This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

SAND2022-5129C

How to Train Your Digital Twin: Practical Deep Learning Approaches to Modeling As-built Components

Cari Martinez

April 26, 2022

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.

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.

How do we build digital twins?

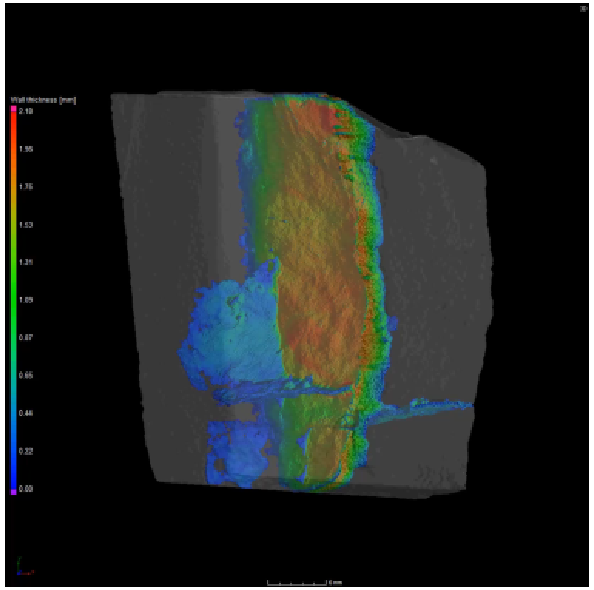

Challenges:

- Idealized models fail to capture impact of defects.
- Overwhelming amounts of data cannot be manually analyzed.
- Uncertainty lurks everywhere.
- Limited or incomplete data availability constrains AI approaches.

Practical approaches:

- Nondestructive multimodal data collection can characterize systems.
- Deep learning models can ingest all the data to predict a range of system properties.
- Domain knowledge can be incorporated to improve models and interpret predictions.

Emerging capabilities enable digital twins

Advances in computer vision
- Feature-based Anomaly Detection System (FADS)
- Volumetric segmentation

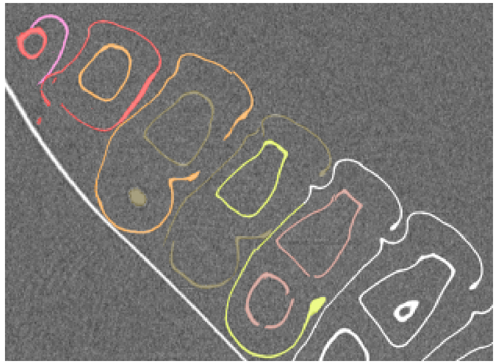

Uncertainty quantification for image-based simulation
- Efficient Quantification of Uncertainty in Image-based Physics Simulations (EQUIPS) workflow

Beyond Fingerprinting
- Physics-Informed Multimodal Autoencoders (PIMA)

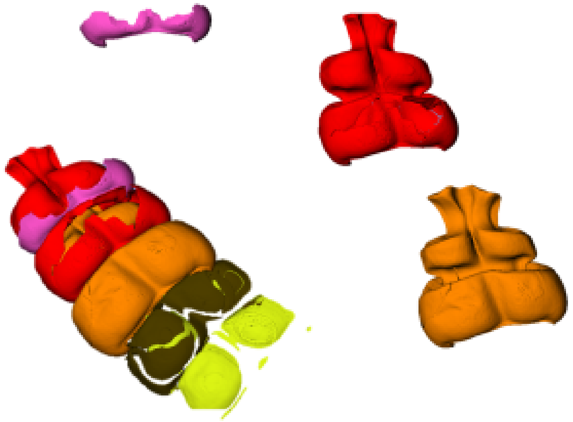

Automatic anomaly detection in high reliability as-built parts from images

Kevin Potter, Anthony Garland

Feature Based Anomaly Detection System (FADS)

Anthony Garland

Extends the deep-one-class classification idea

Works via a **pretrained network** to provide the mapping

- **Record** the activations from the model's convolutional filters
- **Aggregate** each filter to a single value (max, min, mean, etc.)
- Develop a **statistical** model of expected convolutional activations ($\bar{x}$ and σ) based on the nominal images' activations
- At inference time, measure activation of the input image from the same model and **normalize** based on the nominal statistics

Example nominal data (hazelnuts)

