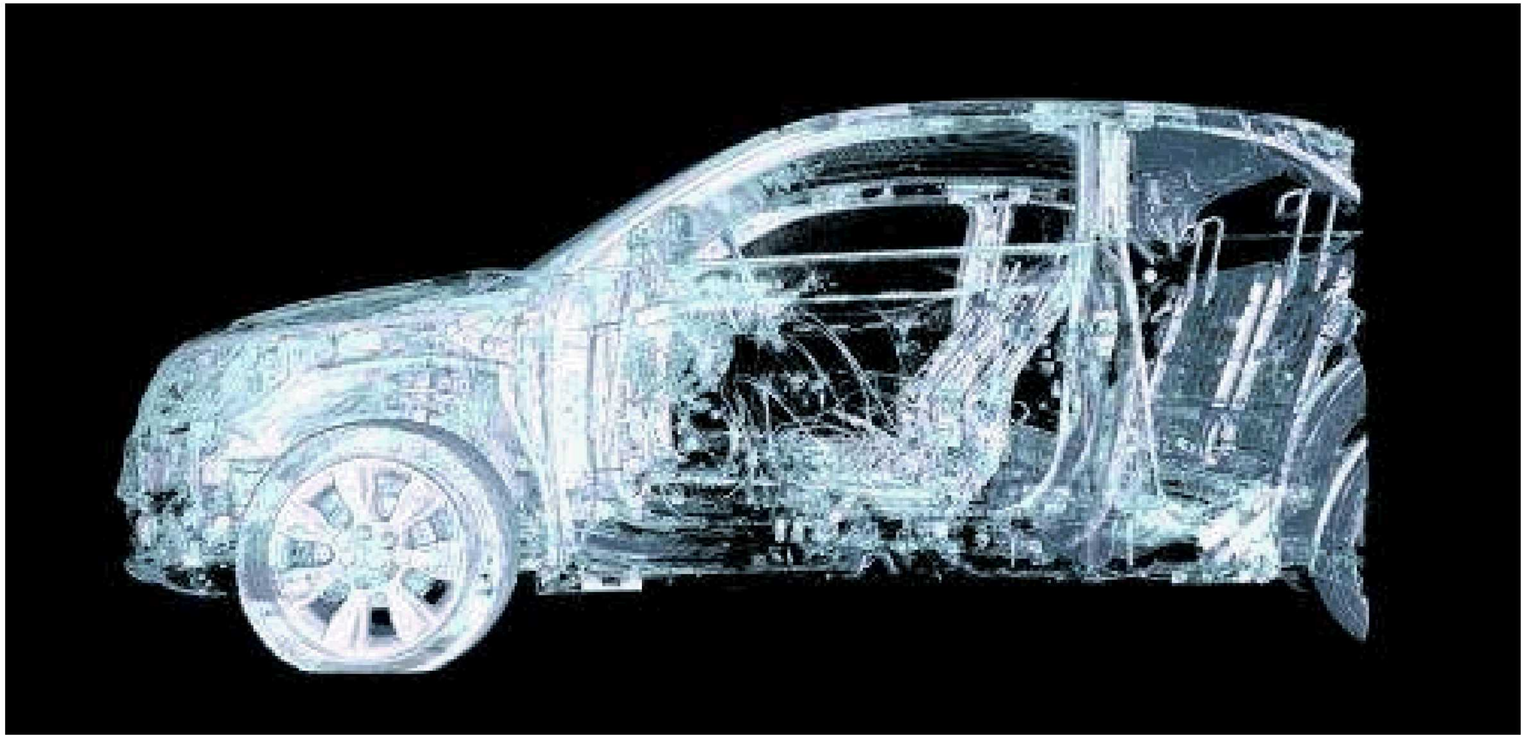

https://phys.org/news/2013-10-scale.html

Modeling an imperfect world

DIGITAL TWINS

Introducing Digital Twins

- Today, we lot test and make a statistical argument.
- With a Computed Tomography (CT) machine, we can scan:
  - After manufacture of a part
  - After each step of assembly
  - After every transport (field or vehicle)
  - In different environmental conditions
  - During motion to determine robustness (CT movie)
- Improves many processes
  - Quality assurance of existing parts
  - Clever use of known faulty parts
  - Better engineering with additional data

https://www.leithcars.com/2017-ford-mustang-raleigh.html

Simulations using digital twins are predictive

- Each component is manually labeled
- High fidelity simulations predict performance of specific units

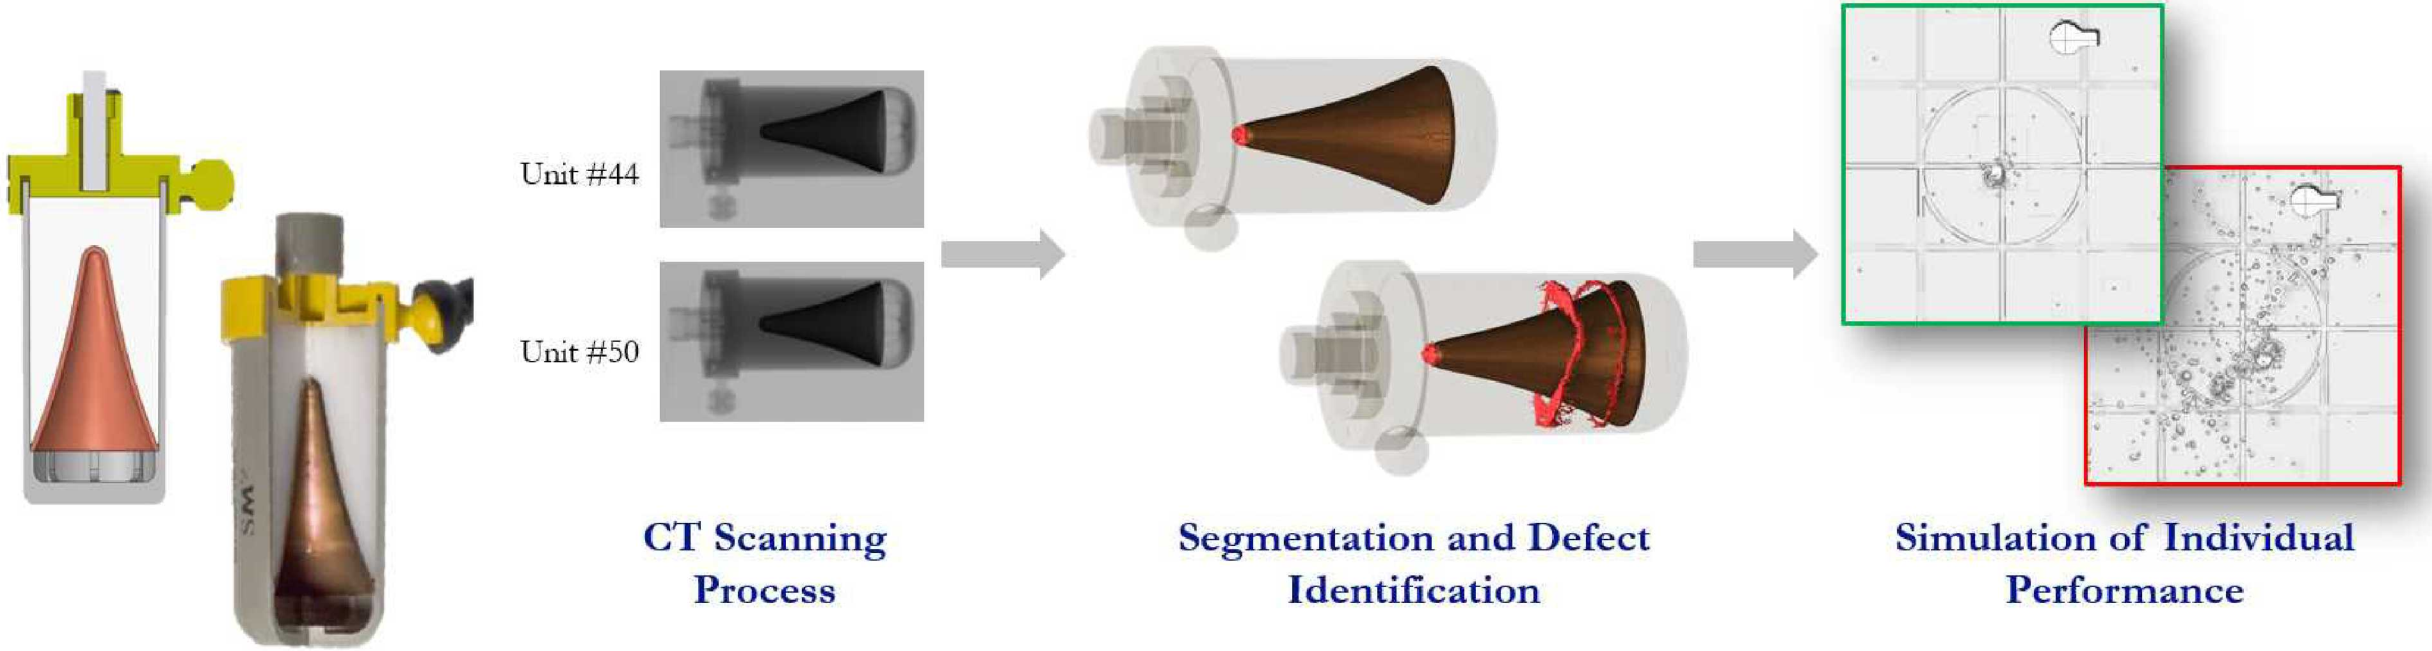

CT Segmentation is hard for humans

- CT scans must be labeled by component for simulations

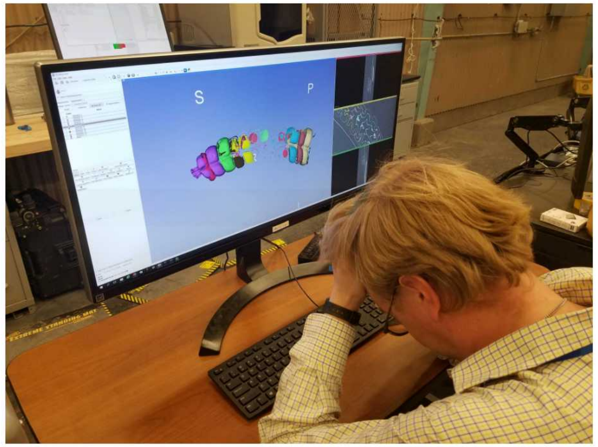

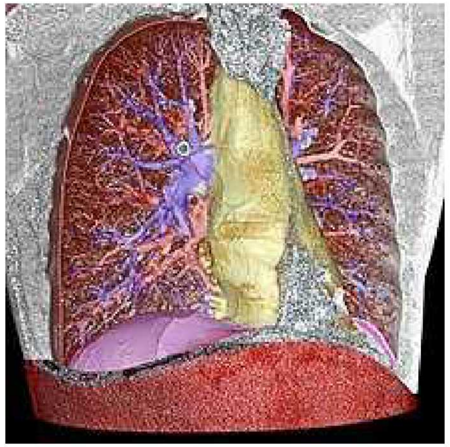

https://en.wikipedia.org/wiki/Image_segmentation

- Labeling by hand does not scale
- Deep learning algorithms
  - Find each component of the shared charge by material
  - Find any defects
  - Prepare for numerical simulations

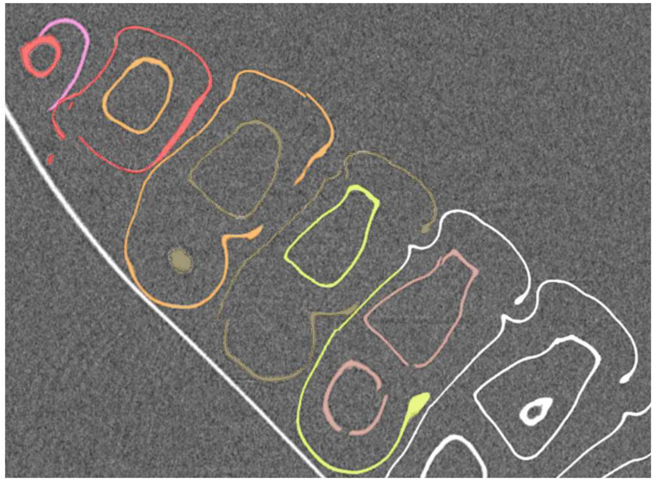

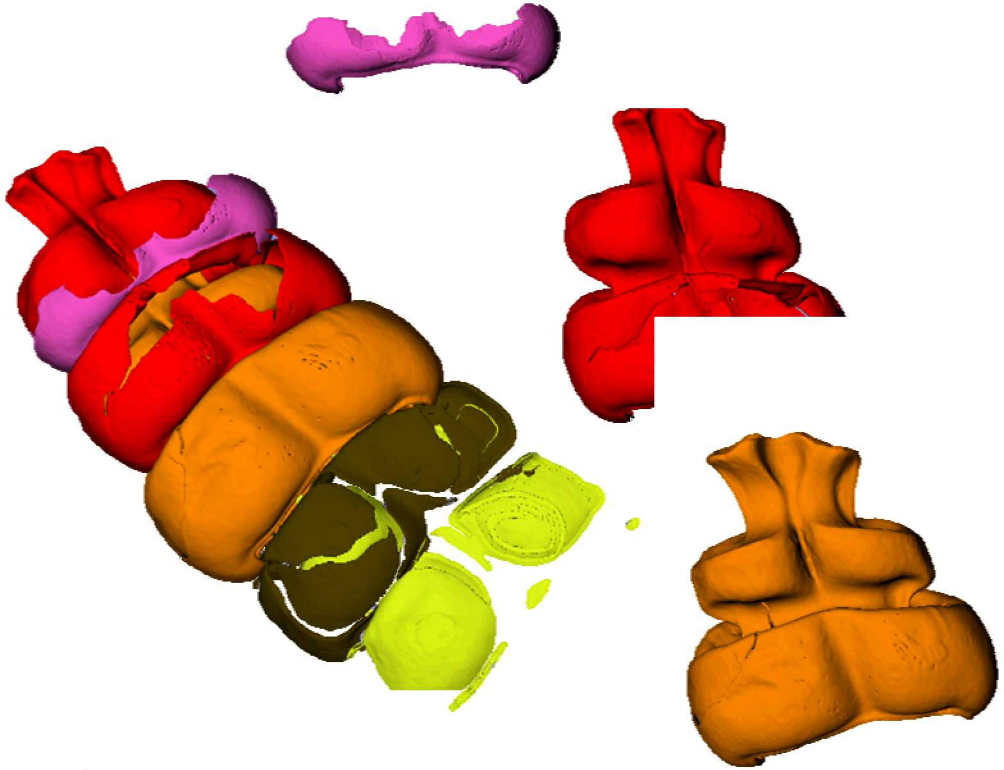

Using ML to save time and effort while improving accuracy

CT SEGMENTATION

Segmentation is a classic computer vision problem

- Image segmentation is well studied
  - Small files
  - Large training sets
- Volumetric segmentation is different
  - Big data
  - Class imbalance (lots of background)
  - Small training sets with “bad” human labels
  - Humans can’t label billions of voxels
- Medical researchers are leading this work toward deep learning solutions

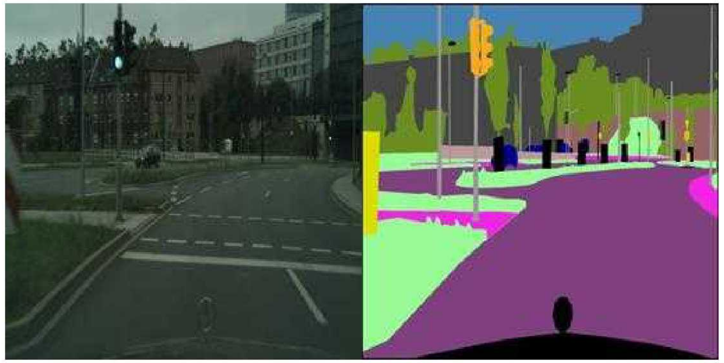

https://www.cityscapes-dataset.com/

Cityscape
(~1e5 pixels)

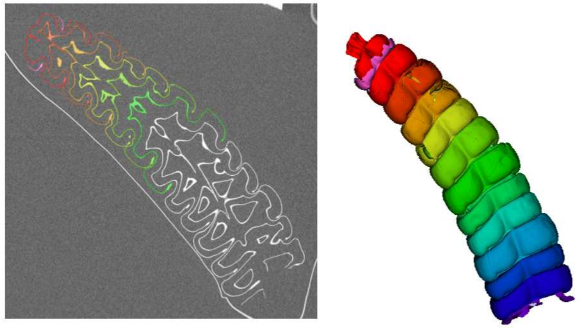

Rattlesnake Tail
(~1e9 voxels)

Mitigating Challenges

- CT scans are large
  - Used chunks of the volume
  - Optimized our model for GPU memory usage on GPU cluster
- Class Imbalance
  - Adjusted loss function that guides training
- Artifacts and noise
  - Selected Convolutional Neural Network (CNN) architecture with strength in shape recognition

