SAND2021-2212C

Sandia National Laboratories

Predicting Mechanical Performance in Additive Manufacturing Components Using Deep Learning

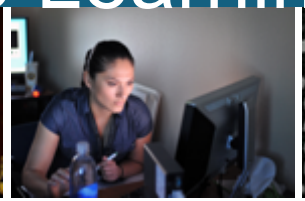

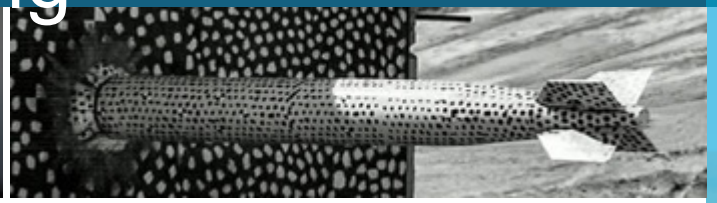

[1]**Kyle Johnson**, [1]Demitri Maestas, [1]John Emery [1]Matthew Smith, [1]Carianne Martinez, [2]Mircea Gregoriu, and [1]Steven Sleder

[1]Sandia National Laboratories, Albuquerque, NM
[2]Cornell University, Ithaca, NY

TMS Annual Meeting
March 15-18, 2021

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.

Vision: Rapid failure prediction based on microstructure, geometry, and loading conditions

Microstructure

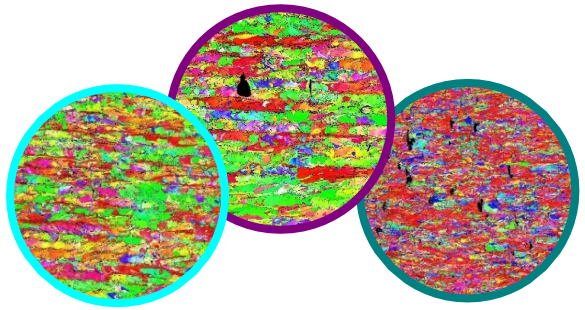

As-built Geometry

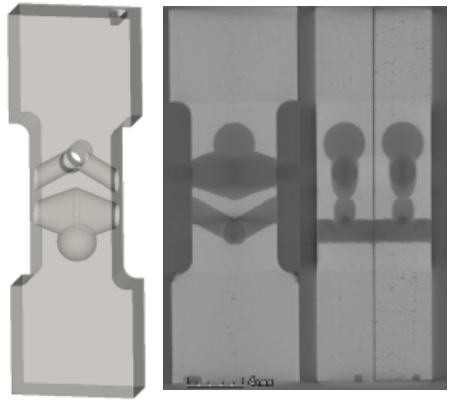

Boundary Conditions

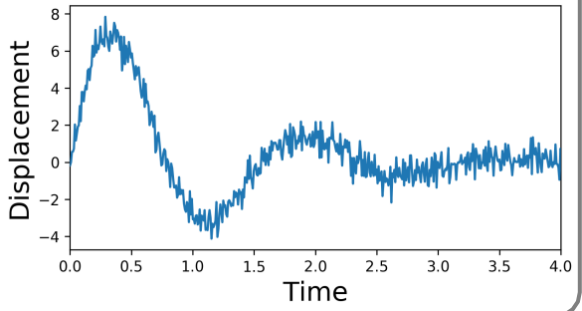

Simulation Code

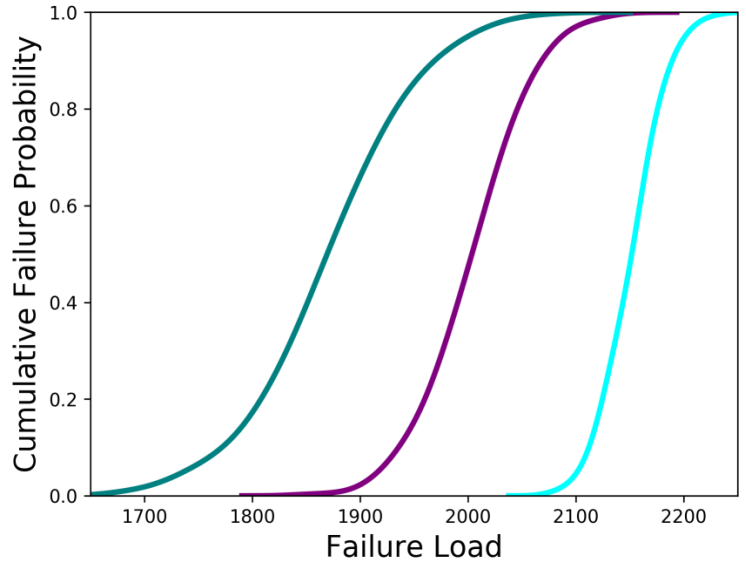

Qualification, Topology Optimization, etc.

Image: oldcomputers.net

Today: Failure Prediction is Difficult and Slow Even With a Team

1. Create Mesh (Days-Weeks)

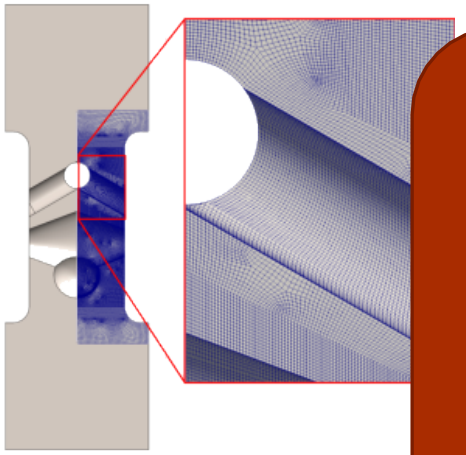

2. Iteratively Calibrate Material Model with Damage (Days-Weeks)

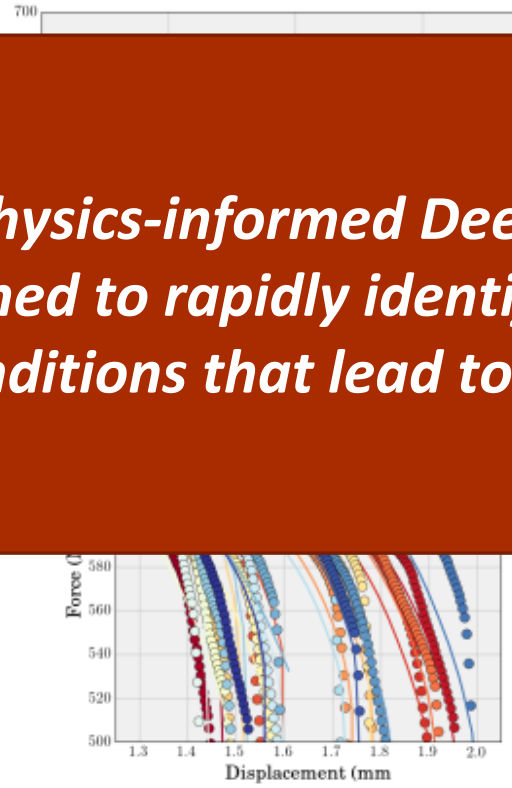

3. Run Many Large Simulations (Days-Weeks-Months)

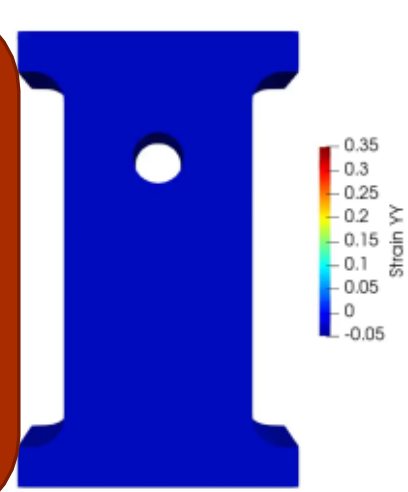

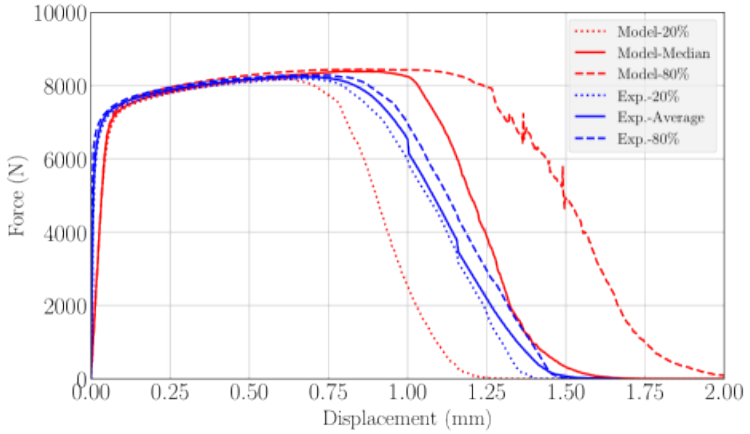

Can physics-informed Deep Learning (DL) algorithms be trained to rapidly identify the initial microstructural conditions that lead to incipient failure initiation?

Johnson et al. *IJF 2019, Kramer, Boyce* et al., *IJF 2019*

Why Machine Learning?

Some problems are difficult to solve with a directly-coded algorithm

- Don’t generalize well
- Have to write a program by hand for each specific task

Machine learning (ML) is good at:

- Recognizing patterns
- Anomaly detection
- Prediction

Slide: Carianne Martinez, Drew Levin, Demitri Maestas, Matthew Smith

Introduction to Machine Learning

Typical computer algorithms:

- have an input,
- do some computation,
- and produce some output

Introduction to Machine Learning

Machine learning algorithms:

- have lots of inputs,
- do some computation,
- and produce a *function*

Introduction to Machine Learning

Layers are functions with ”trainable” parameters
- e.g. y=mx + b

Layering linear functions models non-linear functions

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

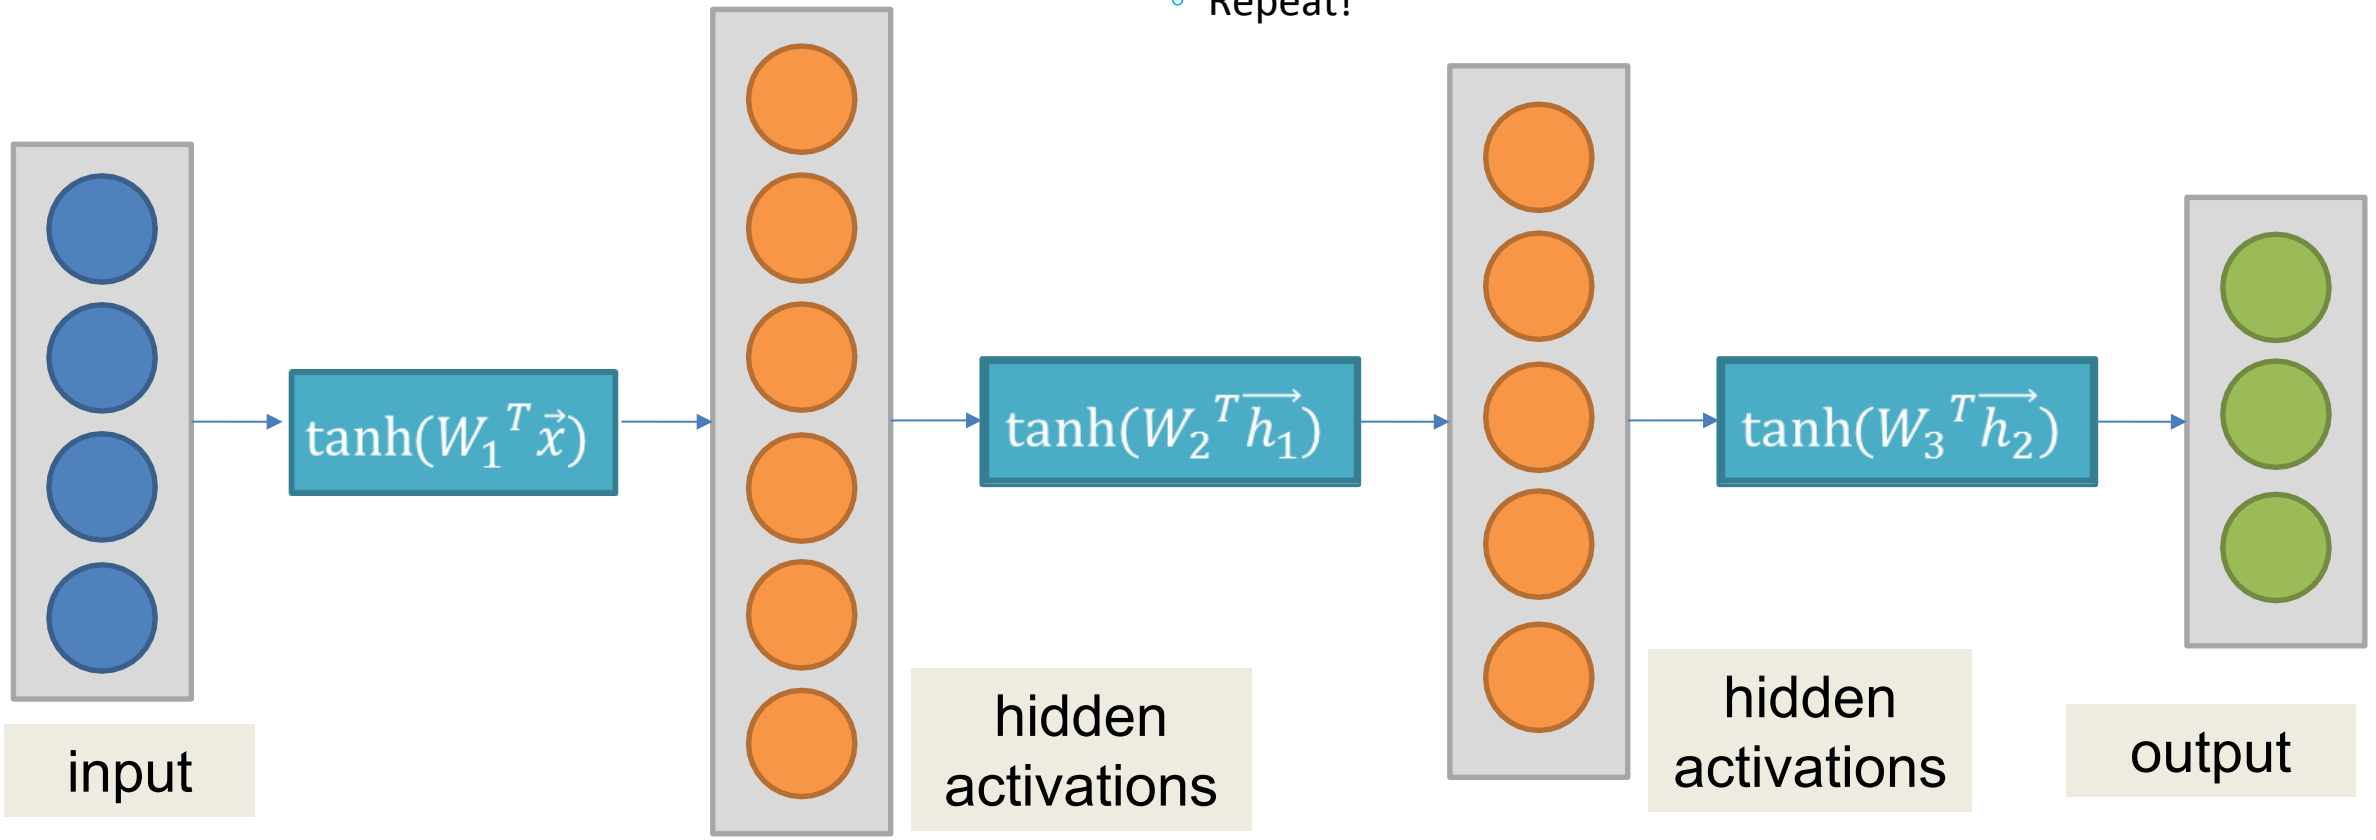

Slide: Carianne Martinez, Drew Levin, Demitri Maestas, Matthew Smith

Deep Learning

Big neural networks

Multiple layers capture “deep” statistical information in data

Require more intensive computing resources

Include:

- Convolutional neural networks
- Recurrent neural networks

Methods:

- Supervised
- Unsupervised
- Semi-supervised

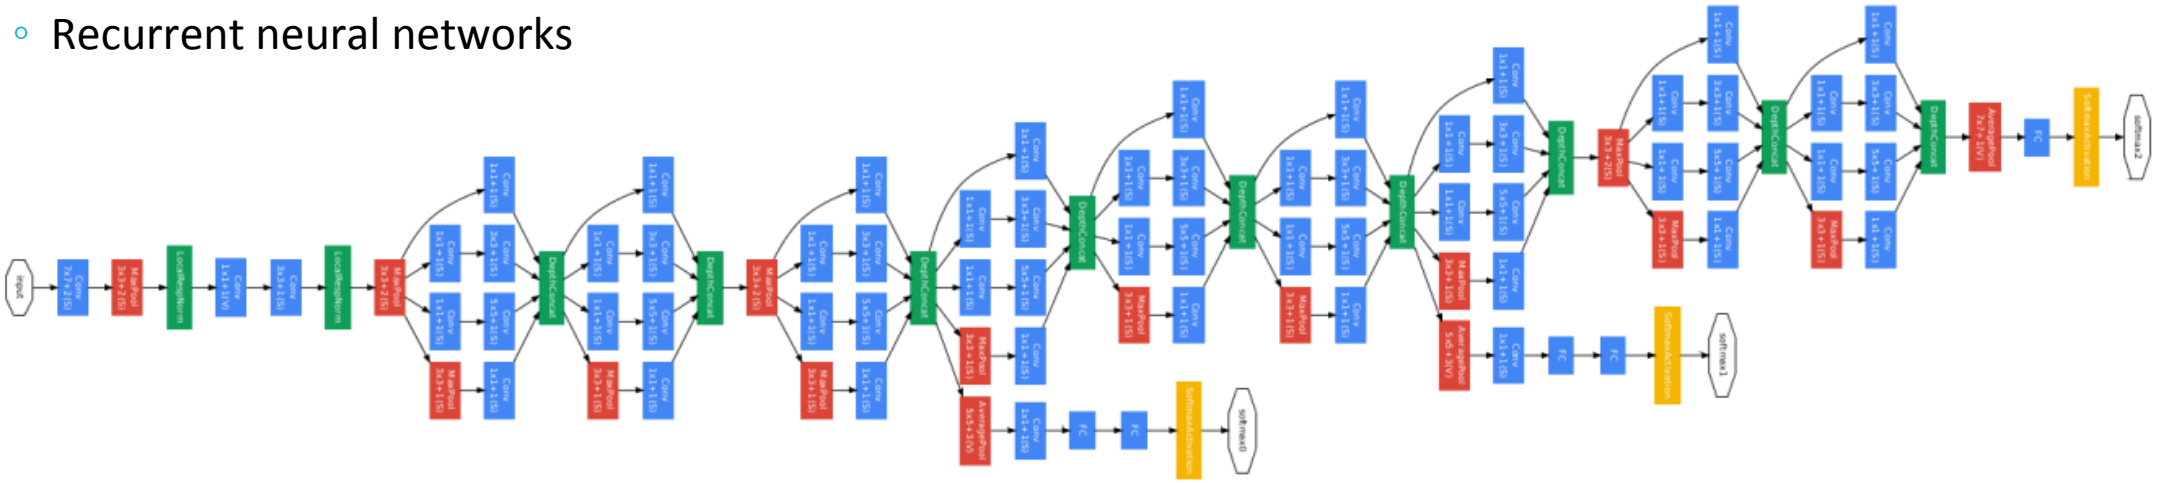

Outline

- Part 1: Predicting AM Mechanical Response
  - Training Data Generation Using Porous AM Material
  - Data Mapping
  - Local Variations in Stress State
  - Deep Learning Algorithm
  - Deep Learning Predictions
  - Conclusions
- Part 2: Current work on predicting failure in AA6061

