The Incorporation of Machine Learning into Scientific Simulations at Lawrence Livermore National Laborator

High Energy Density Science Seminar

July 30, 2020

Katie Lewis, Lawrence Livermore Nation
Advanced Machine Learning Project Lea

LLNL-PRES-808845
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Supercomputing and Computational Physics at Lawrence Livermore National Laboratory

- Lawrence Livermore National Laboratory (LLNL) was founded in 1952
  - Scientific Computing was part of our initial portfolio
- "It is now accepted that in addition to the experimental and theoretical branc there is a third: computer simulation." - David Young, postdoc who worked wi in the late 1960s
- Today, LLNL and the DOE Complex continue to dominate supercomputing for s simulations in support of national security.
- Data Science is increasingly a part of this landscape.

Sierra – 125 pe

www.llnl.gov/about/history
www.llnl.gov/news/berni-alder-pioneer-times
www.top500.org

Data Science potential spans Scientific Computing sp

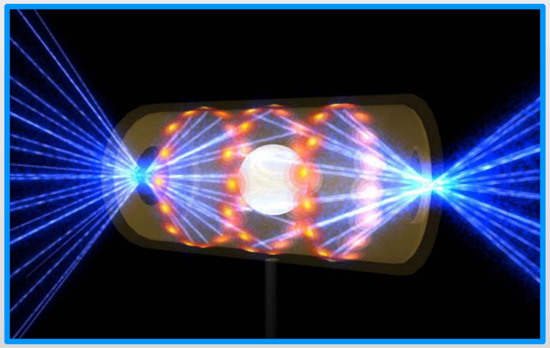

Enhanced Design Workflow

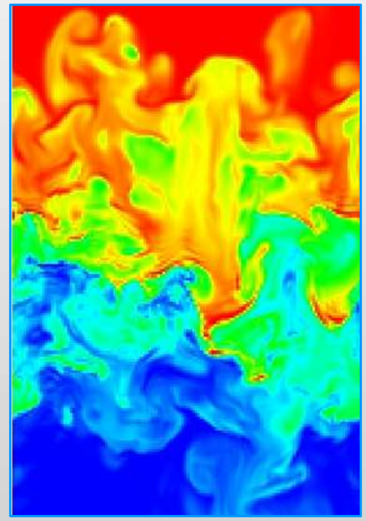