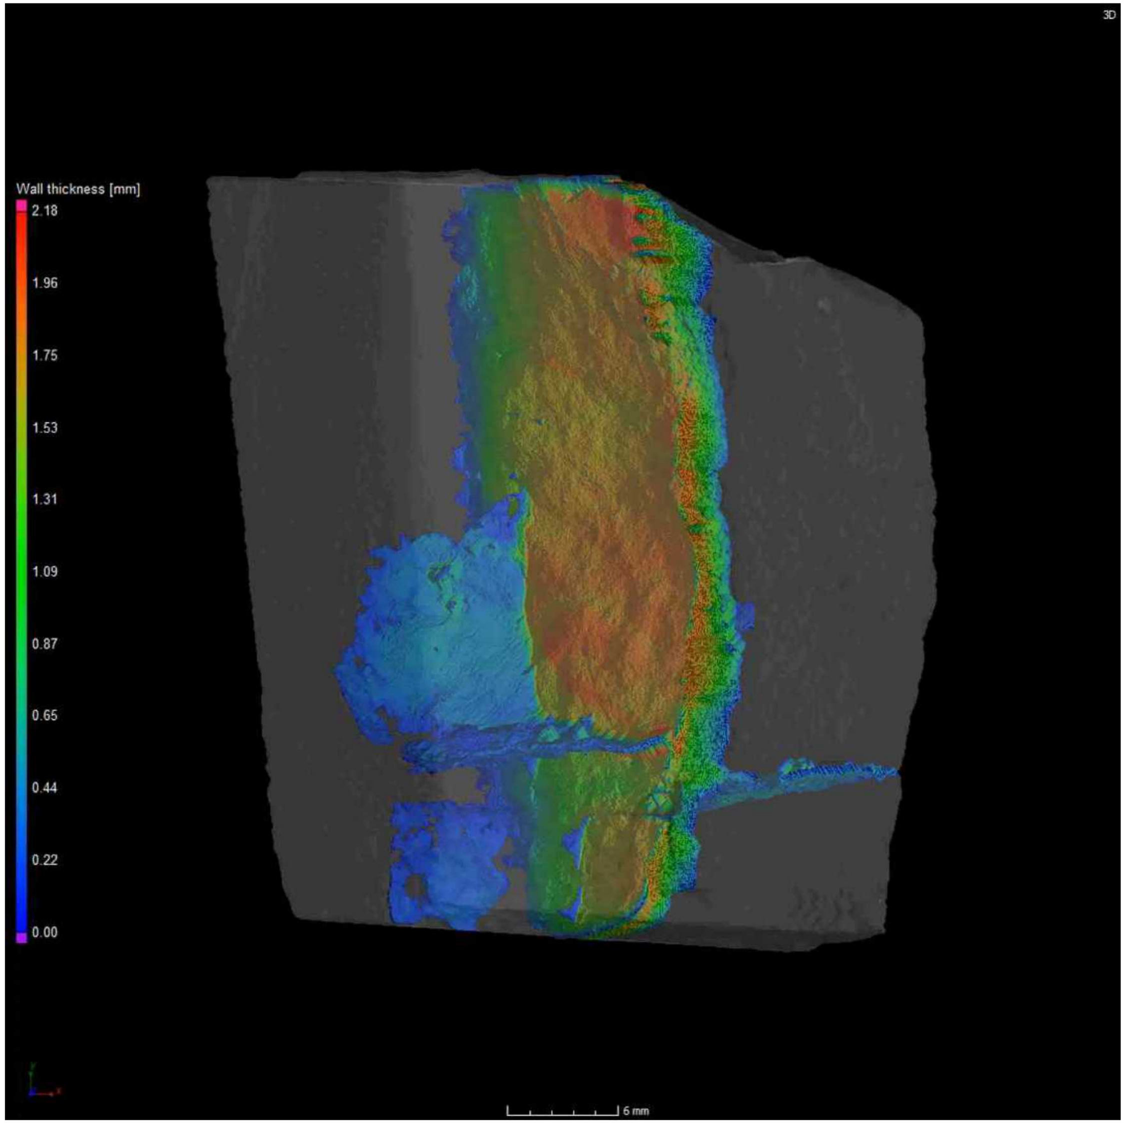

Deep learning is big data and large networks

- *"Deep learning is the first class of algorithms that is scalable... performance just keeps getting better as you feed them more data"*

Andrew Ng, Founder of Google Brain

- U-net, a big advance in biomedical segmentation

  Olaf Ronneberger, Philipp Fischer, Thomas Brox , "U-Net: Convolutional Networks for Biomedical Image Segmentation", in Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer, LNCS, Vol.9351: 234--241, 2015

- V-net follows as a natural extension

  F. Milletari, N. Navab, and S. A. Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in 2016 Fourth International Conference on 3D Vision (3DV), Oct 2016, pp.565–571

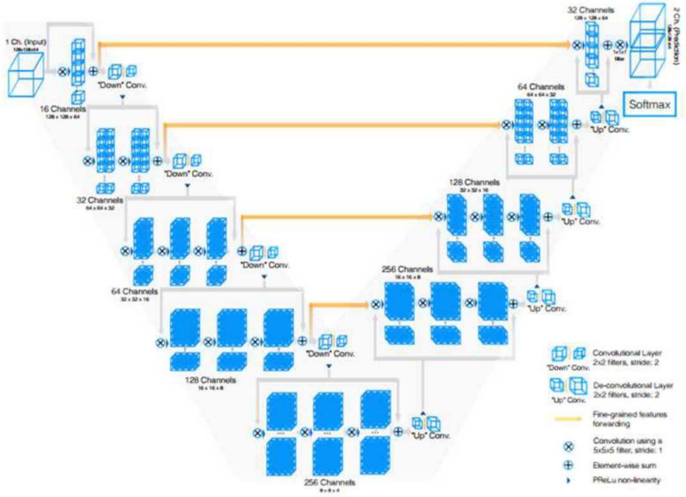

V-Net architecture for segmenting volumetric data (2016)