Part 1 Project Overview: Predicting AM Mechanical Response

Training Data: High fidelity model results from AM dogbones loaded in tension

DL Model: Supervised 3D CNN

Inputs: Porosity and final state equivalent plastic strain

Output: Metric Classification

Submitted to: Additive Manufacturing

Metrics

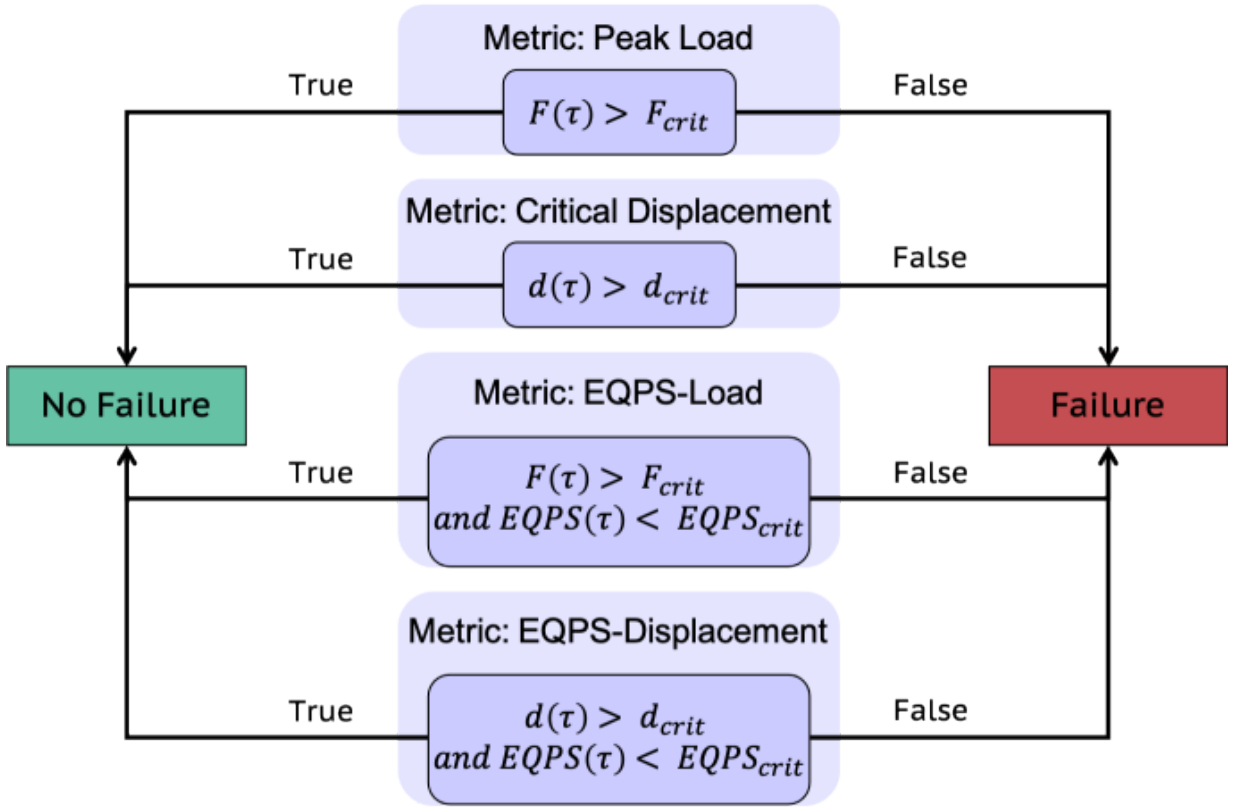

τ = *time corresponding to peak load*

F_{crit} = *required load*

d_{crit} = *required displacement*

$EQPS_{crit}$ = *required equivalent plastic strain*

Experimental CT measurements inform training data meshes

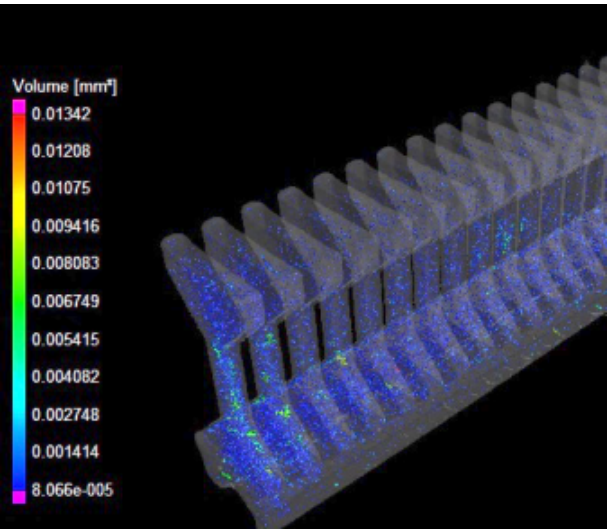

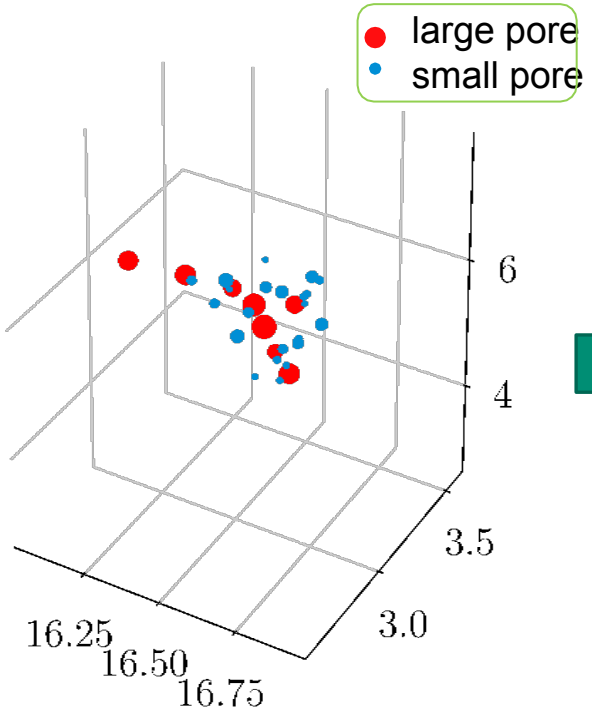

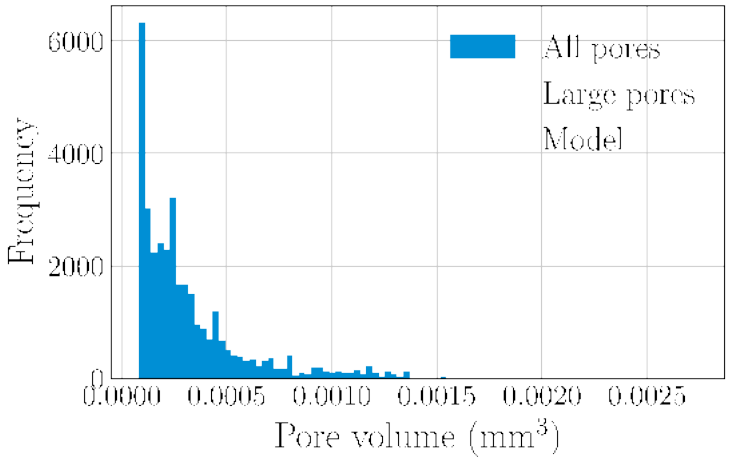

- Training data consists of dogbone gauge sections loaded in tension past peak load

CT Data: Chris Laursen, Jay Carroll

Varying fidelity models test statistical representation of voids

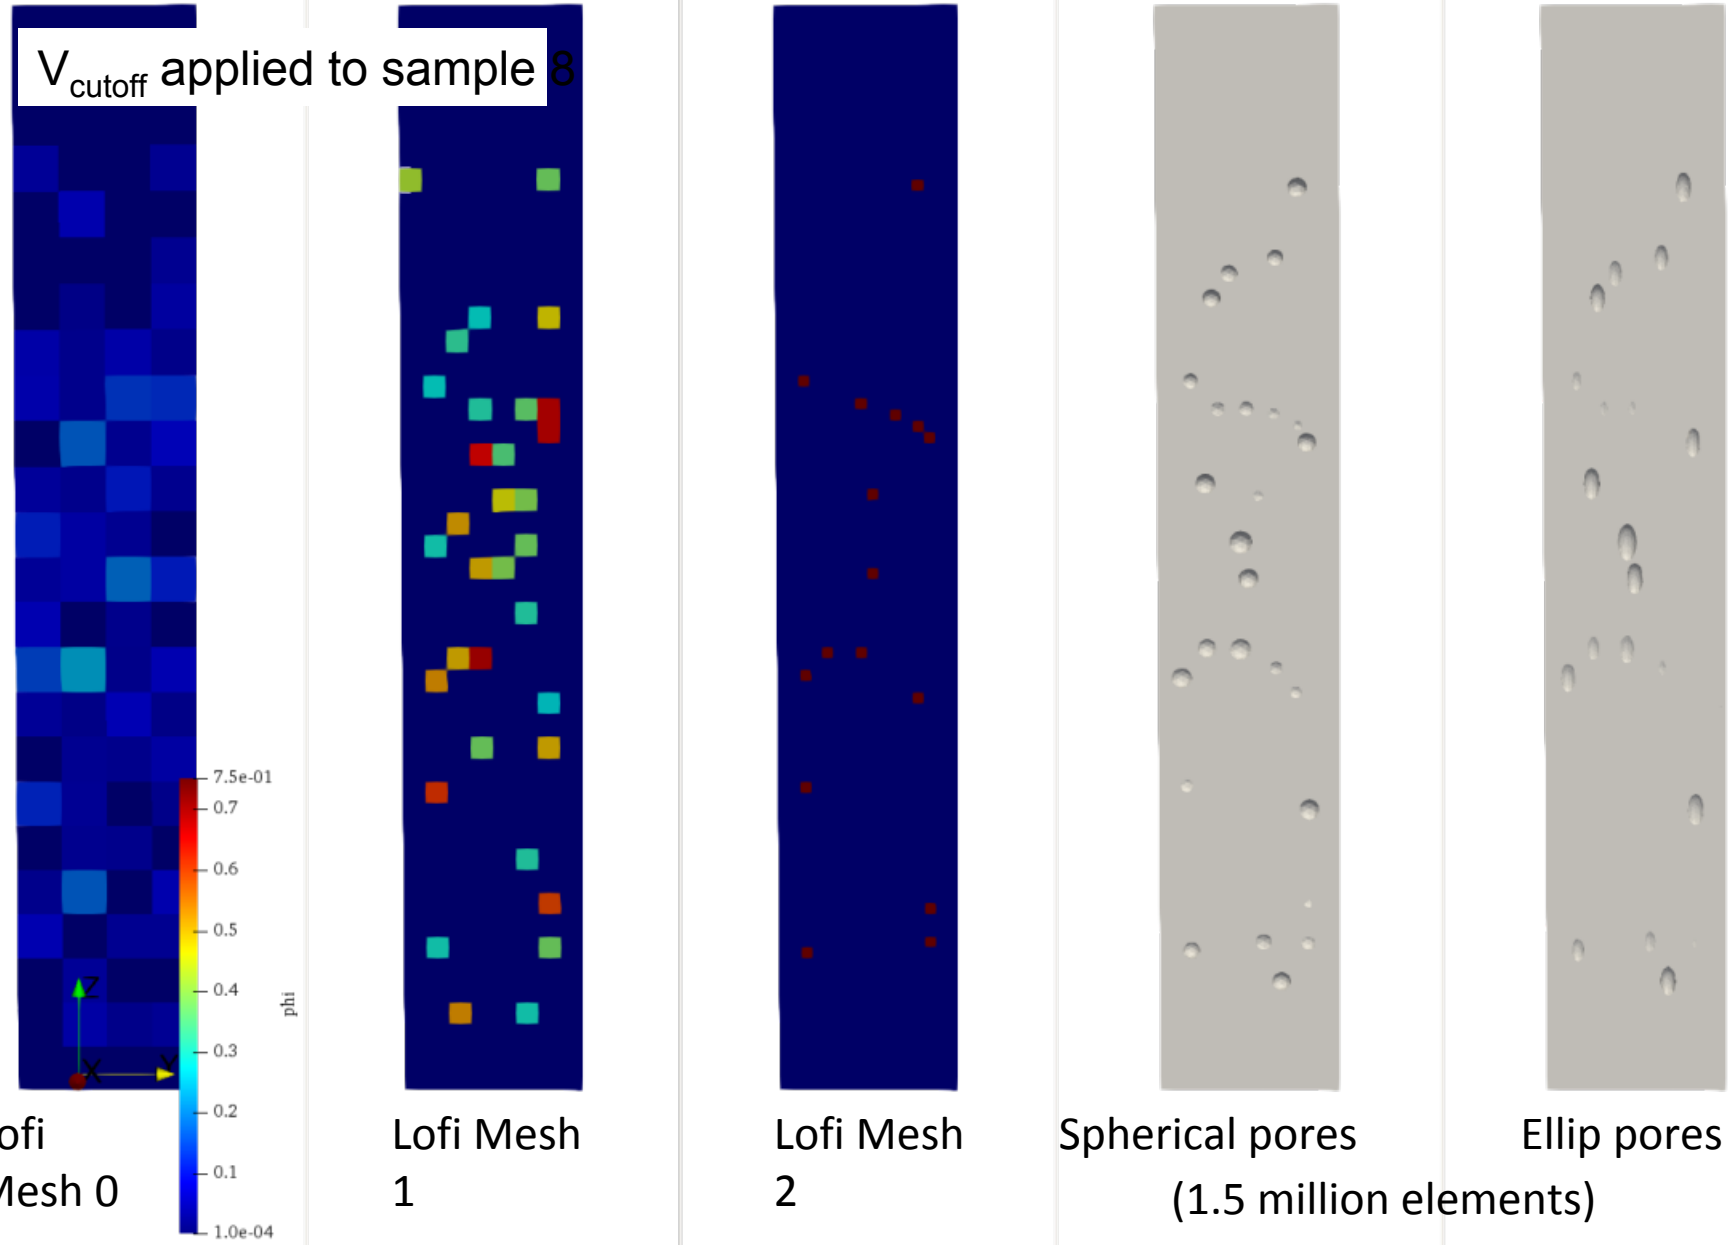

Lofi Mesh 0

Lofi Mesh 1

Lofi Mesh 2

Spherical pores

Ellip pores

(1.5 million elements)