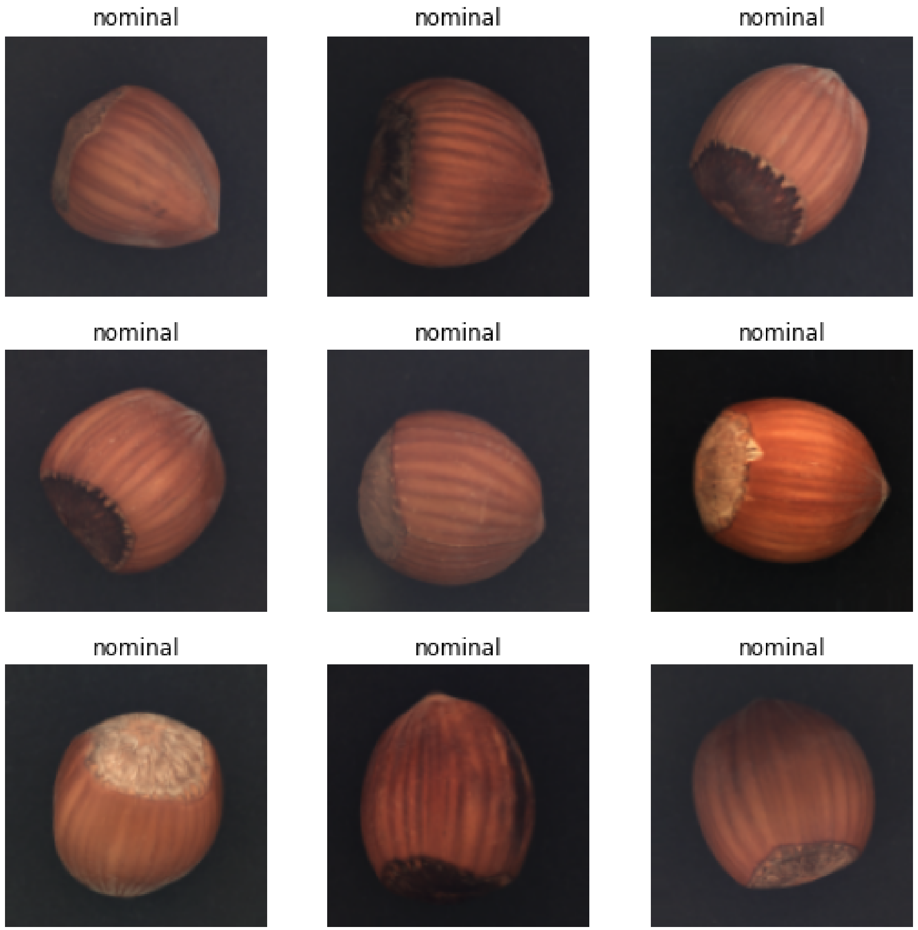

CNN activations

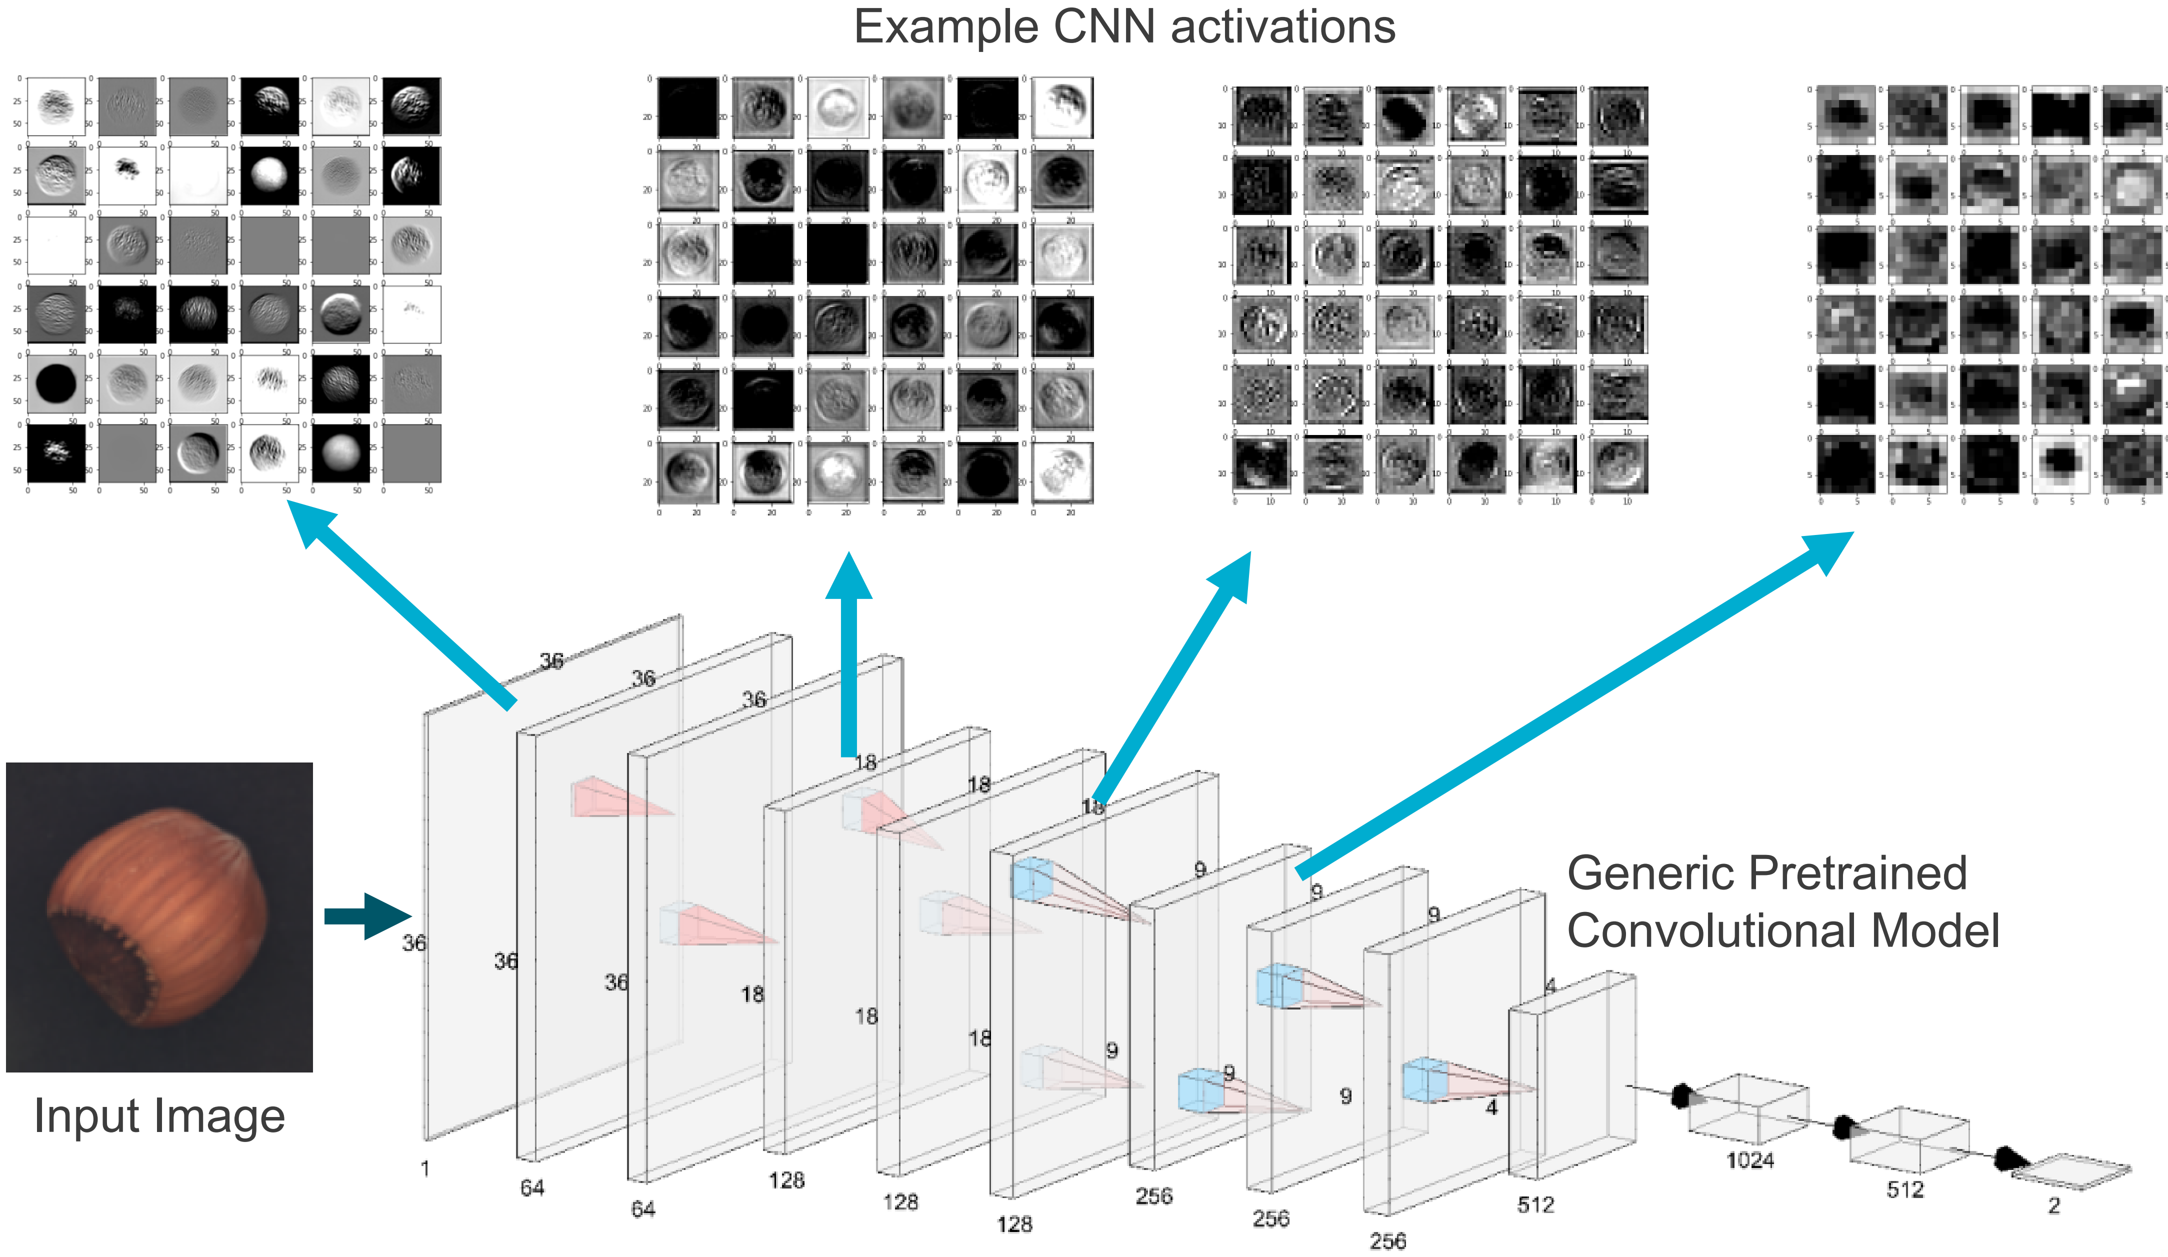

Aggregate activations

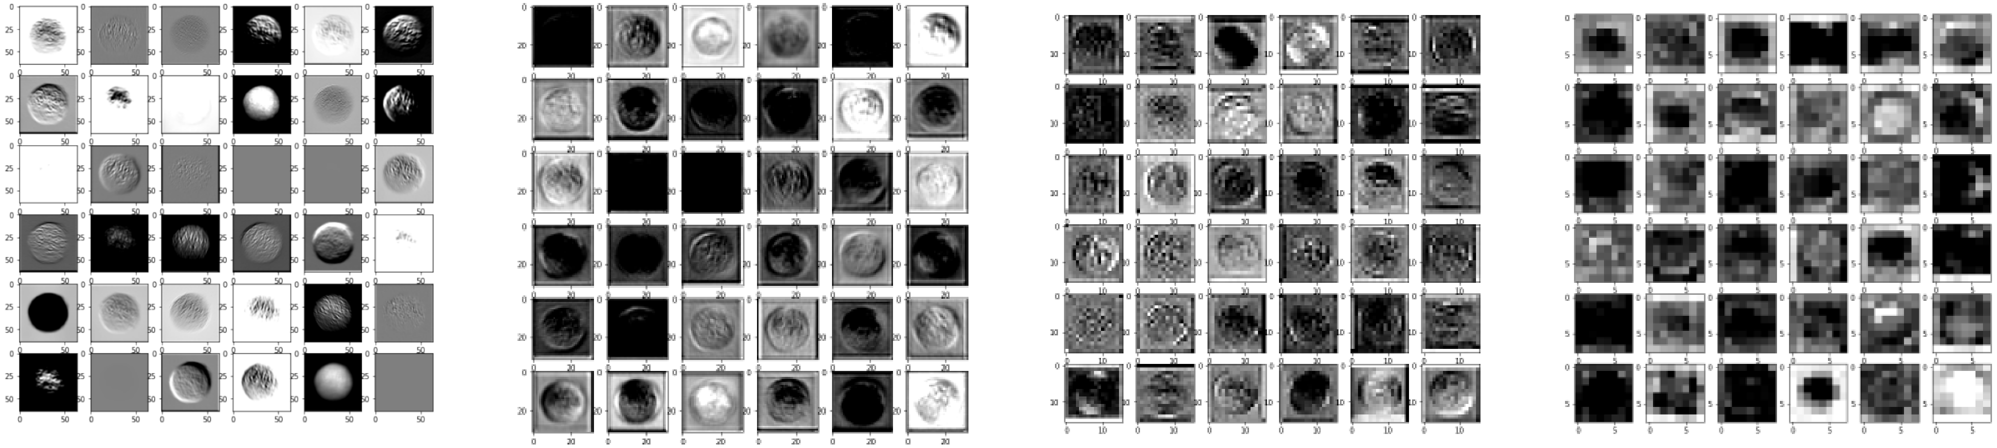

Collapse each filter's activations to a single value and stack

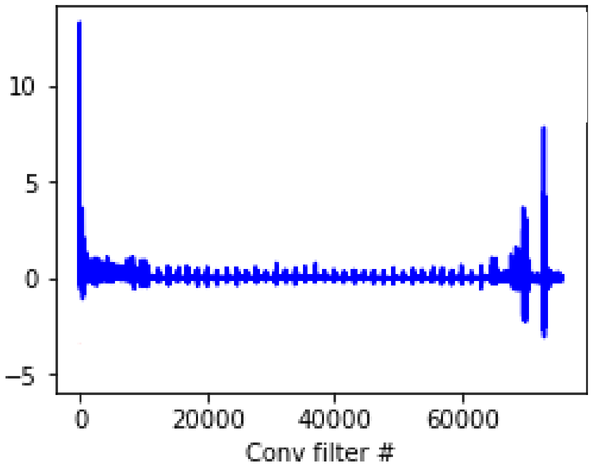

Learn the nominal datasets activation stats for each filter

Nominal images
("training" set)

Generic Pretrained Convolutional Model

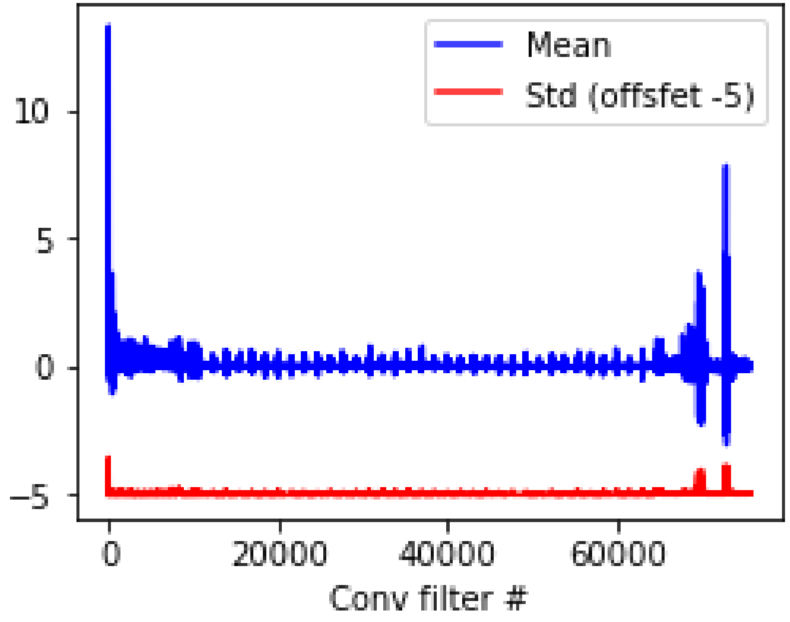

2-D FADS example – inference normalization

Input image

Generic Pretrained Convolutional Model

Feature activations

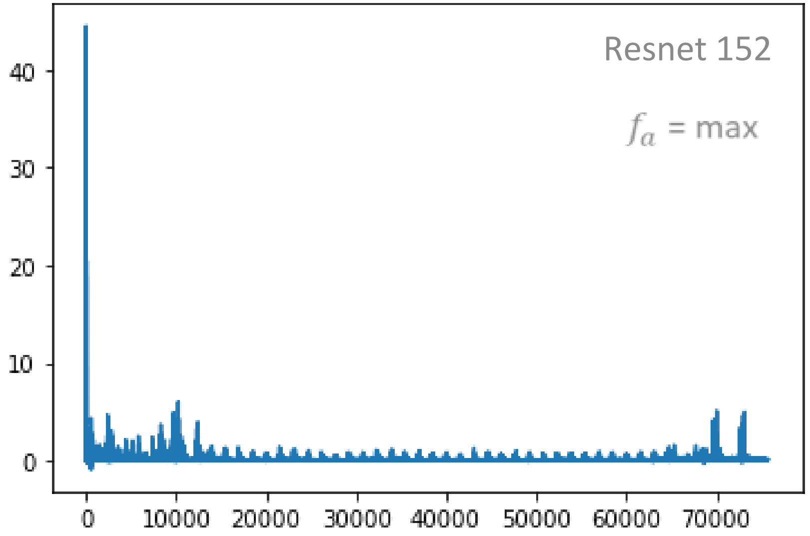

Normalize against "training" set

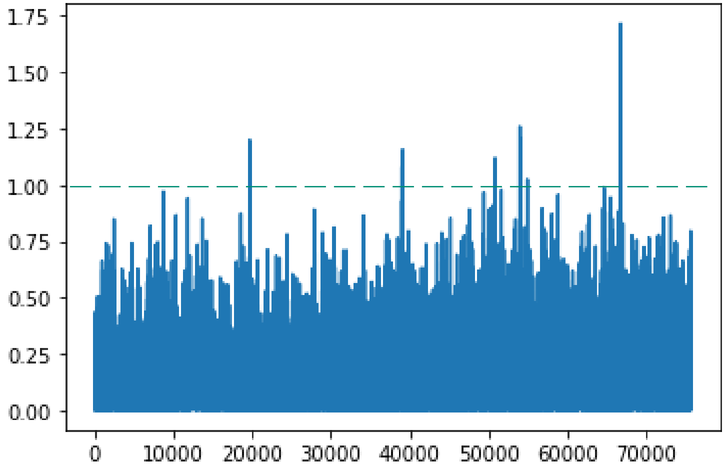

R-vector (standard deviations from nominal mean)

2-D FADS example – converting to a threshold

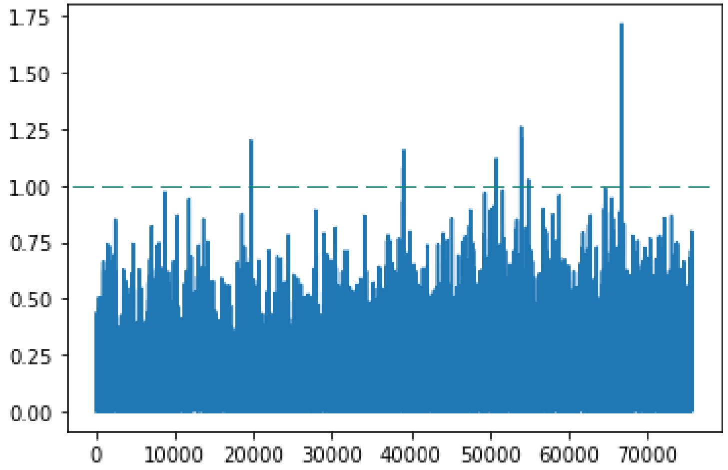

R-vector (standard deviations from nominal mean)

Scoring Function

$$s_m = \max(r)$$

Or

$$S_p = percentile(r, 90)$$

Anomaly Scores

1.75

Or

1.1

FADS on MVTec AD dataset (whole image)

ROC for pill category

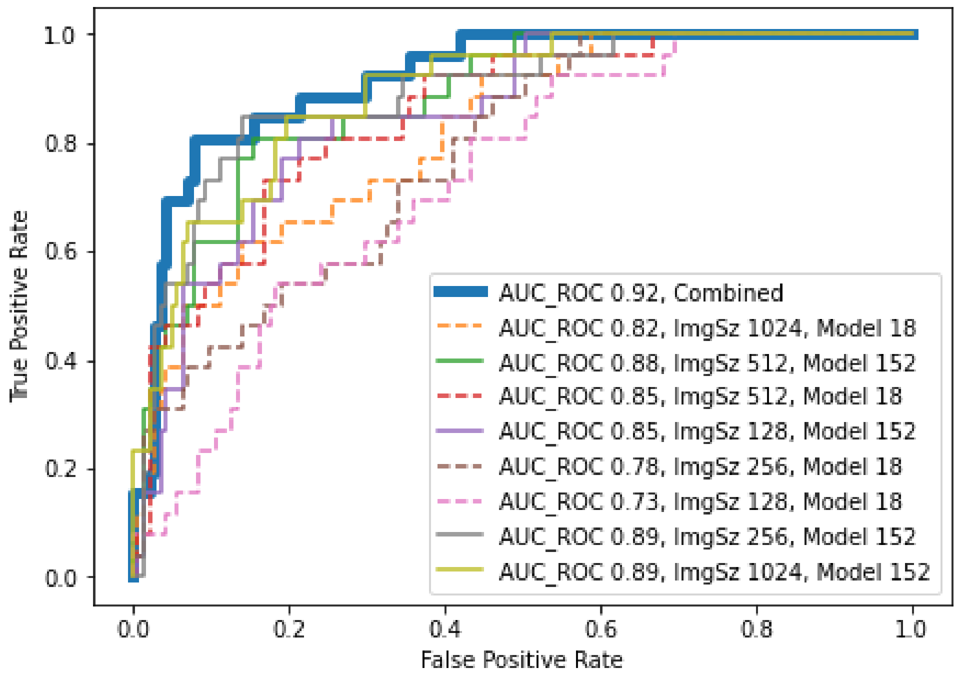

FADs AUC score on each class

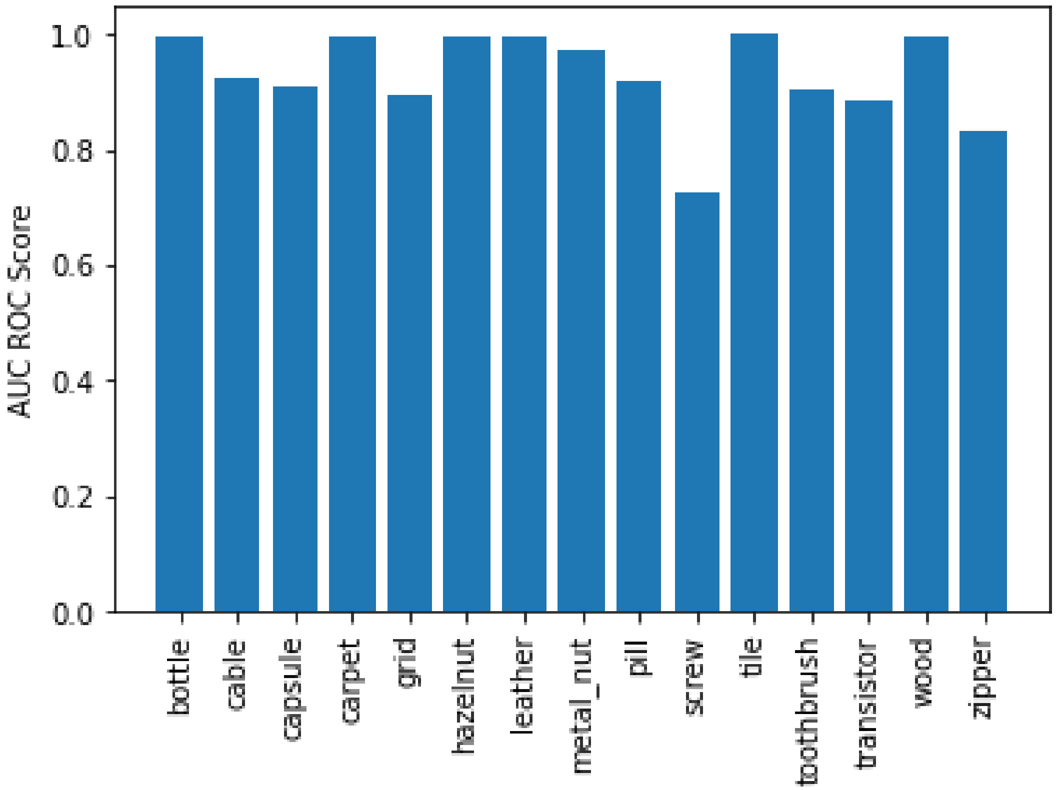

FADs achieves an average AUC of 0.93

FADS can also highlight the anomalies

By taking the gradient to minimize the anomaly score with respect to the input, the pixels that contribute to anomalousness are highlighted

- Very sensitive – wood image picks up a scuff that is barely visible for instance
- Still fast as it uses a single backward pass