We started with a V-Net and made improvements as necessary

Deep Learning produces human quality segmentation

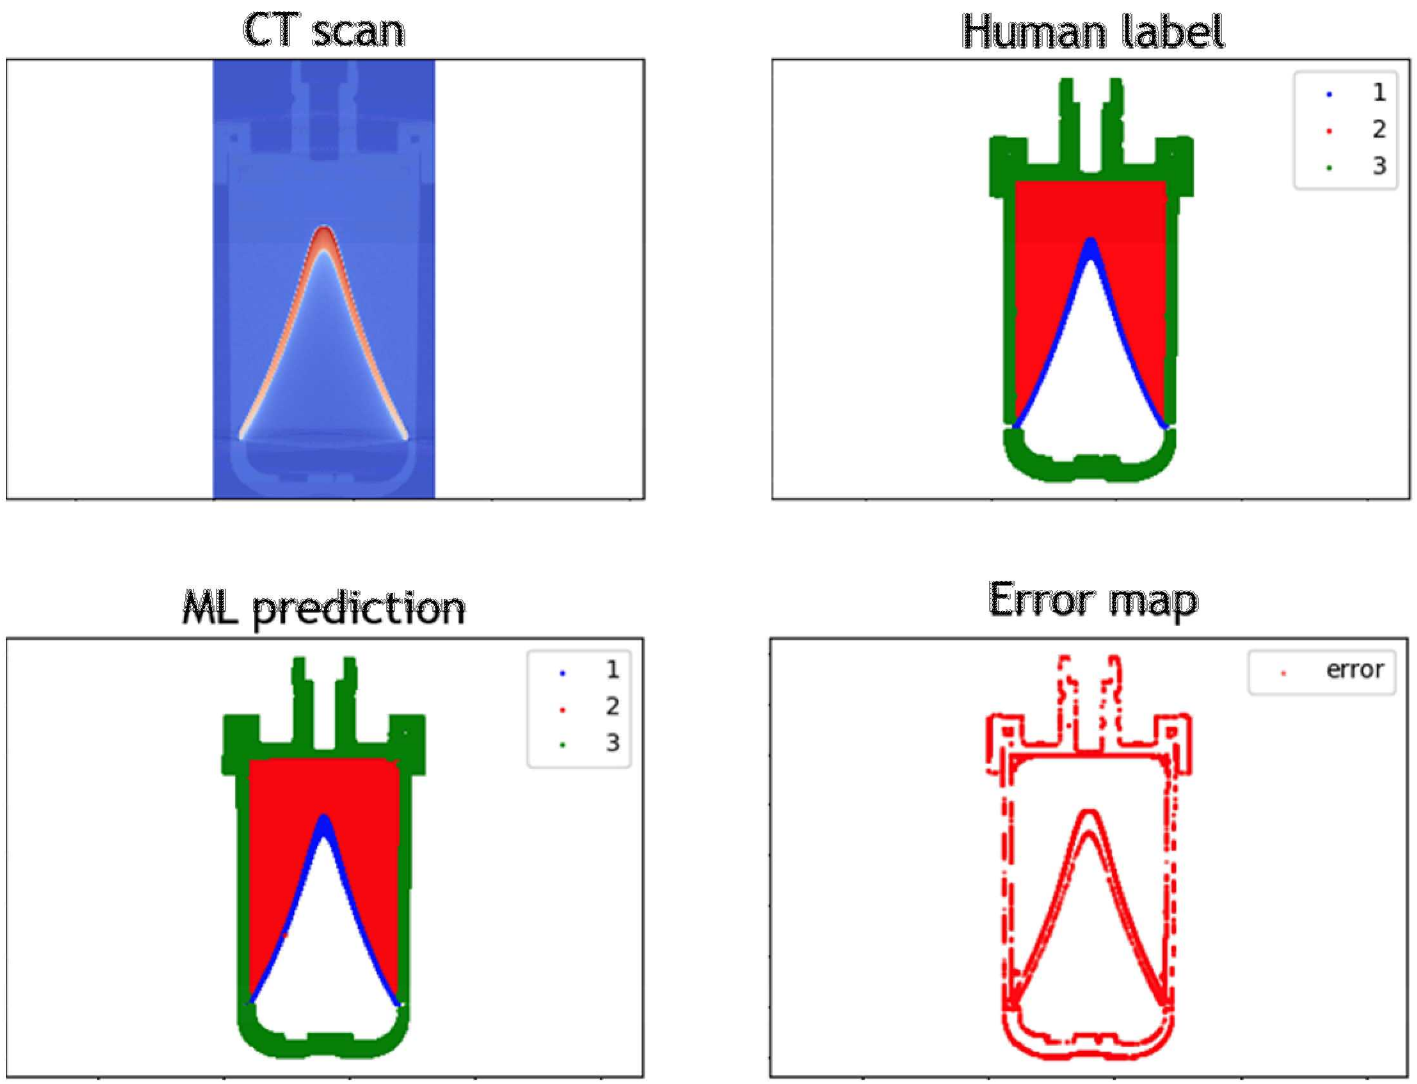

Deep learning model is robust to other domains

Slice of 3D Image

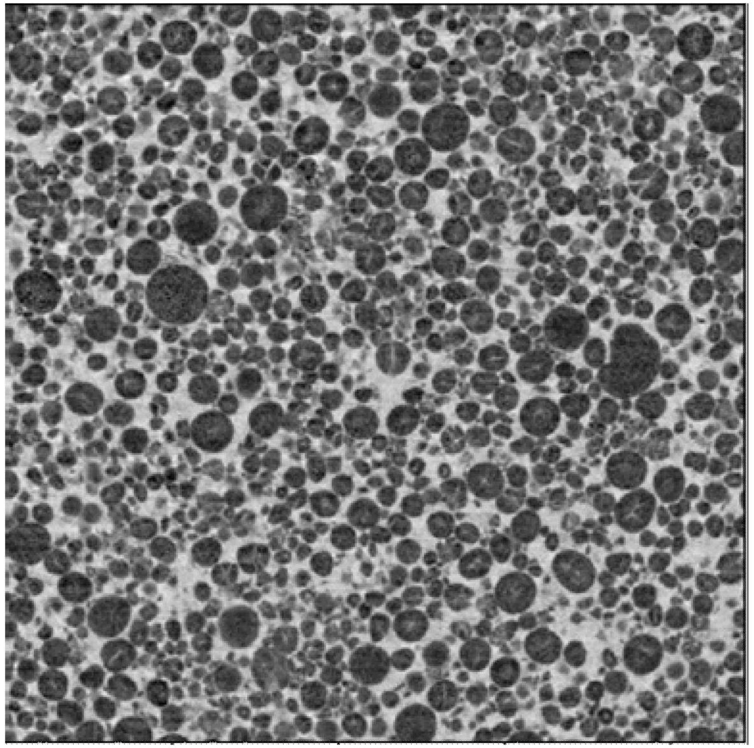

Human label

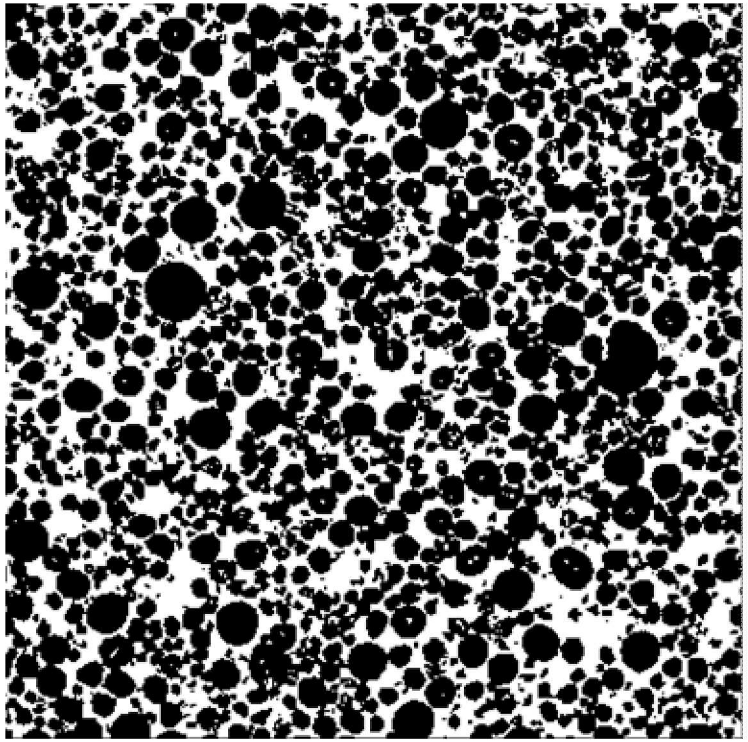

Volumetric battery segmentation achieves high accuracy compared to human labels

Slice of 3D Image

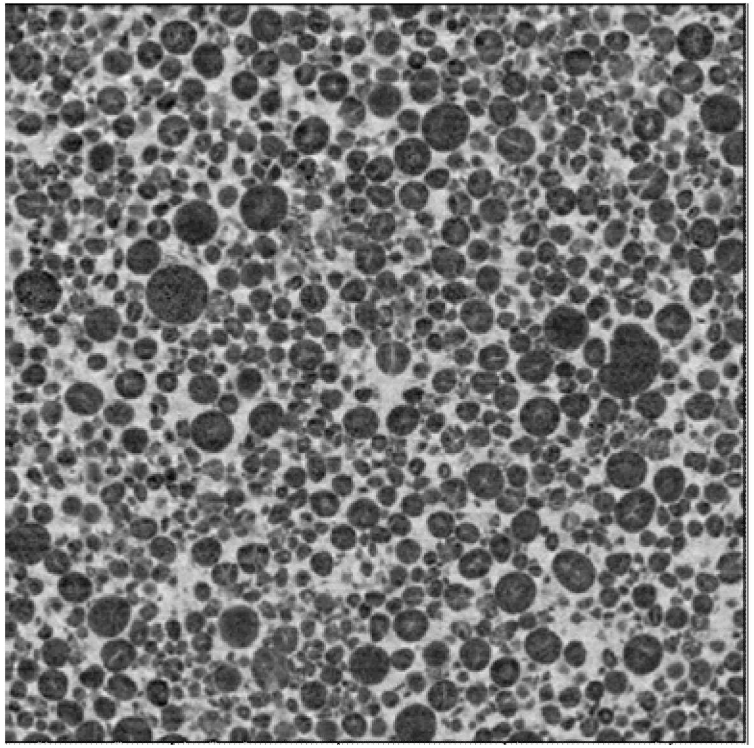

Human label

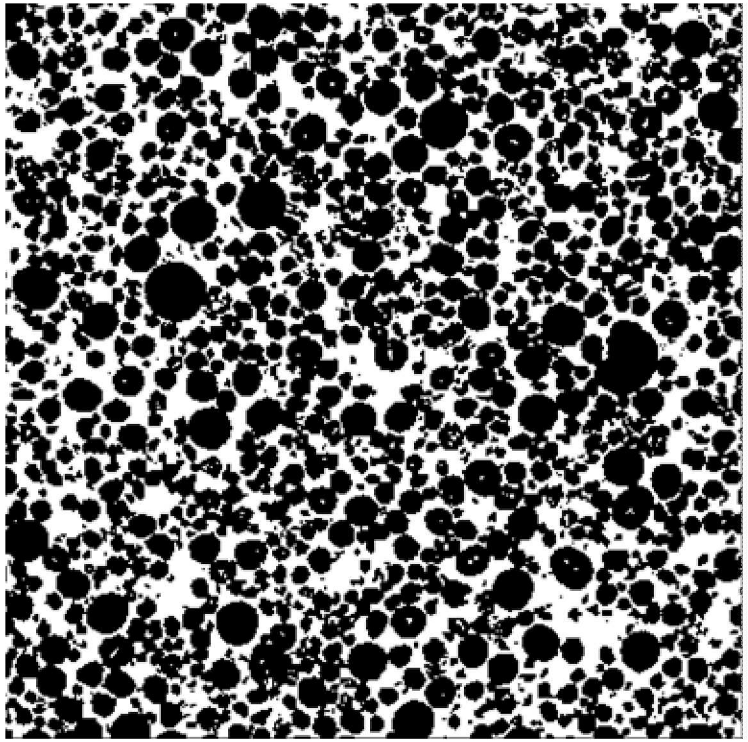

ML prediction

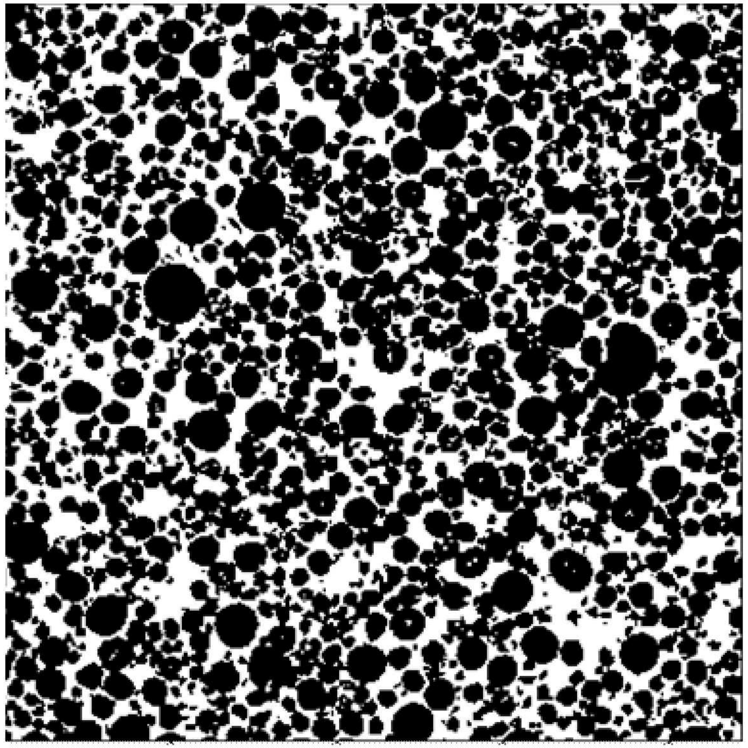

Averaged 99.7% accuracy over held out test set

http://www.cs.ox.ac.uk/people/yarin.gal/website/blog_3d801aa532c1ce.html

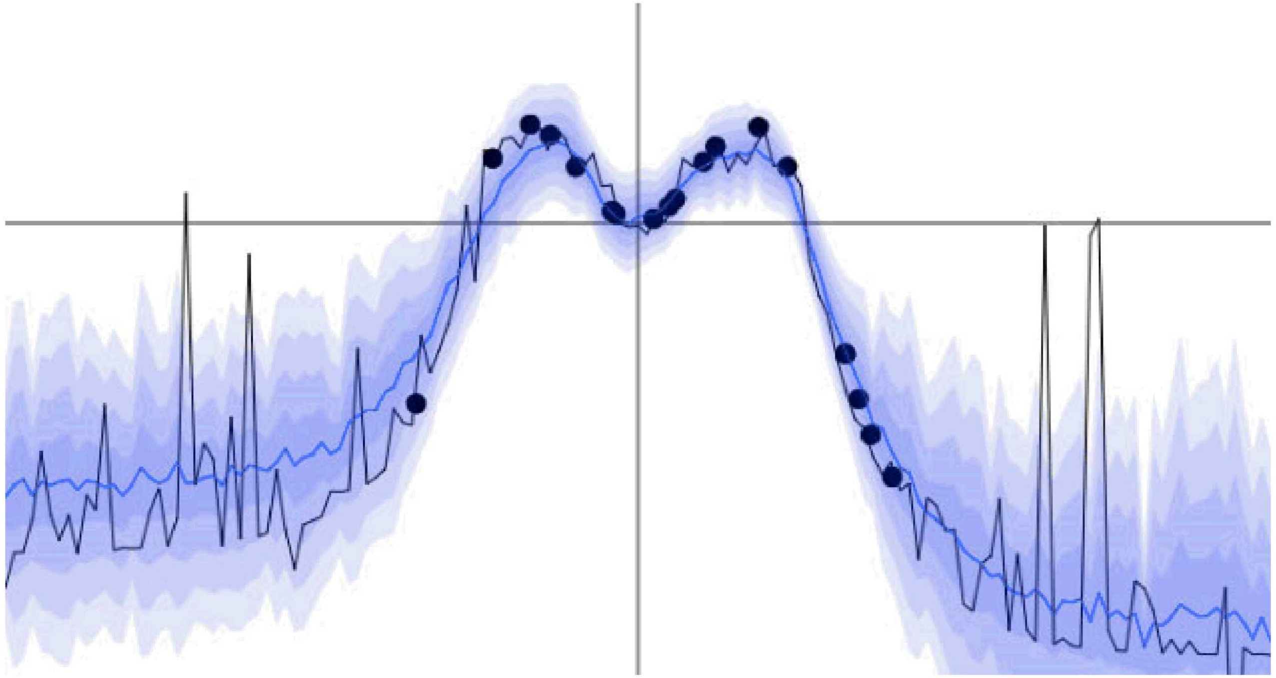

Using dropout to estimate segmentation confidence

UNCERTAINTY QUANTIFICATION

Softmax uncertainty as a baseline

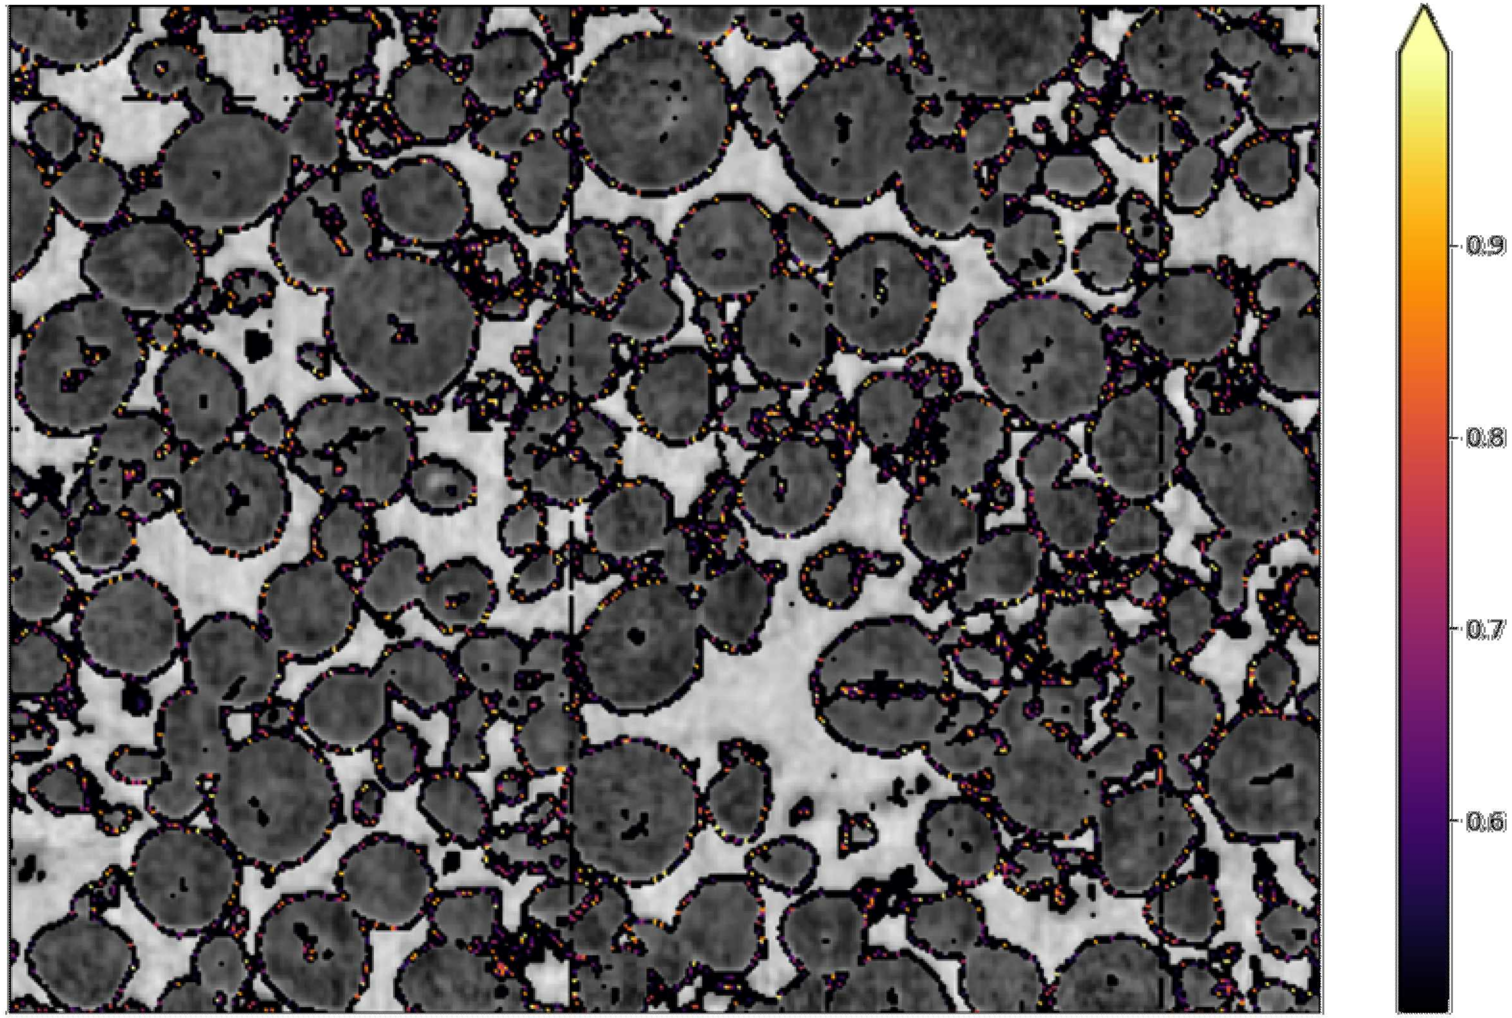

Standard neural network models do not provide error bars

- Softmax output is not always enough to determine model uncertainty

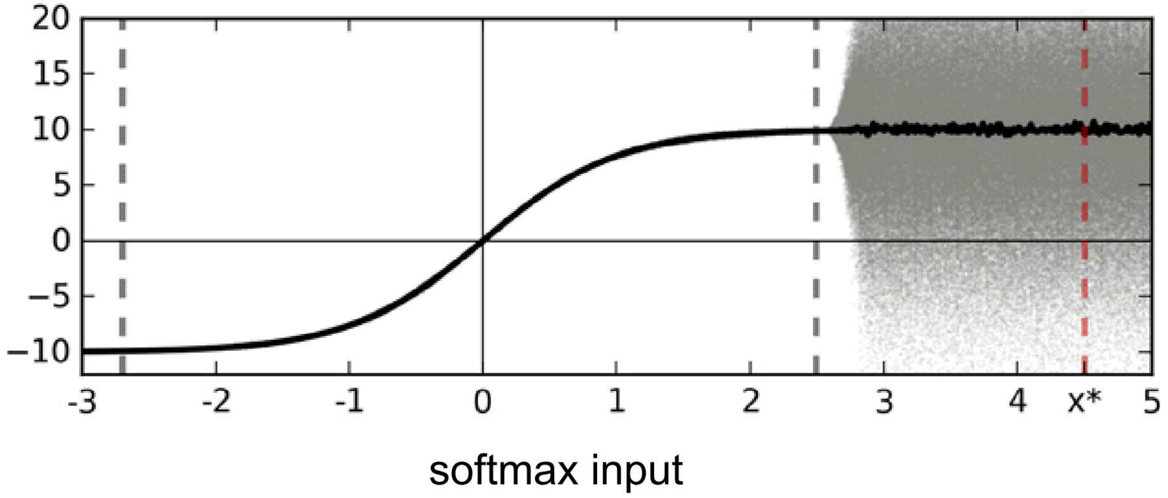

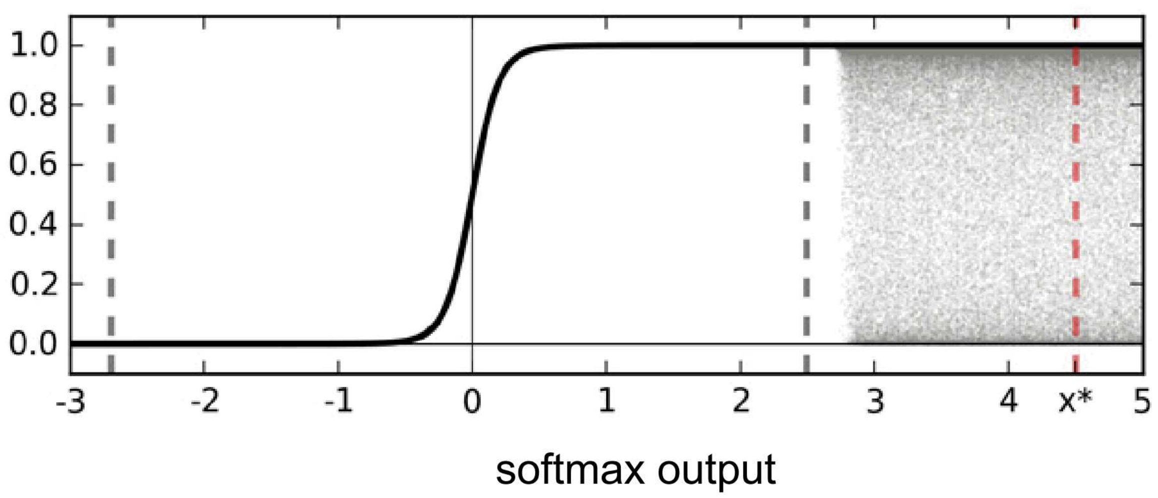

https://www.cs.ox.ac.uk/people/yarin.gal/website/blog_3d801aa532c1ce.html

- We can use dropout at inference time to approximate uncertainty

  Gal, Yarin, and Zoubin Ghahramani. "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning." international conference on machine learning. 2016.

Credible uncertainty in ML can be obtained

- A model trained on pandas should not classify gibbons with certainty

https://www.telegraph.co.uk/news/2016/09/15/pandas-arent-cute-theyre-death-loving-oxygen-thieves-lets-just-e/

https://en.wikipedia.org/wiki/Gibbon

- We can approximate uncertainty with small modifications to deep learning model
  - Gal, Yarin, and Zoubin Ghahramani. "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning." *international conference on machine learning*. 2016.
  - Dropout during inference
  - Sampling inference

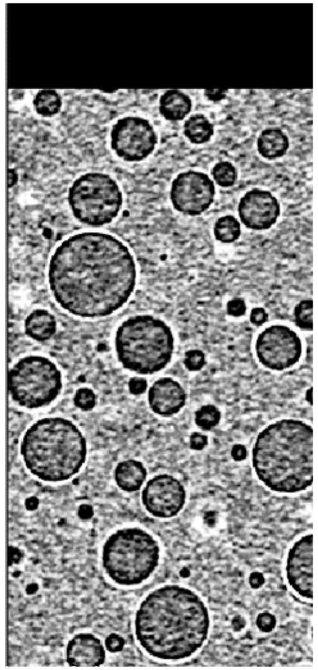

CT image

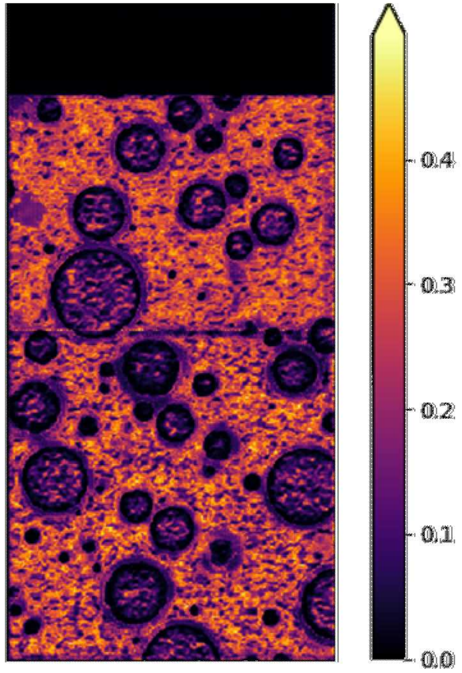

Uncertainty map

Understanding explosive fragment flight

FRAGMENT CHARACTERIZATION

Case Study: Southwest Airlines Engine Explosion

- Southwest flight 1380, April 17th, 2018
  - Engine failure after takeoff from New York LaGuardia
  - Metal fragments from explosion punctured fuselage
  - 1 fatality, several injuries
- How can we understand fragment flight to prevent future safety incidents?

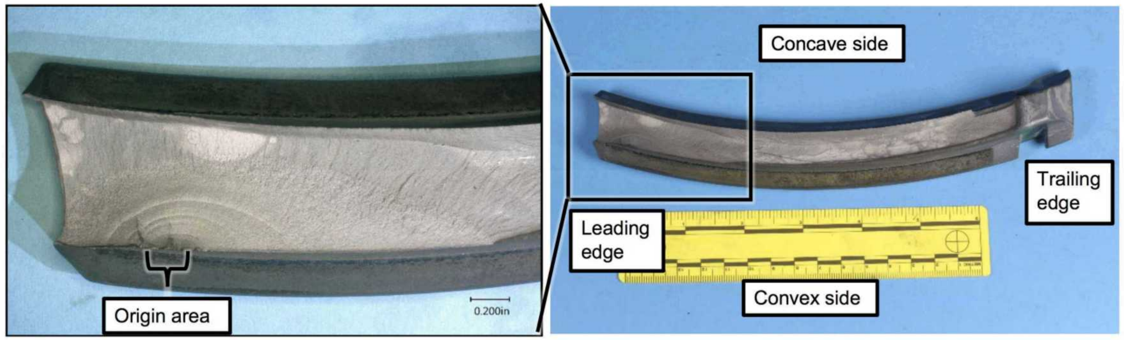

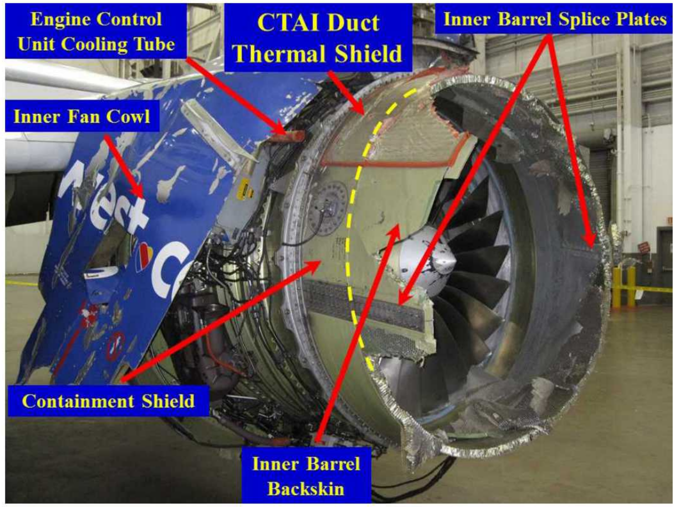

https://www.nytimes.com/2018/04/17/us/southwest-airlines-explosion.html
https://en.wikipedia.org/wiki/Southwest_Airlines_Flight_1380

Problem Statement

- Multi-camera stereo system with synchronized videos
- Fragments from experiment fly across field-of-view
- Using machine learning & computer vision techniques, can we answer:
  - How do fragments form?
  - Where are the fragments located in 3d space?
  - What are their velocities?
  - What are their masses?

