SAND2018-12283PE

Machine Learning for CT Segmentation

Cari Martinez, Kevin Potter

Sandia Interdisciplinary Machine Learning Research (SIMLR) Team

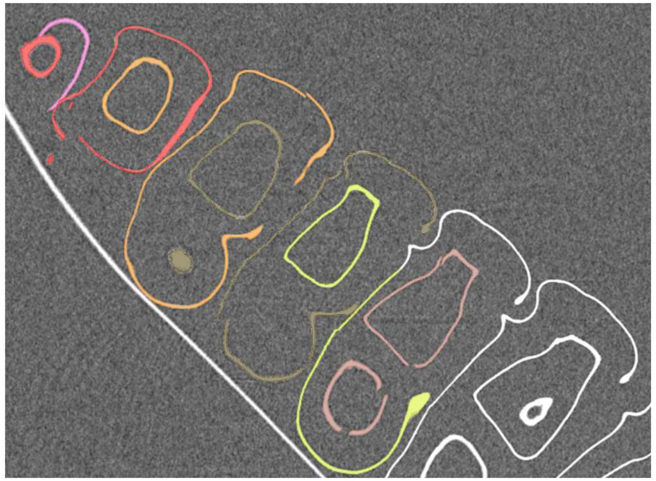

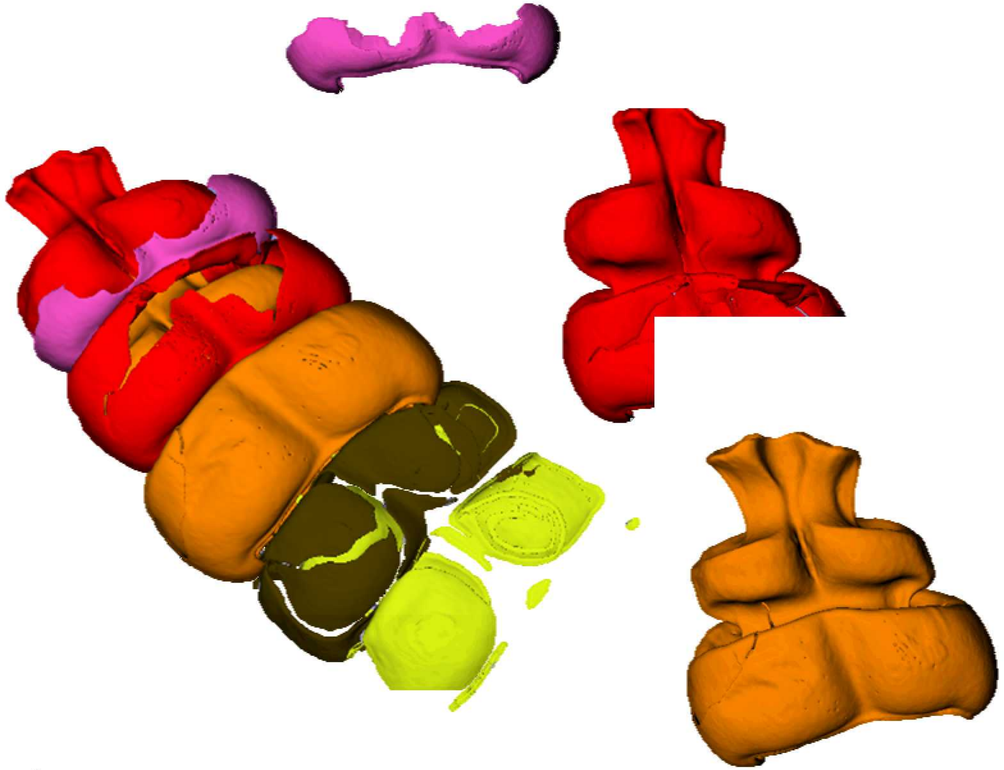

Using ML to save time and effort while improving accuracy

CT SEGMENTATION

Segmentation is a classic computer vision problem

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

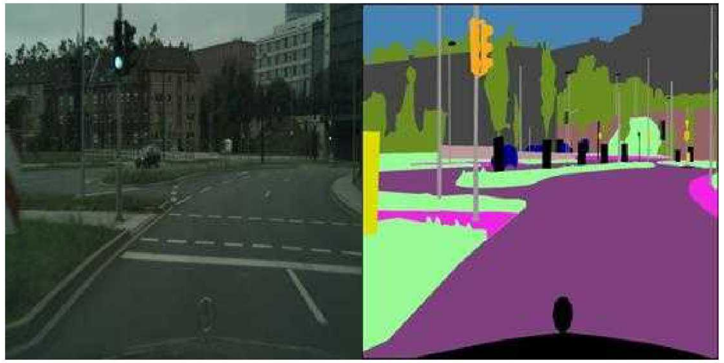

https://www.cityscapes-dataset.com/

Cityscape
(~1e5 pixels)