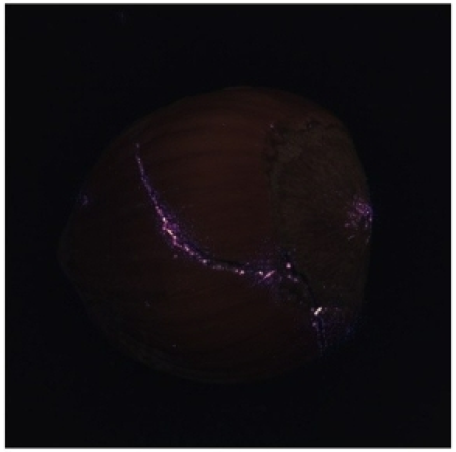
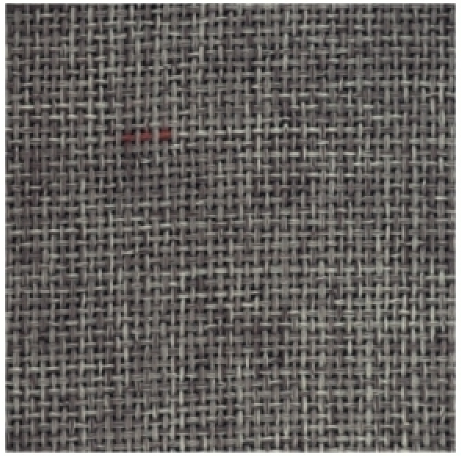
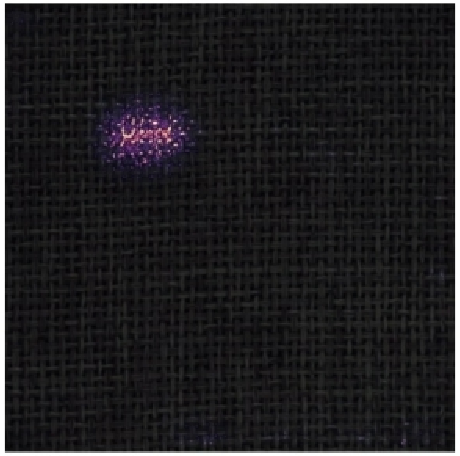

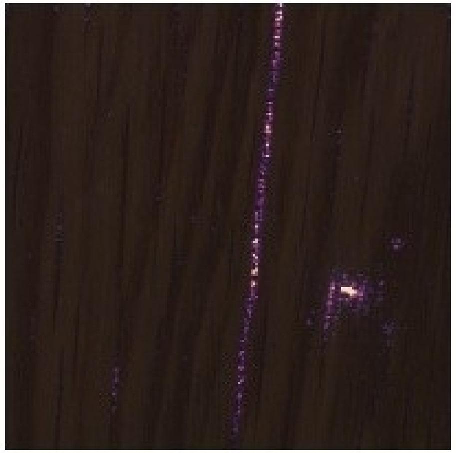

FADS can also highlight the anomalies

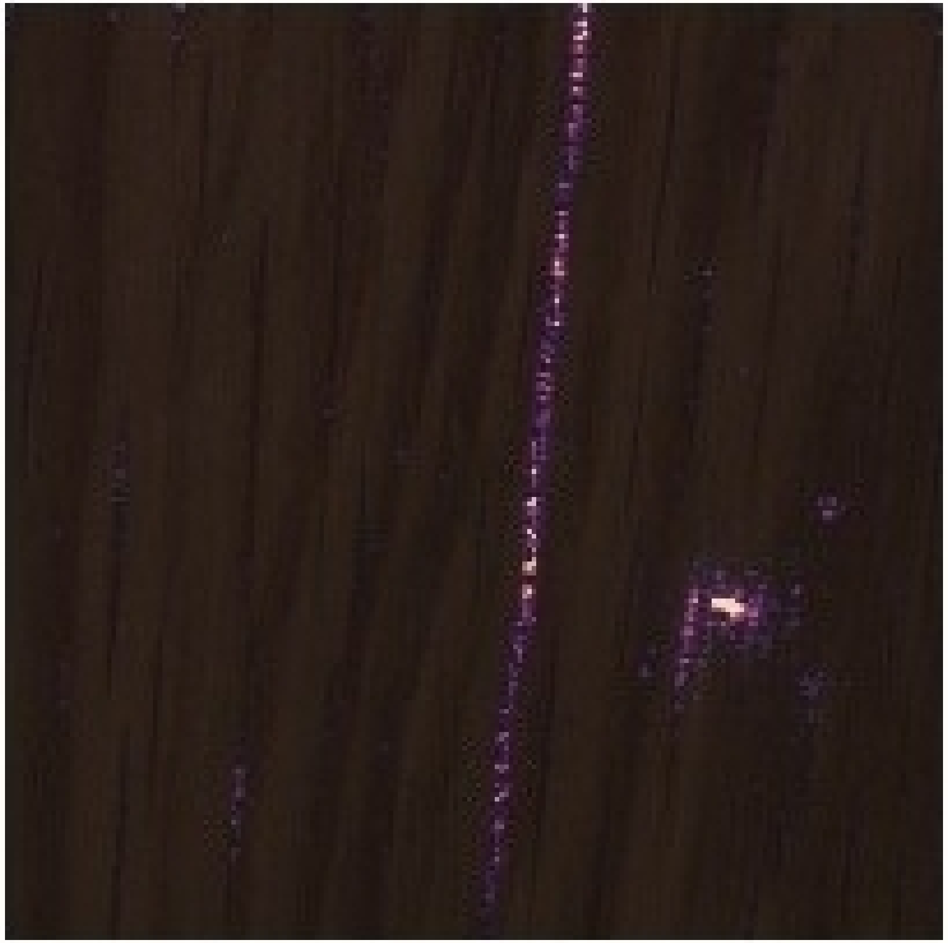

FADS against a real world application

Additive Manufacturing

(right) Localizing flaws in real prints

(below) Dedicated print testing: Using just images, identify defective parts with incorrect print process settings

- "Trained" on 18 lattices
- Result: Avg AUC of 0.99

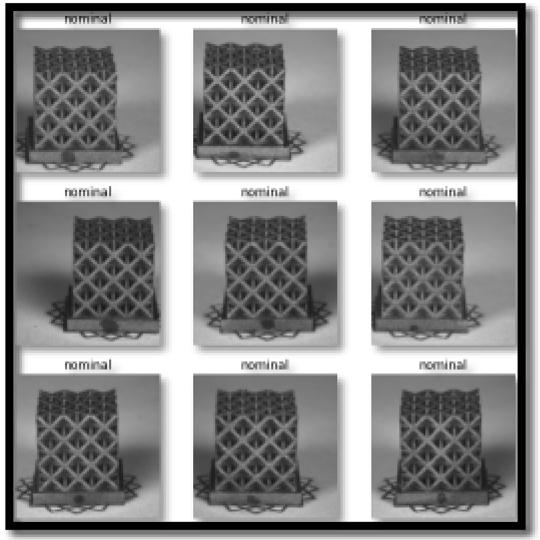

3-D Printed Lattice

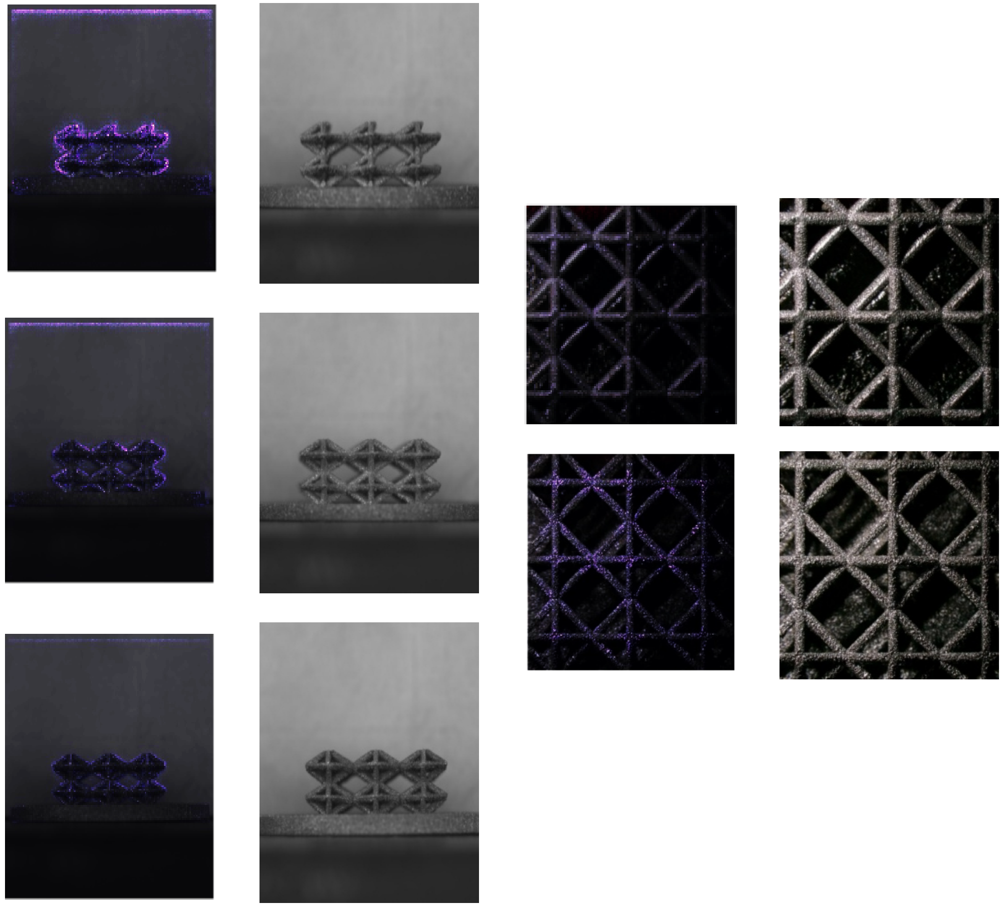

Visualization of the regions causing high anomaly scores

FADS key insights

- Powerful transfer learning from models pretrained on massive, unrelated datasets
- Features relevant to separate normal from abnormal examples highlighted without supervision
- Limited number of training examples adequate for high accuracy

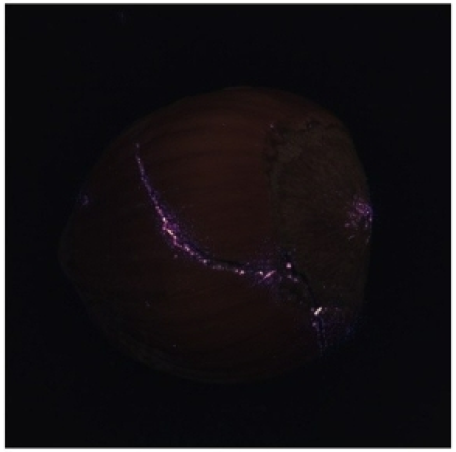

Credible Automated Meshing of Images

Scott Roberts, PI

Credible, Automated Meshing of Images (CAMI) LDRD

Raw greyscale image (XCT)

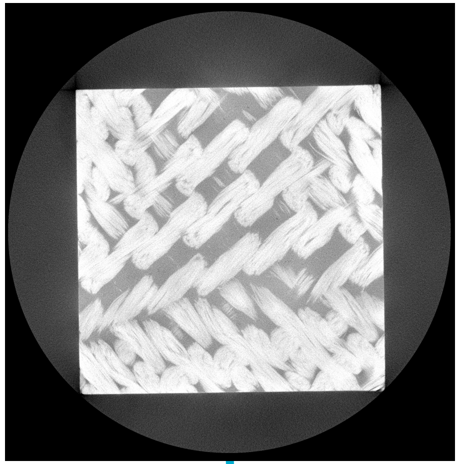

Surface mesh (STL)

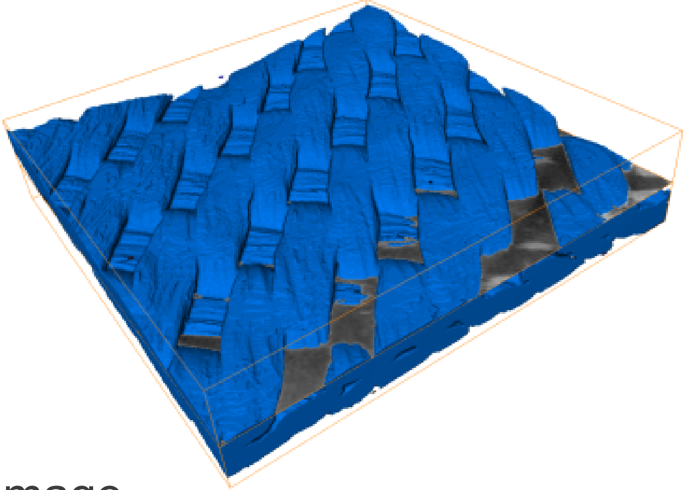

Meshing:
- CDFEM + snap + Emend
- High quality

Physics simulation

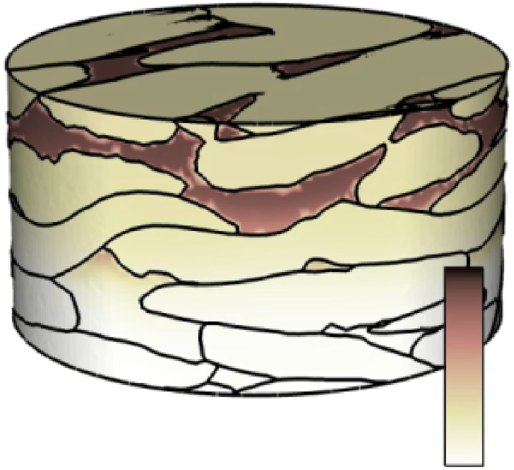

Segmentation:
- Automated: deep learning
- Repeatable

Segmented image

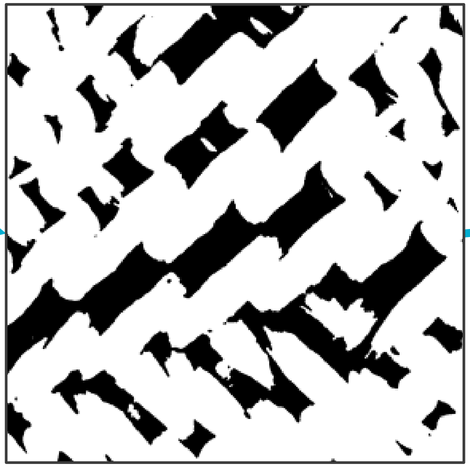

Interface Identification:
- Automated
- Marching cubes on smooth data

Volume mesh

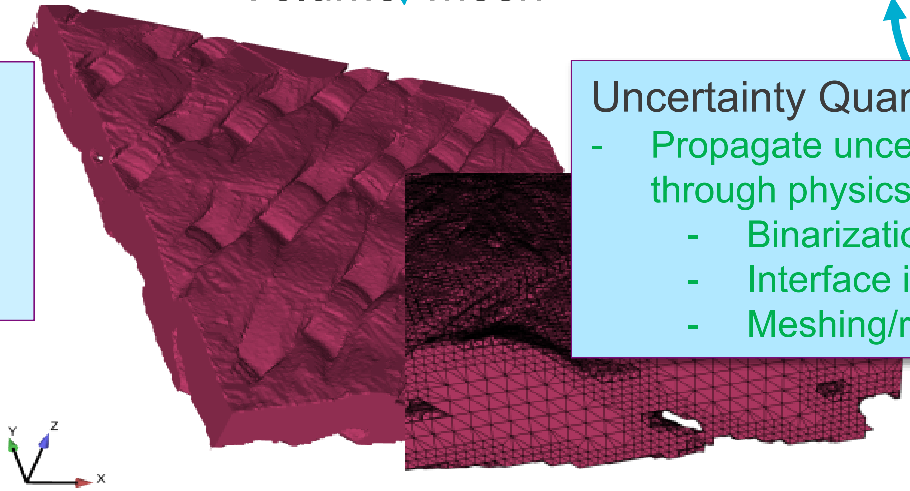

Uncertainty Quantification:
- Propagate uncertainty through physics predictions
  - Binarization
  - Interface identification
  - Meshing/resolution

Credible, automated image-based simulation will revolutionize engineering analysis of as-built parts.