Low fidelity predictions versus full-fidelity prediction

V_{cutoff} applied to sample 8

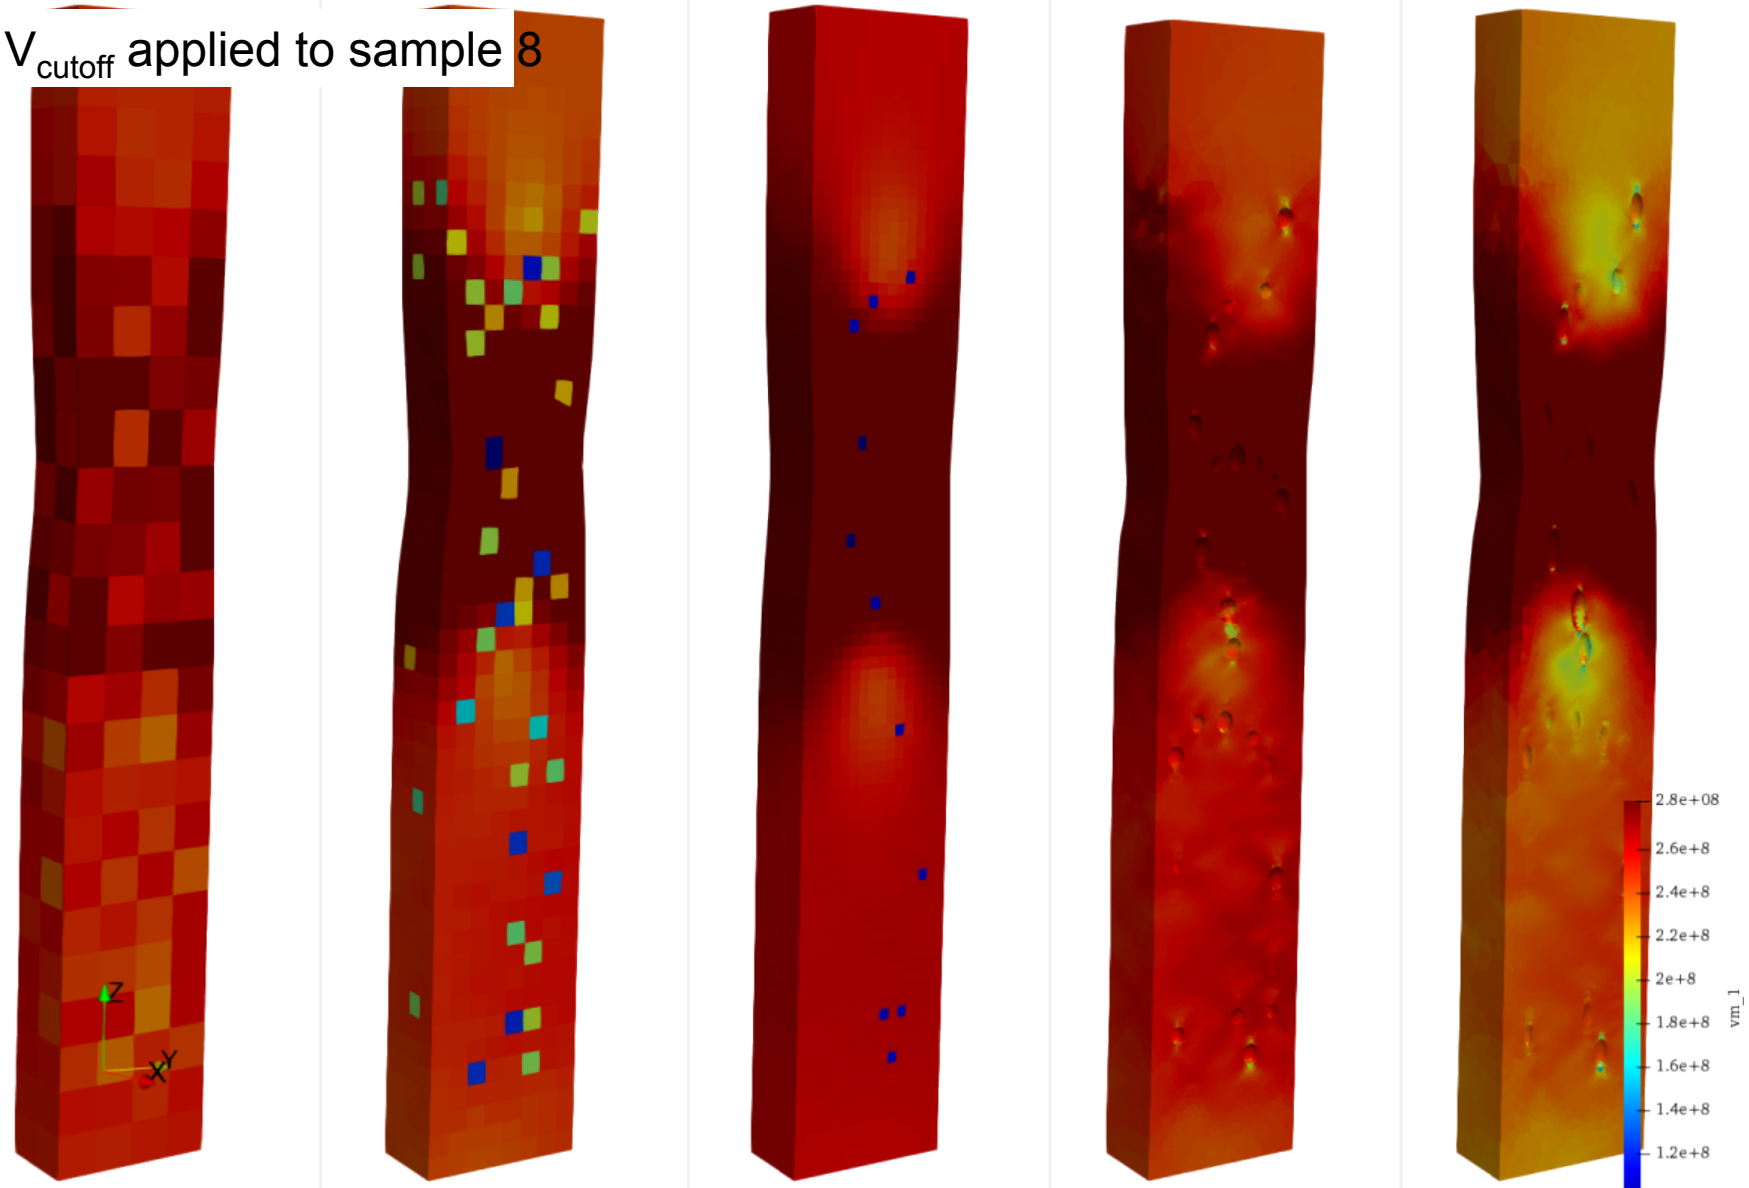

Lofi Mesh 0 | Lofi Mesh 1 | Lofi Mesh 2 | Spherical pores | Ellip pores

Low- and full-fidelity vs. experiment

V_{cutoff} applied to sample 8

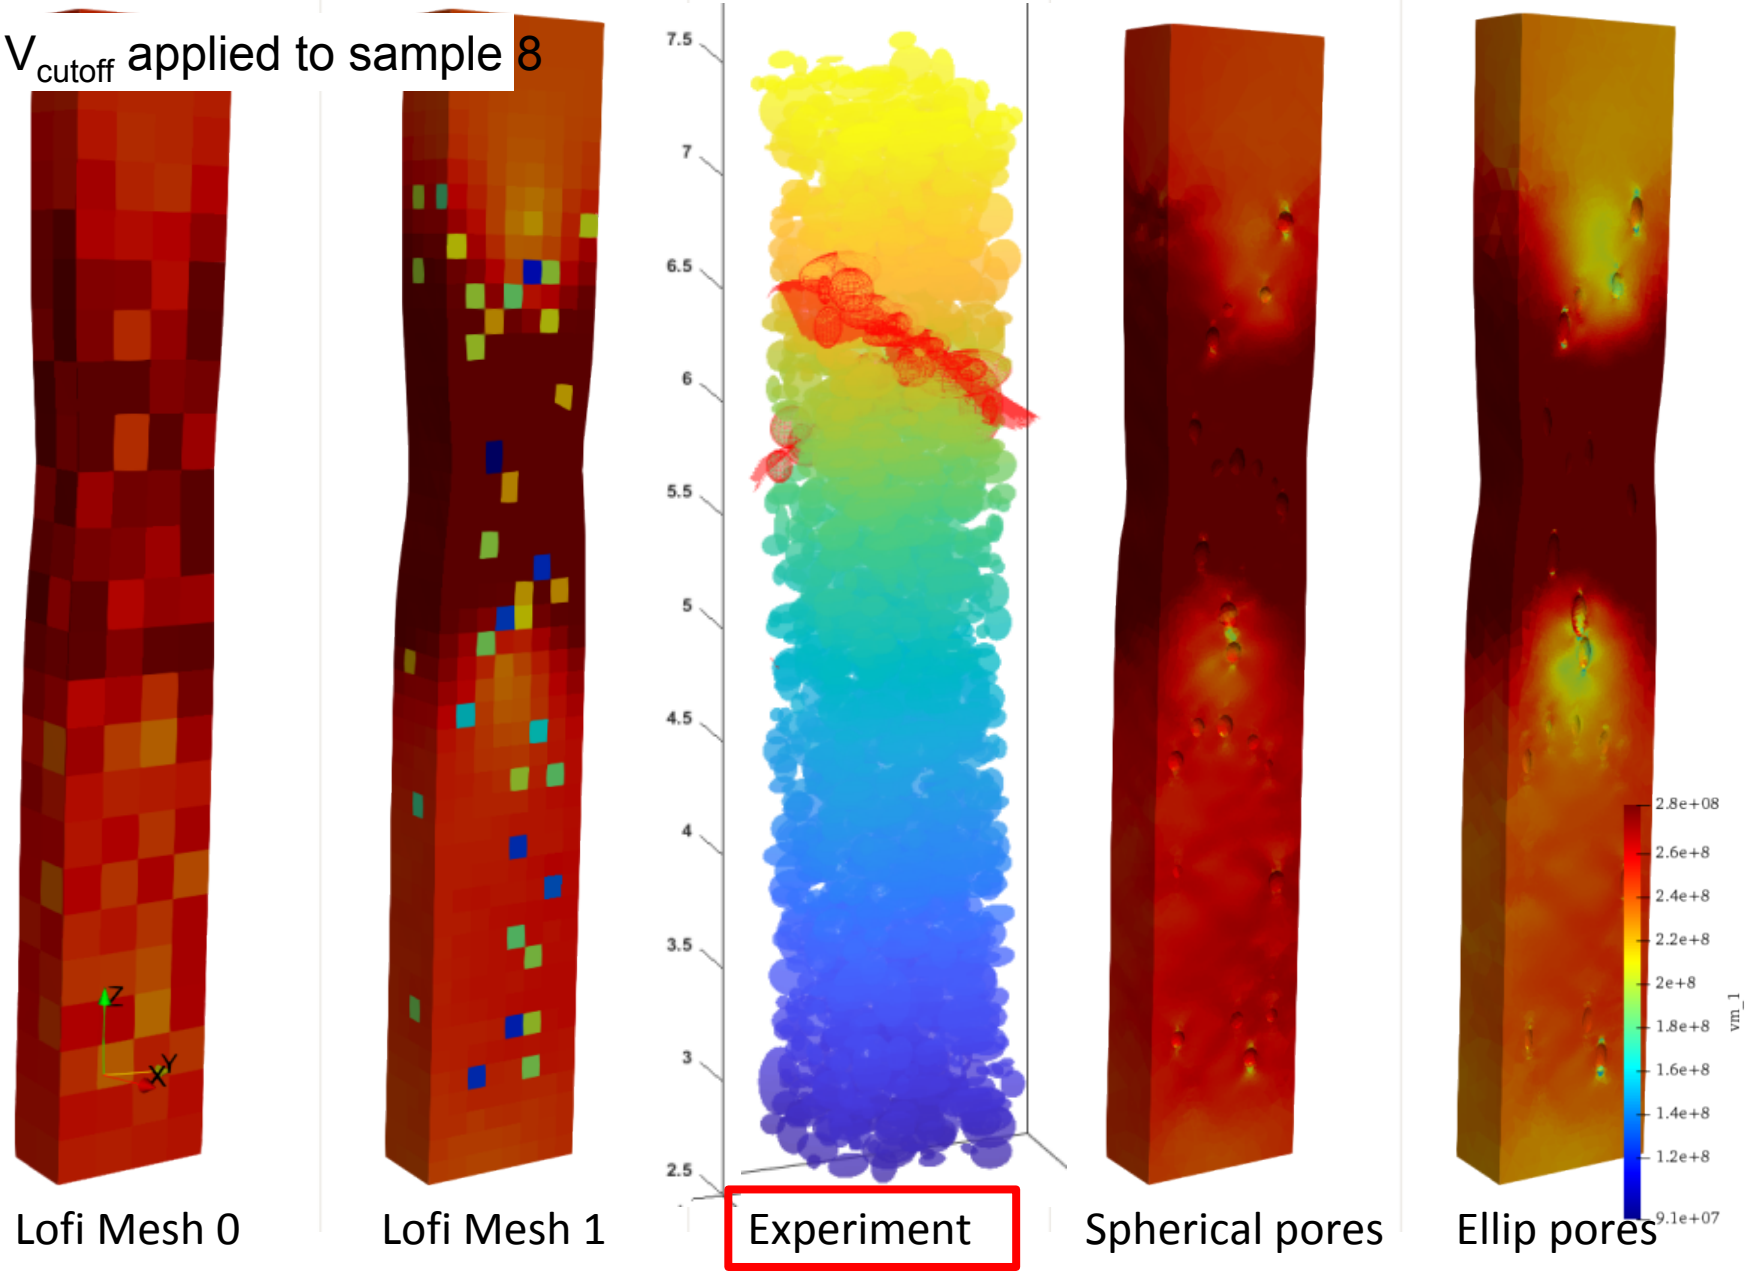

Lofi Mesh 0 | Lofi Mesh 1 | Experiment | Spherical pores | Ellip pores

Calibration to tensile specimens

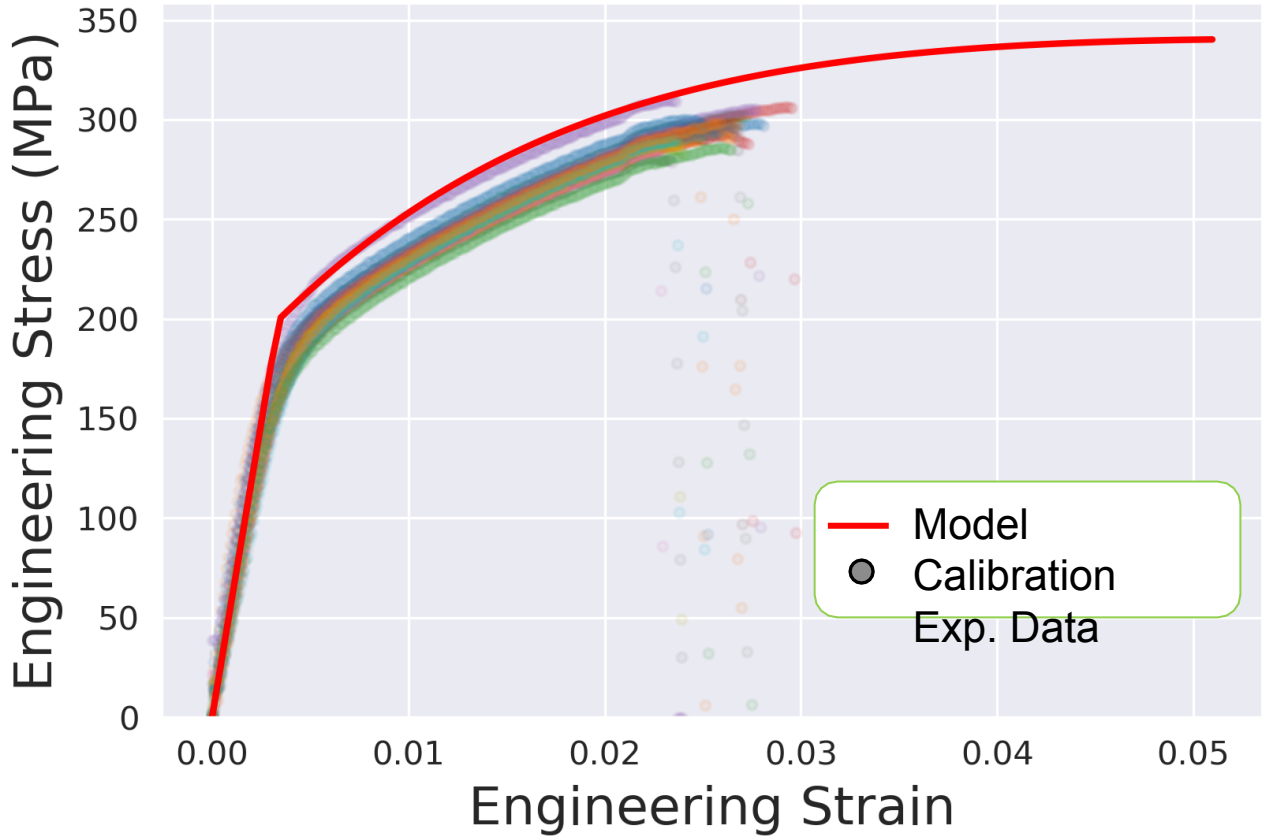

- Plasticity is captured with Voce hardening model

$$\bar{\sigma} = \sigma_y + A(1 - \exp(-n\bar{\varepsilon}^p))$$

- Model calibrated using porous mesh from CT scan – captures “matrix” response

Different porosity samples lead to different local behavior

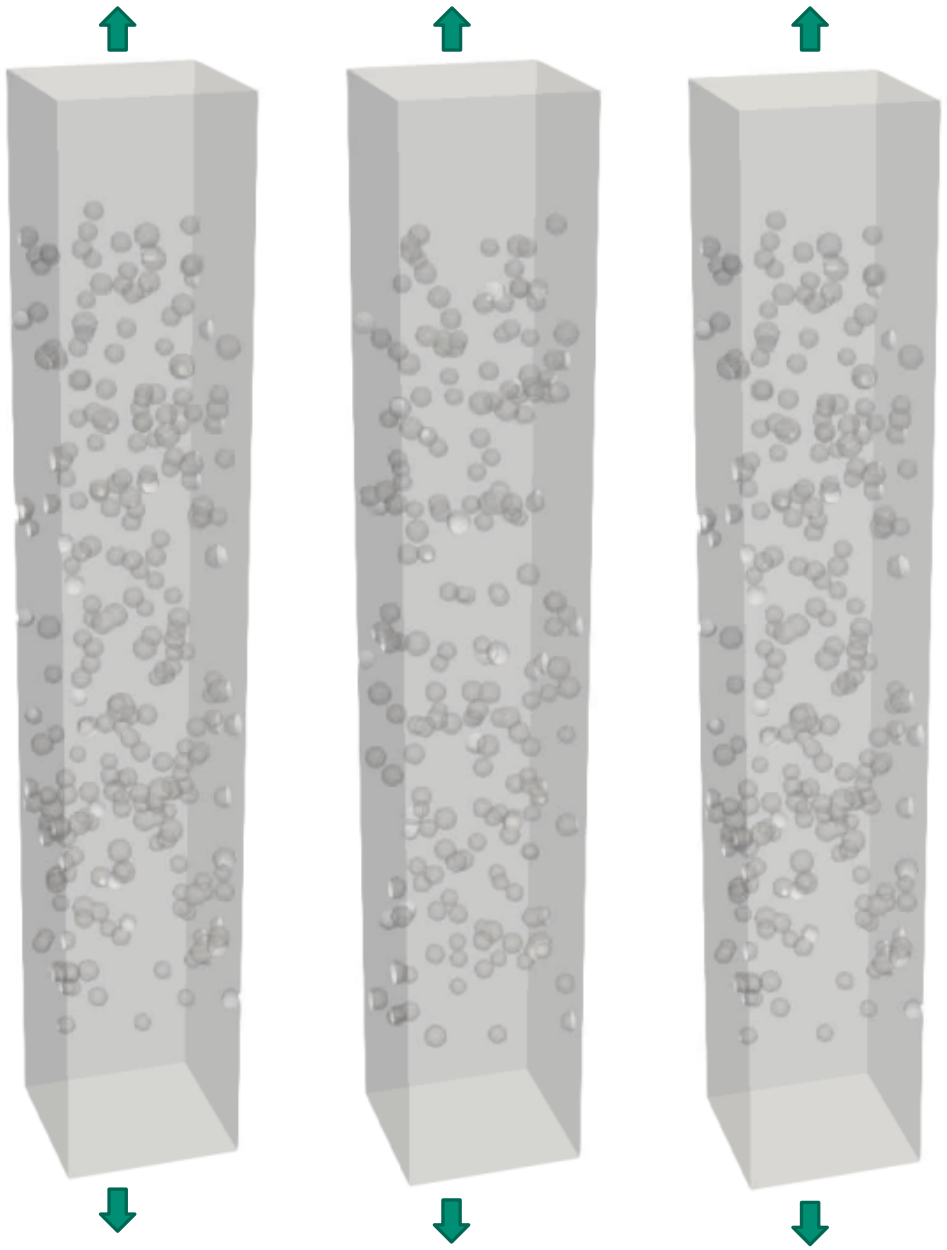