Stereo camera system

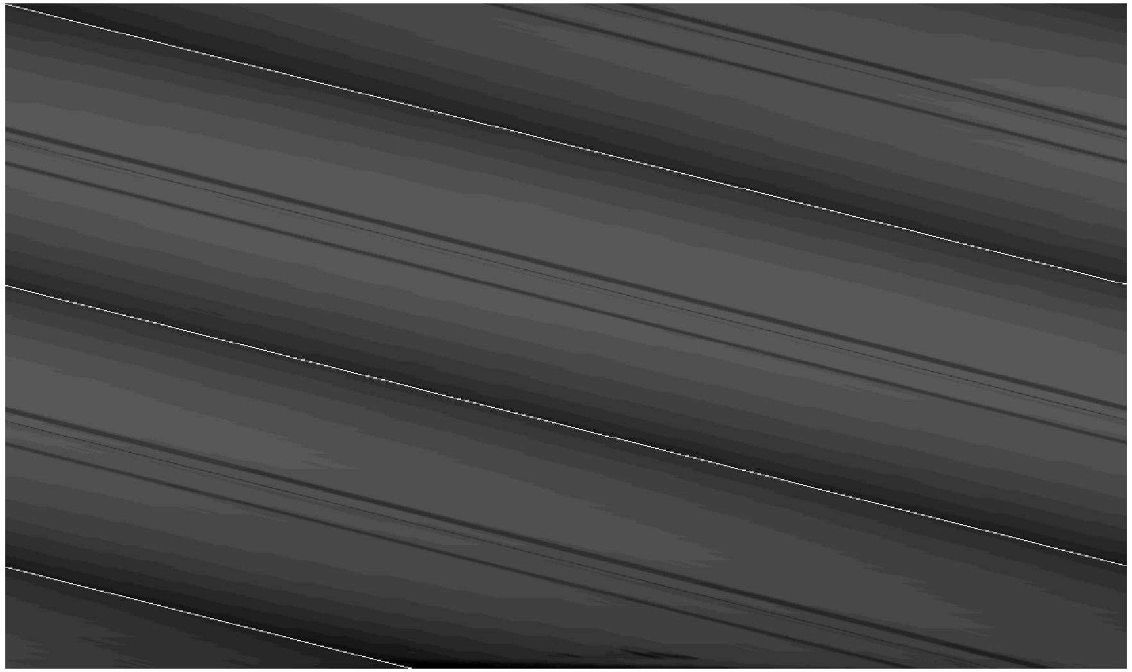

Fragment experiment

http://www.1vision.co.il/Stereo-3D

Simulated experiments

- Simulate an experimental setup with a stereo camera system
  - Verifiable against known ground truth
  - Avoid confusers and noisy data
- Allow any positioning of cameras, informing experiment design
- Can quickly and cheaply simulate a ton of experiments

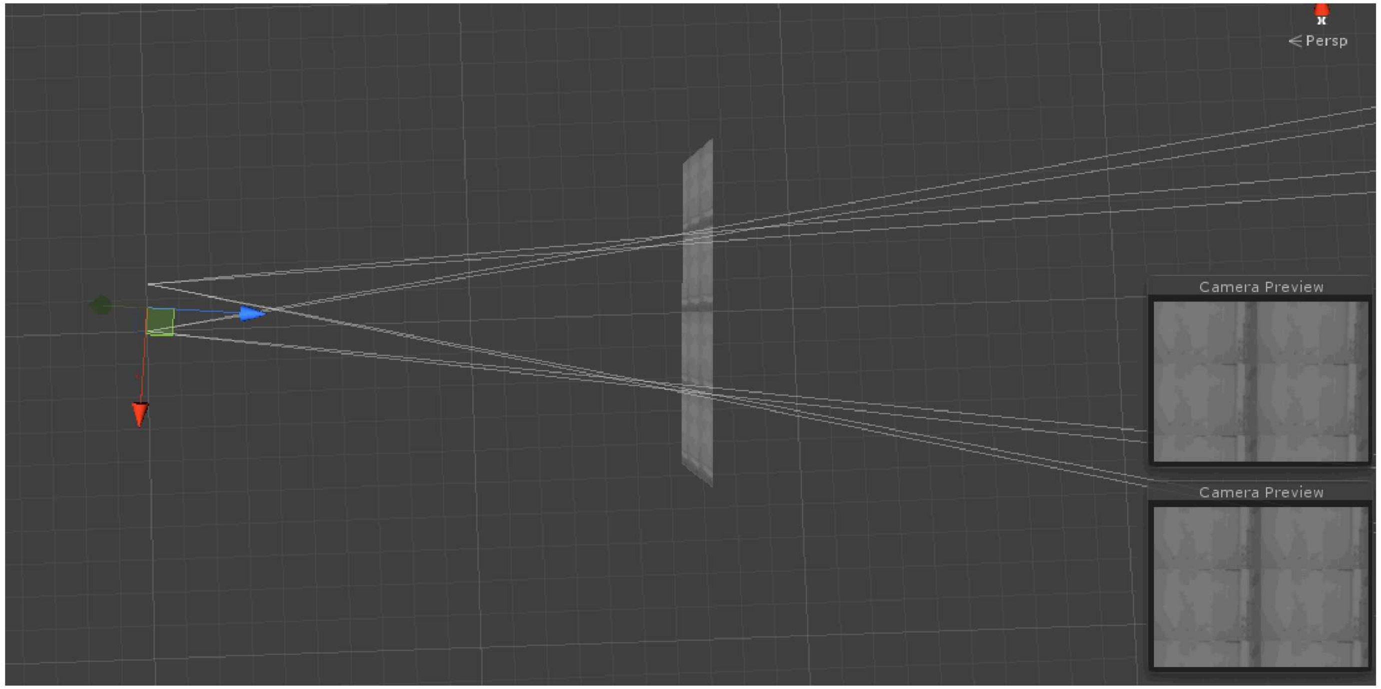

Bird's eye view of simulation environment

Fragment characterization method

- Simulate fragment experiments
- Use deep learning to segment fragment locations
- Track fragments
- Stereo matching between fragments
- Characterize each fragment
  - Positions
  - Velocity
  - 3D reconstruction

Simulated experiment

Generative Adversarial Networks: a game theoretic approach to machine learning

- Generative Adversarial Networks (GANs) pit two competing neural networks against each other
  - **The generator**, tries to mimic real results
  - **The discriminator**, tries to identify mimicked results from real results

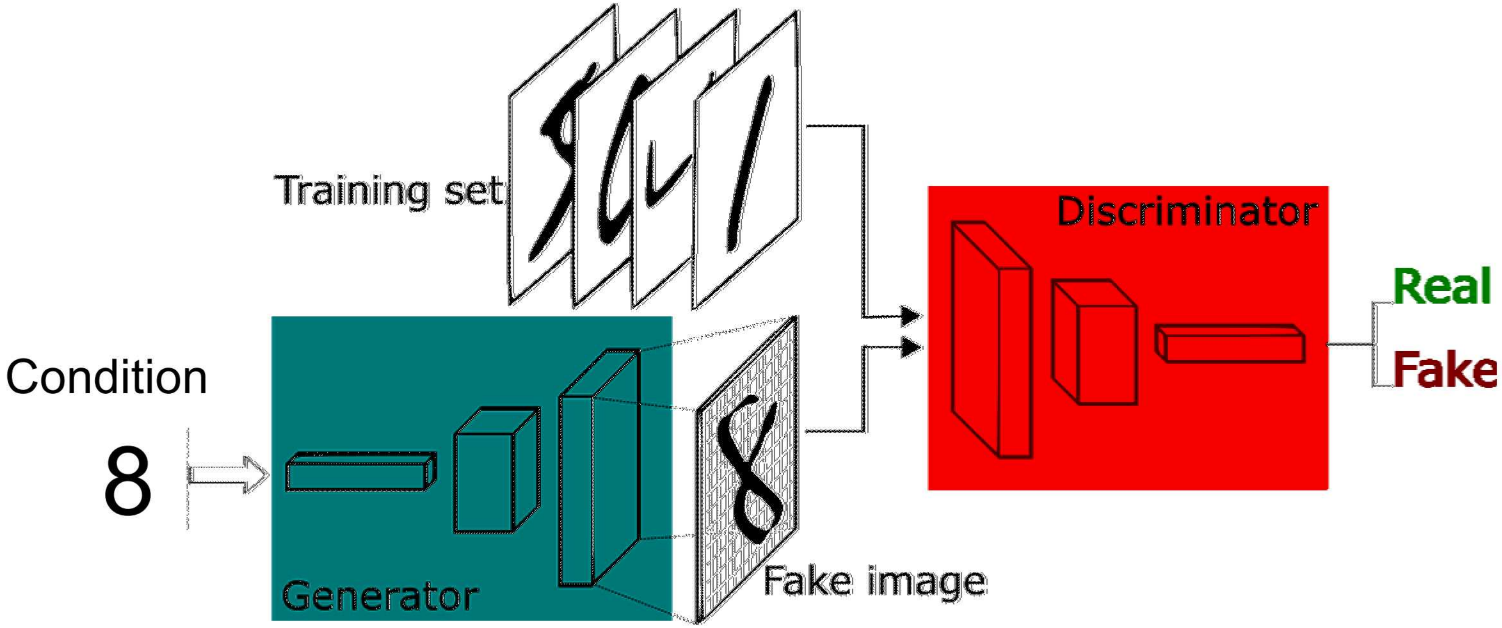

https://deeplearning4j.org/generative-adversarial-network

Pix2Pix and Generative Adversarial Networks

- Pix2pix model
  - Condition is an image instead of a label
    - E.g. color segmentation of a scene
  - GAN has to learn how to fill in segmentations convincingly
  - Training goal is to fool the discriminator

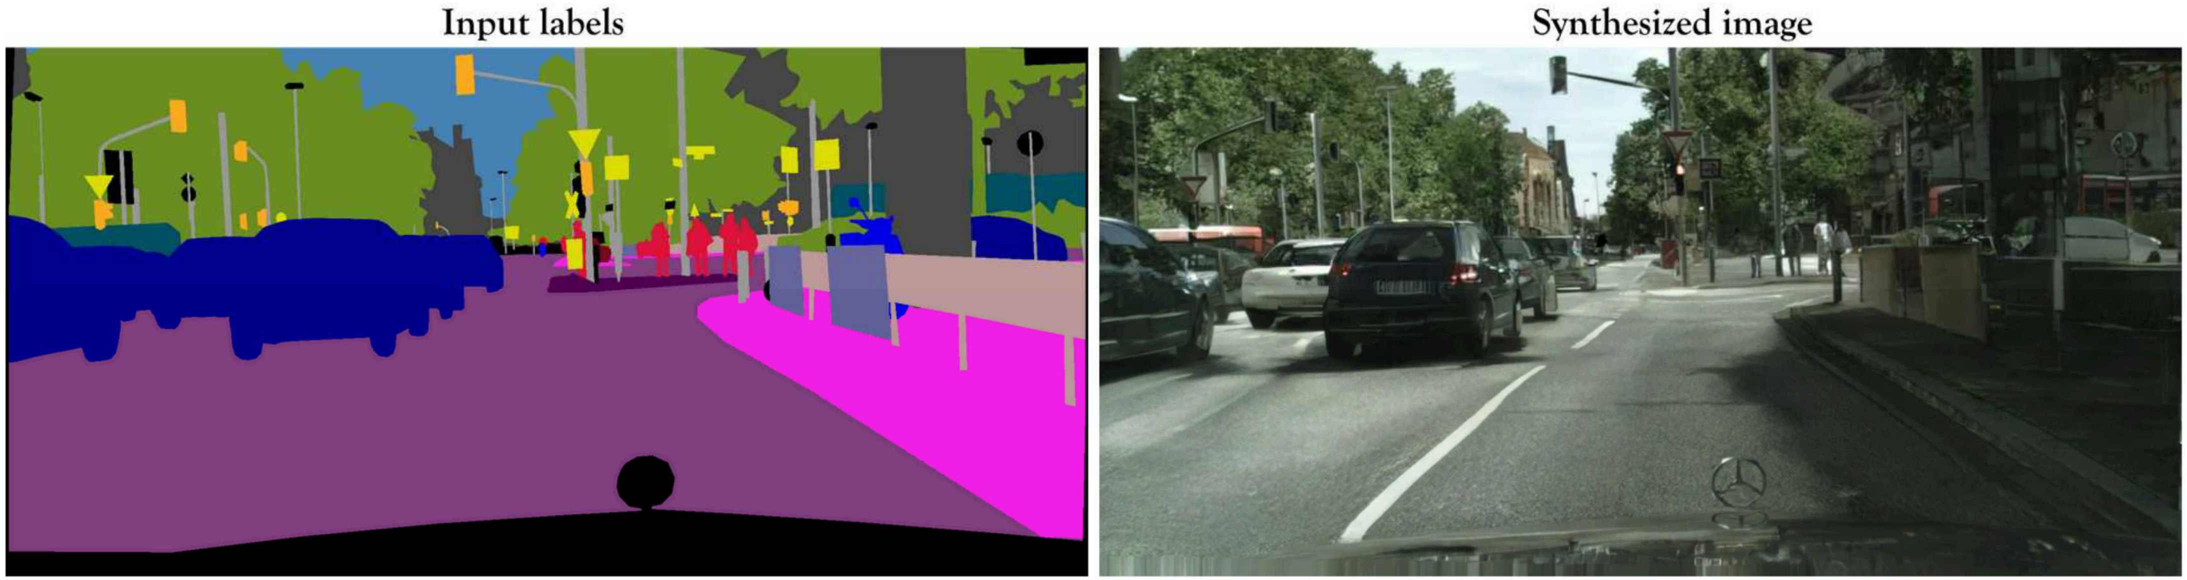

Wang, et. Al. “High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANS”. CVPR 2018

Algorithmic tracking

- Idea: frame the tracking problem as a graph optimization problem
  - Nodes are detection locations for all frames from pix2pix
  - Edges are all possible tracks from fragments in one frame to fragments in next frame
- 3-way matching algorithm
  - Given edges between 3 successive frames
  - Find the edges that make physical sense:
    - Approximately constant velocity
    - Fragment shapes are similar

Fragment correlation

Tracking

- CSRT algorithm:
  - User specifies first bounding box (or automatic segmentation algorithm)
  - Algorithm learns a filter from affine transformations of the image patch
  - Filter correlated with next frame, max response is the detected location
- Future research: improve tracking by incorporating 3-way matching

CSRT algorithm applied to simulated video

Given fragment geometries, can we predict flight?

- Estimate a physics calculation incredibly quickly using machine learning (GPUs)
- Exascale and future platforms double as training data generation of previously unseen physics

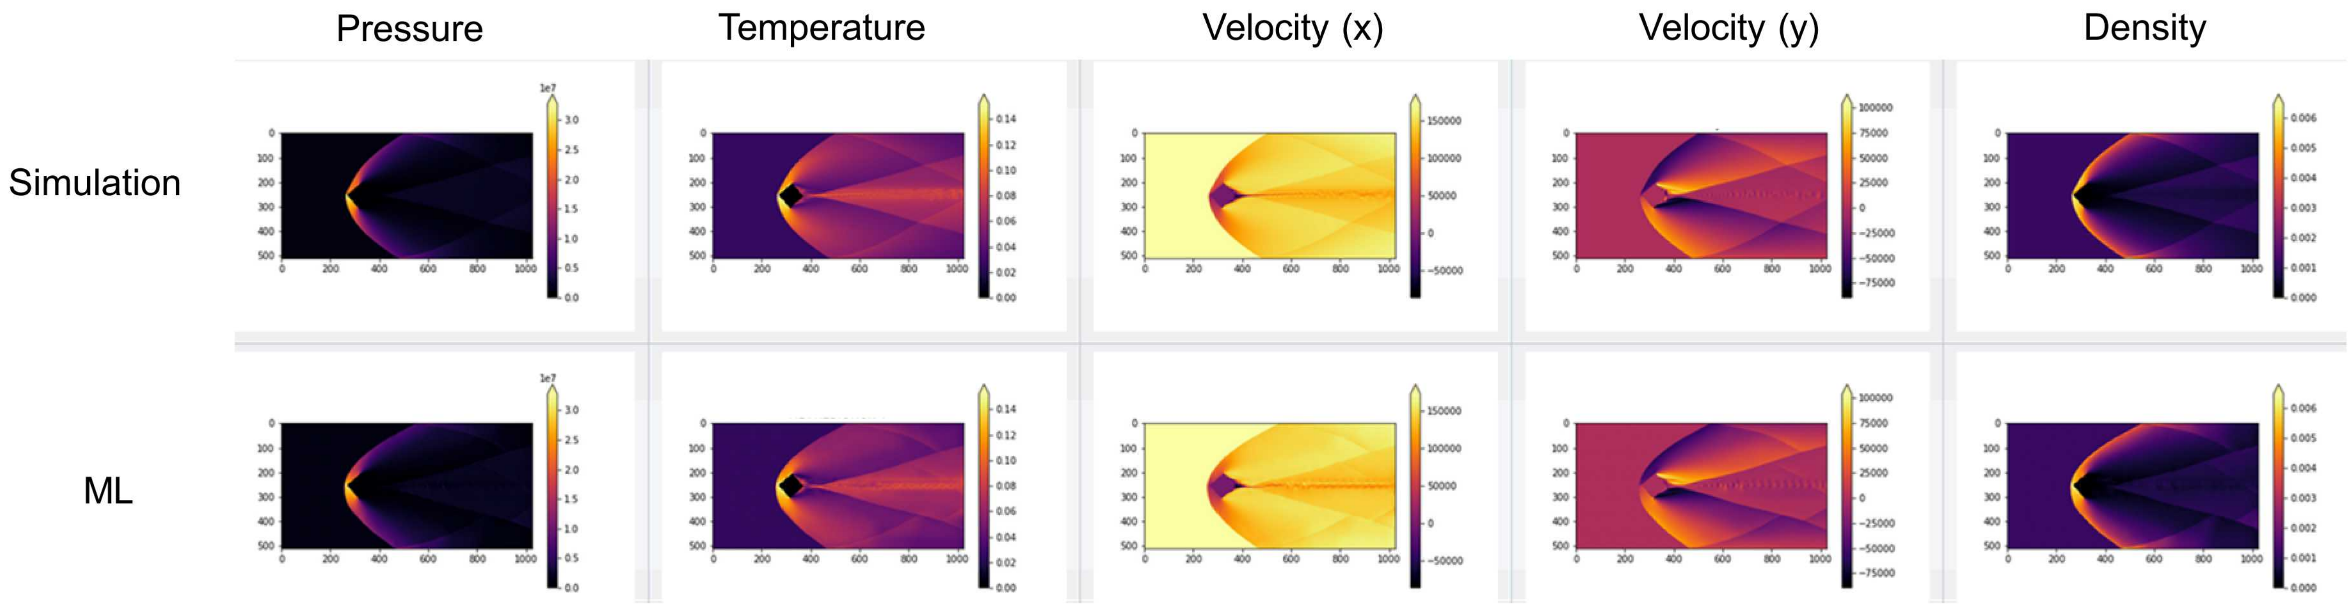

Conditional GAN mimics flow field solutions

- Does well qualitatively, not so well quantitatively
- Adding physics to loss:
  - More accurately represents shockwave
  - Better around the edges of the object
- Types of physics + qualitative loss combination remain to be studied

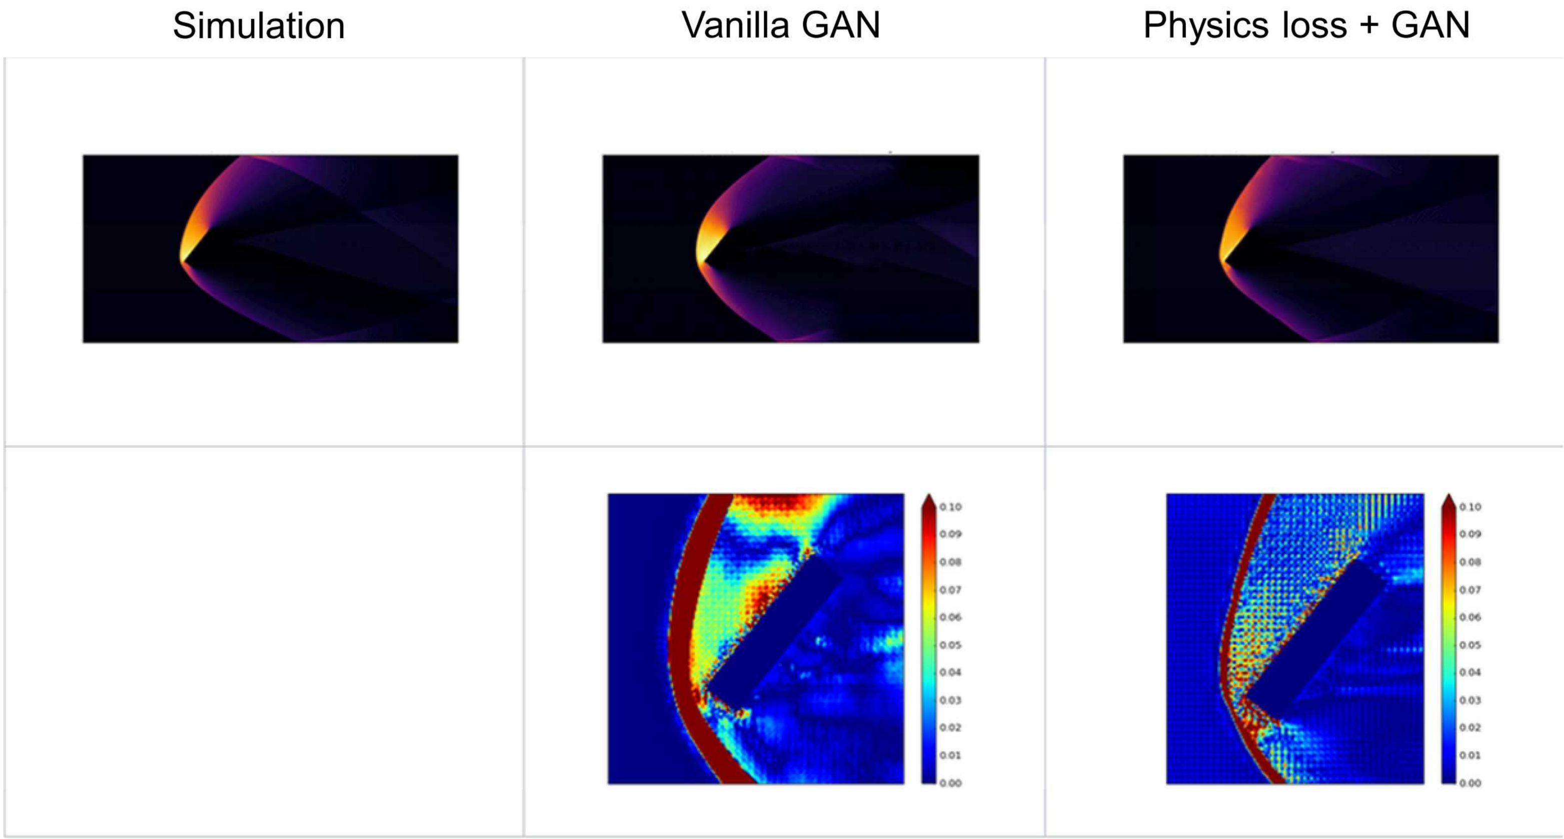

Quantities of interest are close

- Quantities of interest (Added to loss function):
  - Drag
  - Lift
  - Torque
- Exciting but preliminary
  - Experiment involves only rectangles
  - Results for 2D only. 3D in progress.
  - Fixed at Mach 5
- Time speed up
  - 2 min 30 s (HPC codes) -> ~250 ms (ML)
  - >100x faster