Segmentation is a classic computer vision problem

Image segmentation is well studied
- Small files
- Large training sets

CT segmentation is different
- Volumetric; larger files
- Class imbalance (lots of background)
- Noise/artifacts in scans
- Small training sets with “bad” human labels
- Inconsistent scan quality (domain shift)

Medical researchers are leading this work toward Deep Learning solutions

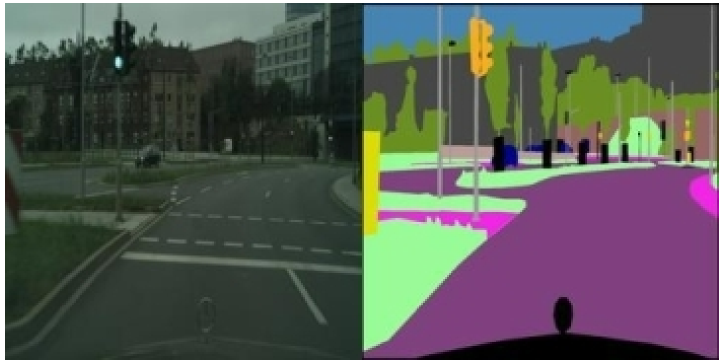

https://www.cityscapes-dataset.com/

Cityscape
(~1e5 pixels)

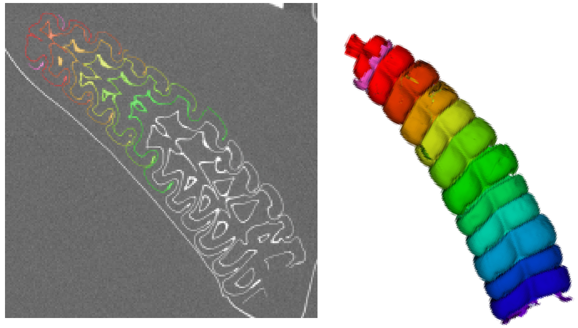

Rattlesnake Tail
(~1e9 voxels)

Mitigating challenges

Volumetric; larger files
- Train with random subvolumes

Class imbalance (lots of background)
- Loss function sets weights inversely with class fraction for each subvolume
- Normalization methods can separate foreground from background

Noise/artifacts in scans
- Convolutional Neural Network (CNN) architecture learns to recognize shapes

Small training sets with “bad” human labels
- 1-3 volumetric training examples is often sufficient
- Errors in labels are overcome if the errors are inconsistent

Inconsistent scan quality (domain shift)
- Use UQ to drive corrections to predictions
- Cari’s dissertation!

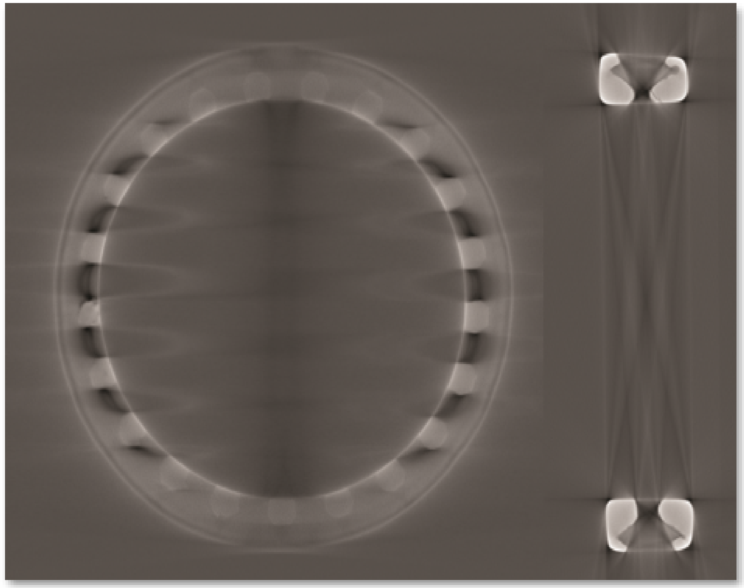

Supervised: Encoder-decoder network with skip connections

Encoder learns features at different resolutions

Decoder uses encoded features passed via skip connections for segmentation

U-net: significant advance for 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 to handle 3D images

- 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

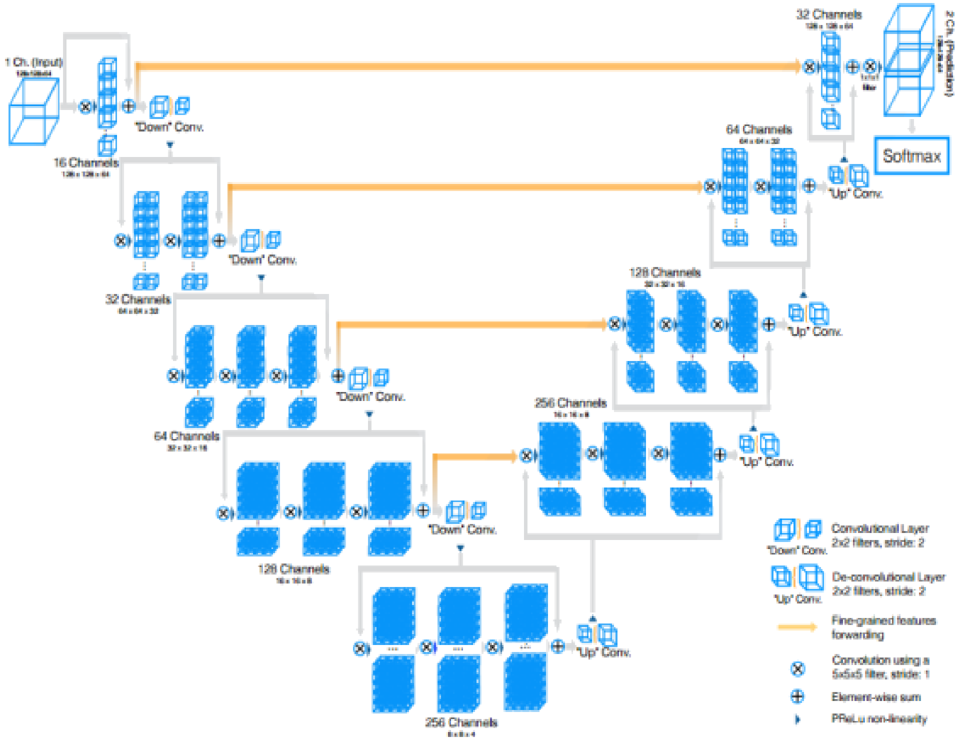

V-Net architecture for segmenting volumetric data (2016)

Example: Train V-Net to segment batteries

DOMAIN NAME	ACCURACY
E35	0.984
Tesla	0.973
Litarion	0.966
25R6	0.955
Electrode_I_1	0.948
Electrode_III_1	0.945
GCA400	0.928
Electrode_IV_1	0.917
Electrode_II_2	0.902
GCA2000	0.900
Electrode_I_2	0.892
Electrode_III_2	0.773
Electrode_IV_3	0.748
Electrode_IV_2	0.745
Electrode_II_3	0.699
Electrode_III_3	0.668
Mean	**0.8714375**

TRAINING SET

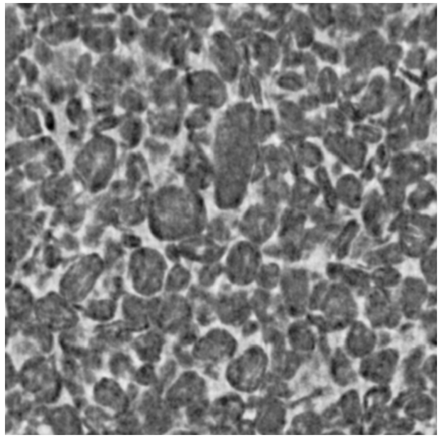
Litarion

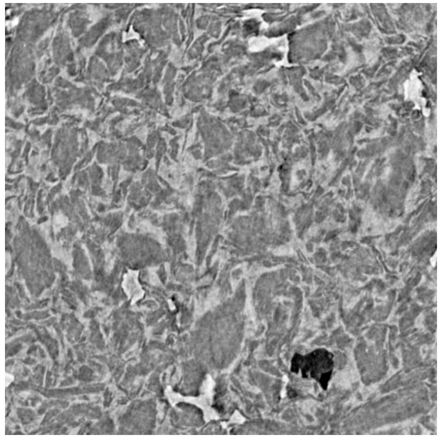
Electrode IV_1

TEST SET

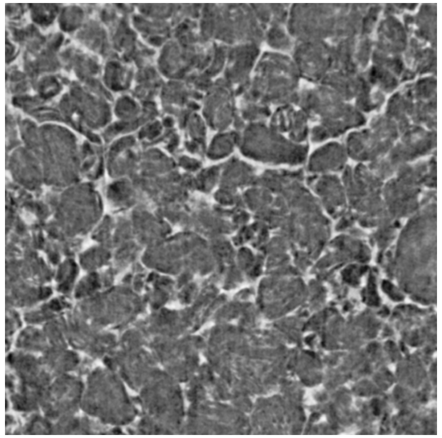
E35

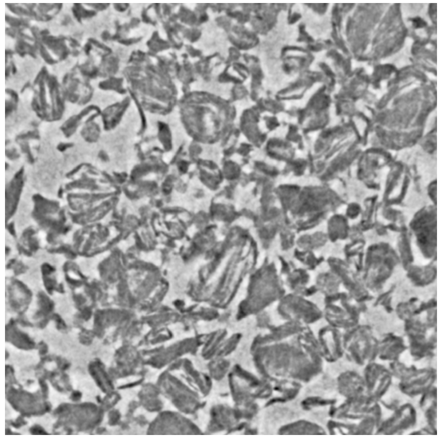
GCA400

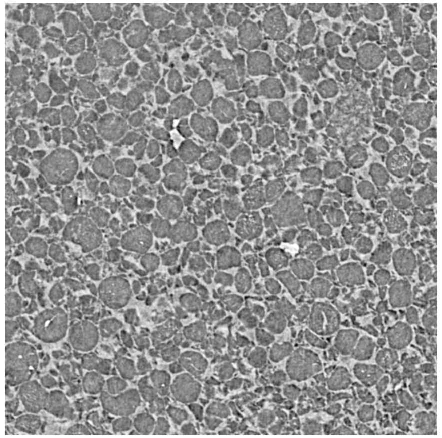
Electrode II_3

Inference results in training domain are as expected

Litarion CT scan slice

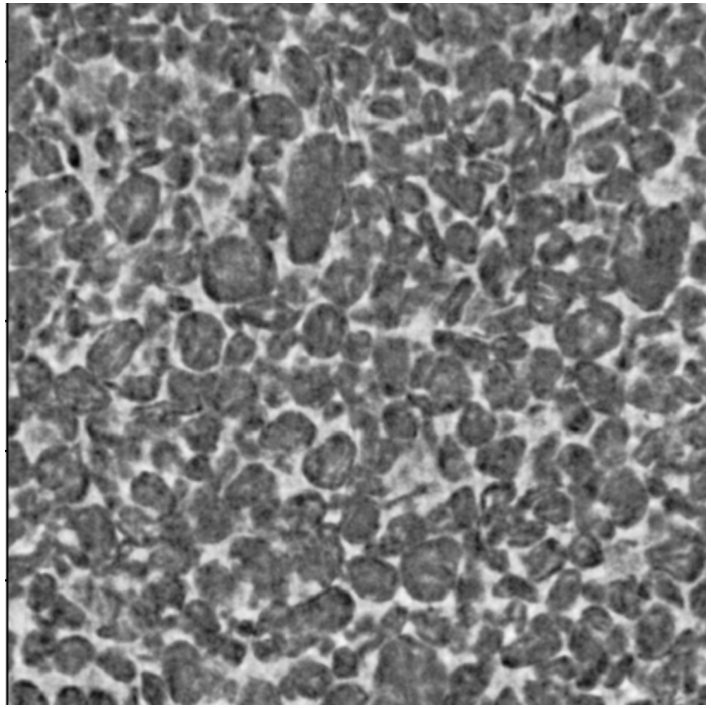

Human label

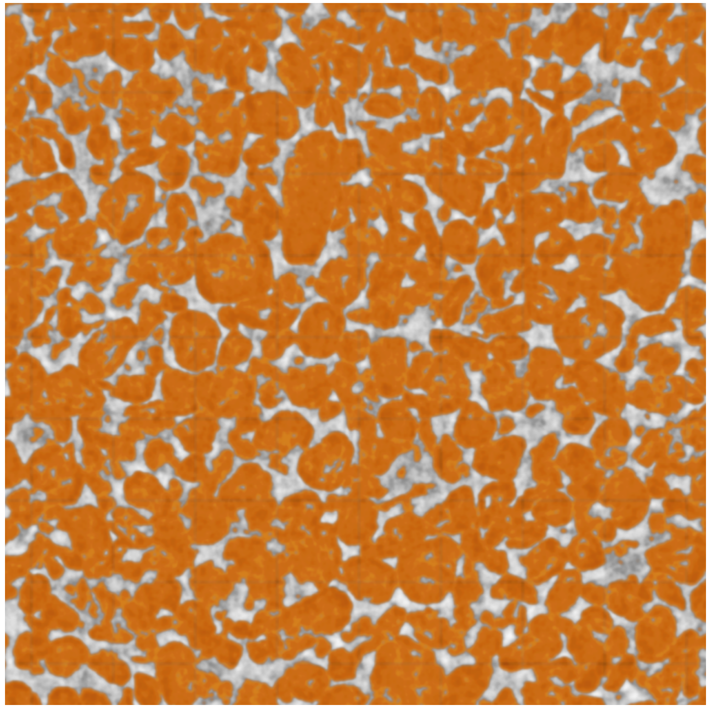

ML prediction

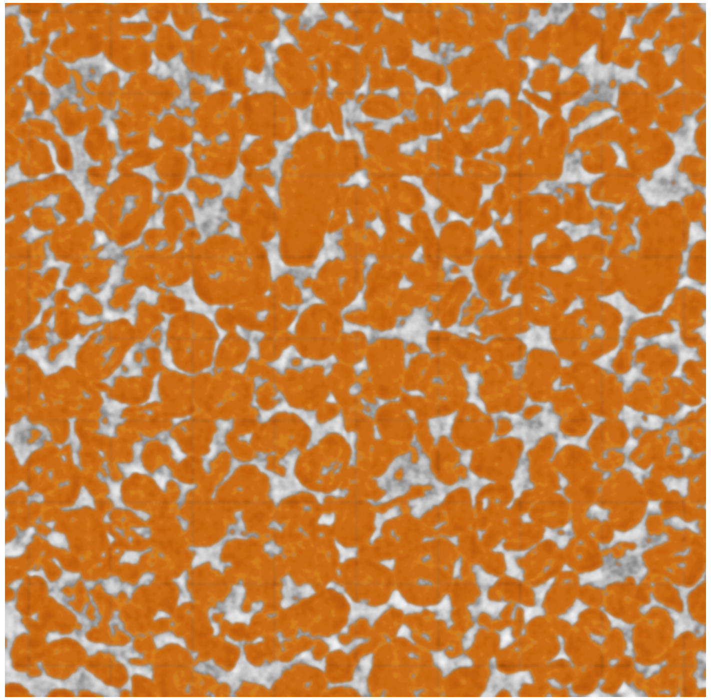

ML segmentation is 96.6% accurate to the human label

Inference results outside the training domain are qualitatively better than accuracy measurements indicate