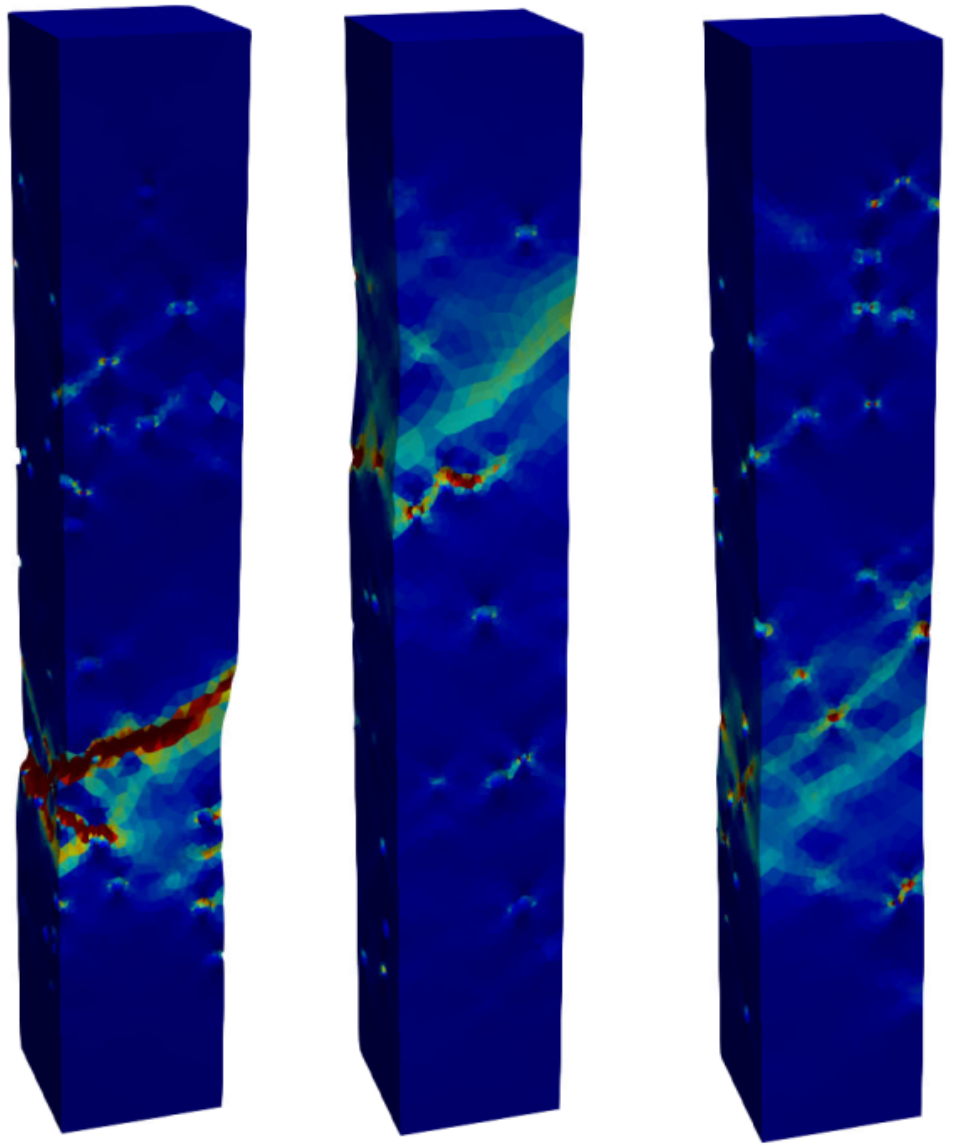

Force-displacement and max EQPS show large variations due to pore structures

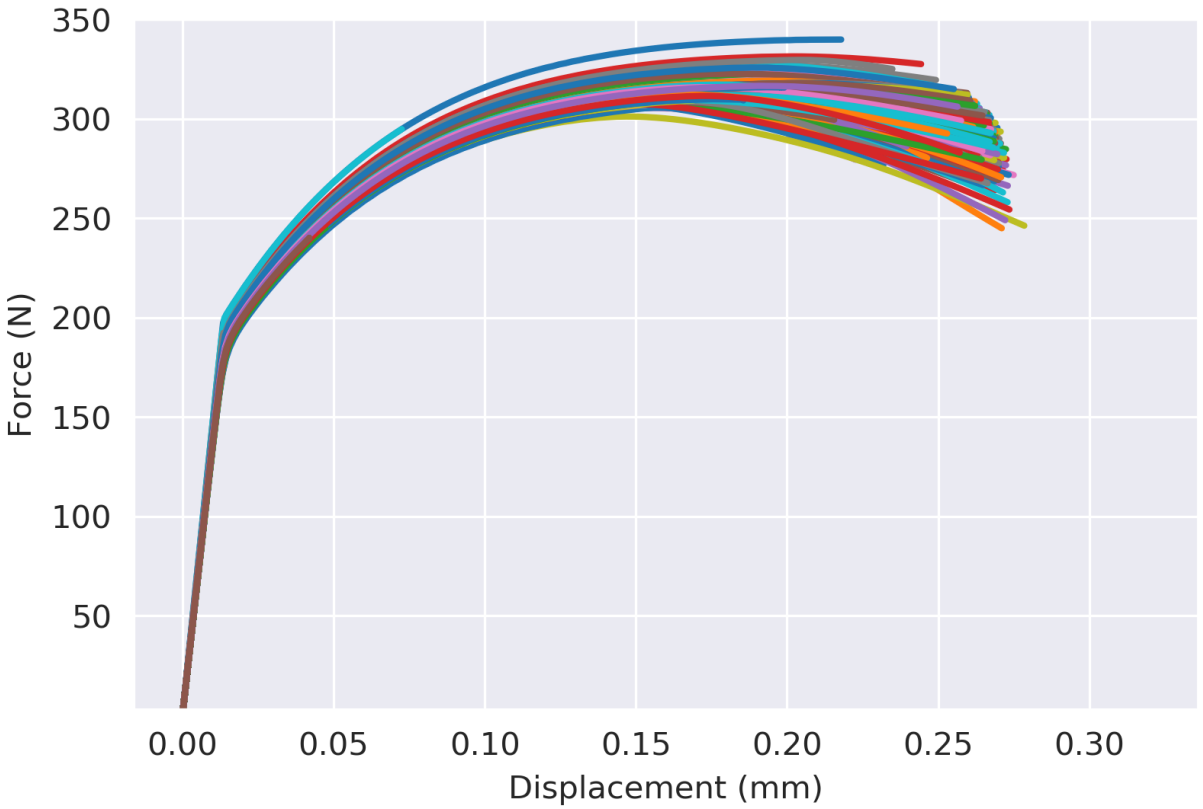

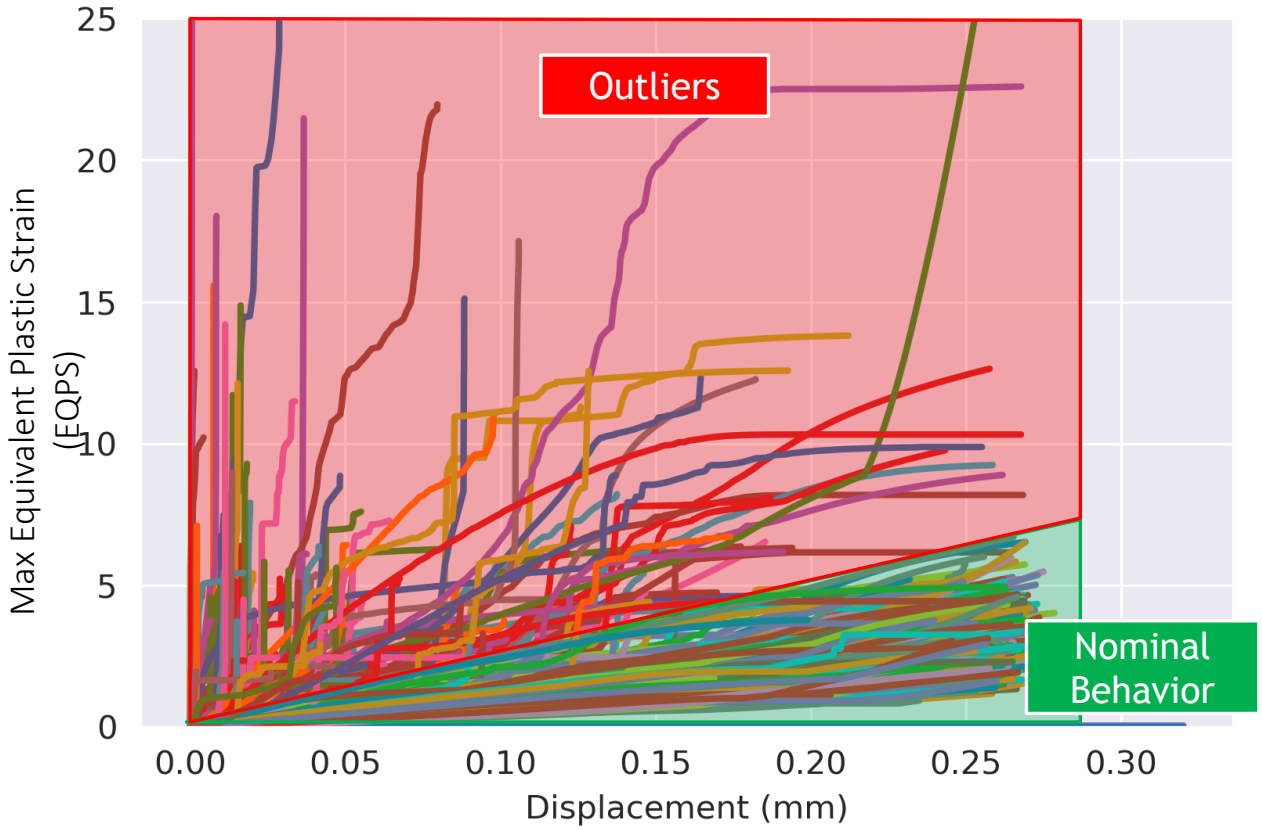

Deep Learning algorithm requires uniform voxel (3D Pixel) data format

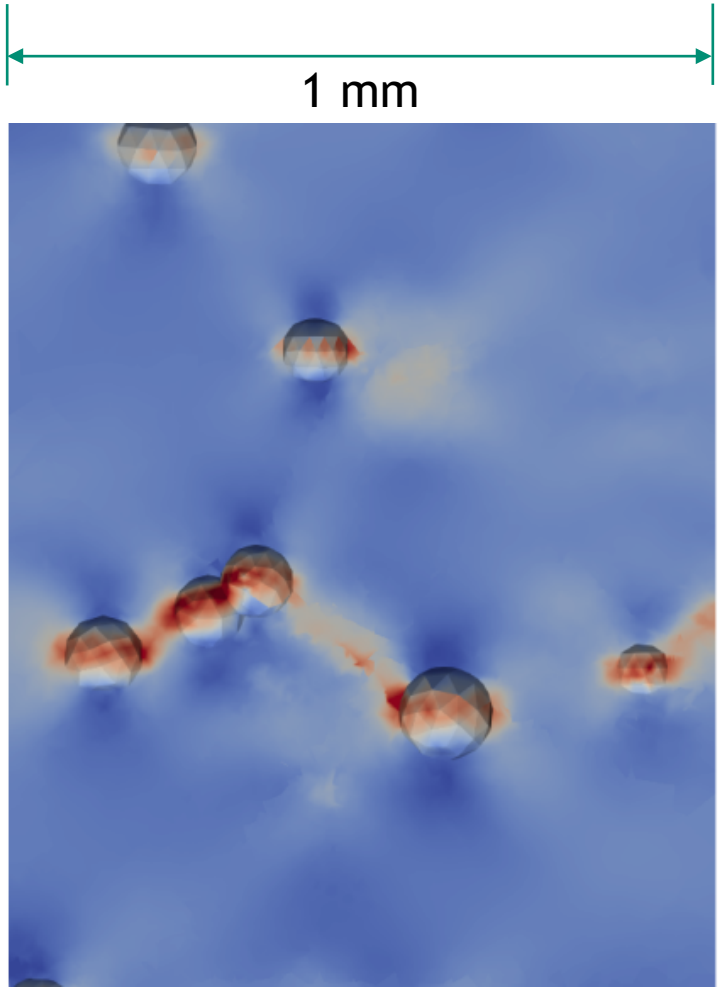

Sierra simulations results from tetrahedral mesh

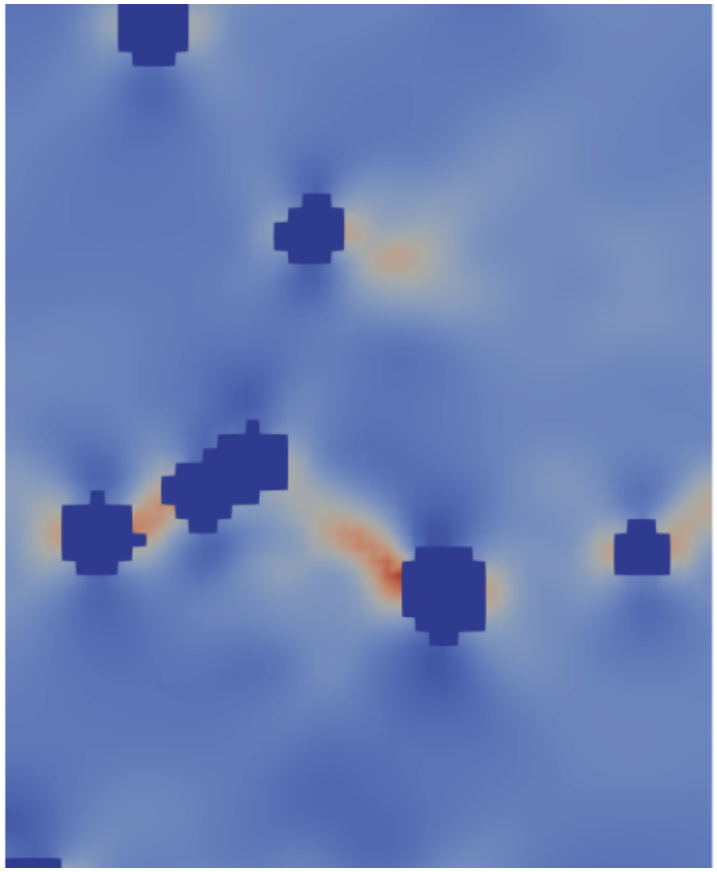

Results mapped to uniform hexahedral mesh

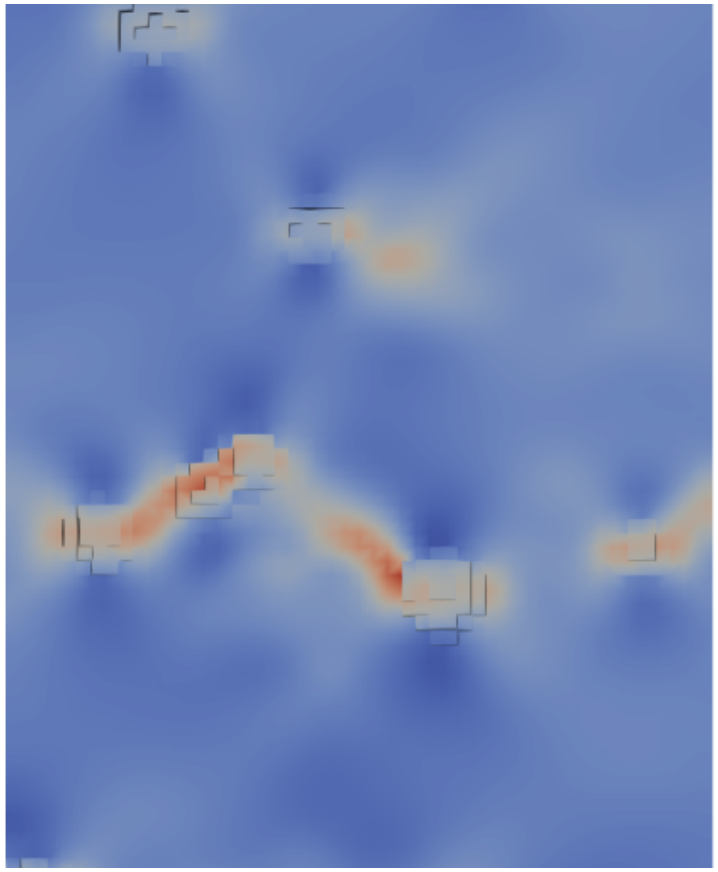

Hex elements in pores removed

Statistics of original data EQPS vs. mapped data EQPS illustrates smoothing

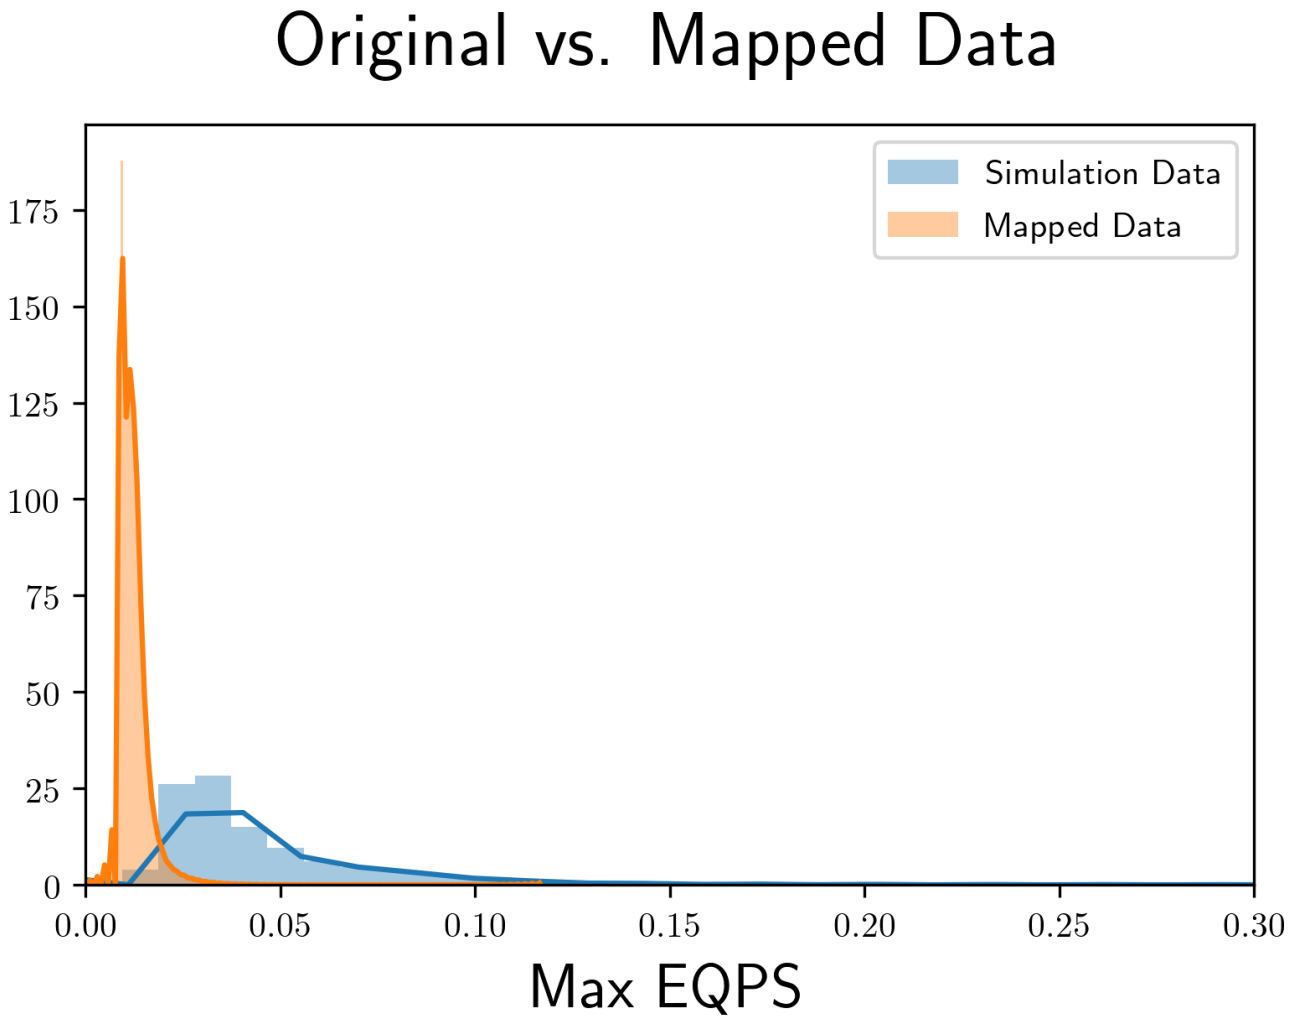

Specimens show large local stress state variations even in nominal uniaxial tension simulations > Reduces risk of extrapolation