	Avg Relative Error
Drag	1.87%
Lift	5.63%
Torque	2.29%

Full 3D is showing the same early promise as 2D

- Preliminary results
- Have not enabled physics loss or force loss yet
- Not calculating 1 lift force and 2 torques yet

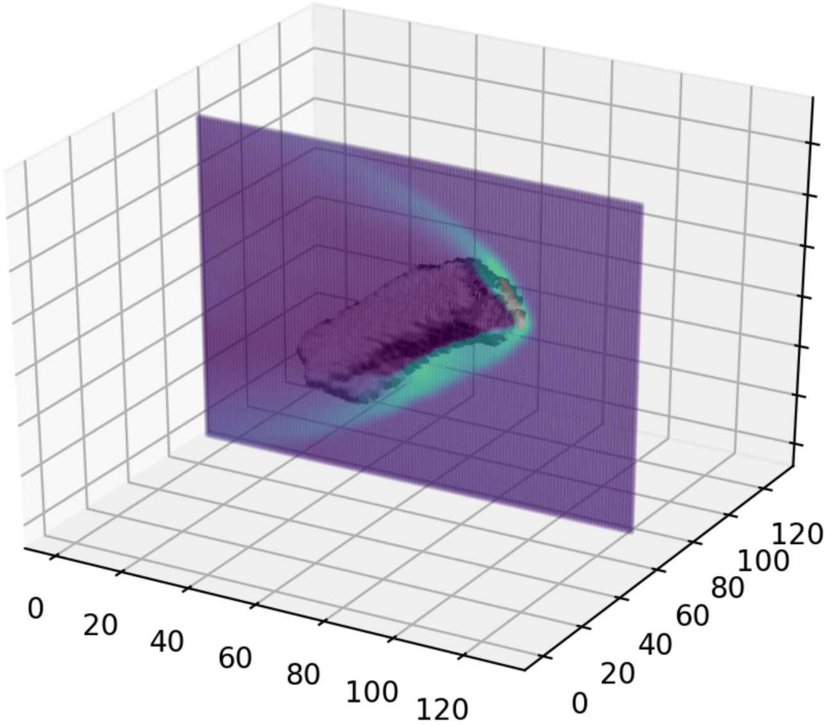

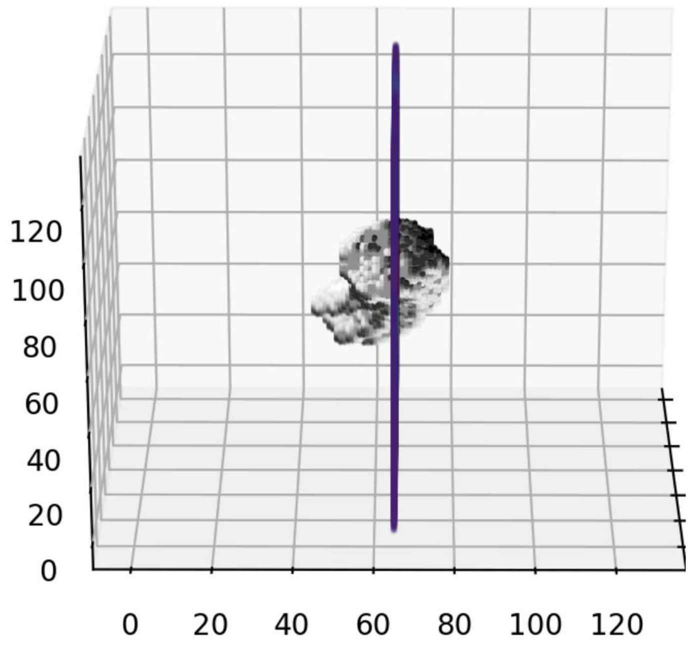

	Avg Relative Error
Drag	6.8%
Lift (Z)	29%
Torque (Z)	13%

Machine learning applied to structural dynamics

- A simple three degree-of-freedom (3DOF) system was devised to study the applicability of the LSTM to problems with intermittent contact.
- The system was excited with a random vibration input at one end while the other end was kept fixed. A small gap between M_2 and M_3 enabled intermittent contact to occur.

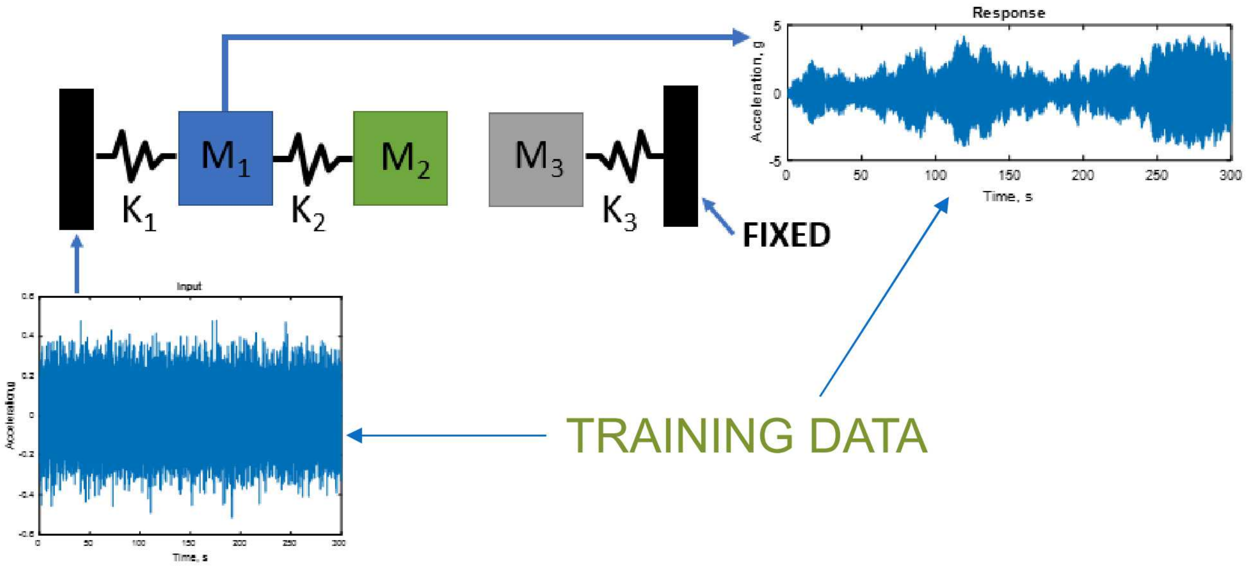

LSTM can predict system responses

- The LSTM was then used to predict the system response to a new time history (of similar statistics) by feeding it only a portion of the output.

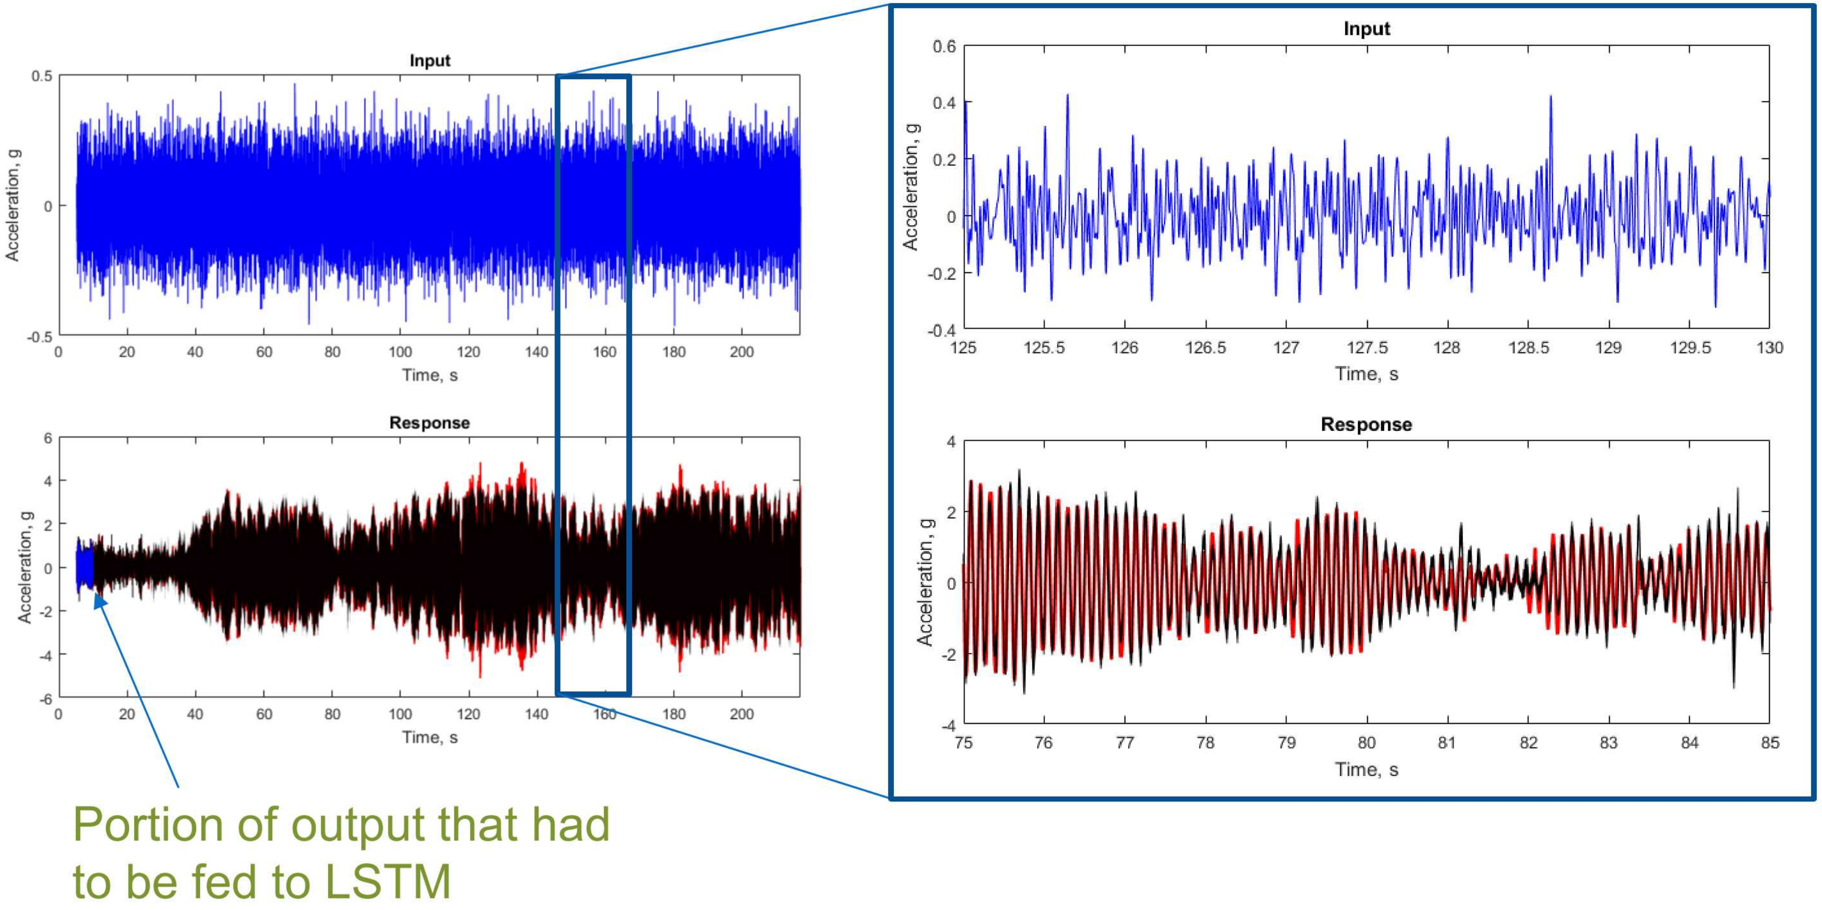

Next steps: Multiphysics analysis via machine learning

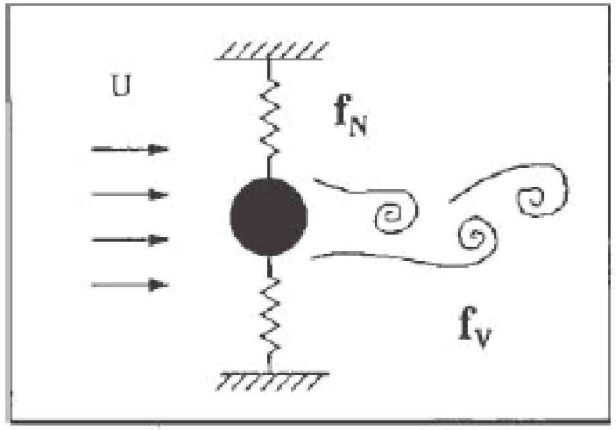

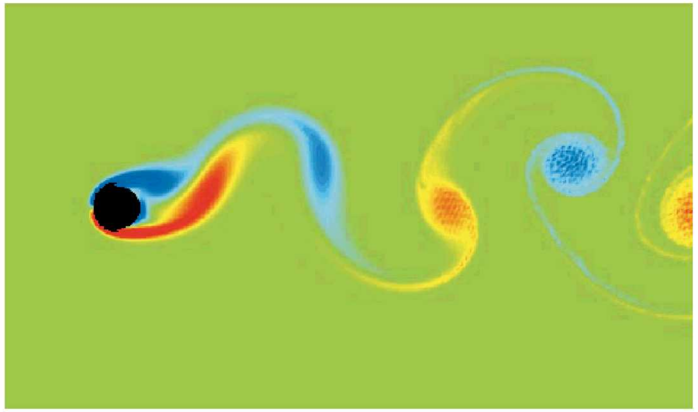

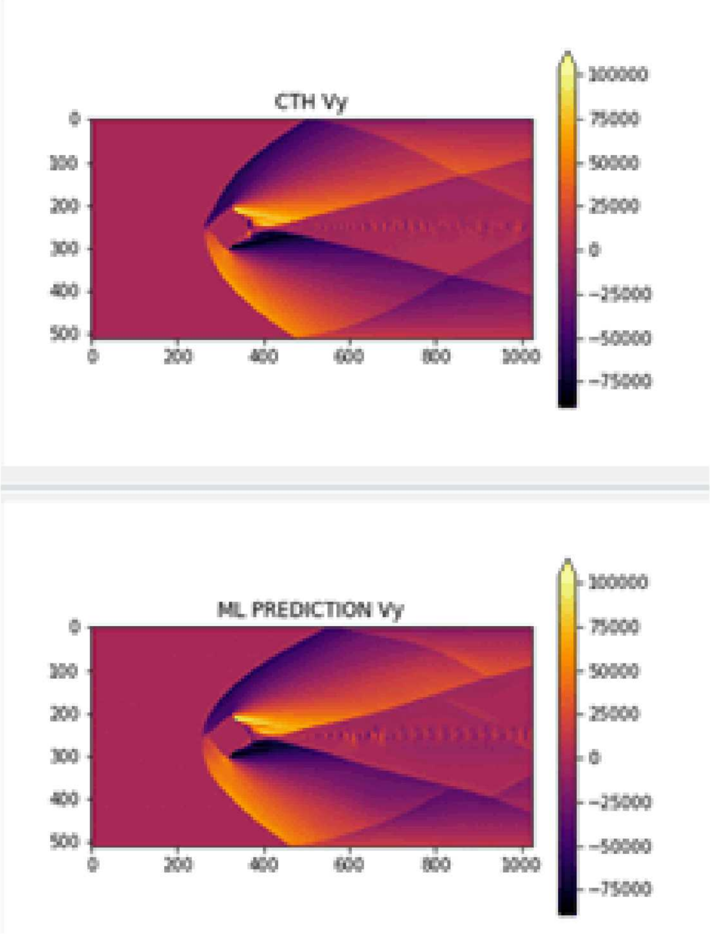

Can multiple machine learning models communicate effectively to predict complex system dynamics?

Machine Learning augments all Sandia missions

- Machine Learning - give a computer or algorithm learning capabilities
- We work with teams across the Labs
  - Physics modeling
  - Material discovery and design
  - Object recognition
  - 3D reconstruction
  - Natural language translation
  - Image/Volume Segmentation
  - Many more

Sandia Interdisciplinary Machine Learning Research Team

simlr@sandia.gov

?

BONUS SLIDES

What can ML do for you?