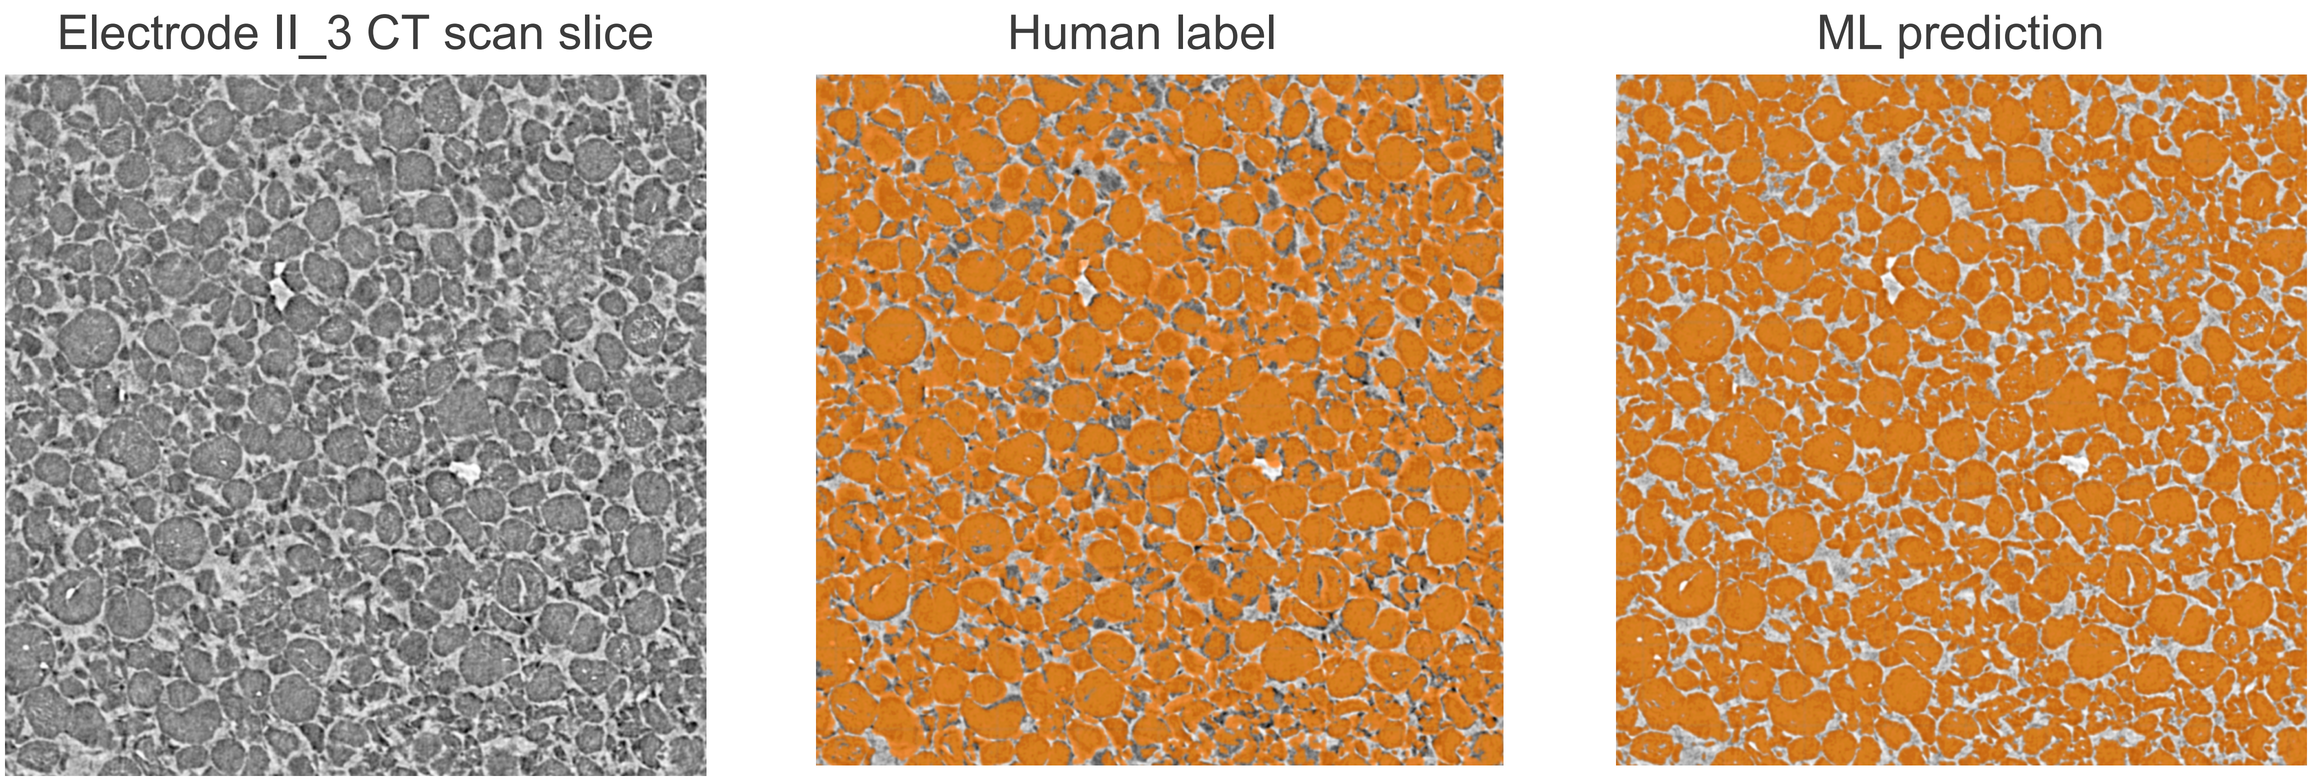

ML segmentation is 69.9% accurate to the human label…but looks qualitatively better

Alternative approach: 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.

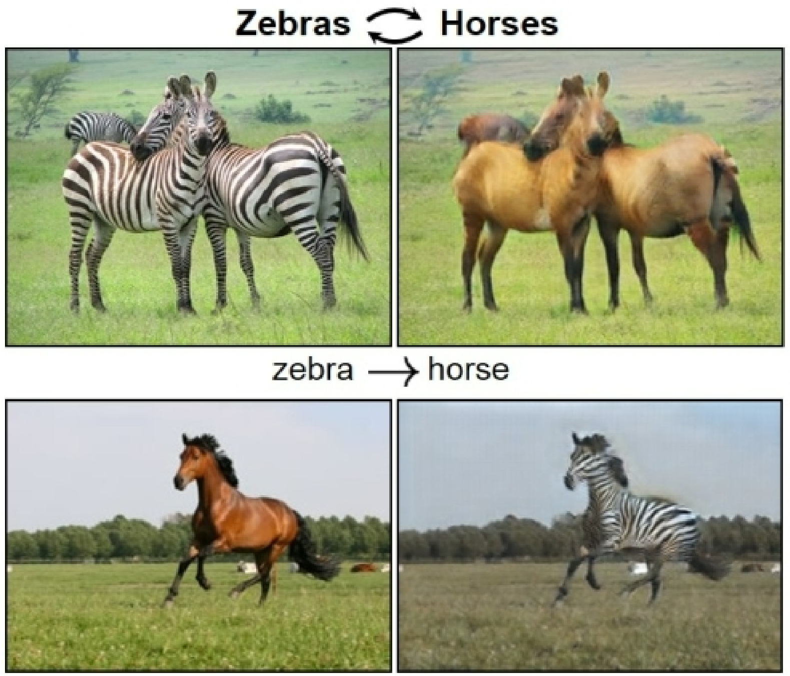

Learns two functions:

F(x) = Horse to zebra

G(x) = Zebra to horse

Cycles back to starting point to learn without paired examples

F(G(x)) = x

CycleGAN translates between material domains capturing relevant features

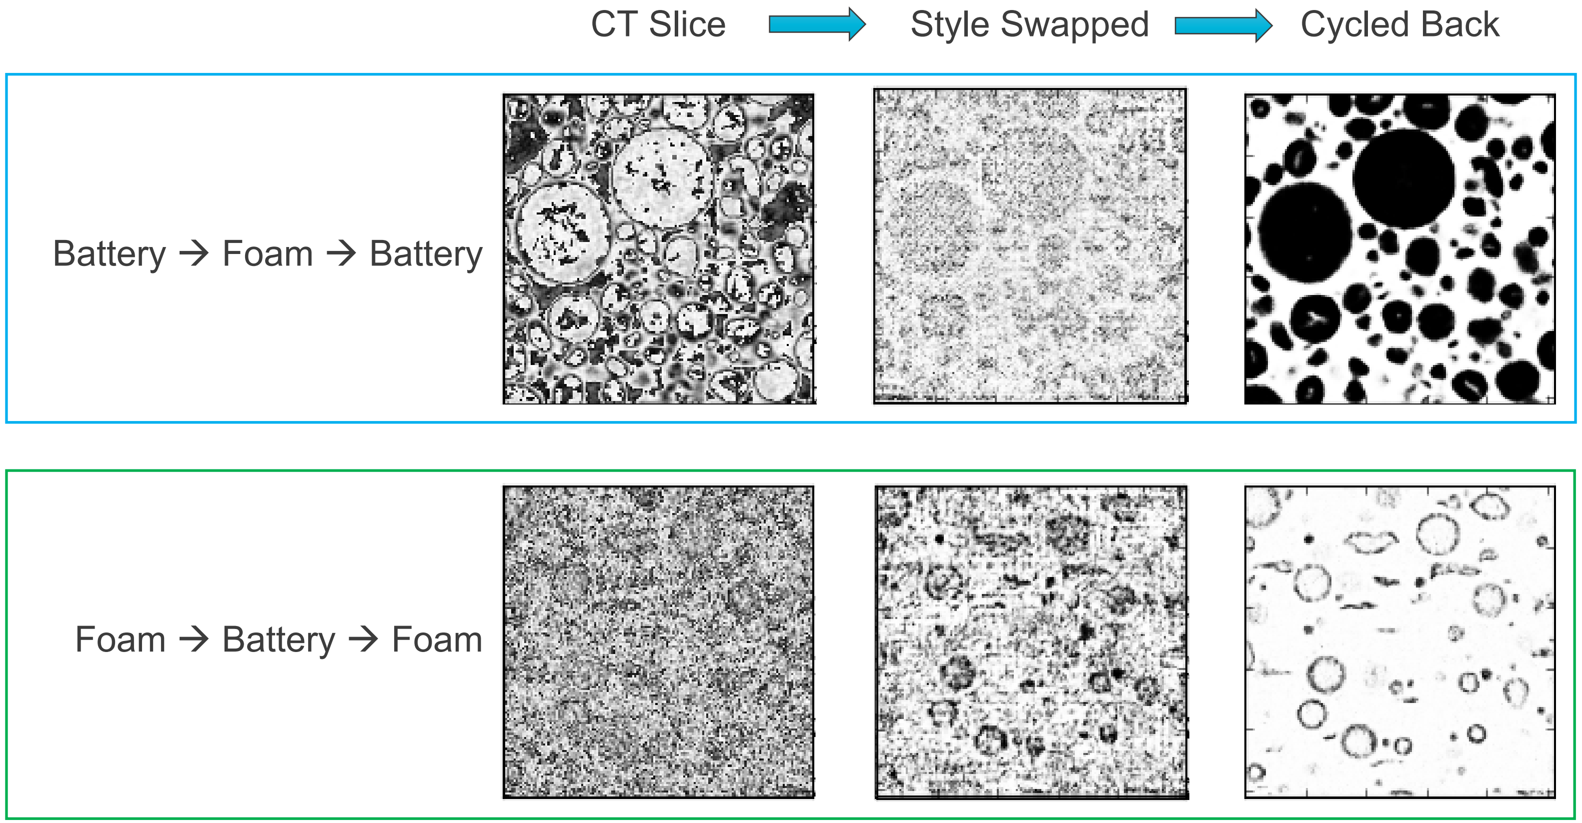

CycleGAN provides a rough segmentation of both battery and foam

Semi-supervised: Domain adaptation can reduce supervised labeling cost

- Repurpose labels from one domain (battery) to another domain (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

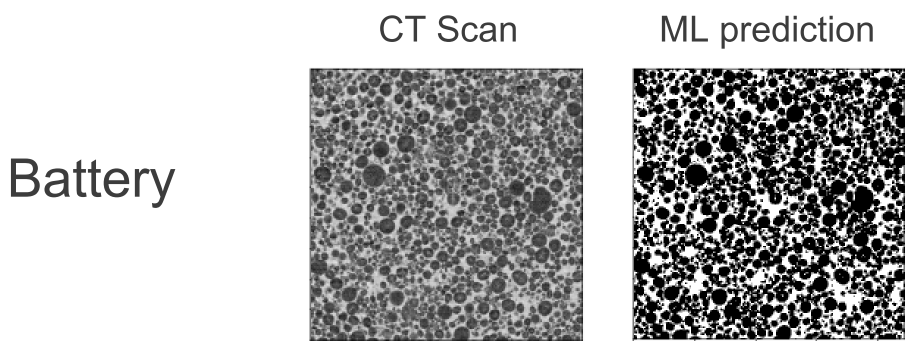

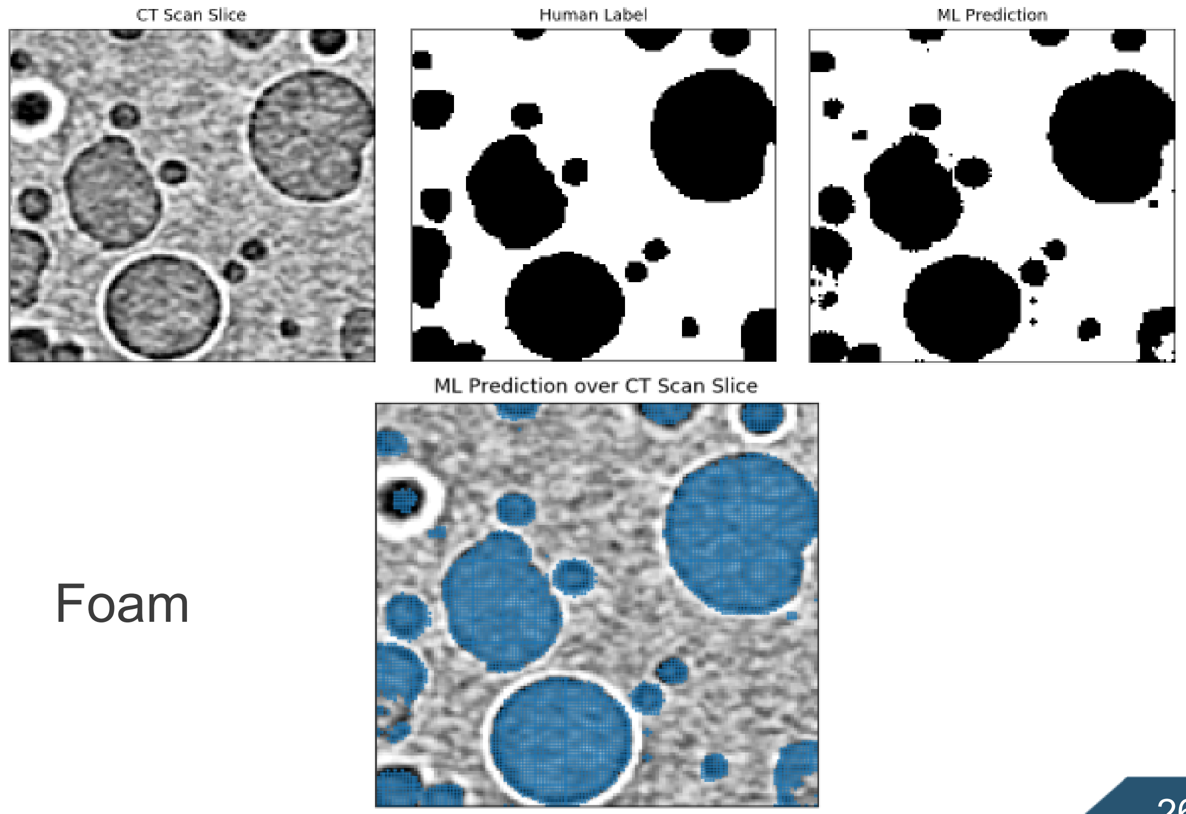

Data preprocessing and augmentation

Making the problem as easy as possible for the DL algorithm with preprocessing can dramatically improve results
- Choose normalization function that helps to separate challenging classes
- Manual inspection is important for selecting the best methods
- Example: Using log function might push most of the background to negative values