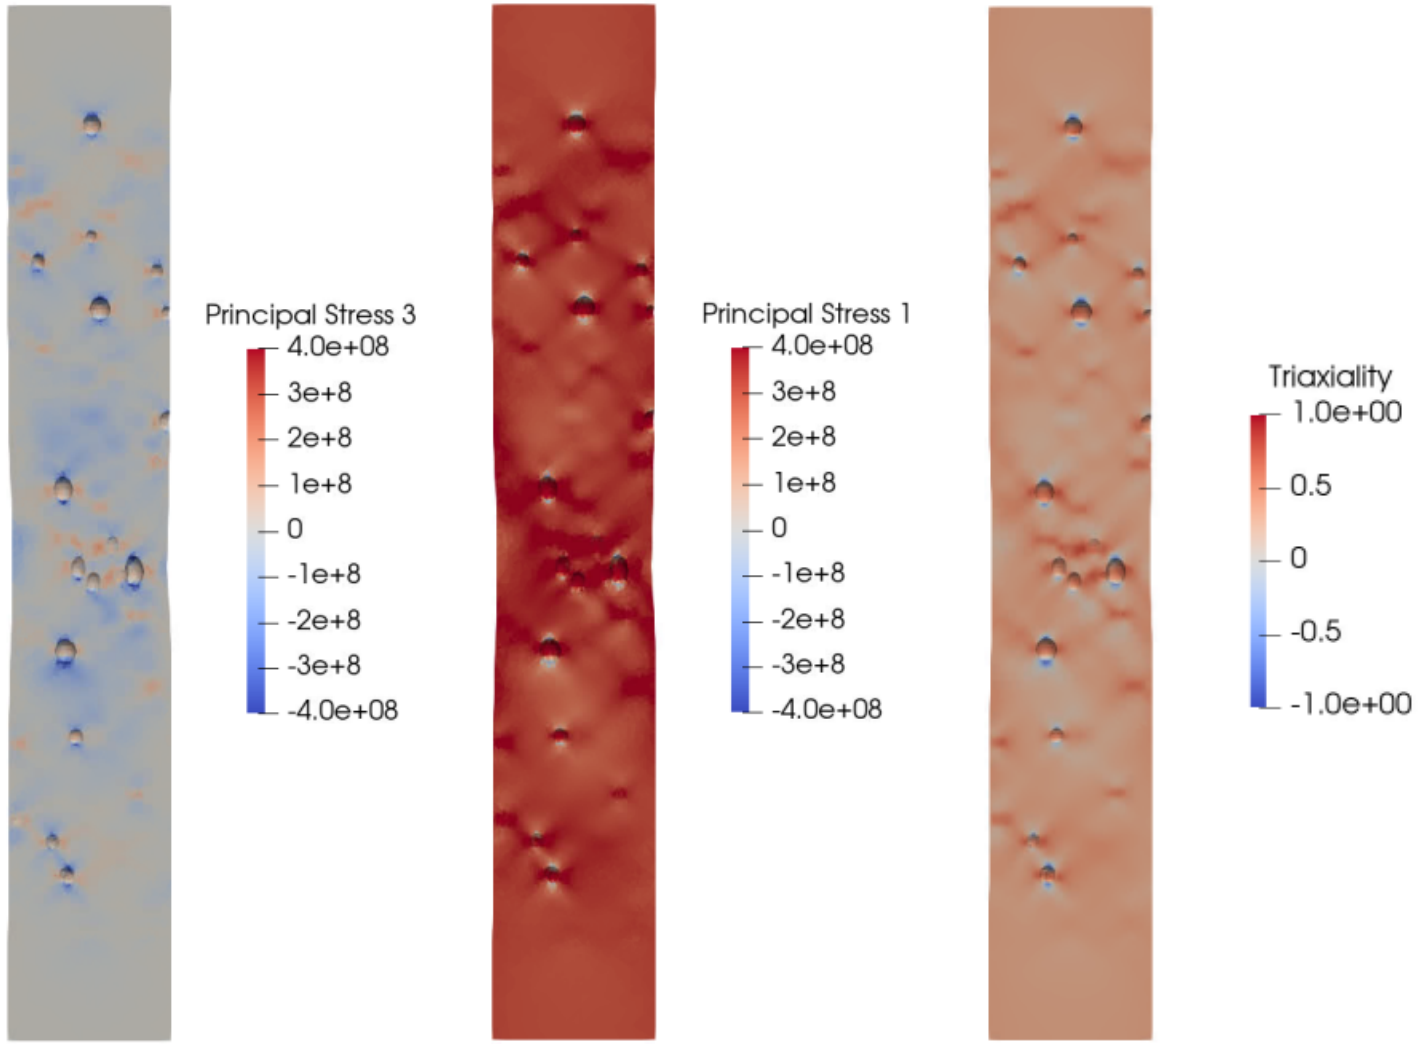

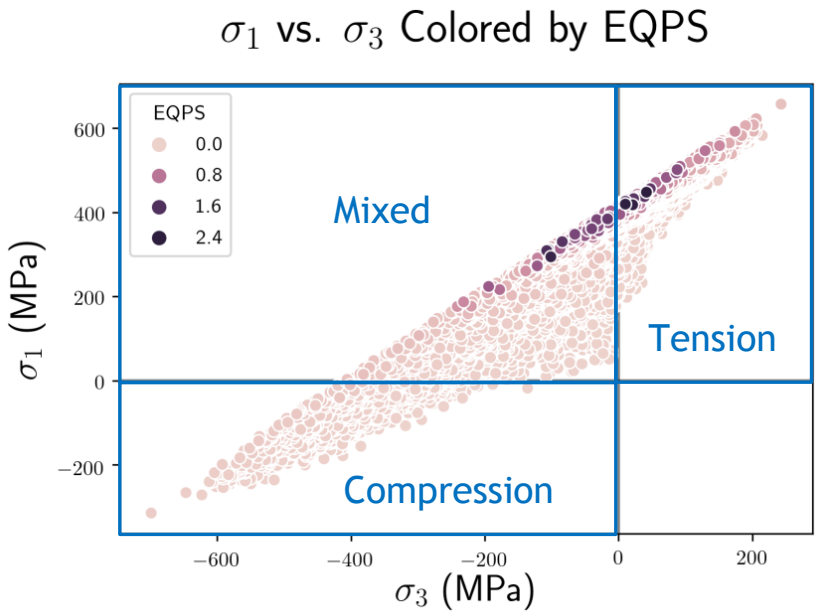

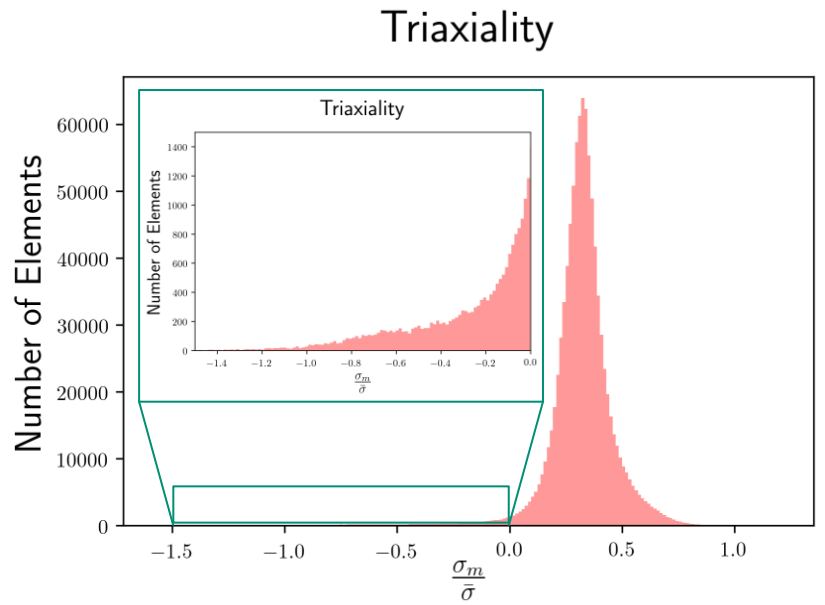

Specimens show large local stress state variations even in nominal uniaxial compression simulations > Reduces risk of extrapolation

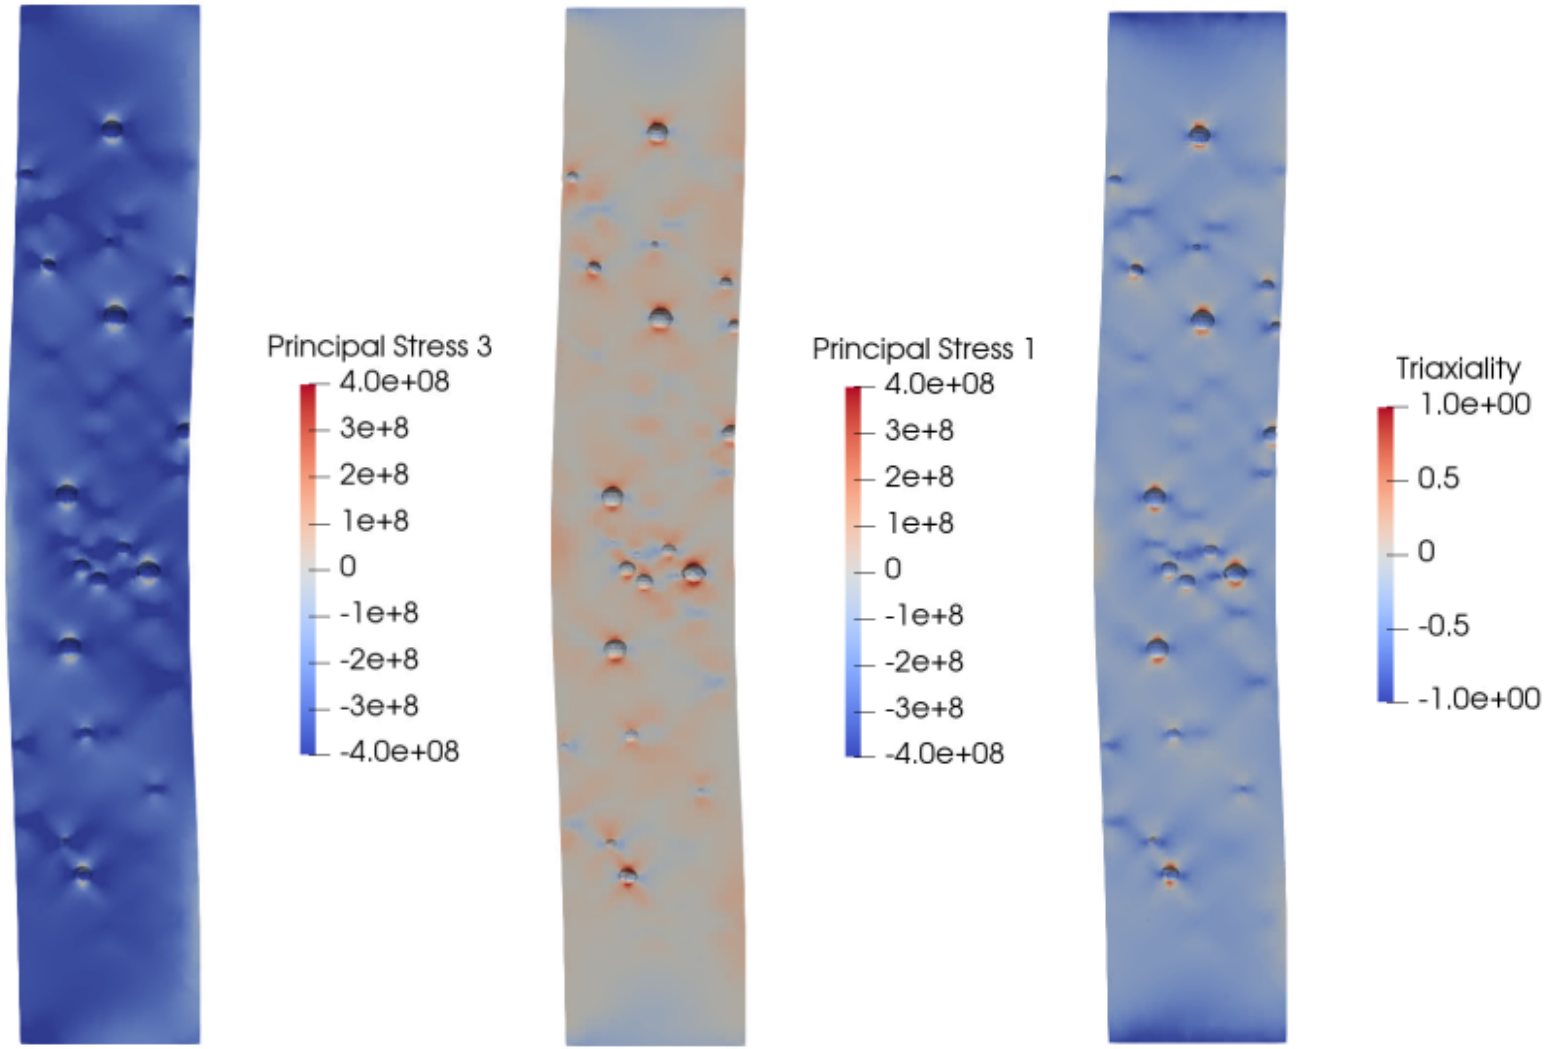

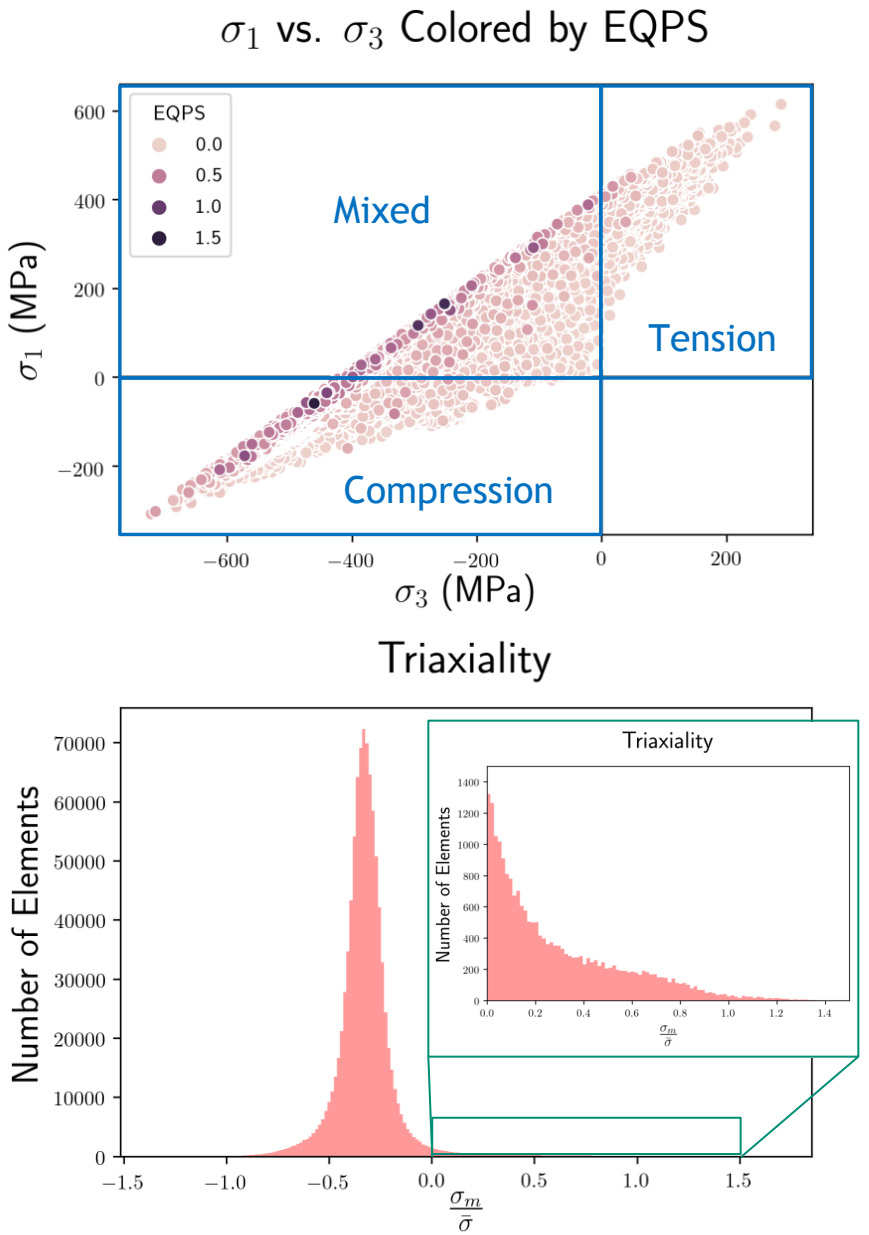

Deep Learning algorithm architecture

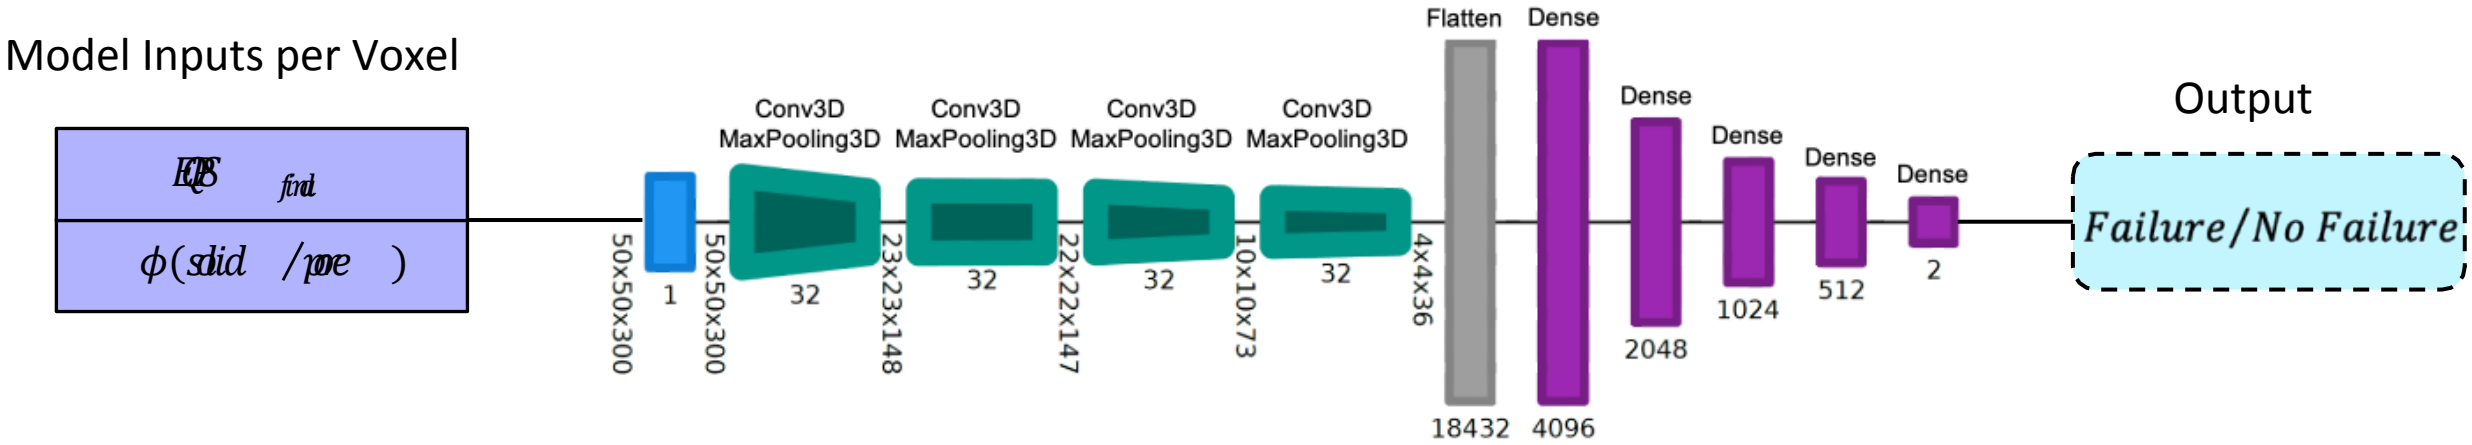

- Model architecture based on Huang et al. *Front. Neurosci*. 2019
- Output is classification – pass/fail for failure metric