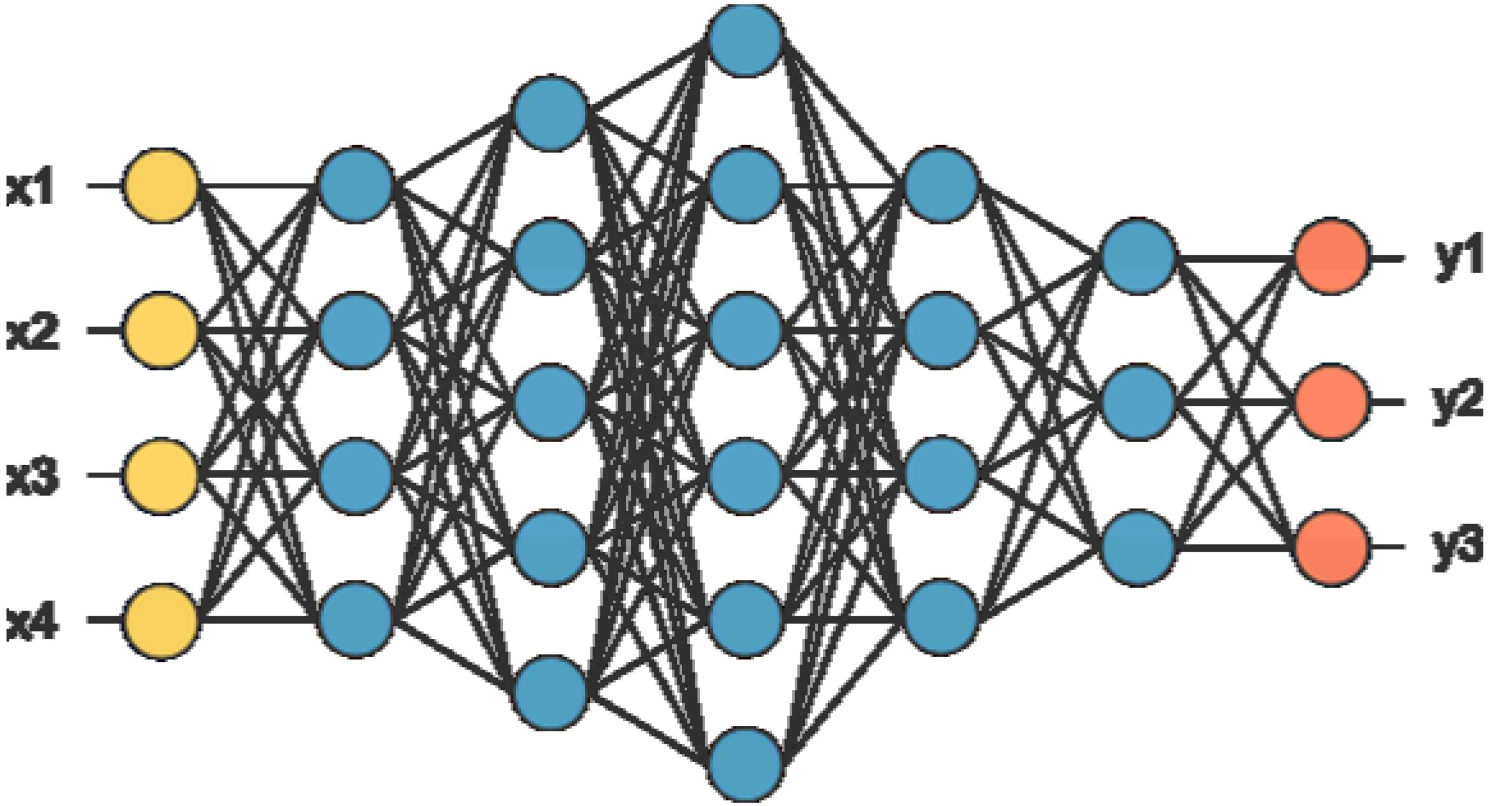

https://www.neuraldesigner.com/

A general introduction to the field

MACHINE LEARNING & COMPUTER VISION

Why Machine Learning?

https://finance.yahoo.com/quote/BTC-USD

- Typical computer algorithms:
  - have an input,
  - do some computation,
  - and produce some output

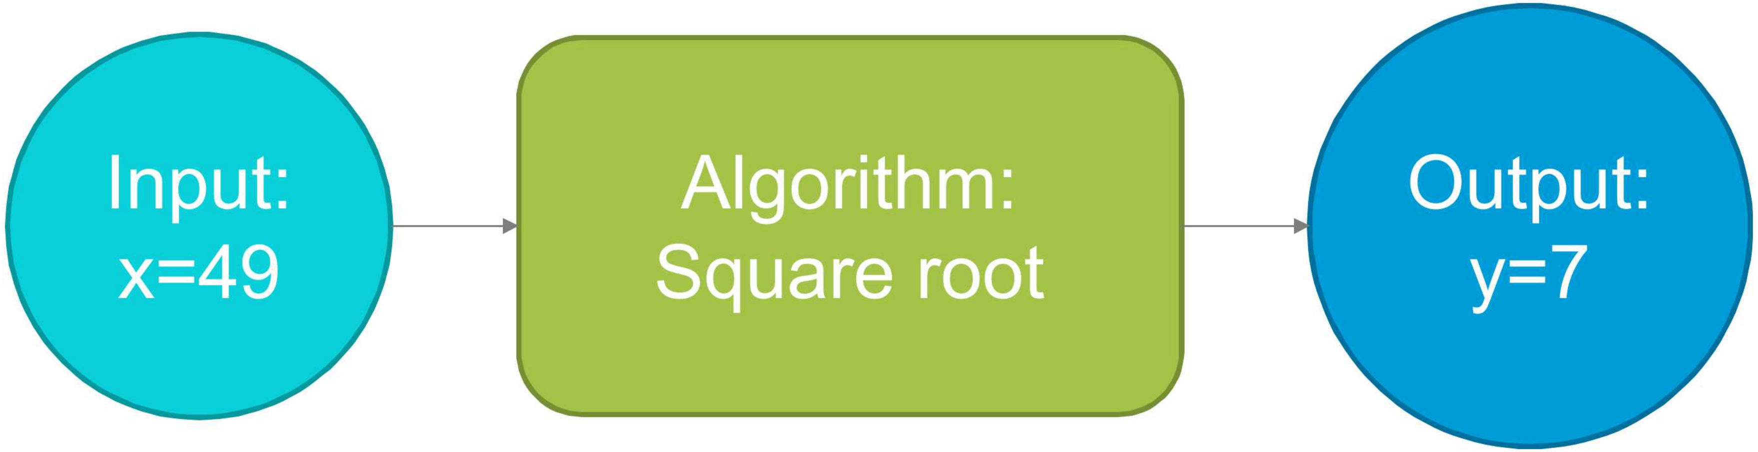

- Machine learning algorithms:
  - have lots of inputs,
  - do some computation,
  - and produce a *function*

- The input to an ML algorithm is a dataset
  - X's: inputs to the function you want to learn
  - Y's: desired outputs to the function
- The ML output is a "trained classifier/regressor"
  - Given an unseen input, it will (hopefully) produce the correct output
  - Types:
    - Binary
    - Classification
    - Regression

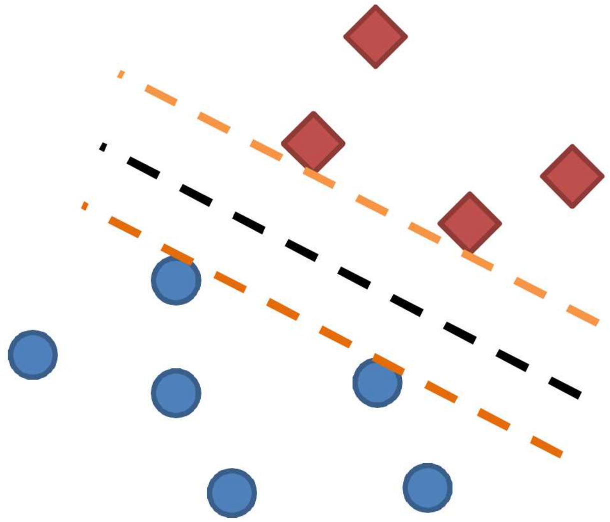

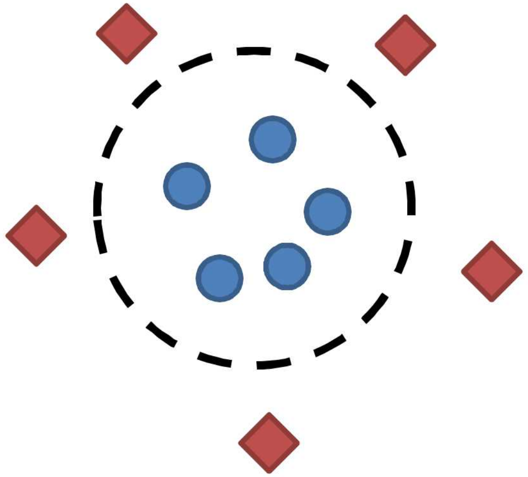

Neural Networks

- Layers are functions with “trainable” parameters
  - e.g. y=mx + b
- Layering linear functions models non-linear functions

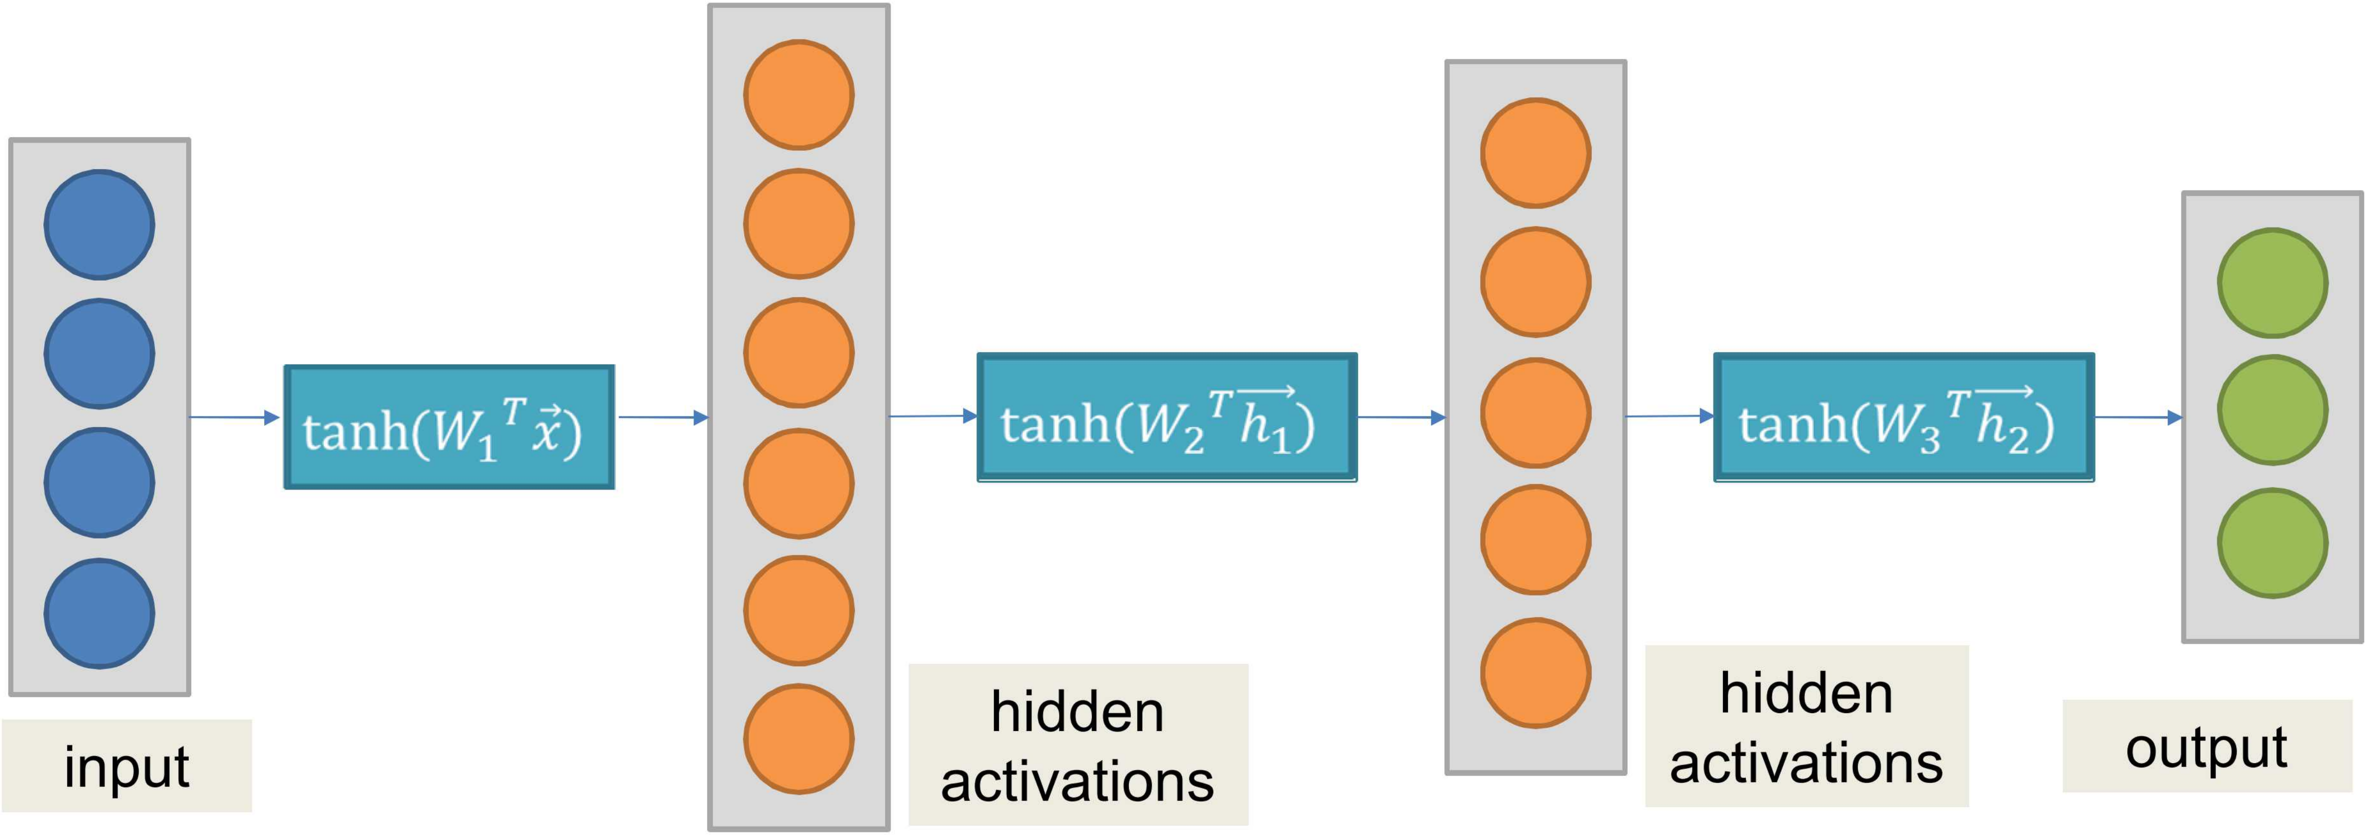

Training Neural Networks

- Training step:
  - Send an input, x, through the network
  - Compare the output, y, to the known output
  - Update layer parameters
- Repeat!

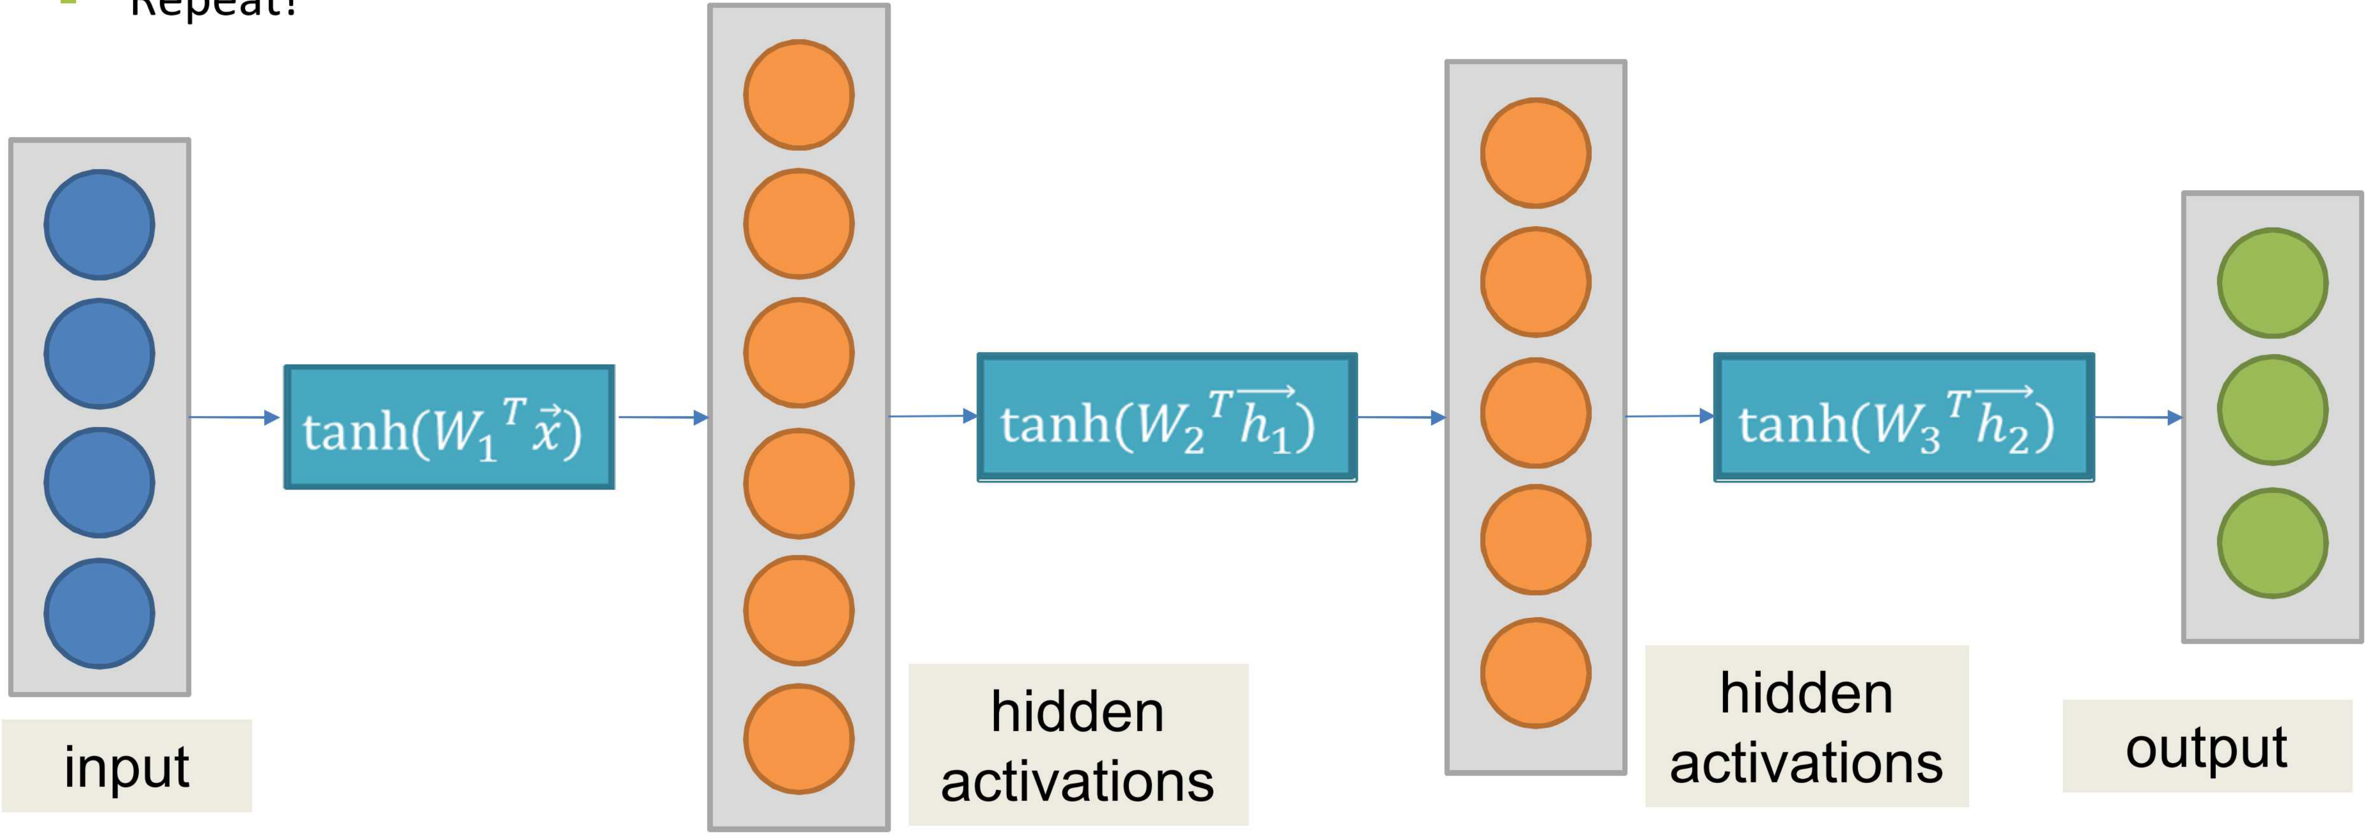

Convolutional neural networks

- Instead of a linear function (i.e. a matrix multiply) use convolutional layers
- Convolutions exploit spatial locality in pixels
- Early layers in the network learn simple features
- Later layers learn more abstract features