- 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

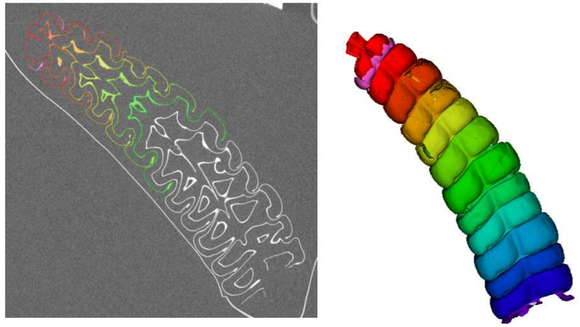

Rattlesnake Tail
(~1e9 voxels)

- Medical researchers are leading this work toward Deep Learning solutions

Digital Twins are a bold concept

Collaborators – John Korbin & David Peterson

- Takes human expert 8 hours to segment
- Deep Learning can play a large role across the full pipeline

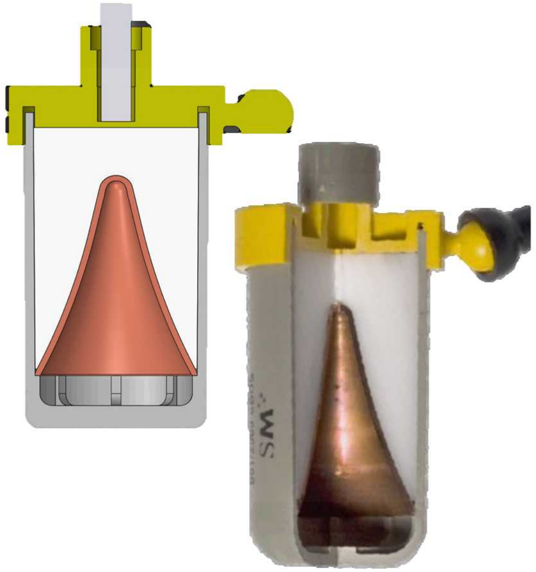

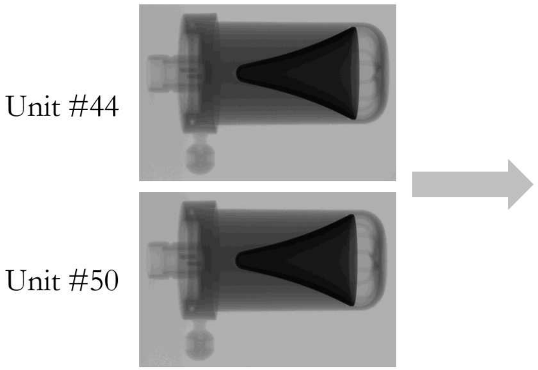

CT Scanning Process

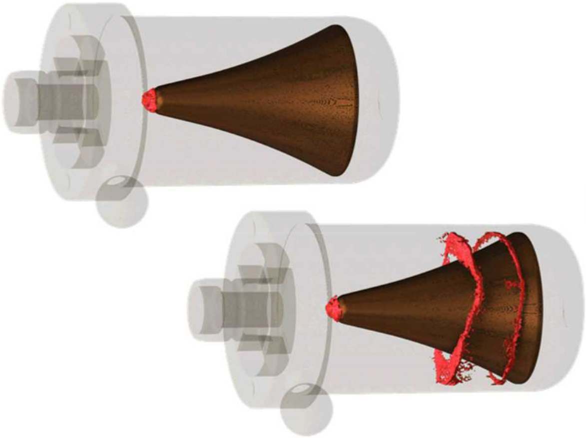

Segmentation and Defect Identification

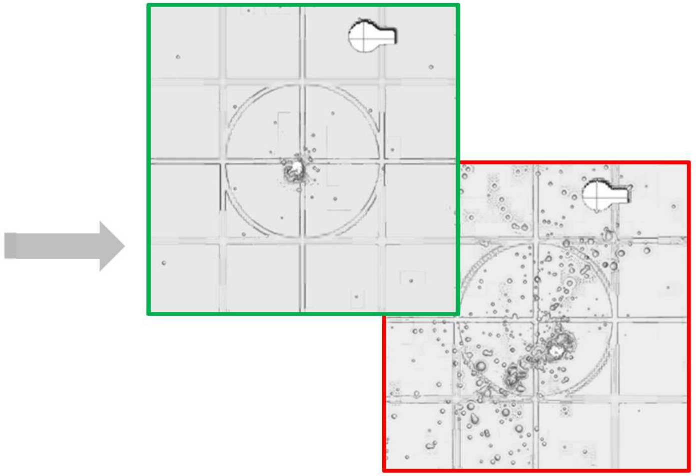

Simulation of Individual Performance

CT Segmentation is hard for humans

- CT scans must be labeled by component for simulations

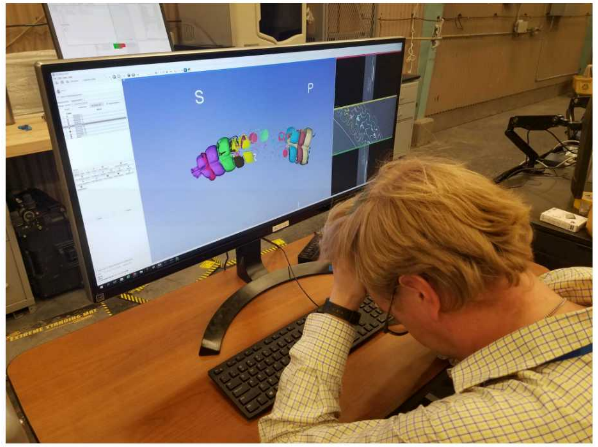