Data augmentation methods
- Flip along each axis
- Generate synthetic data from perfect numerical models: ASTRA toolbox https://www.astra-toolbox.com/

Numerical model of material

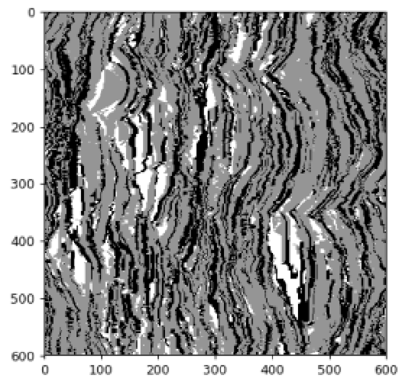

Slice of synthetic CT scan

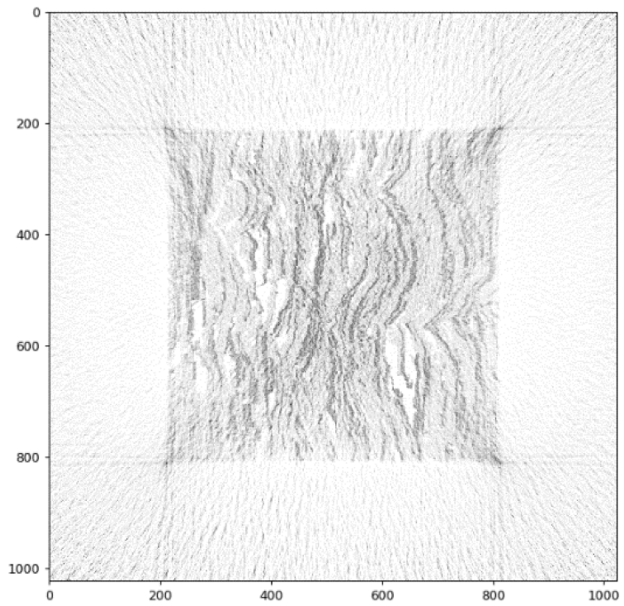

Martinez, Carianne, et al. "Automated Segmentation of Porous Thermal Spray Material CT Scans with Geometric Uncertainty Estimation." SAND2020-9099

Image segmentation uncertainty

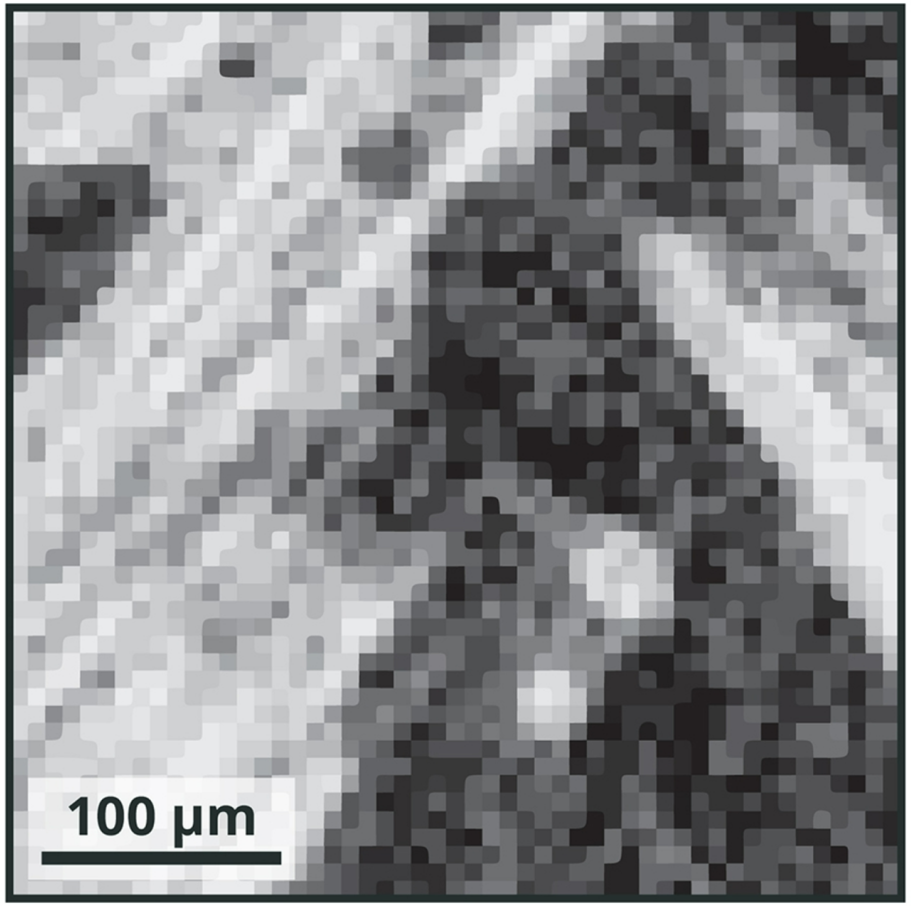

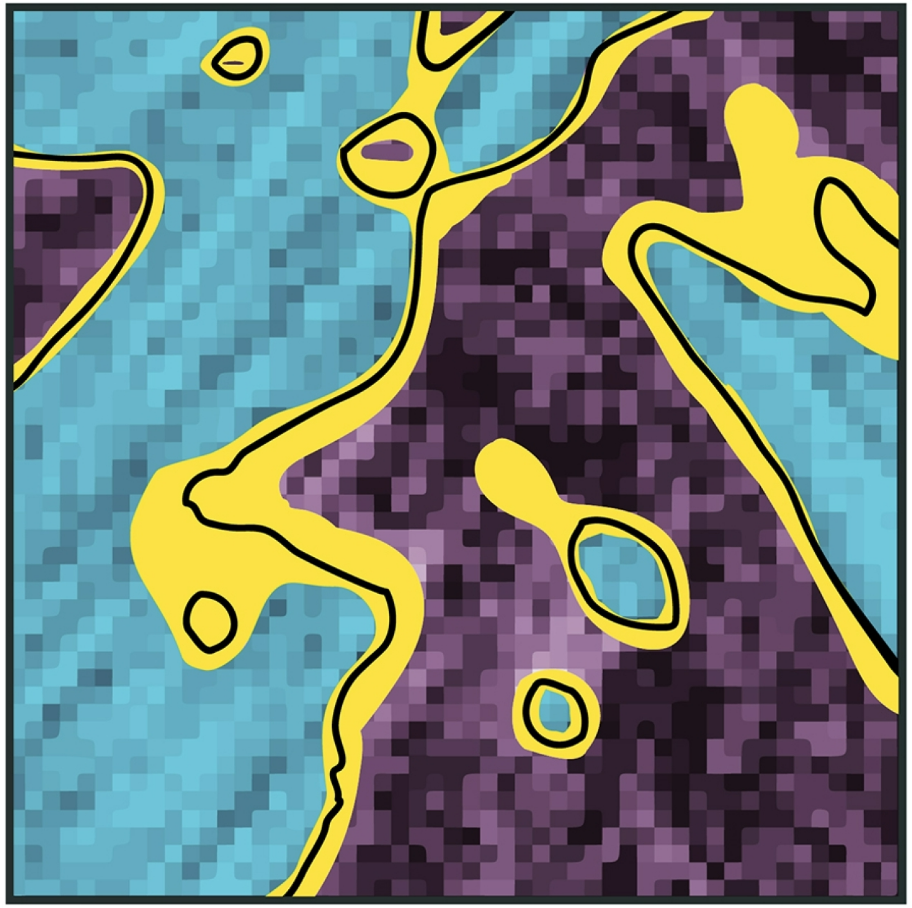

Uncertainty can be used to inform segmentation

Neural networks measure per-voxel segmentation uncertainty

Provides a measure of the model’s credibility on a particular task

Enables neural networks to overcome domain shift

- This additional information offers an insight into the model
- New ways to mitigate common problems

Slice of scan of woven composite material

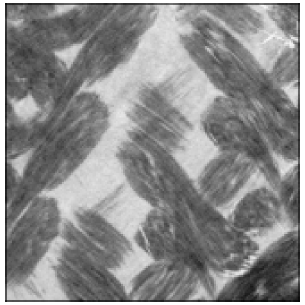

Uncertainty map - brighter pixel values indicate higher uncertainty

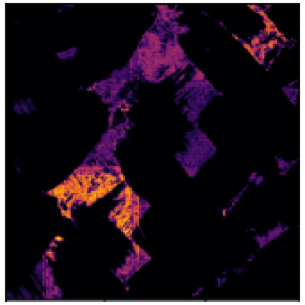

High uncertainty indicates this model should not be trusted in this domain

We leverage uncertainty maps to enable generalization of a trained model to shifted domains

CT slice from shifted domain

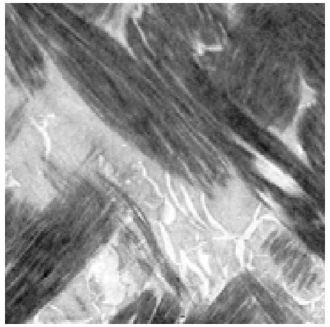

Predict segmentation using model trained on original domain

Unusable segmentation

+

Uncertainty map

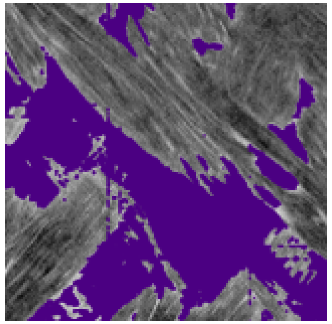

Apply advanced uncertainty based refinement method

Refined segmentation

Martinez, C., et al. (2019). Segmentation certainty through uncertainty: Uncertainty-refined binary volumetric segmentation under multifactor domain shift. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (pp. 0-0).

Happy accident leads to anomaly detection algorithm (Kyle Karlson)

Training the neural network on the wrong labels resulted in poor segmentations but high uncertainty around interesting features

Preliminary result – requires further research and validation

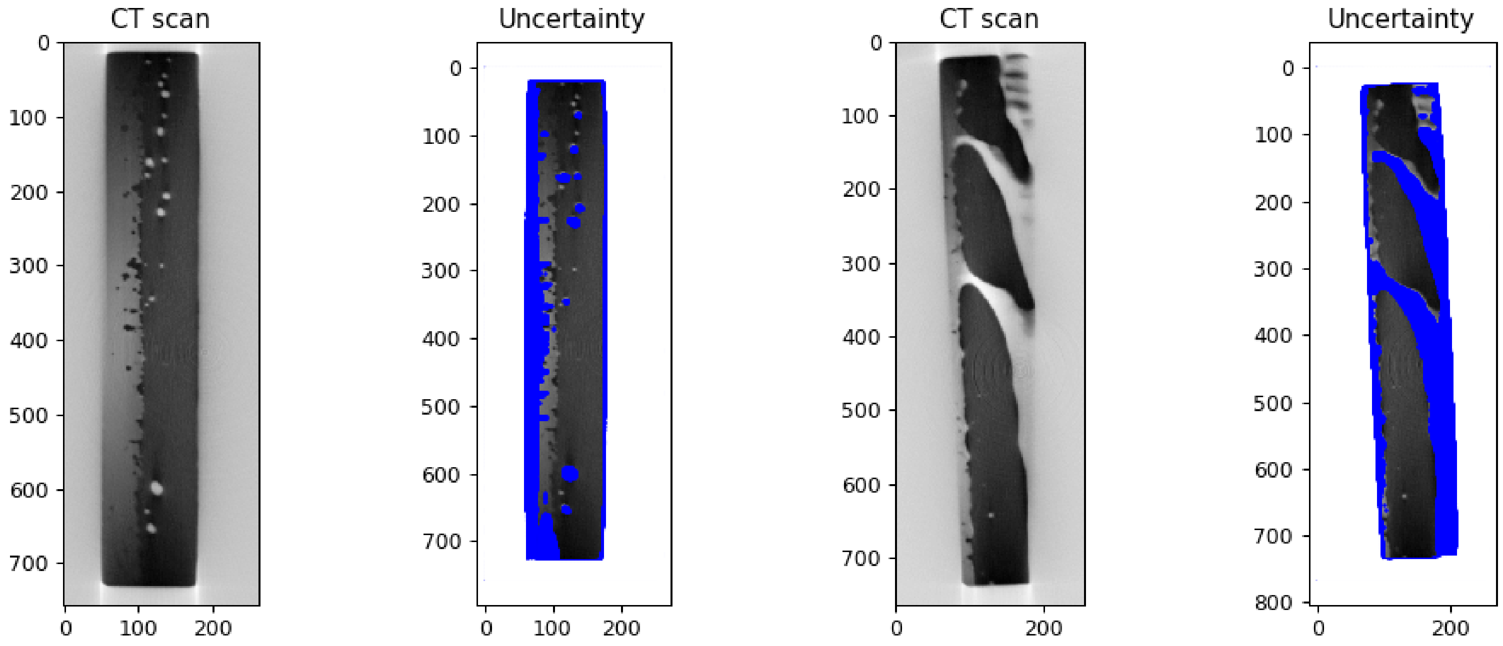

Potential Impact: Uncertainty can highlight anomalous regions

Efficient Quantification of Uncertainty in Image-based Physics Simulations (EQUIPS)