Failure prediction results in test sets for network trained only on tension

Failure Metric: Sample failed to reach a required load before onset of strain localization.
FEA Simulation Time: 88 minutes on 216 cpus
DL Network Inference Time: 0.02 s on 2 GPUs
Speedup: 264000x

Square Tension

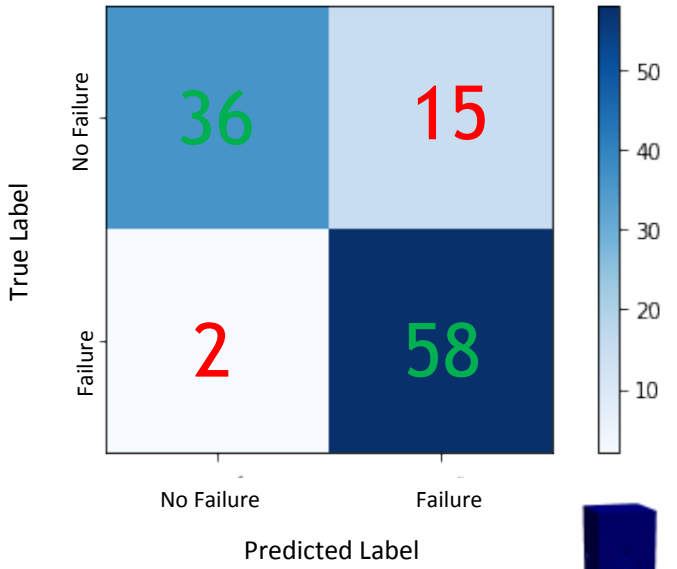

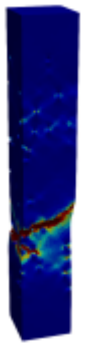

Test Accuracy: **84.7%**

Cylindrical Tension

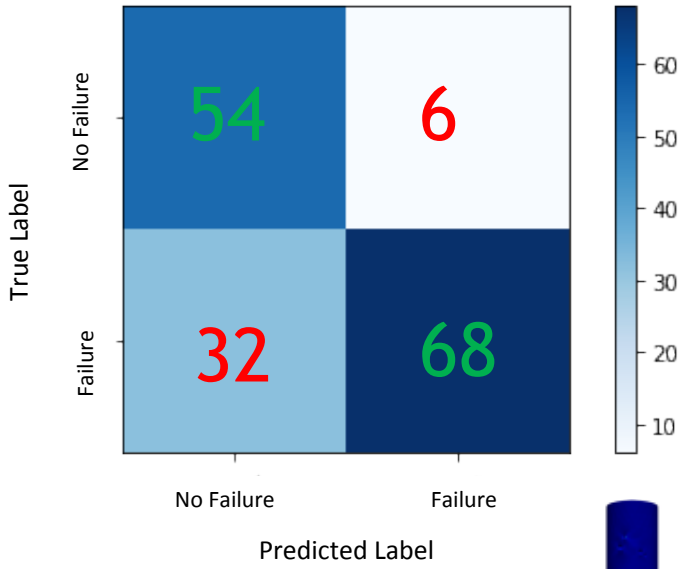

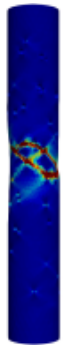

Test Accuracy: **76.2%**

Square Compression

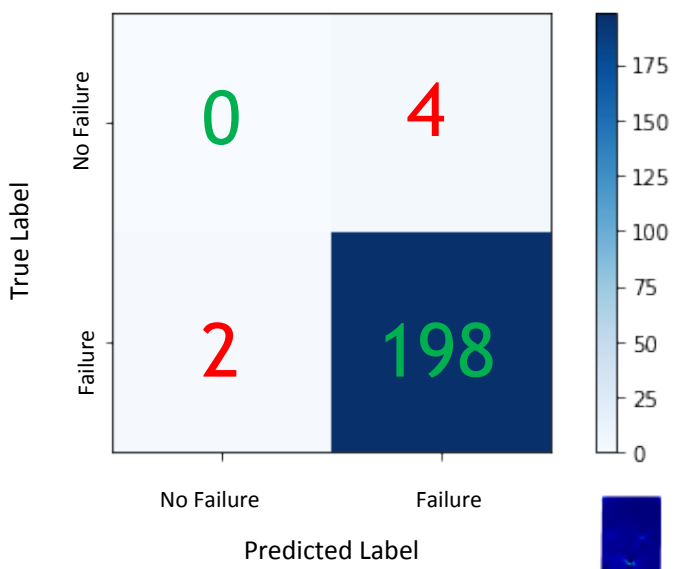

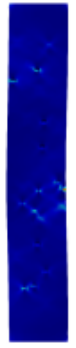

Test Accuracy: **97.1%**

Failure prediction results in test sets for network trained only on tension

Failure Metric: Sample Max EQPS remained below critical value and displacement reached a required value at onset of localization.

FEA Simulation Time: 88 minutes on 216 cpus

DL Network Inference Time: 0.02 s on 2 GPUs

Speedup: 264000x

Square Tension

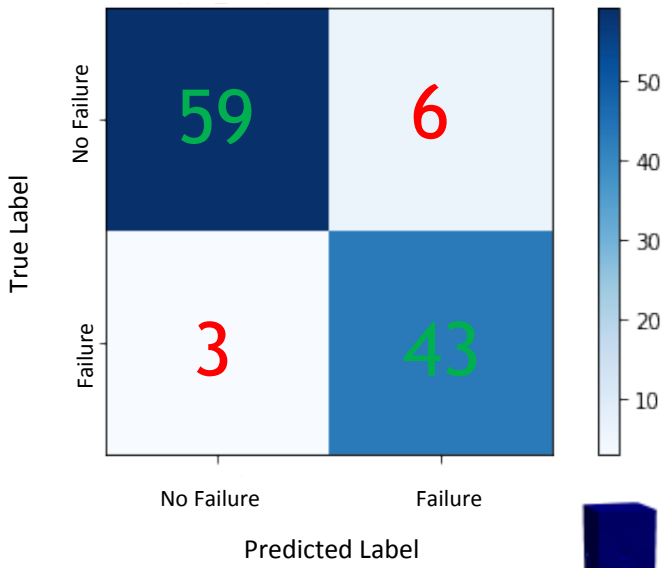

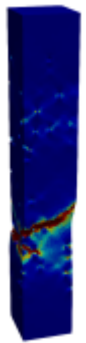

Test: Accuracy: **91.9%**

Cylindrical Tension

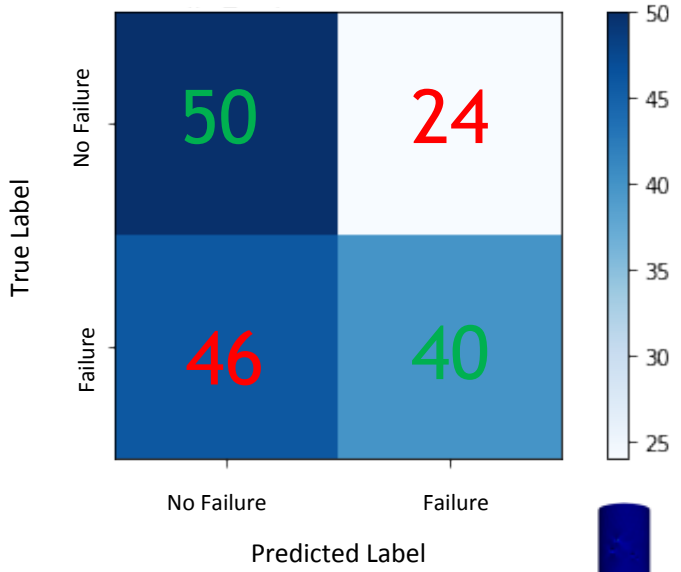

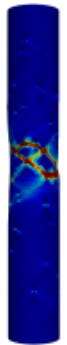

Test: Accuracy: **56.2%**

Square Compression

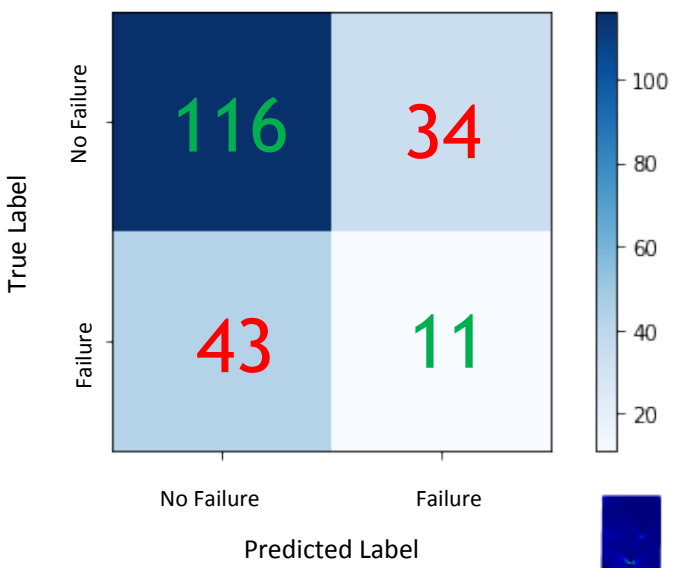

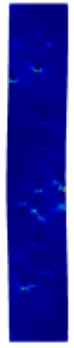

Test: Accuracy: **62.3%**

Failure prediction results in test sets for network trained only on tension

Failure Metric: Sample Max EQPS remained below critical value and force reached a required value at onset of localization.

FEA Simulation Time: 88 minutes on 216 cpus

DL Network Inference Time: 0.02 s on 2 GPUs

Speedup: 264000x

Square Tension

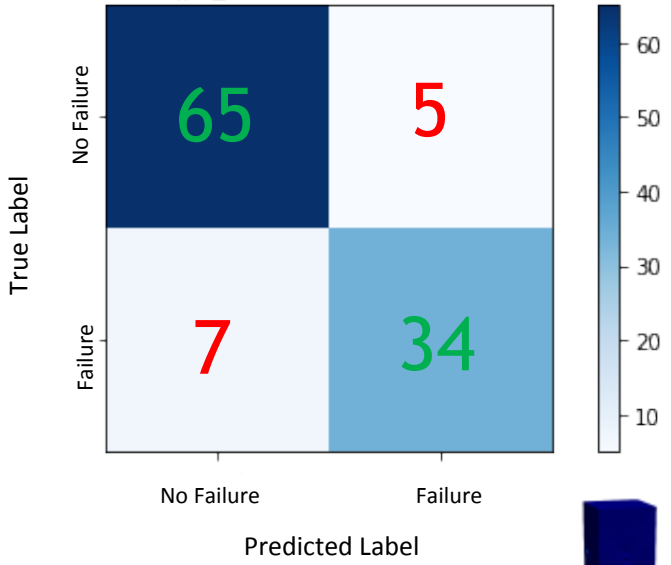

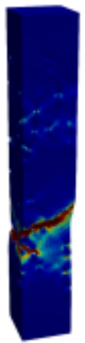

Test Accuracy: **89.2%**

Cylindrical Tension

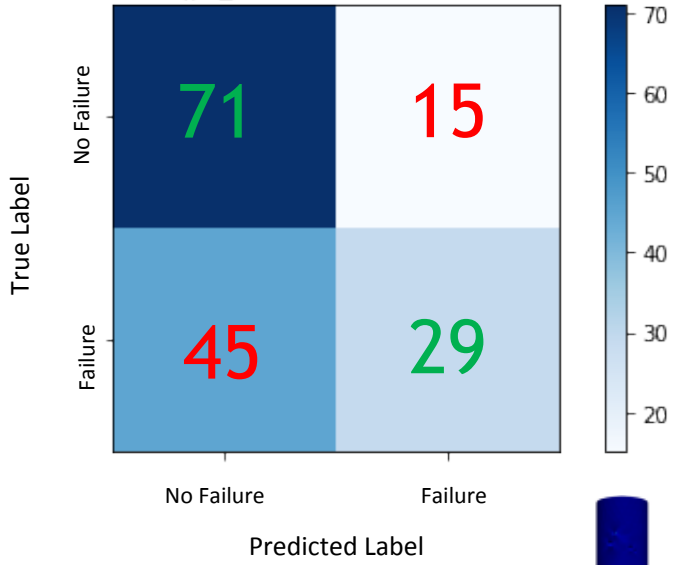

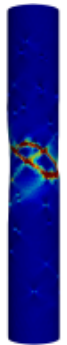

Test Accuracy: **62.5%**

Square Compression

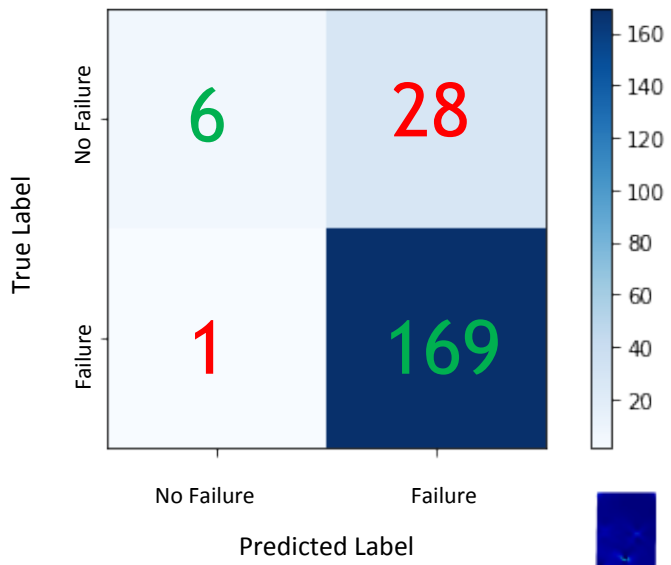

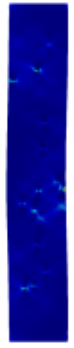

Test Accuracy: **85.8%**

Failure prediction results in test sets for network trained only on tension

Failure Metric: Sample displacement reached a required value at onset of localization.
FEA Simulation Time: 88 minutes on 216 cpus
DL Network Inference Time: 0.02 s on 2 GPUs
Speedup: 264000x

Square Tension

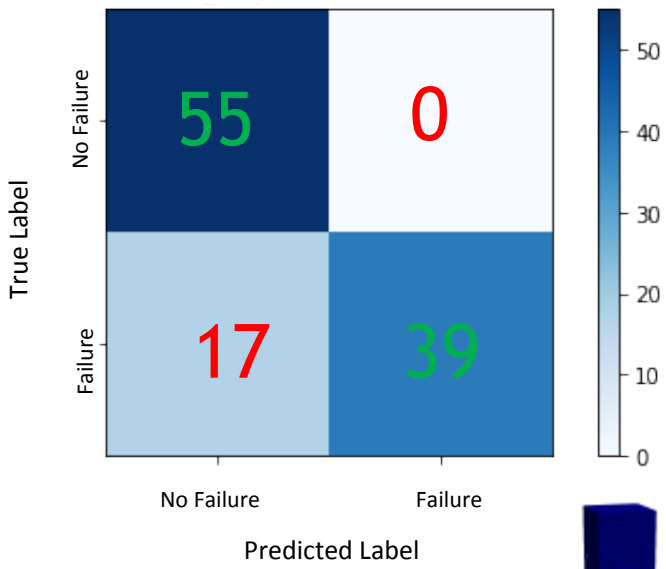

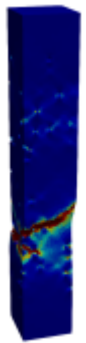

Test Accuracy: **84.7%**

Cylindrical Tension

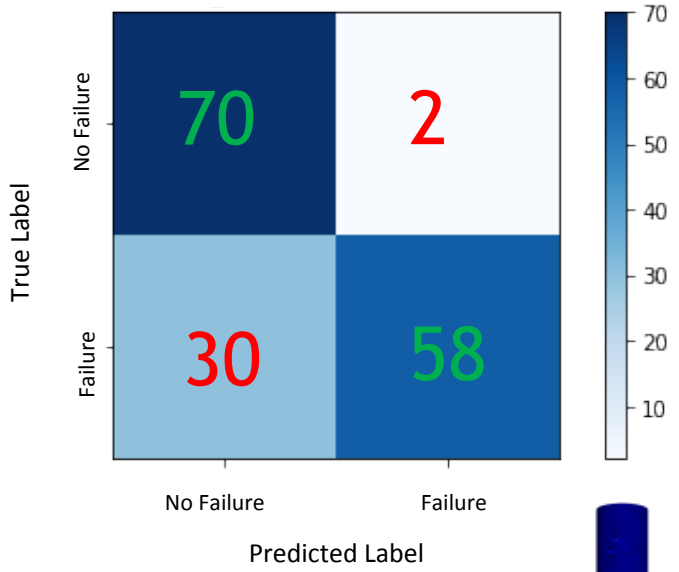

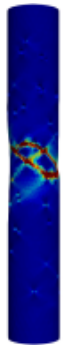

Test Accuracy: **80.0%**

Square Compression

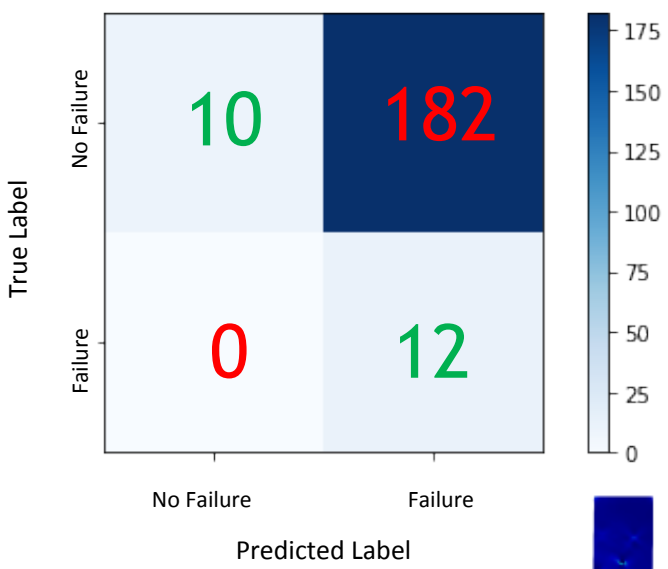

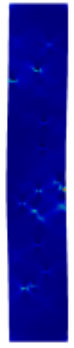

Test Accuracy: **10.8%**

Memory management becomes an issue with large datasets

Current 3D training samples contain 750k voxels – approaching memory limits on GPU.

Multiple components of the DL process live in GPU memory simultaneously.