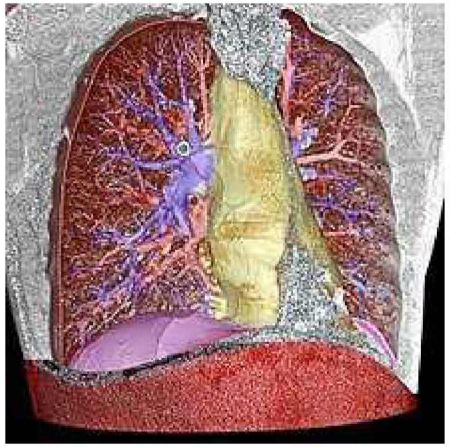

https://en.wikipedia.org/wiki/Image_segmentation

Labeling by hand does not scale

Deep Learning algorithms
- Find each component of the shaped charge by material
- Find any defects
- Pass this to a usable form for numerical simulations

Automated Segmentation presents challenges

- CT scans are large
  - Medical literature: 128x128x64
  - Ours: 800x800x1600 (~1000x more)
  - Soon: 2000x2000x2000
- Class Imbalance
  - Empty space dominates the scan
  - Defects ~0.6% of volume
- Artifacts and noise
  - Difficult to separate materials of similar density
  - Shadow effects

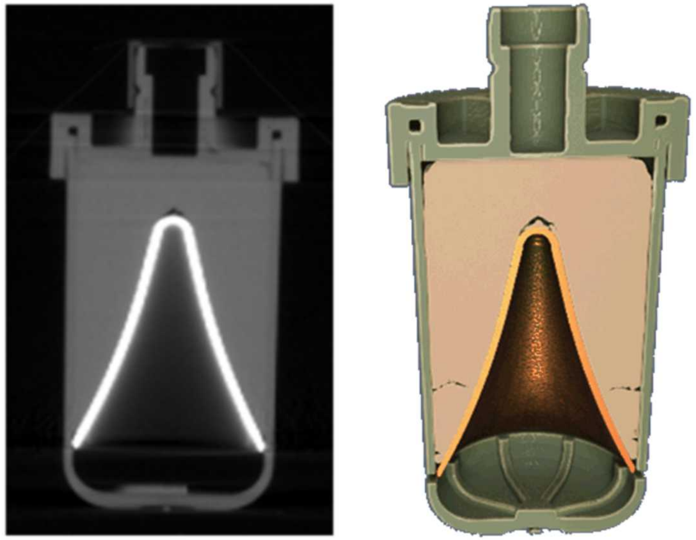

SM33 Shaped Charge

Mitigating Challenges

- CT scans are large
  - Used 240x240x240 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