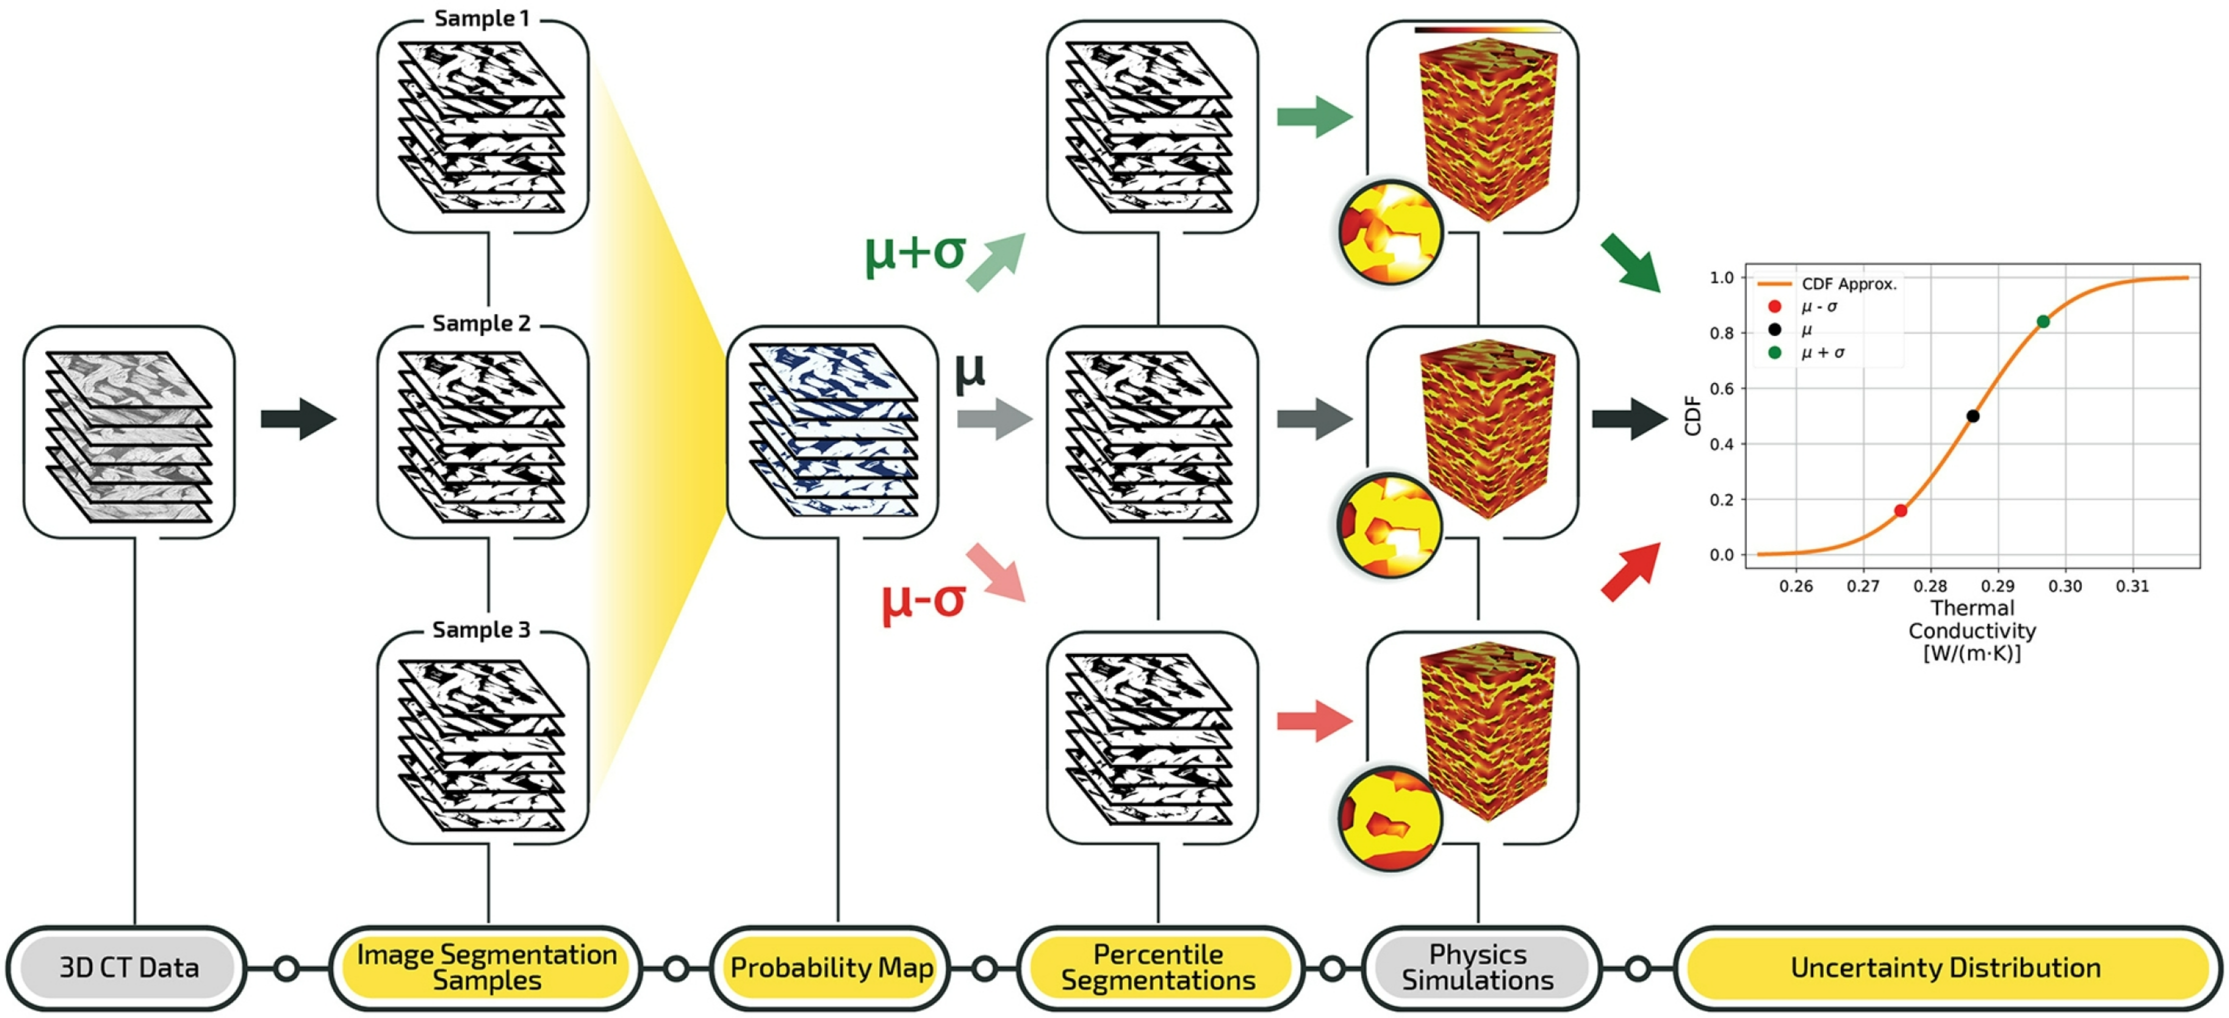

Krygier, M. C., LaBonte, T., Martinez, C., Norris, C., Sharma, K., Collins, L. N., ... & Roberts, S. A. (2021). Quantifying the unknown impact of segmentation uncertainty on image-based simulations. *Nature communications*, *12*(1), 1-11.

Exemplar: Thermal protection system (TPS) materials

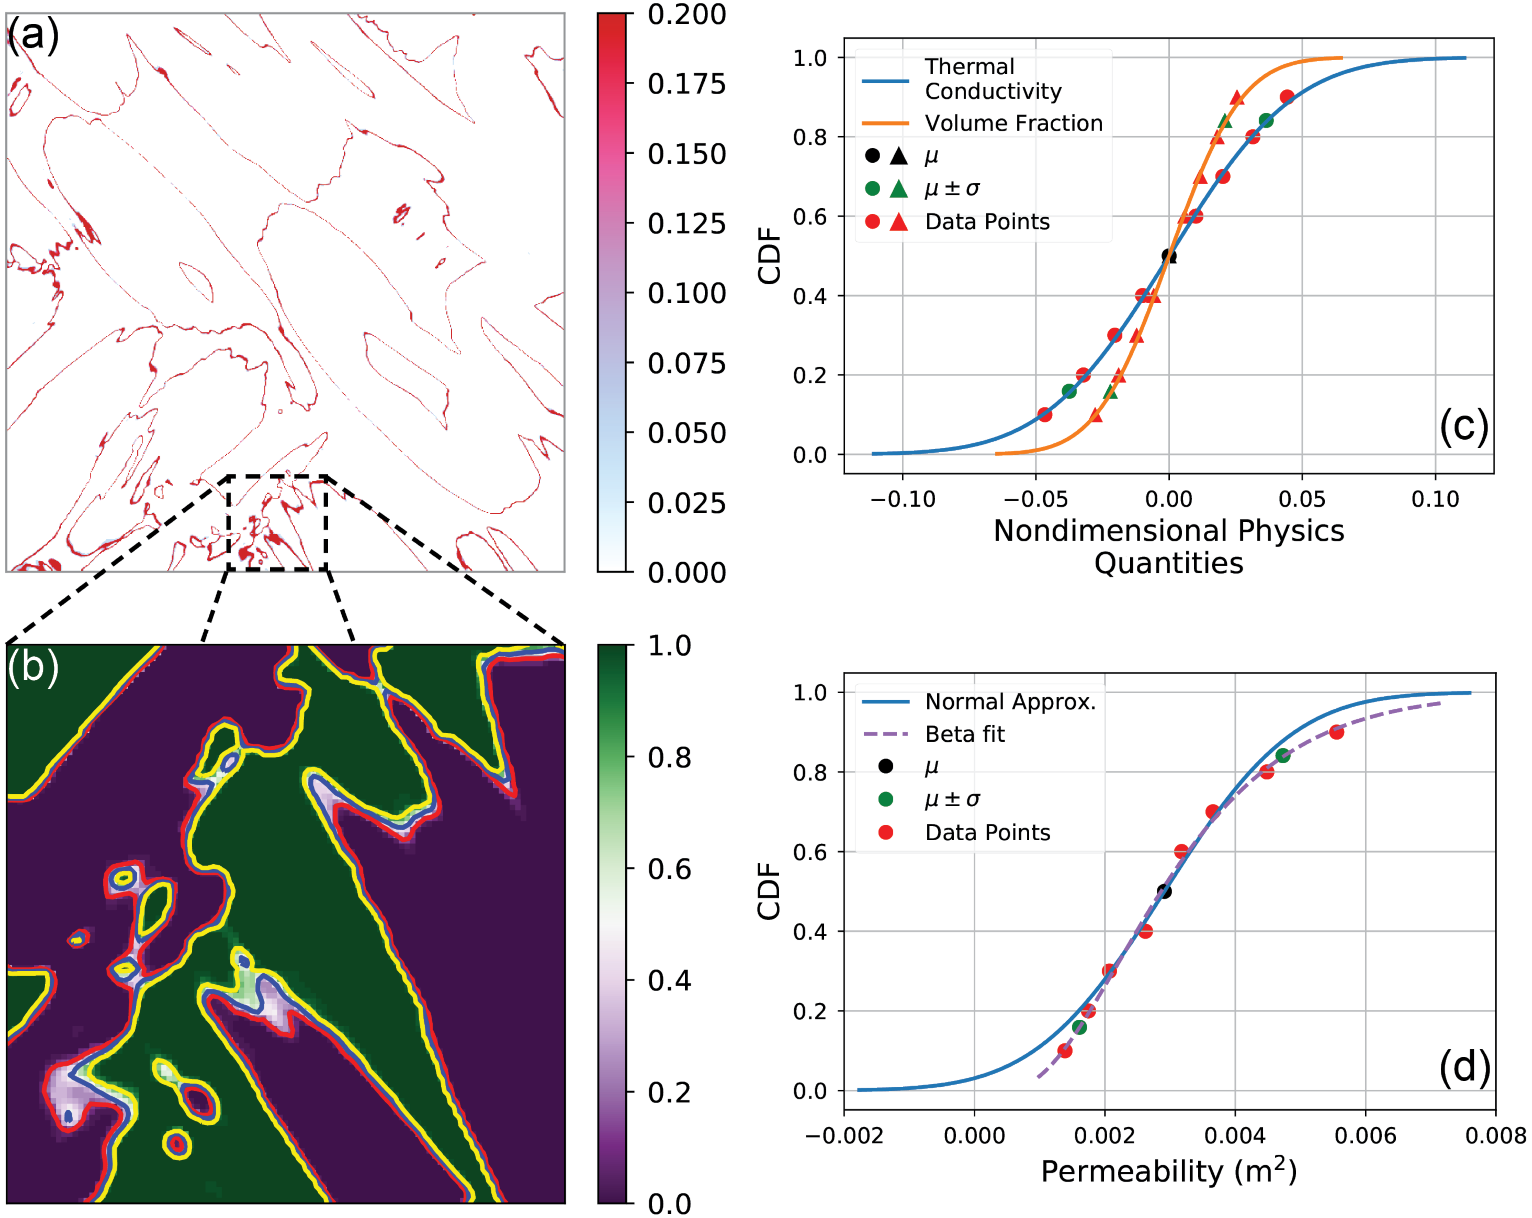

CAMI key insights

- Volumetric segmentation of materials can be credibly automated with few labeled examples.
- Image-based simulations can be sensitive to small changes in geometries.
- Deep learning models can interpret images into geometries with uncertainty.
- A subset of simulations can characterize expected system properties.