Enhanced Modeling

Improved Systen Performance

Physics Constrained Predictions

Interpretable Predictions

Community Engagement

Many research topics are already being investigated

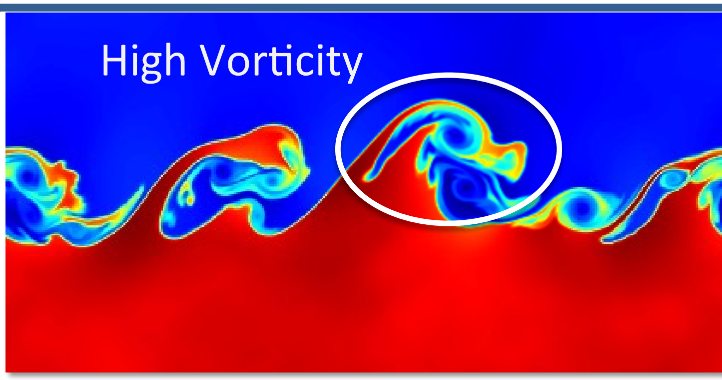

Improved Design Workflows

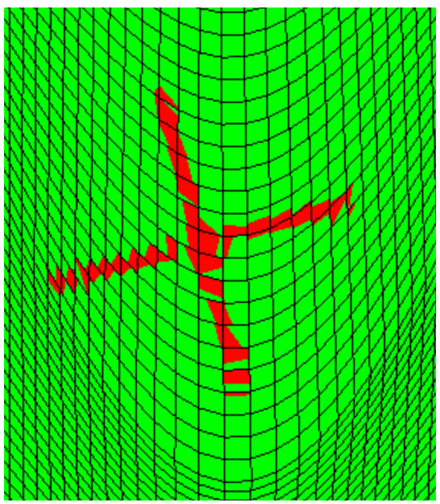

Improved Material Interface Reconstruction

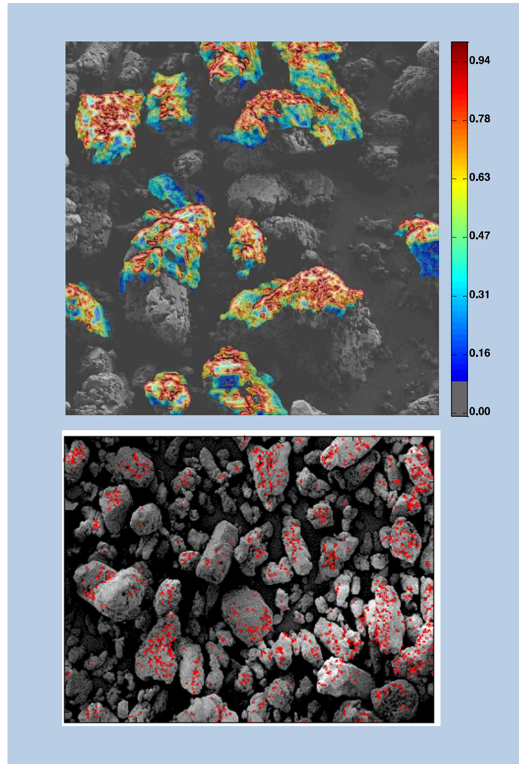

Material Discovery

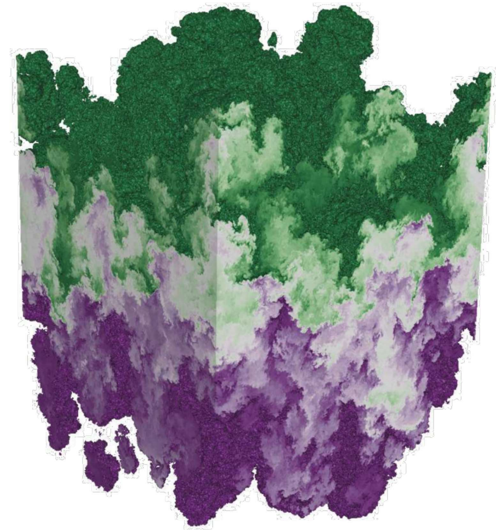

Augmented Turbulence Modeling

Su

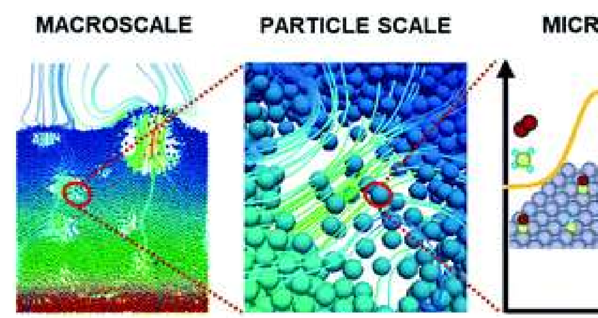

Multi-scale Coupling

Journal for Reaction Chemistry & Engine

Terminology

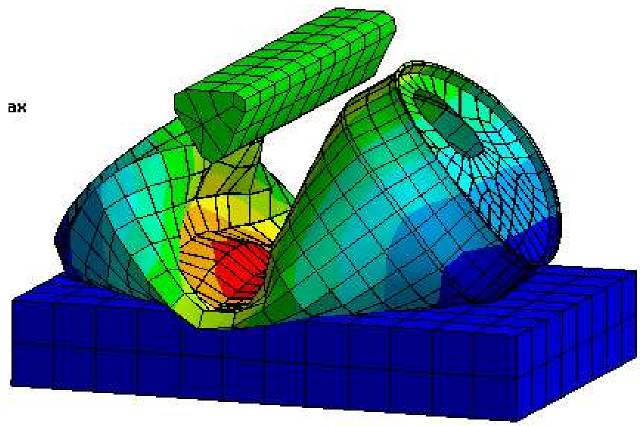

3D Lagrangian simulation
faculty.washington.edu

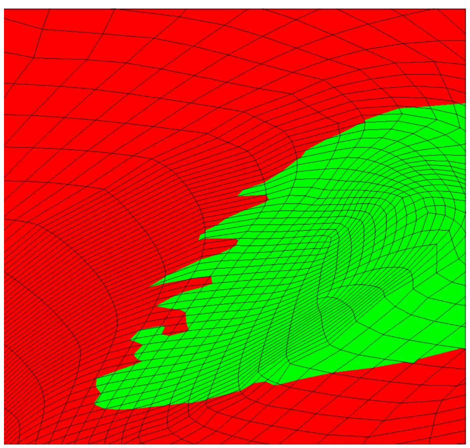

2D ALE Simulation

- Geometry is discretized using a "mesh" or "
  discretized into timesteps.
- Fields (like density, velocity, or temperature
  calculated at the mesh points or "zones".
- In Eulerian simulations, the mesh is static a
  materials move through it.
- In Lagrangian simulations, the mesh moves
  material.
- In Arbitrary Lagrangian-Eulerian (ALE) simul
  mesh moves, but not necessarily with the m