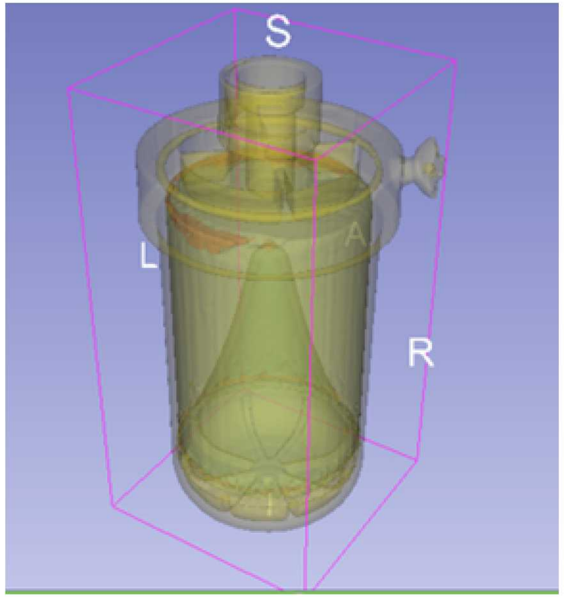

Deep learning is big data and large networks

- *"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

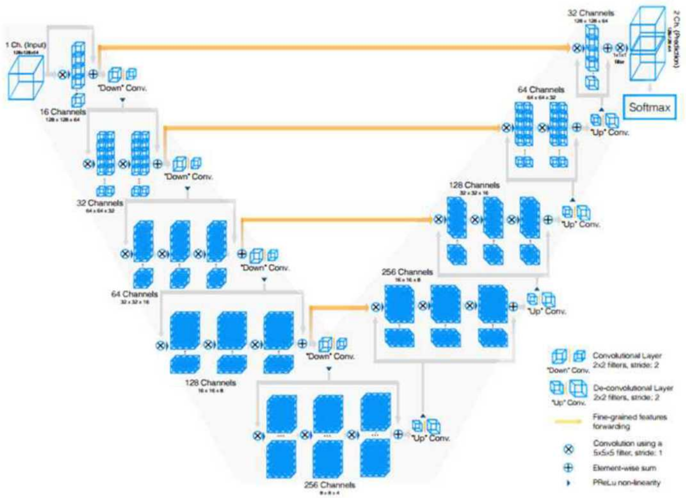

V-Net architecture for segmenting volumetric data (2016)

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

Results on the SM33 shaped charge

- Training time on DGX-1: ~2 days
- Inference time on DGX-1: ~2 minutes
- Error: ~1.02% on average for held out test set