- DL model weights
- Activation function values
- Backpropagation update values
- 3D element values (batch of inputs)

Scaling options

- Batch size limitations
- Smaller/simpler architecture

- Loading different architecture layers across GPUs
- Loading single architecture layer across GPUs

- Patch-wise CNNs
- Physics Informed Neural Networks – use physical insights to perform computations

Increasing Complexity

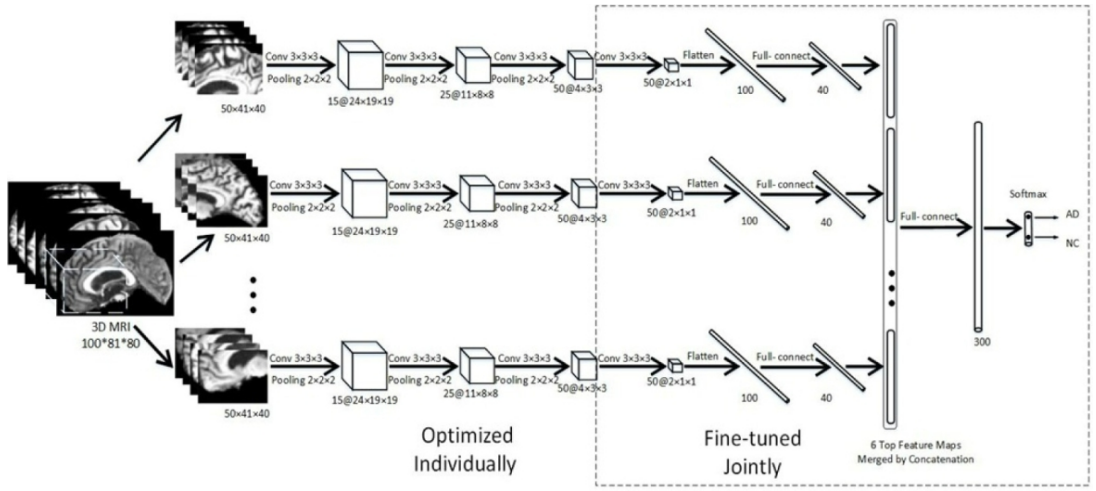

Patch-level CNN Classification (Cheng et al., ICDIP Proceedings 2017)

Using GANs to augment CT images of AM material – Collaboration with Amir Farimani and Francis Ogoke (CMU)

- Generative Adversarial Networks (GANs) produce **new samples** from a training set while **preserving the underlying statistics**.

- GANs are trained to minimize the distance between the distribution of the training data and the generated samples.

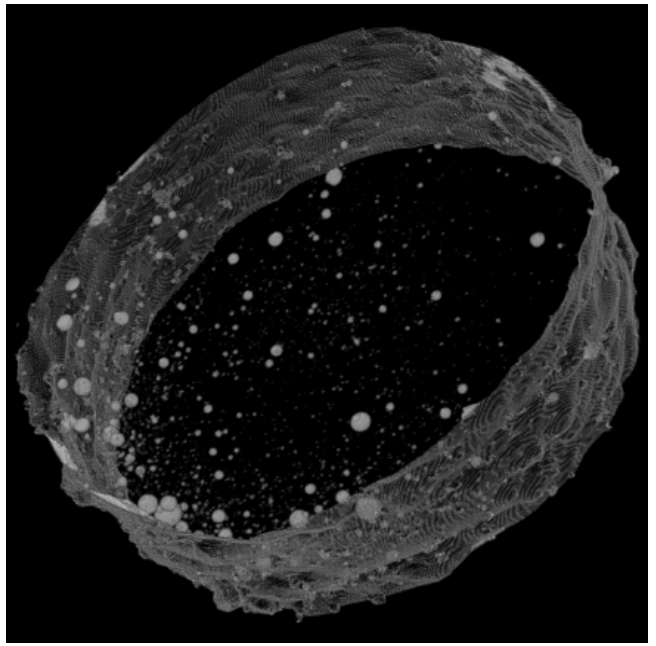

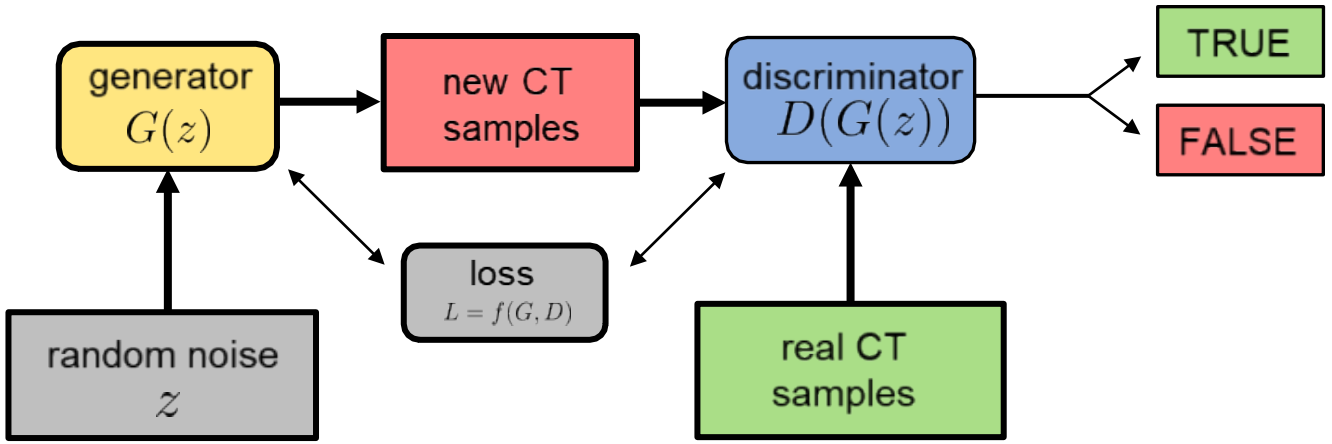

GAN network schematic

Metric	Notes
Volume Distribution	
Nearest Neighbor Distances	
Location Distribution	
Ellipticity	$\sqrt{\frac{a^2-c^2}{a^2}}$
Moment of Inertia	
Surface Area	
Mallat Scattering Transform	$\|\|x * \psi_{\lambda_p}\| * \psi_{\lambda_{p'}}\| * \phi_J$

Conclusions

- Deep Learning was able to predict performance variation due to microstructural features 5 orders of magnitude faster than FEA (nearly instantaneous).
- Model maintained predictiveness in different part geometries and stress states.
- DL is able to pick up on patterns that subject matter experts cannot. Prior to this work we explored looking at stress measures to predict ductility with little success.
- Using large datasets for DL training, such as volumetric data used here, on GPUs is a challenge.

Vision: Rapid failure prediction based on microstructure, geometry, and loading conditions enabled by Deep Learning

Microstructure

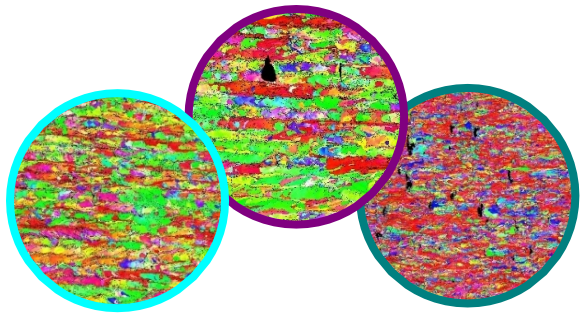

As-built Geometry

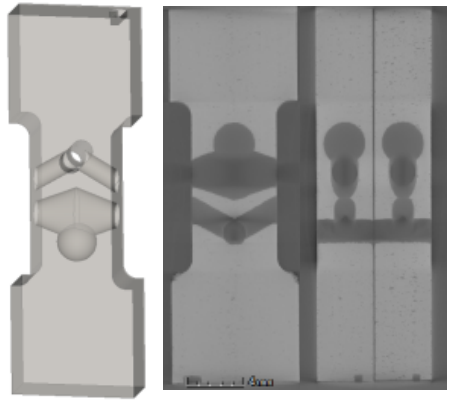

Boundary Conditions

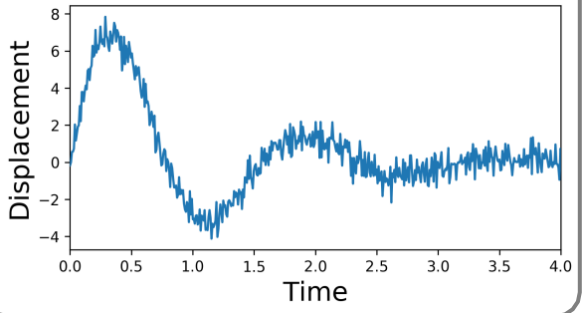

Simulation Code

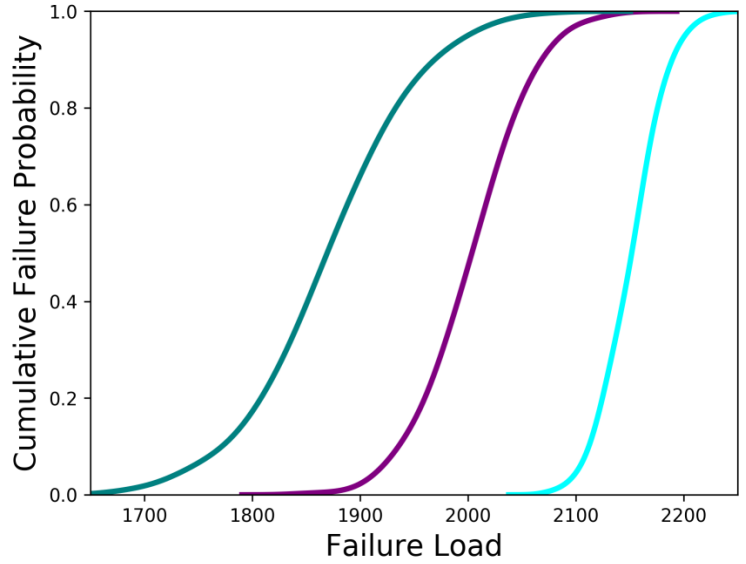

Image: oldcomputers.net

New 3-year project: Failure Prediction Using Deep Learning in AA6061

- Extending approach to include Diffraction Contrast Tomography (DCT), in situ µCT, Digital Volume Correlation (DVC), crystal plasticity, and continuum damage modeling.
- Collaboration with Prof. Mike Sangid (Purdue) for HEDM
- Collaboration with Billy Oates and Brandon Krick (FAMU/FSU) for TEM

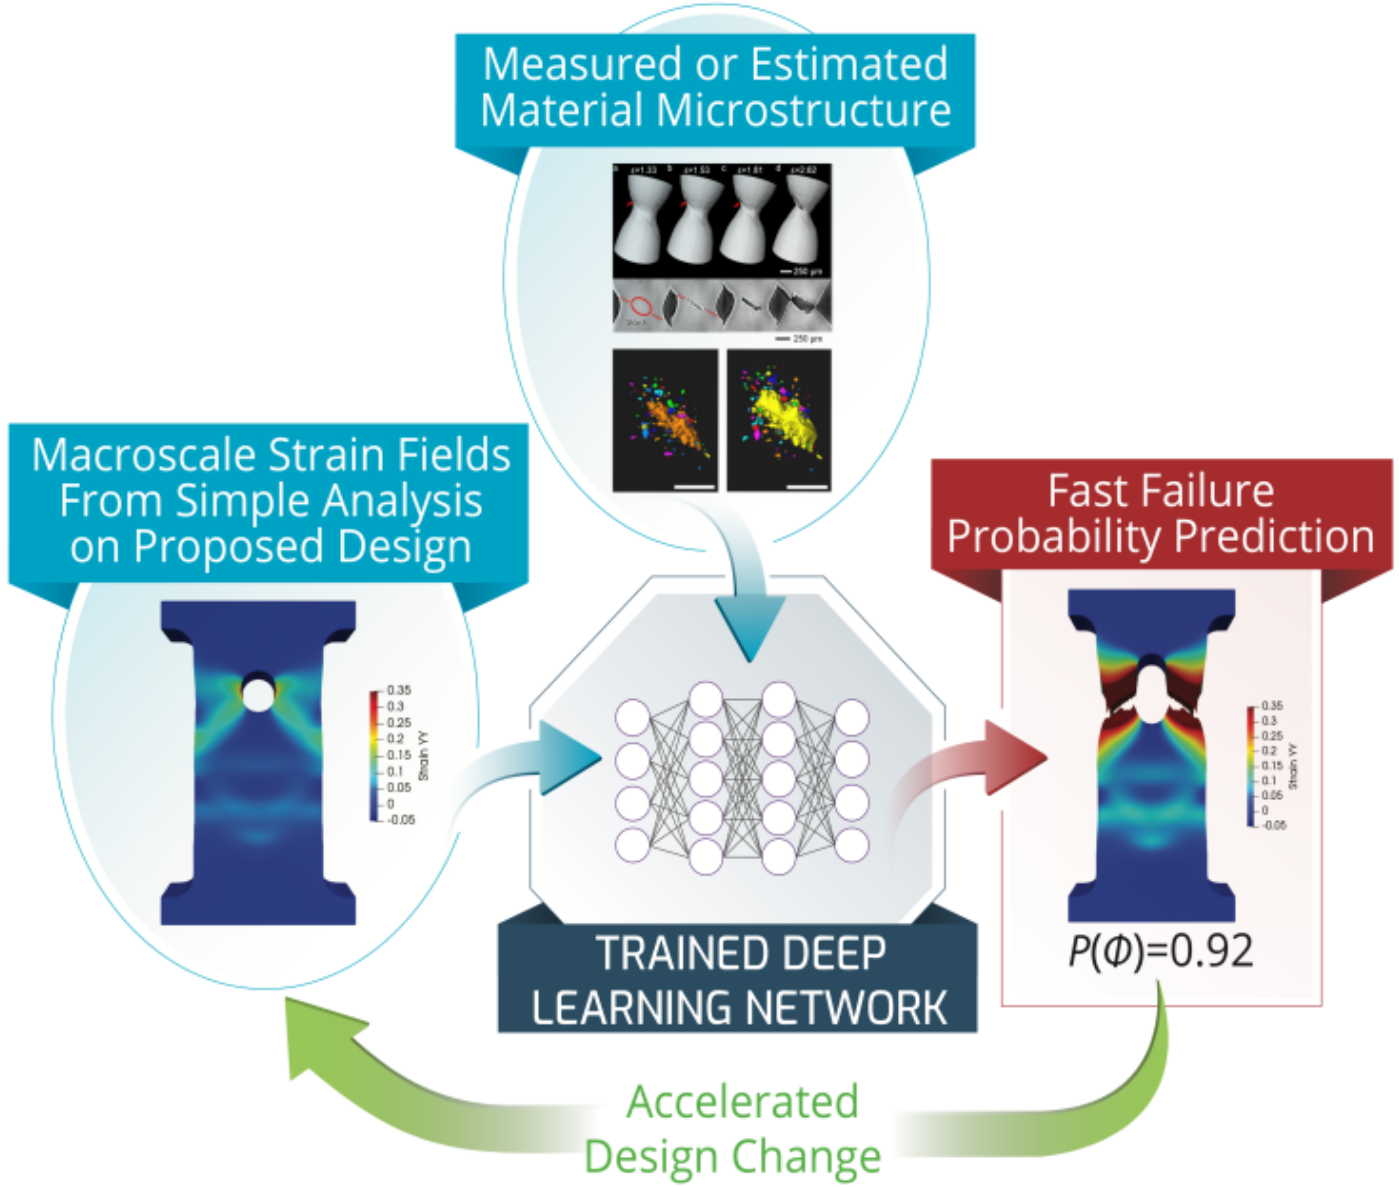

Diffraction Contrast Tomography (DCT) to Crystal Plasticity workflow enables simulation of real microstructures

- ~156 grains, 2.5 um voxel size
- Removed all grains < 10 voxels
- 15,410,688 data points (254x237x254)