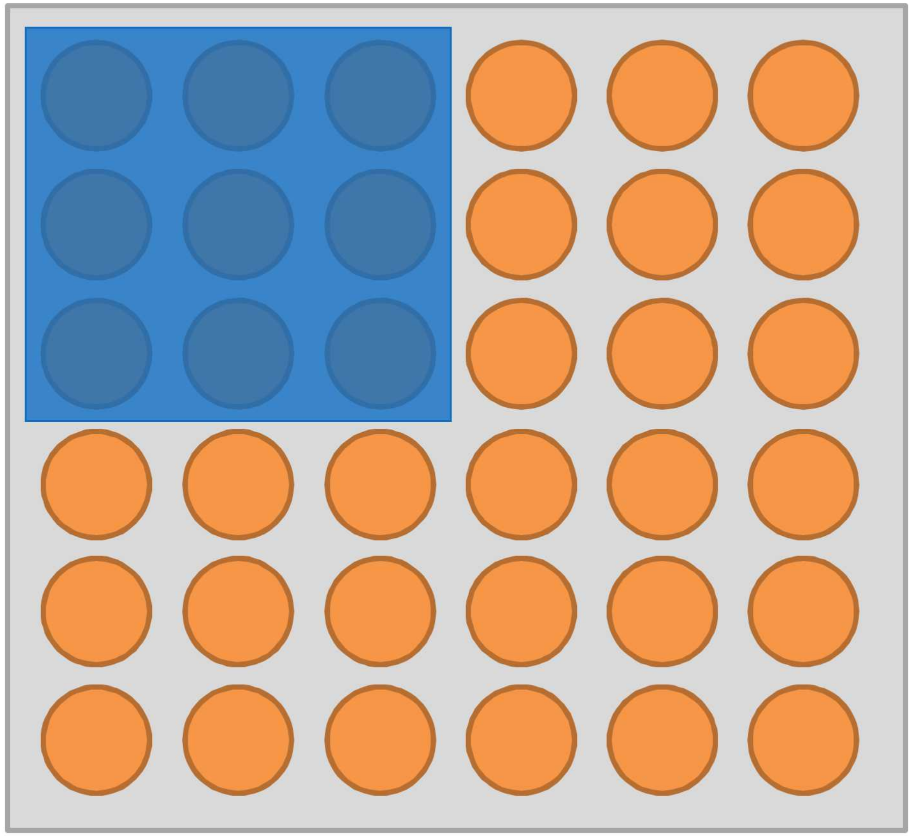

Convolutional neural networks

- Instead of a linear function (i.e. a matrix multiply) use convolutional layers
- Convolutions exploit spatial locality in pixels
- Early layers in the network learn simple features
- Later layers learn more abstract features

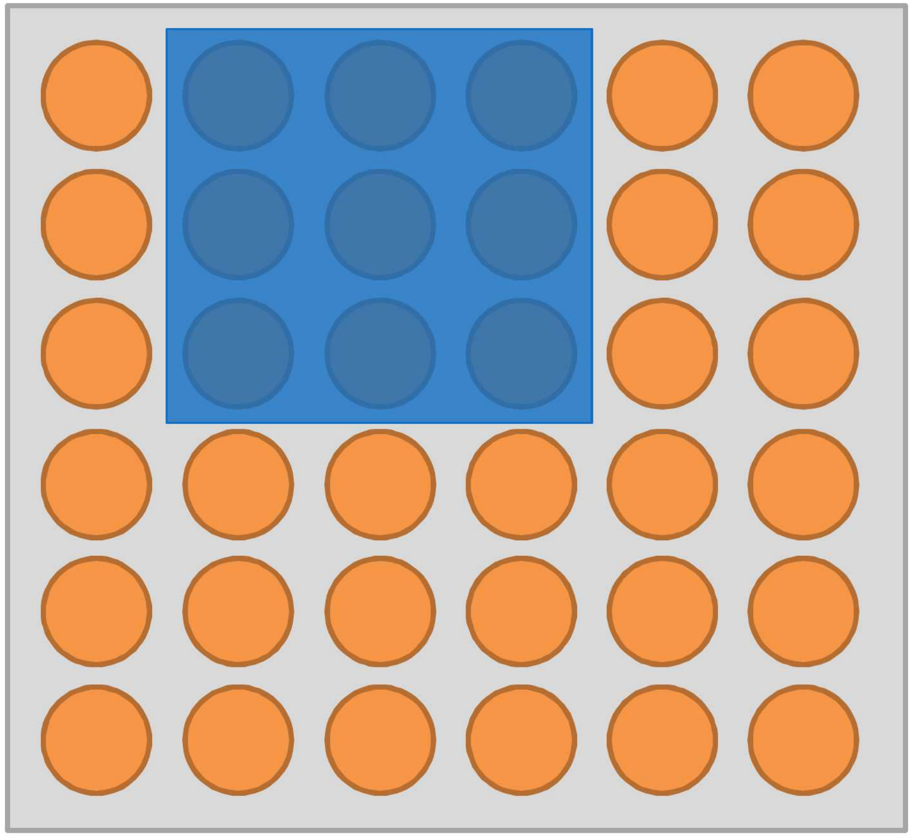

Convolutional neural networks

- Instead of a linear function (i.e. a matrix multiply) use convolutional layers
- Convolutions exploit spatial locality in pixels
- Early layers in the network learn simple features
- Later layers learn more abstract features

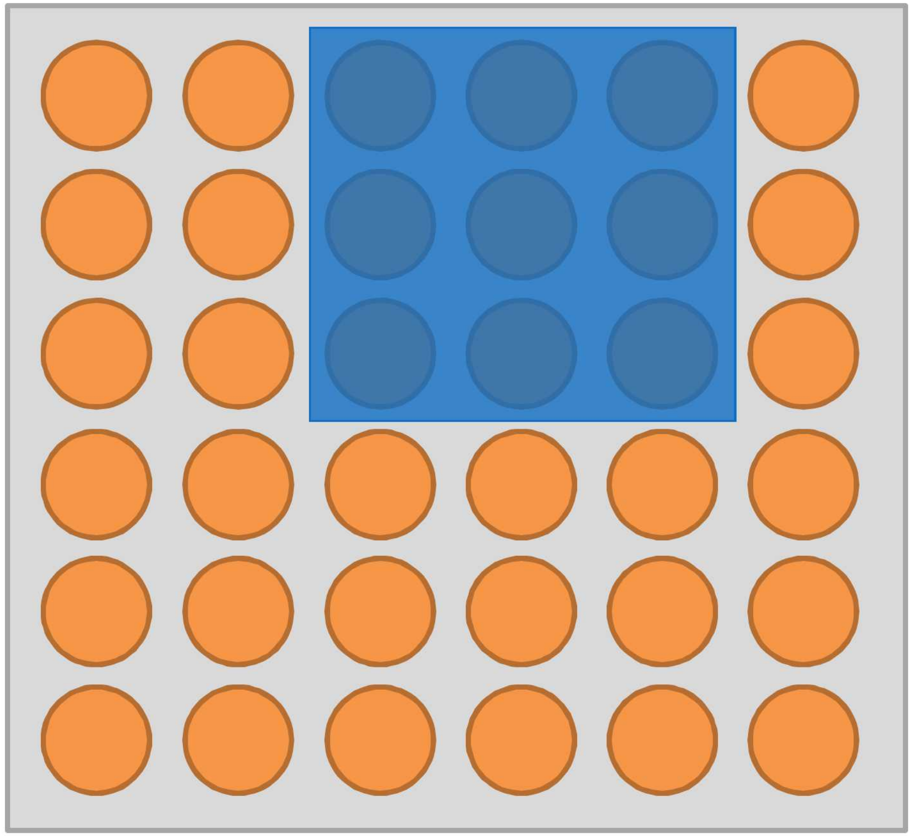

Convolutional Neural Network Layer Visualization

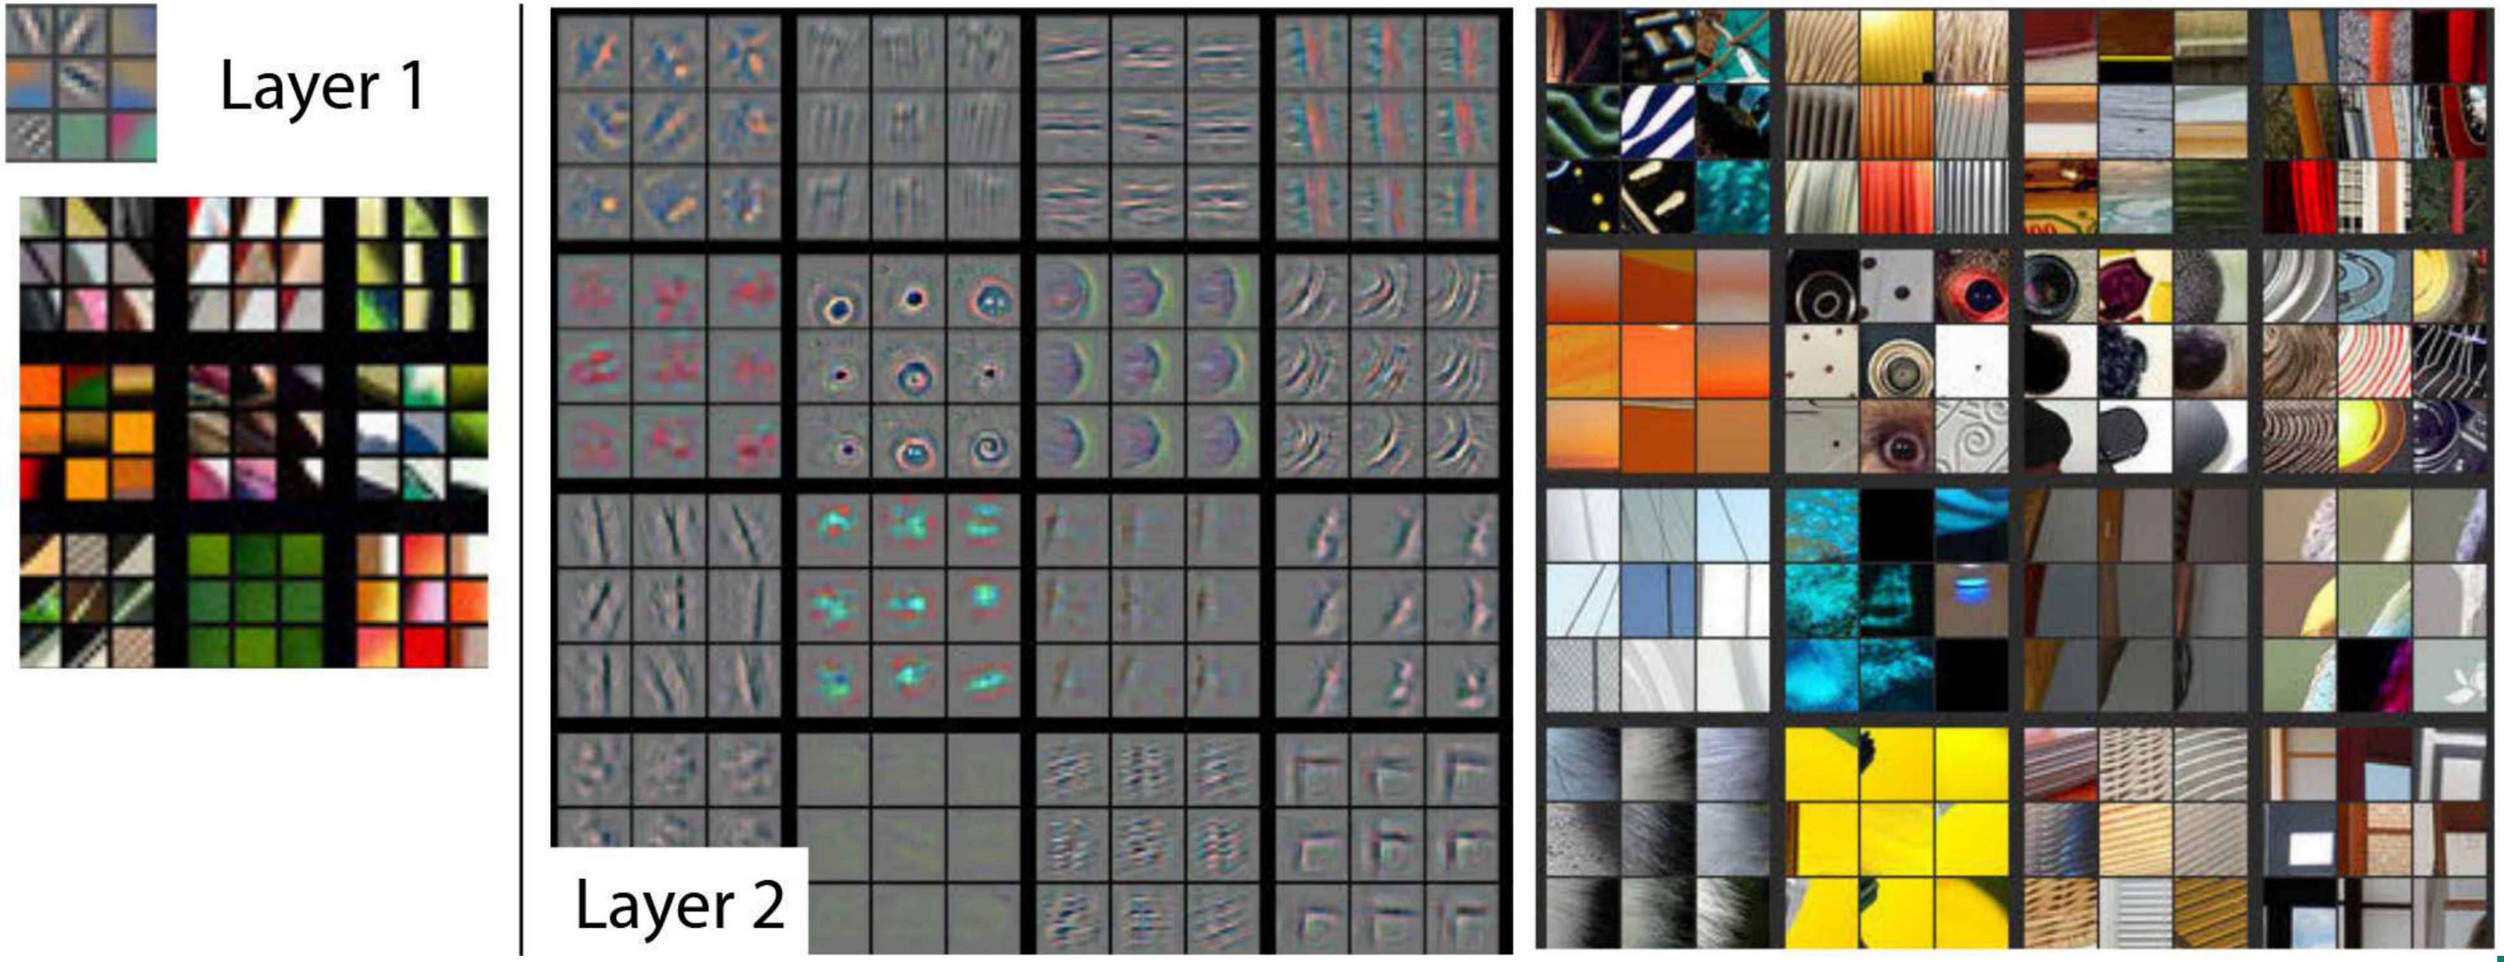

Zeiler, Matthew D., and Rob Fergus. "Visualizing and Understanding Convolutional Networks." Computer Vision – ECCV 2014 Lecture Notes in Computer Science, 2014, pp. 818–833., doi:10.1007/978-3-319-10590-1_53.