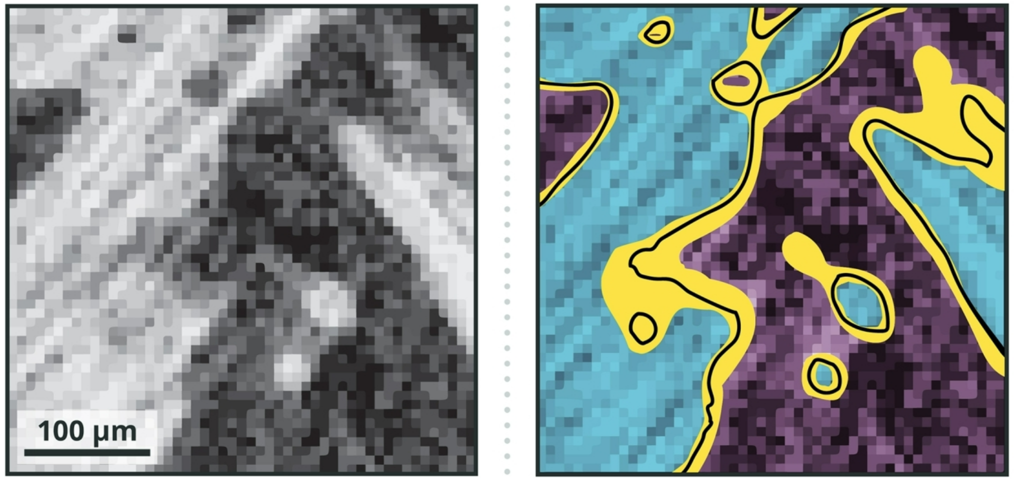

Beyond Fingerprinting Grand Challenge LDRD

Brad Boyce and Remi Dingreville, Co-PIs

Goal: AI-enabled high-throughput materials co-design

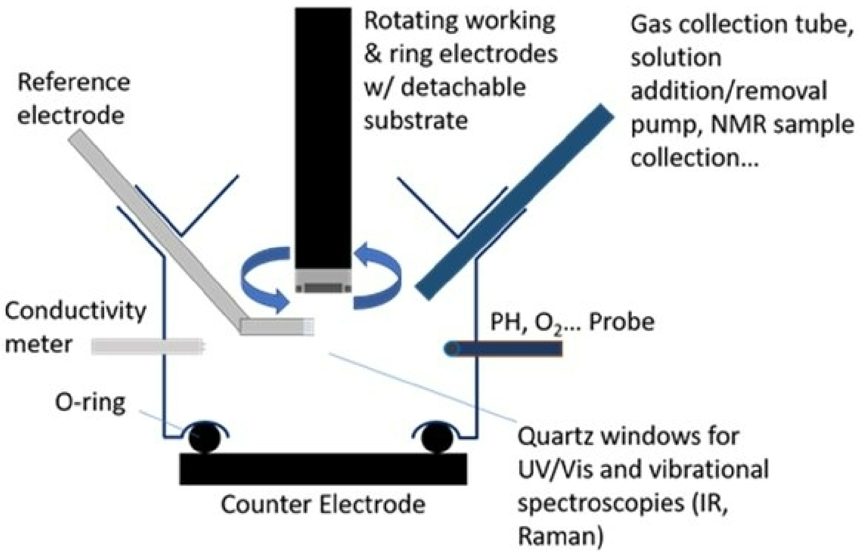

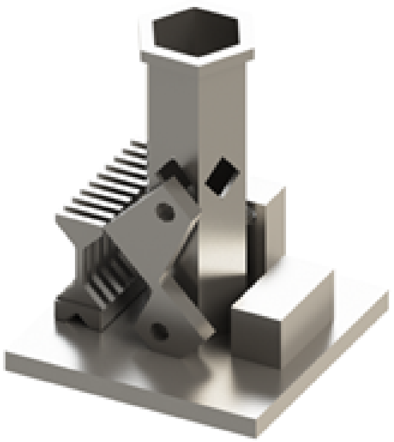

Figure 7. An example witness artifact that enables rapid measurement of 11 material properties in a compact footprint.

High-throughput fabrication

High-throughput characterization

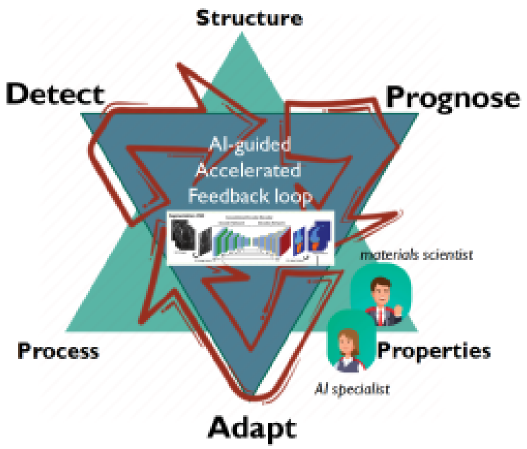

Fabrication: thin films (laser powder bed fusion, electroplating, physical vapor deposition), semiconductor components, integrated lasers and silicon photonics
Process: high-dimensional control parameter space
Structure: molecular/mesoscale description of material
Property: targeted mechanical/electromagnetic property

Hybrid-informed multilayer algorithm (Himulaya)

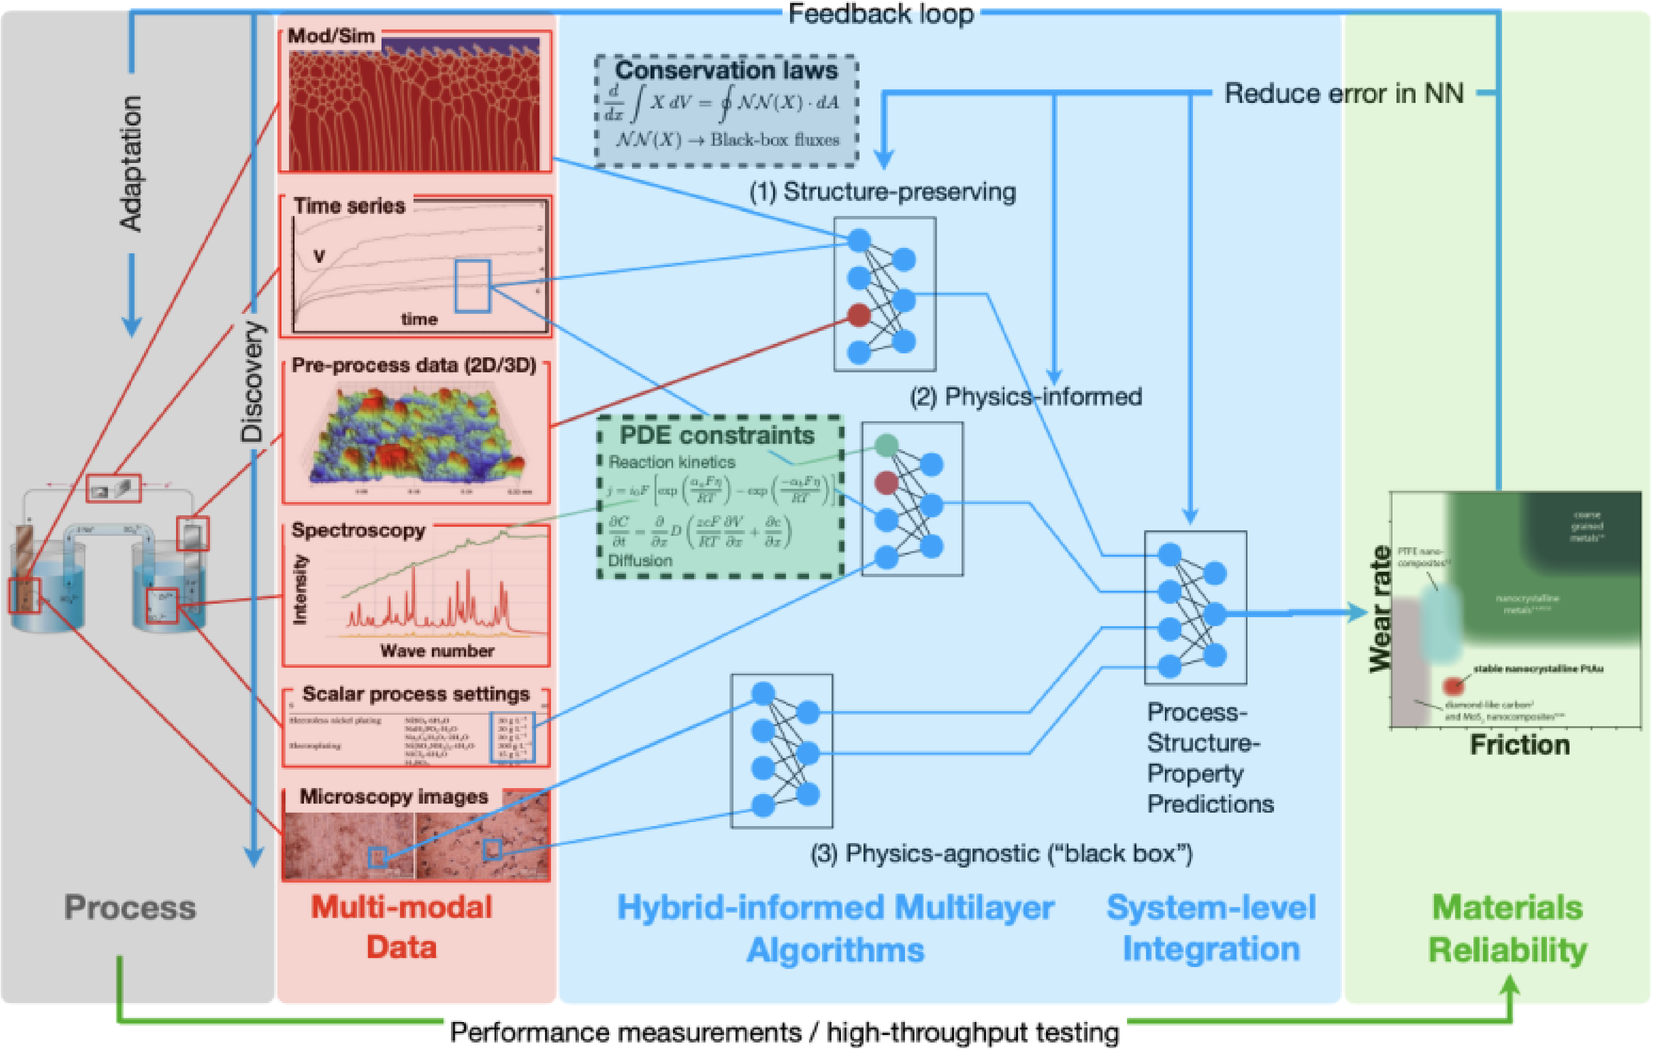

Objective: Exploit multimodality spanning process-structure-property gap, embed physical modeling expertise, **learn fingerprints to detect, prognose + adapt**

PIMA – physics-informed multimodal autoencoder (Nat Trask)

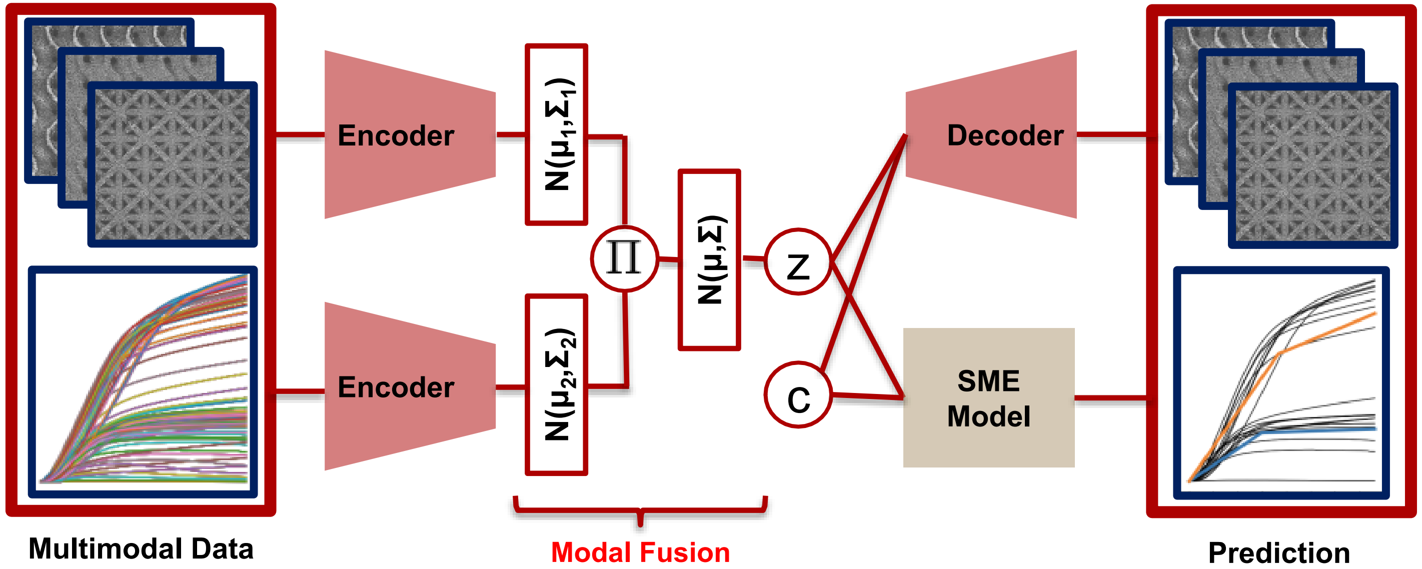

Trask, N., Martinez, C., Lee, K., & Boyce, B. (2022). Unsupervised physics-informed disentanglement of multimodal data for high-throughput scientific discovery. *arXiv preprint arXiv:2202.03242*.

Informal Idea:
We discover a shared latent representation of data providing a Rosetta stone for across modalities <u>w/ uncertainty</u>

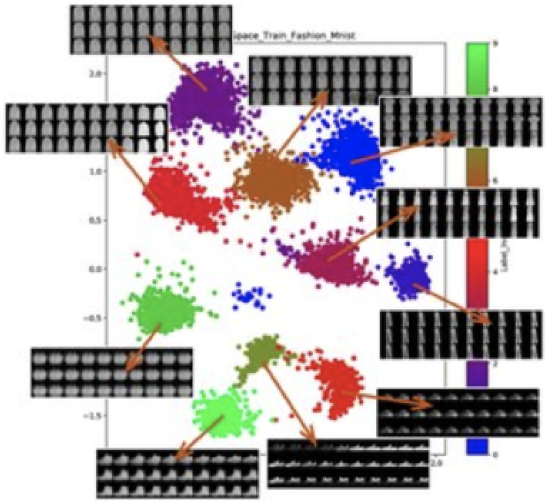

Formal Idea:

- **Gaussian product distribution gives deep posterior embedding for each modality**
- **Gauss mixture prior in latent space identifies populations in data across modalities**
- ***Closed form expressions* for loss – no Monte Carlo**
- **Supports Bayesian inference across modalities**

Conclusion

Digital twin building blocks:

- Advances in automated processing of multimodal data
- Leveraging transfer learning
- Uncertainty quantification for each workflow step
- Incorporation of domain knowledge
- Integrated interdisciplinary teams

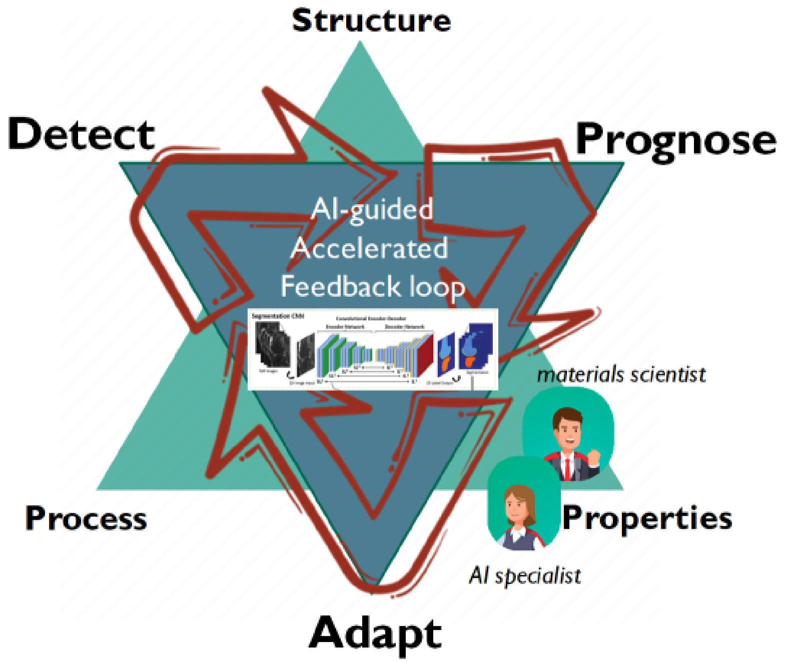

Questions?

cmarti5@sandia.gov