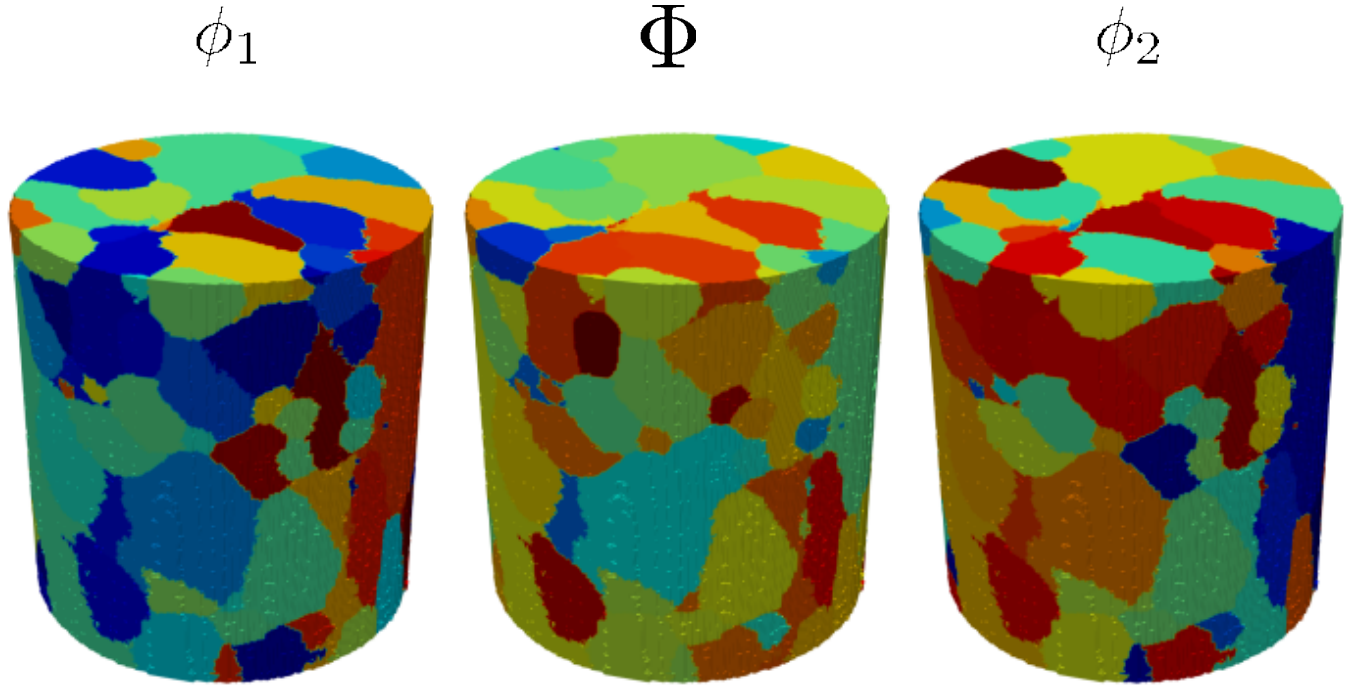

DCT video of tensile gauge section: Philip Noell
Crystal Plasticity: Hojun Lim

Microstructure information

Full resolution (no defeature)
156 grains
9,284,343 elements

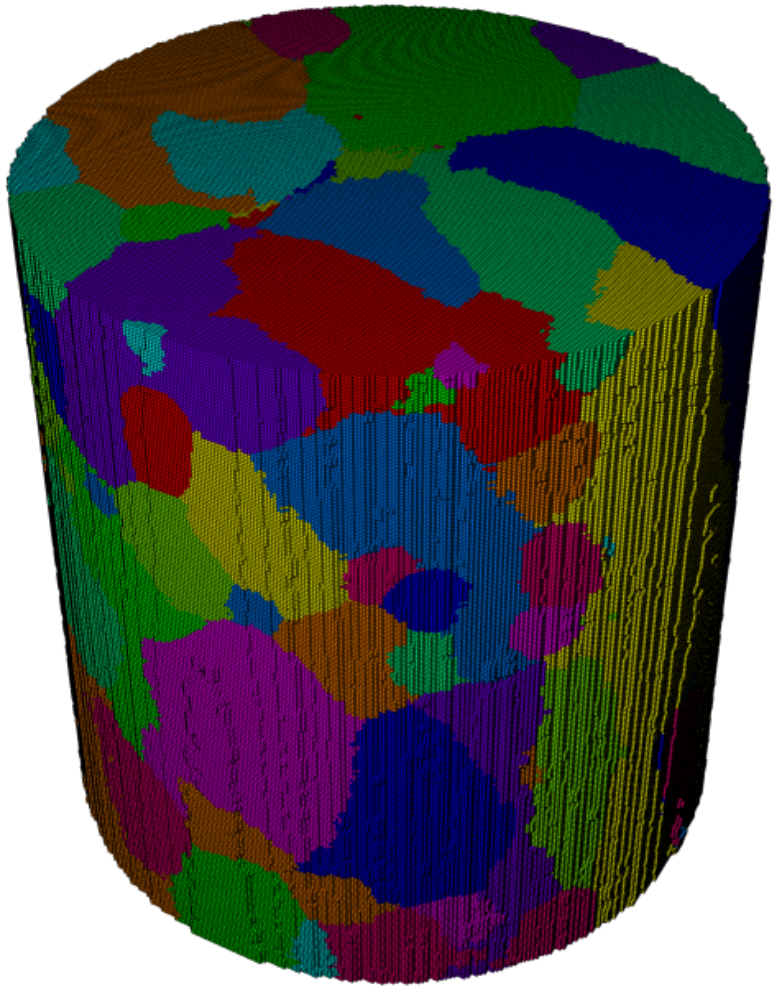

Grain volume distribution

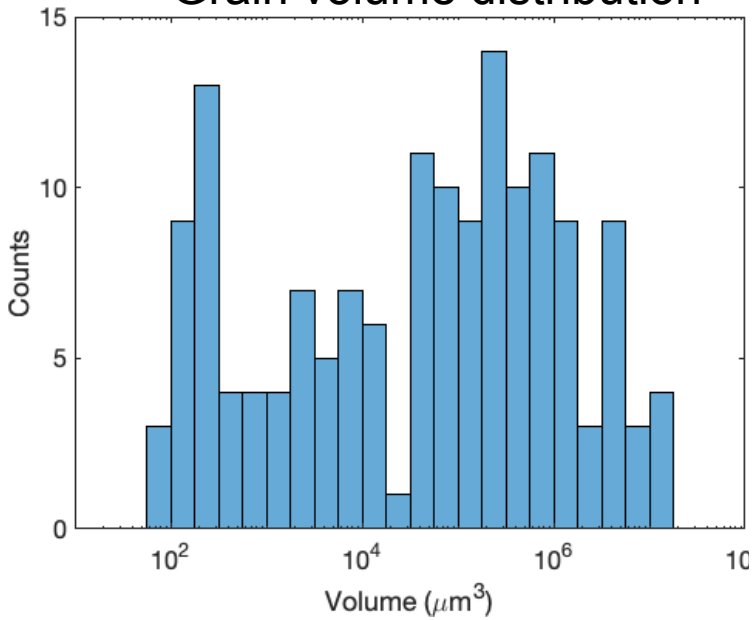

Grain volume distribution (voxels)

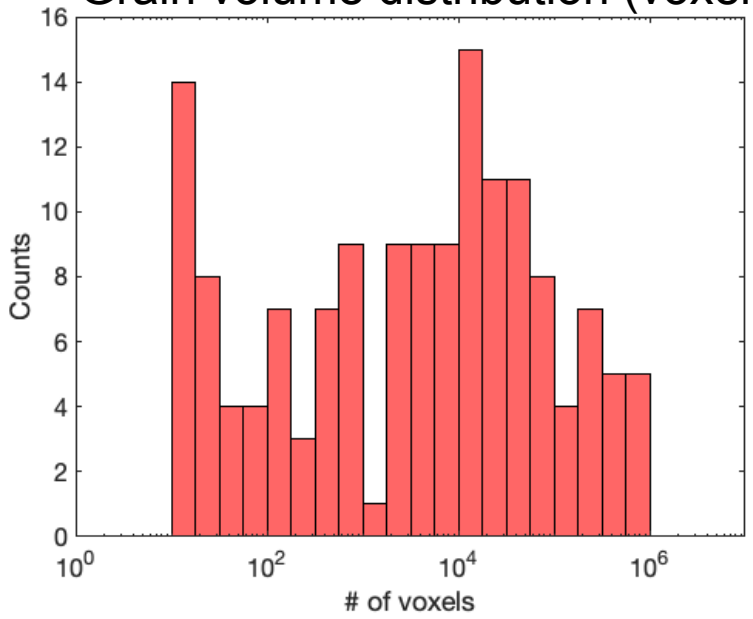

Initial texture

(100)

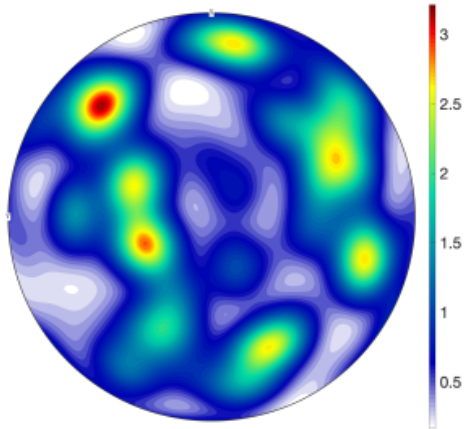

(110)

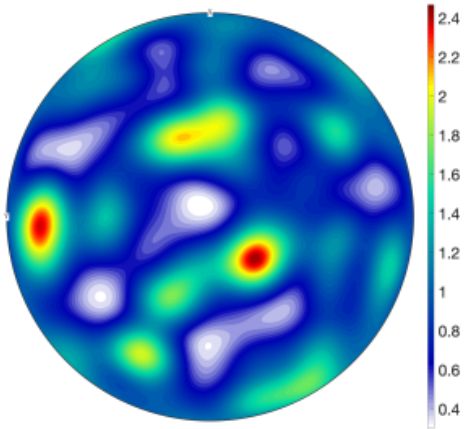

(111)

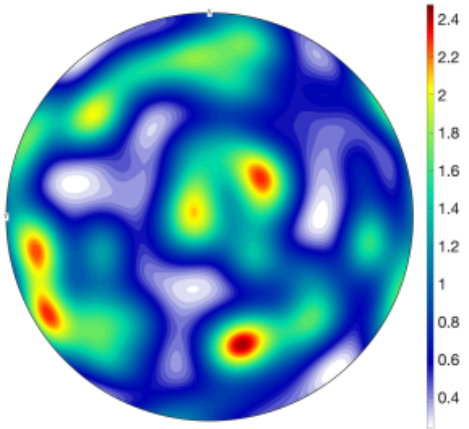

Crystal plasticity simulations – 10% uniaxial tension

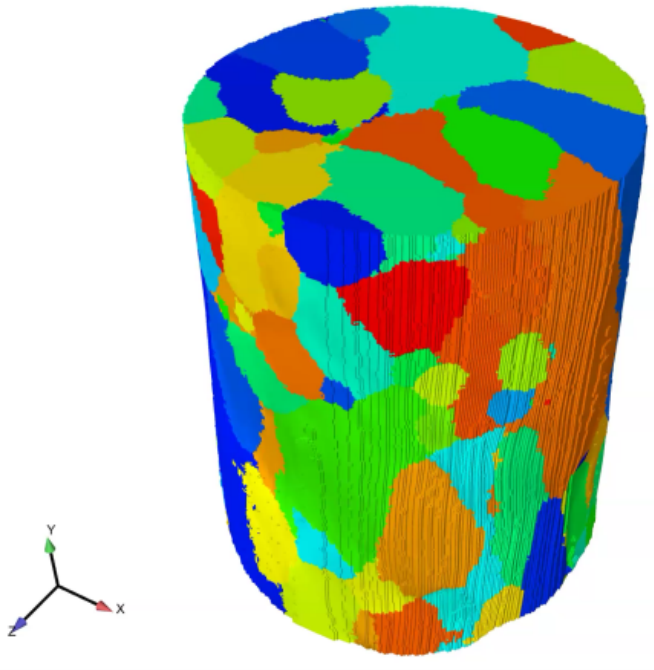

Grain IDs

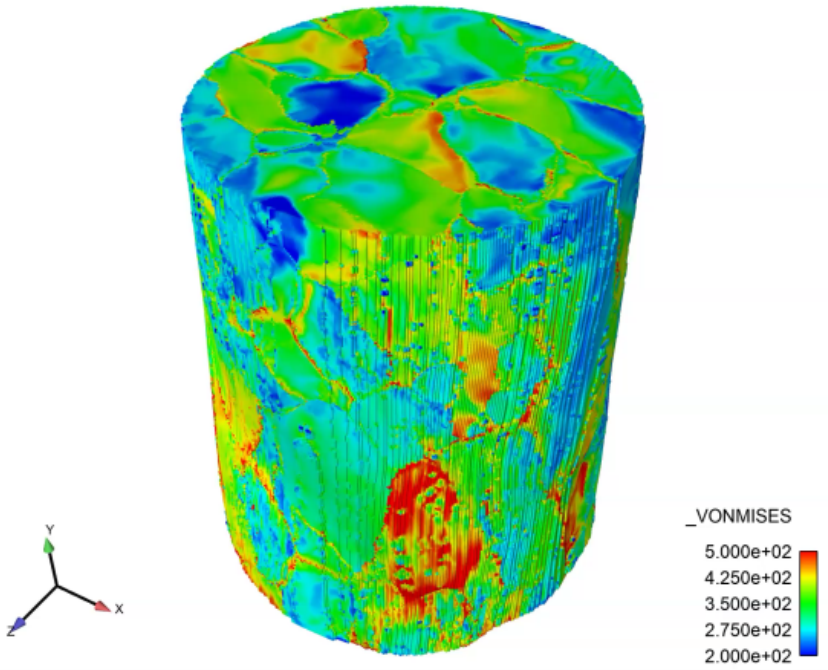

Von Mises stress (MPa)

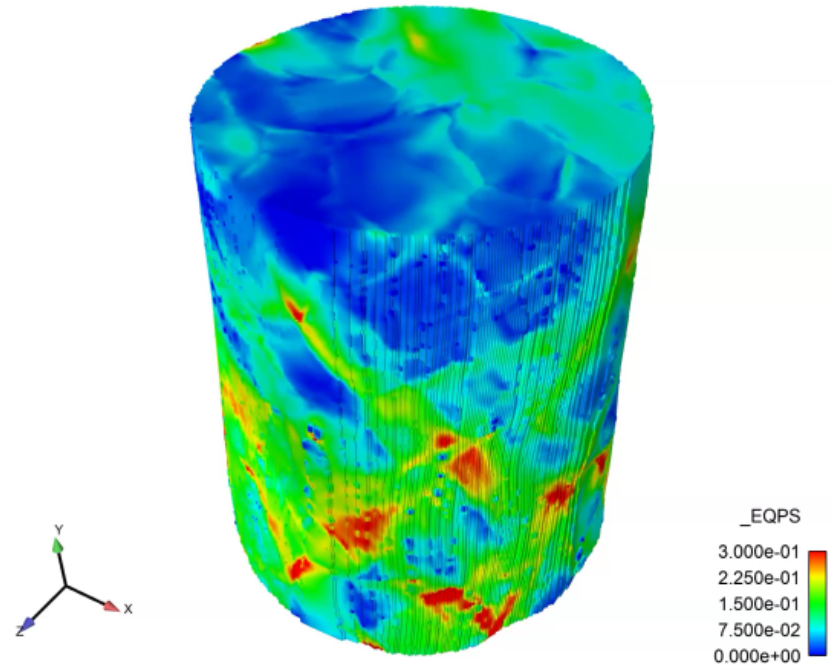

Equivalent plastic strain

Von Mises stresses at 10% deformation – 4 largest grains

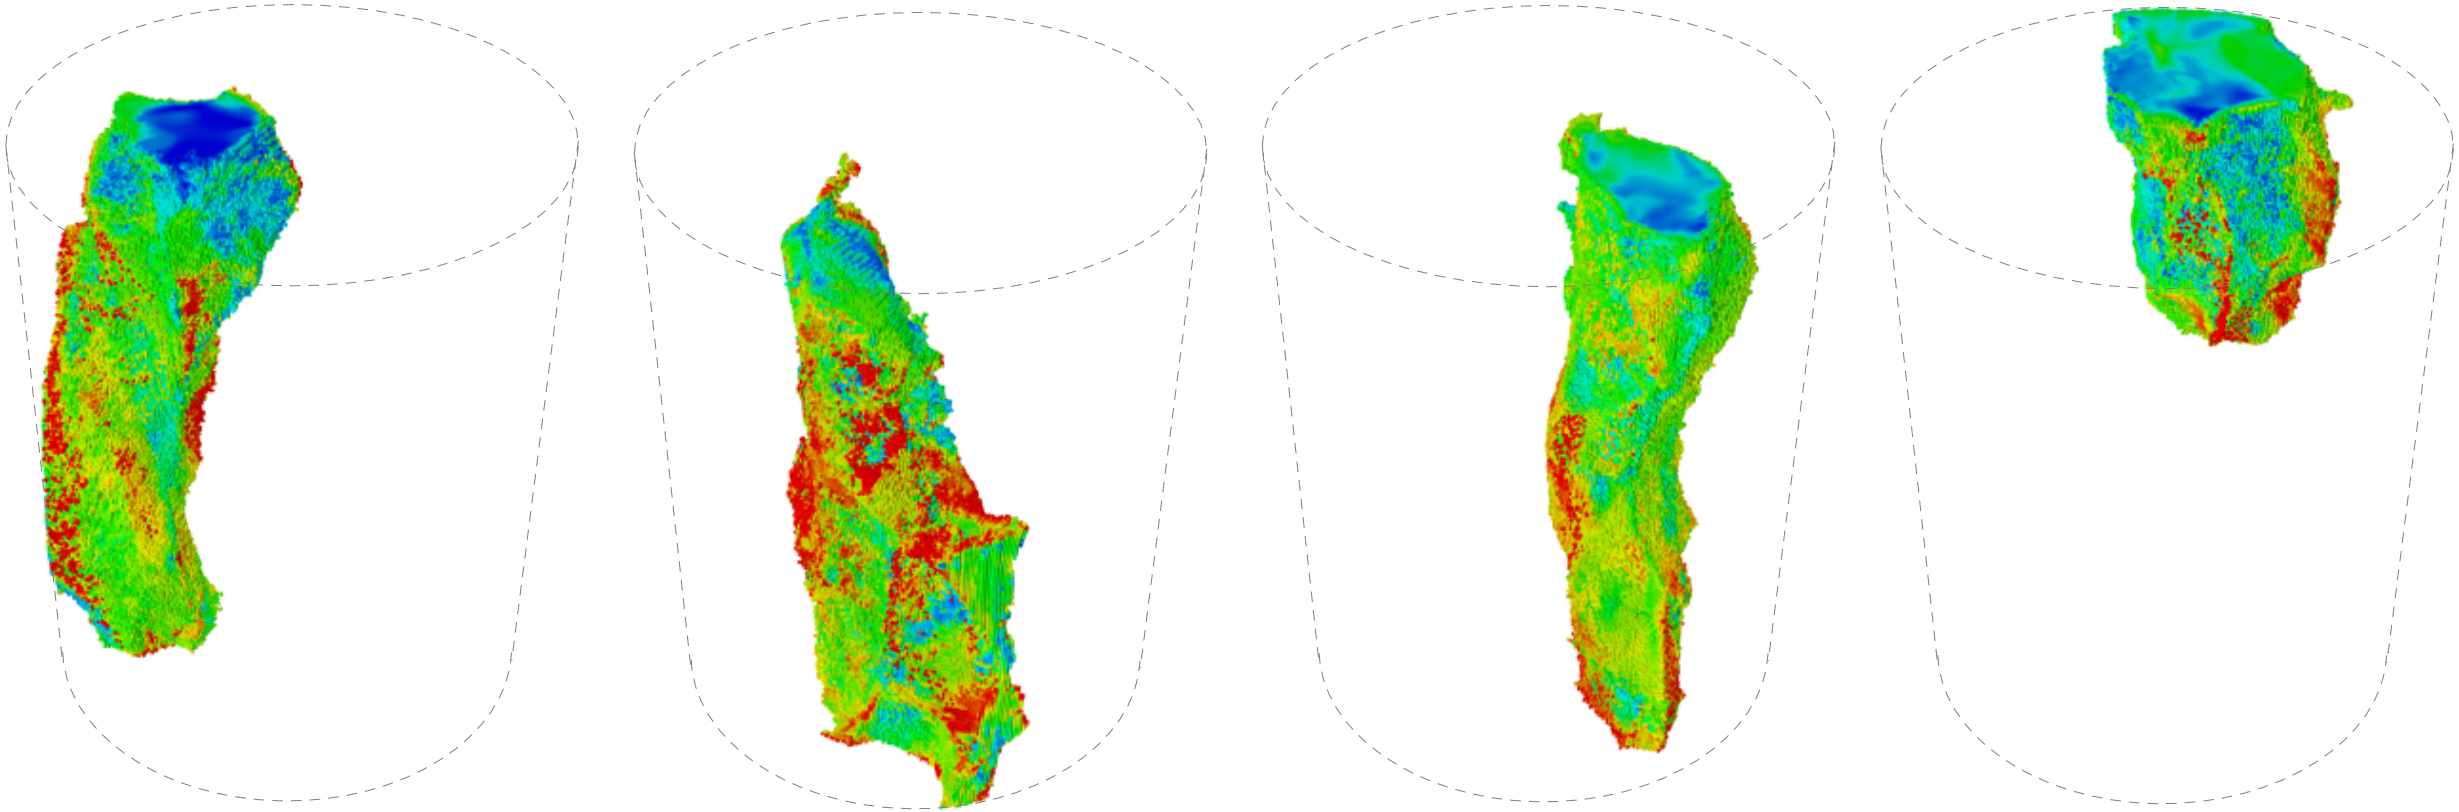

_VONMISES
5.000e+02
4.250e+02
3.500e+02
2.750e+02
2.000e+02

Grain #100
V_f =9.5%

Grain #119
V_f =7.6%

Grain #33
V_f=7.5%

Grain #75
V_f =7.2%

Future Work

- Multiscale coupling – mesoscale CP to macroscale continuum damage simulations
- Digital Volume Correlation (DVC) testing
- In situ micron-scale CT testing
- High Energy Diffraction Microscopy (Prof. Mike Sangid)
- TEM for failure initiation mechanisms (Profs. Billy Oates and Brandon Krick)
- TriBeam characterization on deformed DCT+CP simulated sample
- Combining all of the above in an ML/DL model for failure predictions

Questions?

kyljohn@sandia.gov