Deep Learning

- Big neural networks
- Multiple layers capture “deep” statistical information in data
- Generally require more intensive computing resources
- Include:
  - Convolutional neural networks
  - Recurrent neural networks

- Methods:
  - Supervised
  - Unsupervised
  - Semi-supervised

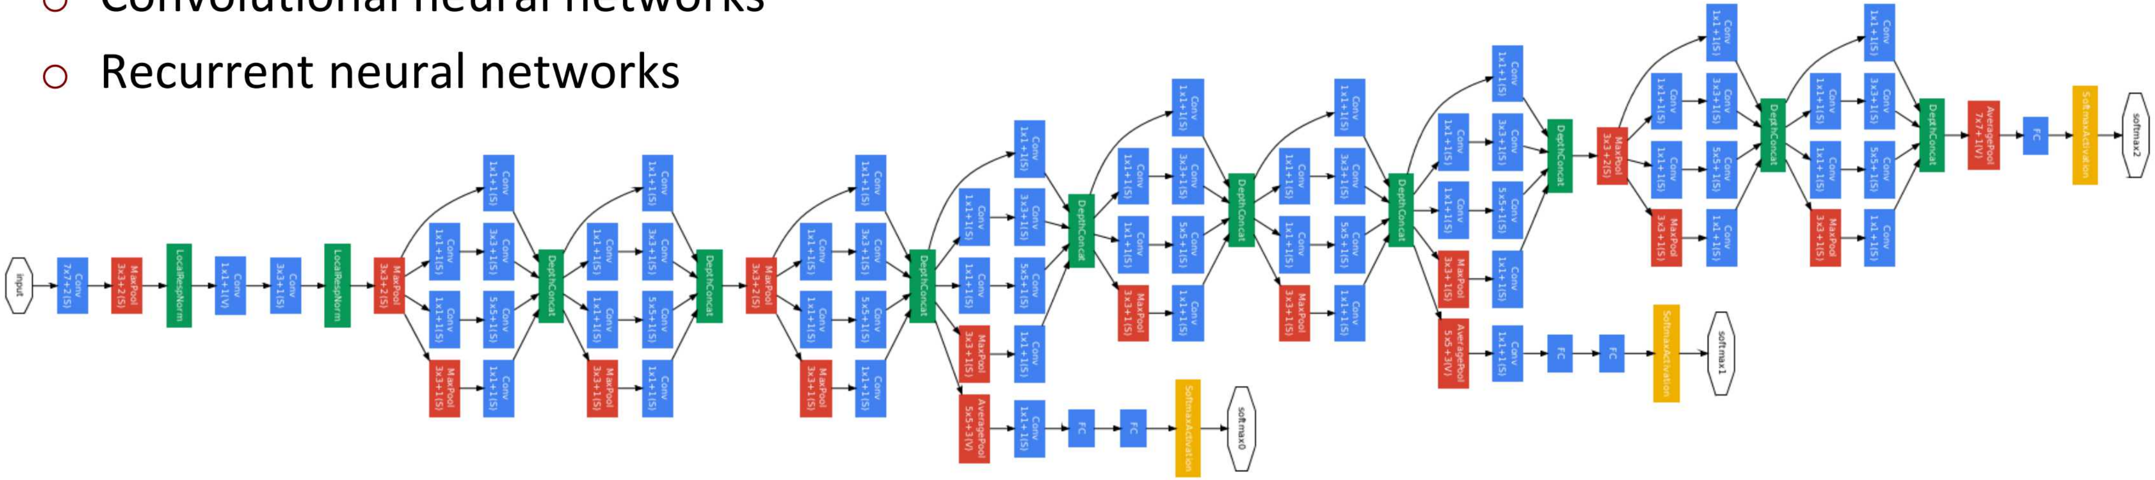

https://medium.com/coinmonks/paper-review-of-googlenet-inception-v1-winner-of-ilsvlc-2014-image-classification-c2b3565a64e7

Simple description of GANs

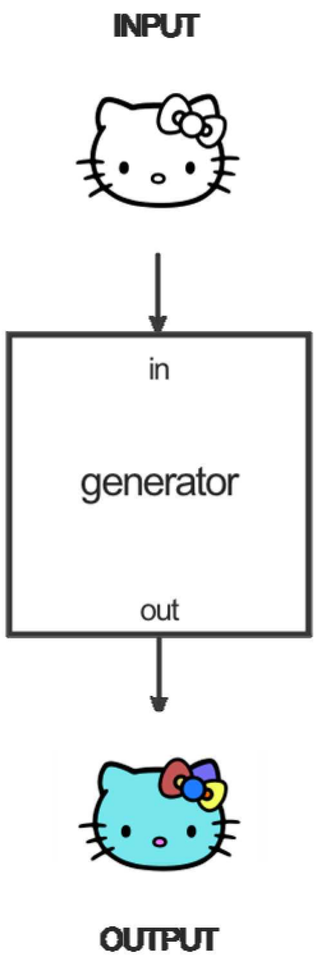

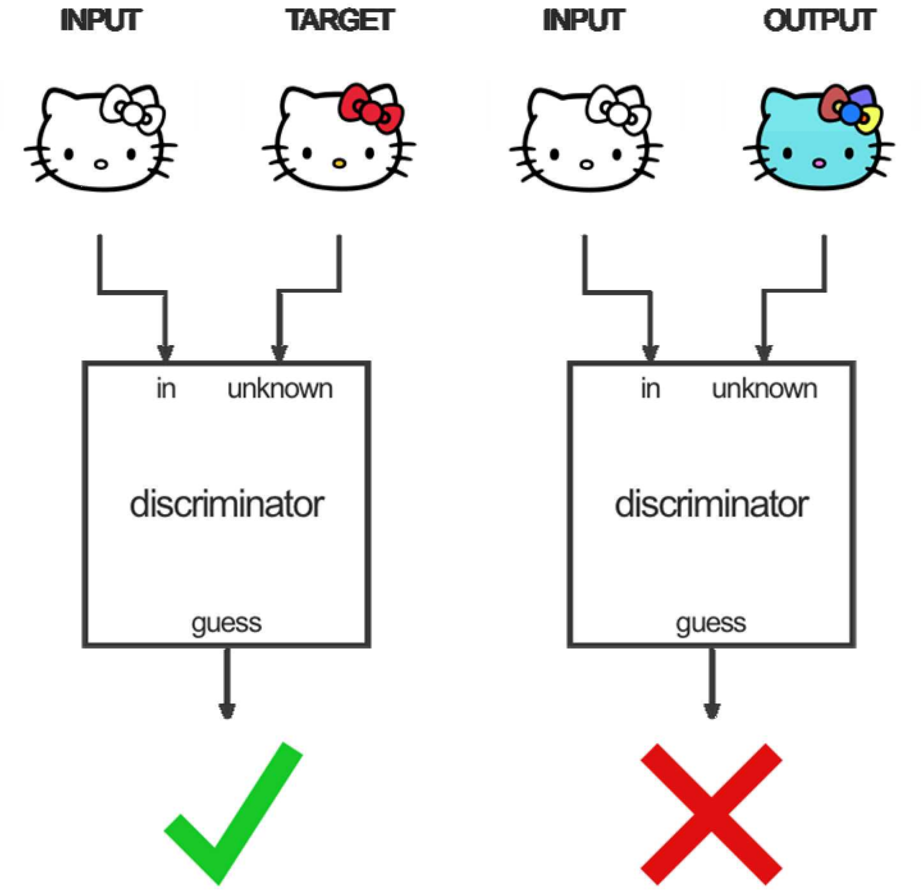

https://affinelayer.com/pix2pix/

Computer Vision

- Research field that seeks answers from digital images:
  - What are the objects in a photograph?
  - How do objects move in videos?
- Traditional techniques:
  - Hand-engineered features
    - Edge detection
    - Corner detection
  - Panoramic stitching
  - Stereo reconstruction

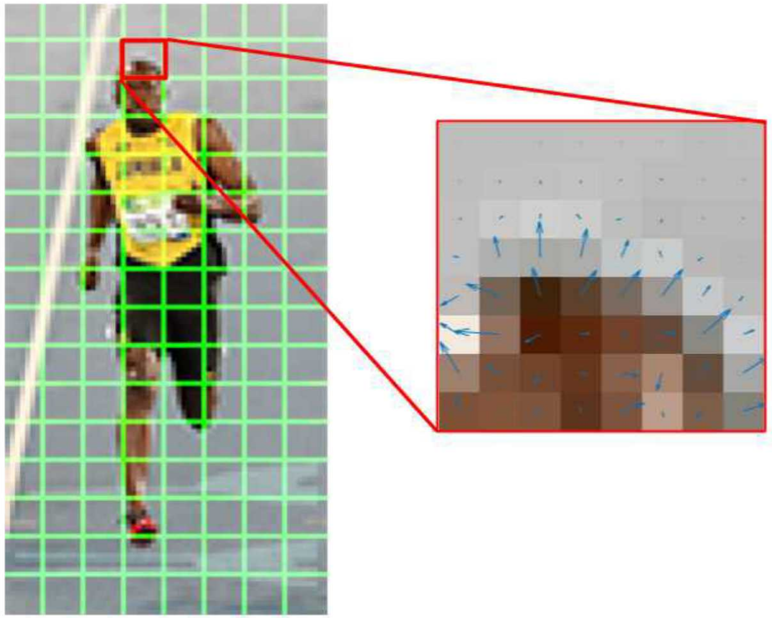

2	3	4	4	3	4	2	2
5	11	17	13	7	9	3	4
11	21	23	27	22	17	4	6
23	99	165	135	85	32	26	2
91	155	133	136	144	152	57	28
98	196	76	38	26	60	170	51
165	60	60	27	77	85	43	136
71	13	34	23	108	27	48	110

Gradient Magnitude

80	36	5	10	0	64	90	73
37	9	9	179	78	27	169	166
87	136	173	39	102	163	152	176
76	13	1	168	159	22	125	143
120	70	14	150	145	144	145	143
58	86	119	98	100	101	133	113
30	65	157	75	78	165	145	124
11	170	91	4	110	17	133	110

Gradient Direction

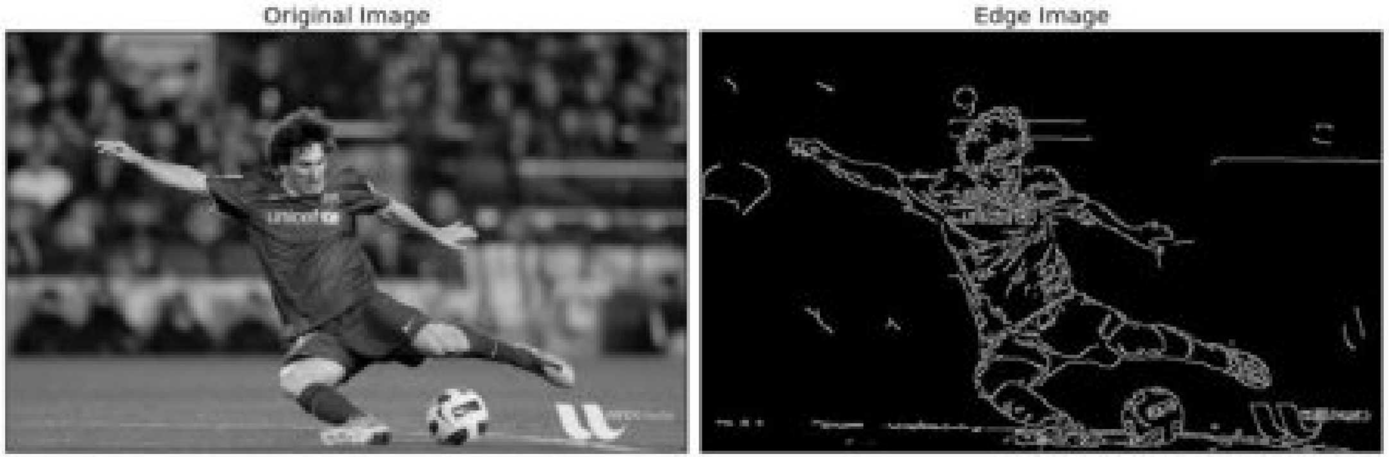

https://www.learnopencv.com/histogram-of-oriented-gradients/
https://docs.opencv.org/3.4.3/da/d22/tutorial_py_canny.html

Computer Vision & Deep Learning

- The emergence of deep learning has radically changed the field of computer vision
- Convolutional neural networks:
  - Replace manual feature engineering
  - Dramatically increase accuracy
  - Solve more difficult problems than traditional CV can answer

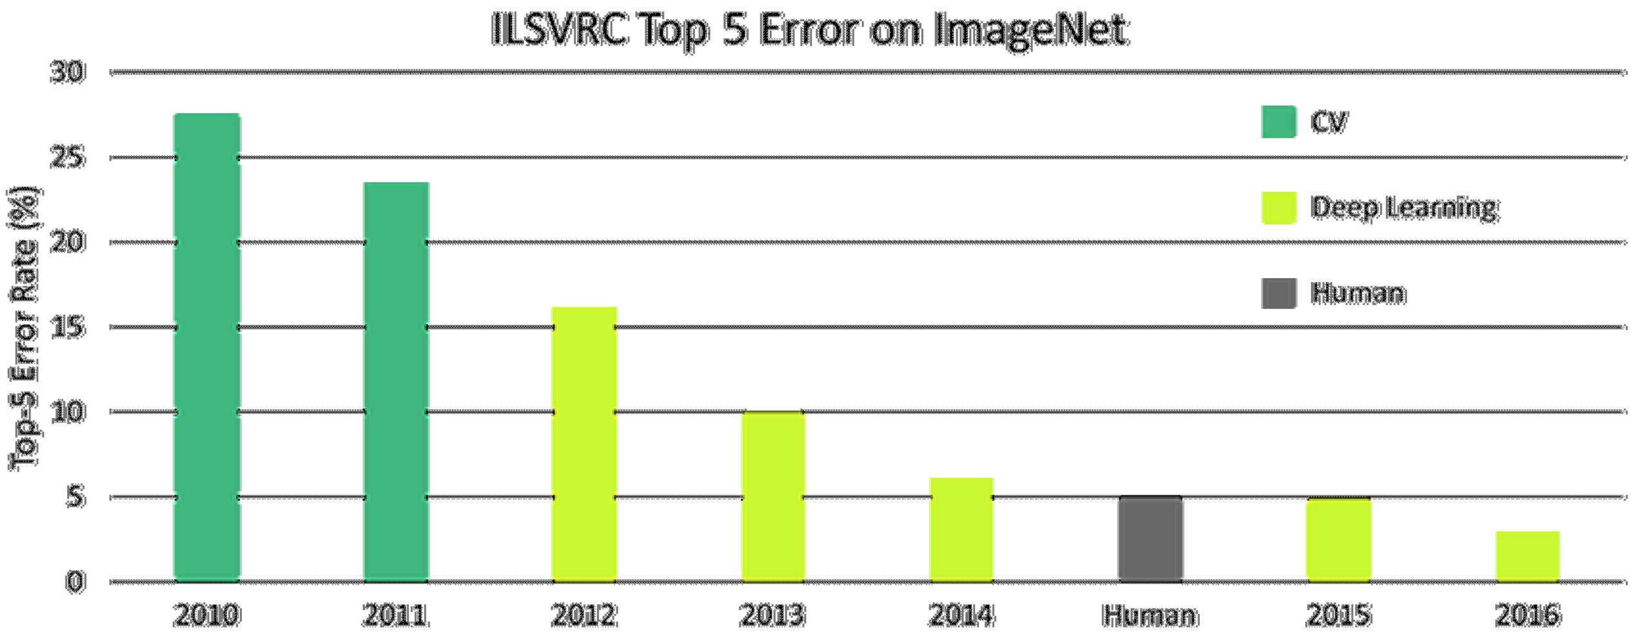

https://www.dsiac.org/resources/journals/dsiac/winter-2017-volume-4-number-1/real-time-situ-intelligent-video-analytics

CycleGAN translates images between domains

Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. arXiv preprint.

https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix

Learns two functions:
F(x) = Horse to zebra
G(x) = Zebra to horse

Cycles back to starting point to learn without paired examples

$$F(G(x)) = x$$

Generative adversarial networks are powerful new models that are widely applicable

Batteries to Foam — Leveraging existing datasets via CycleGAN

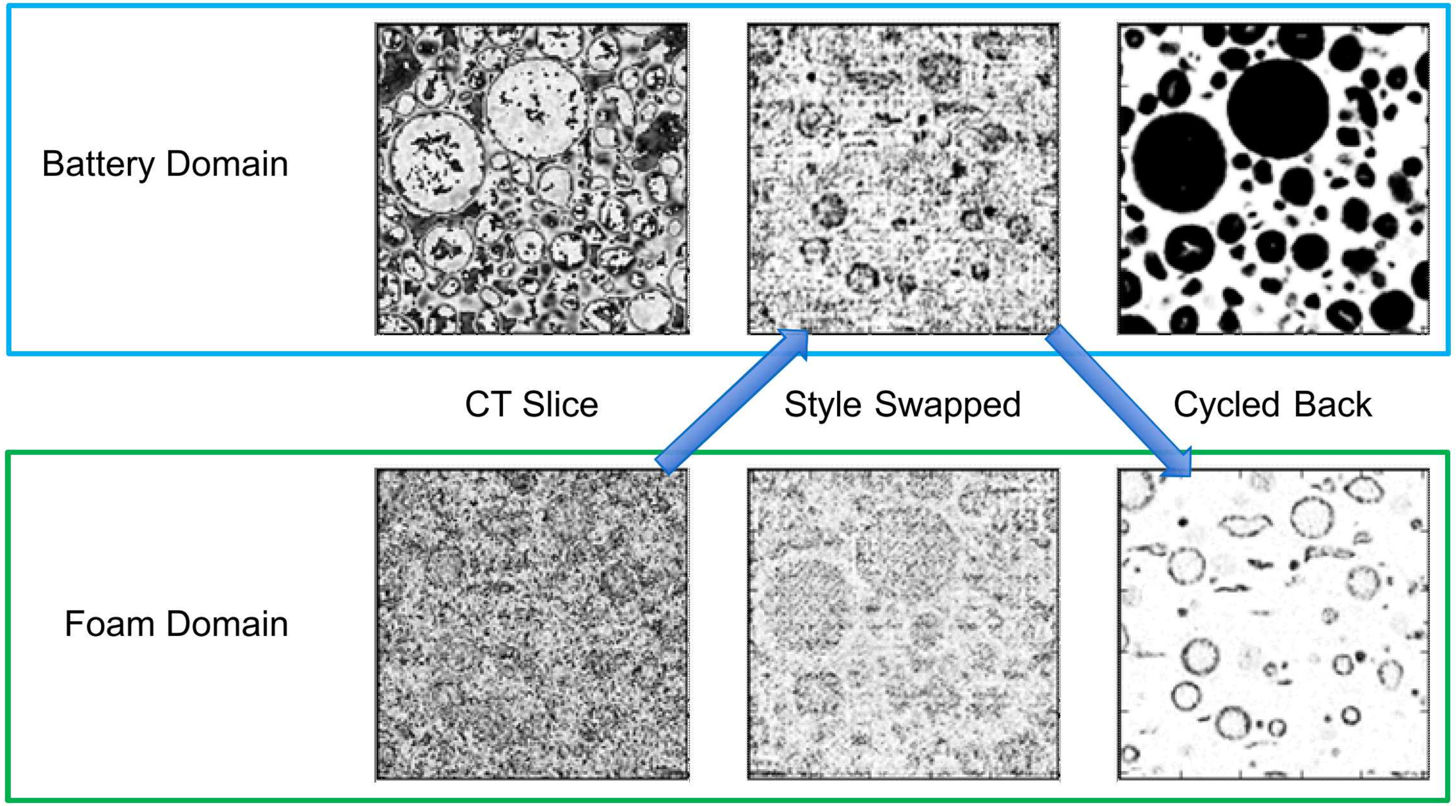

Domain Adaptation could reduce supervised labeling cost

Repurpose labels from one domain (battery) to another domain (glass micro balloon foam)

- CycleGAN transforms foam CTs into the "style" of battery labels
- Semi-supervised

- Hand labeled small slices from 7 CT scans of foam
- Used 2 labels to select stopping point
- Inferred over remaining 5 volumes
- Post-process (fill in gaps) with standard CV methods
- Average 94.8% accuracy when compared with human labeled slices