- Defects remain challenging -

  Believed to be related to CT artifacts

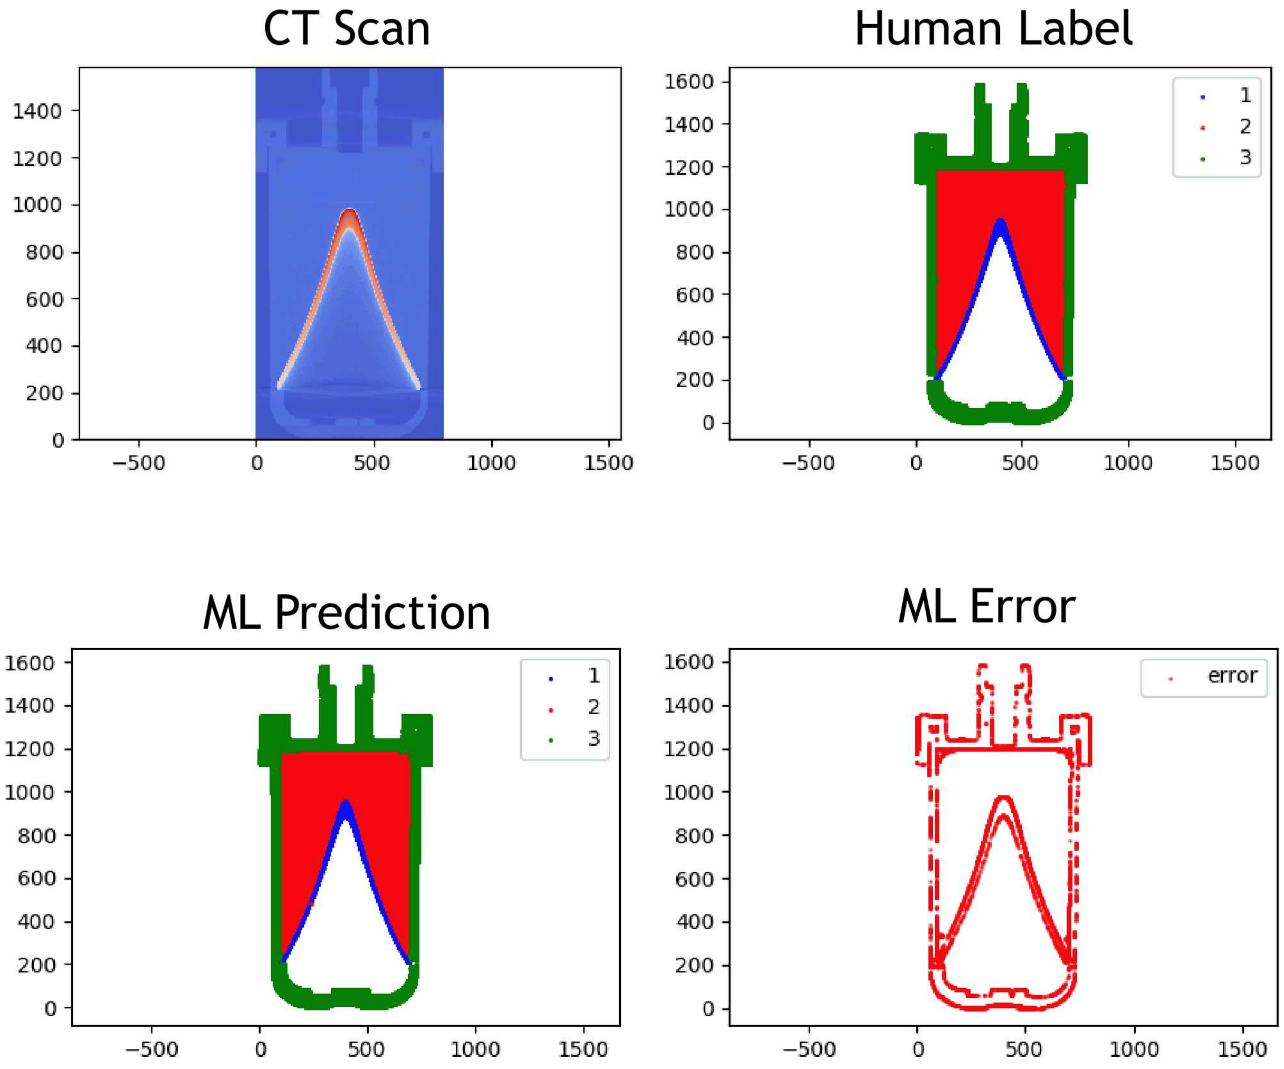

Defect detection

- Our plan
  - Mask explosive using our prediction
  - Solve binary classification problem: “defect or no defect”
- Problems
  - Massive class imbalance
  - Defects are extremely difficult to see => human labels are unreliable
- New Plan: Generate better images

Using CycleGAN for automatic segmentation

SEMI-SUPERVISED LEARNING

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.

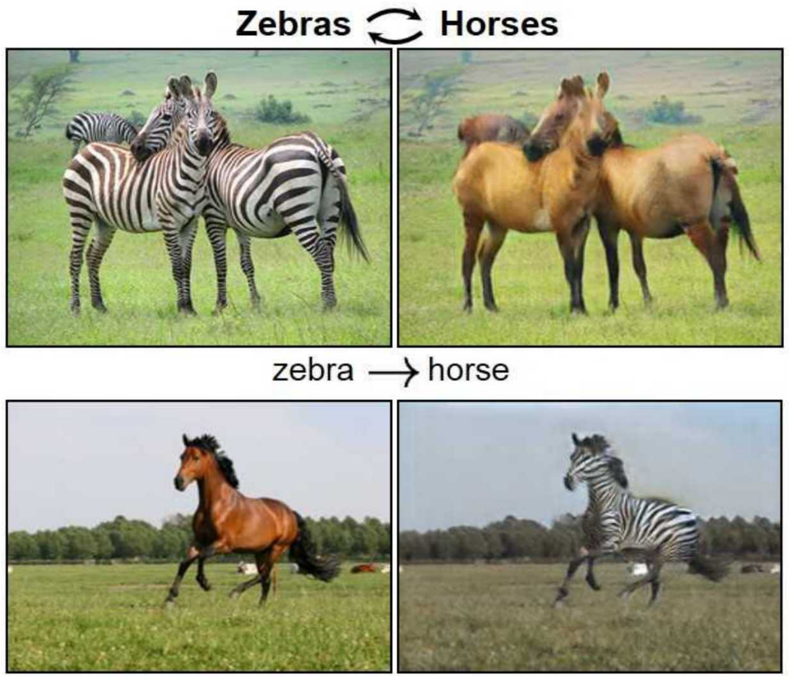

Learns two functions:

$F(x)$ = Horse to zebra

$G(x)$ = Zebra to horse

Cycles back to starting point to learn without paired examples

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

Generative adversarial networks are powerful new models that are widely applicable

Batteries to Foam — Leveraging existing datasets via CycleGAN

Collaborators – Kevin Long & Dan Bolintineanu

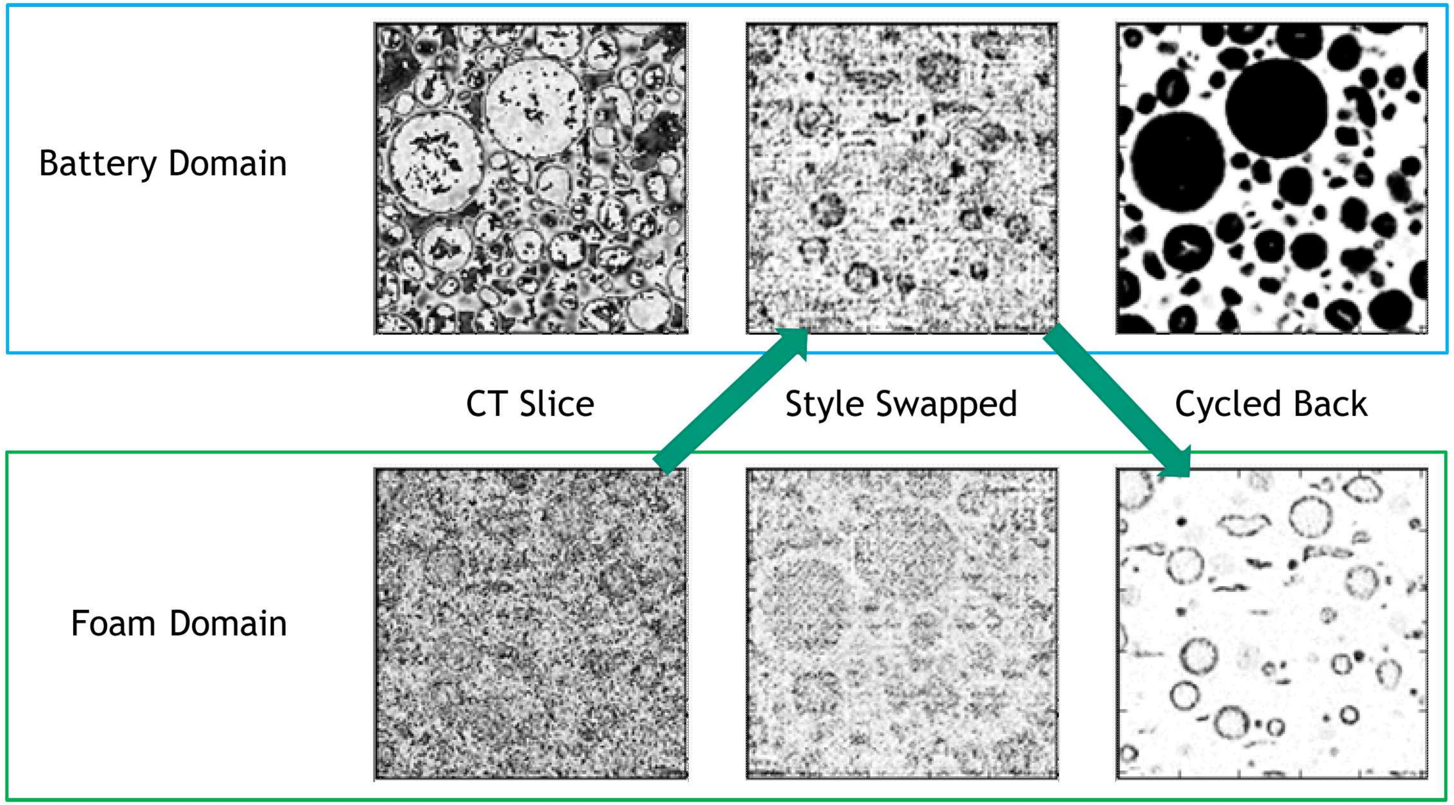

Domain Adaptation could reduce supervised labeling cost

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