Application: ML to control Arbitrary Lagrangian-Eulerian

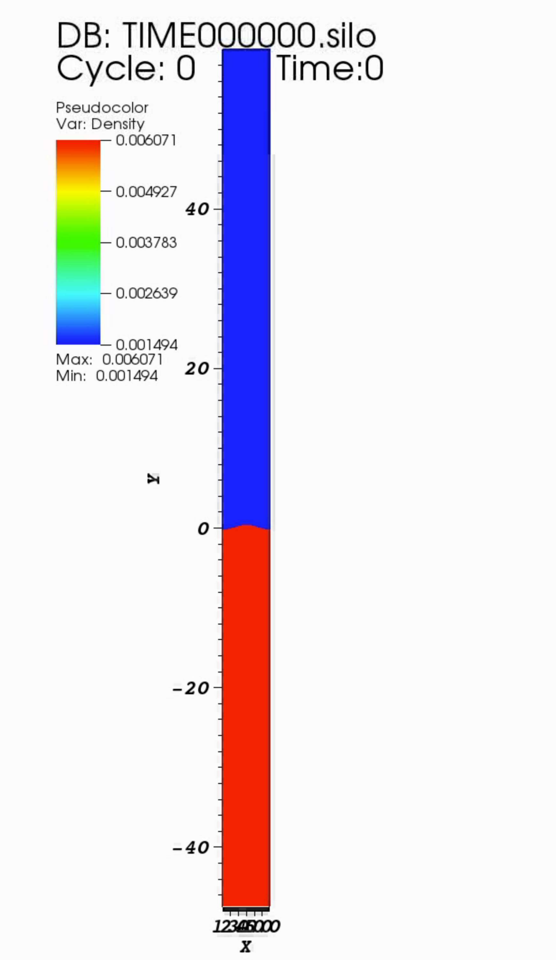

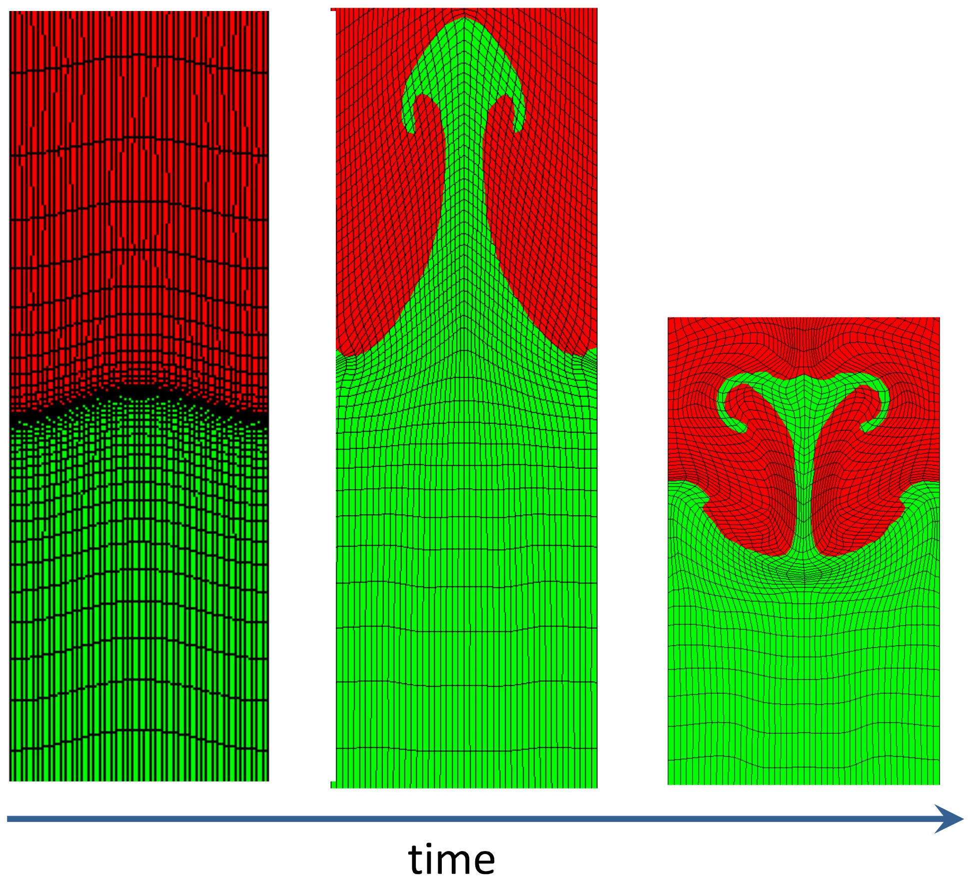

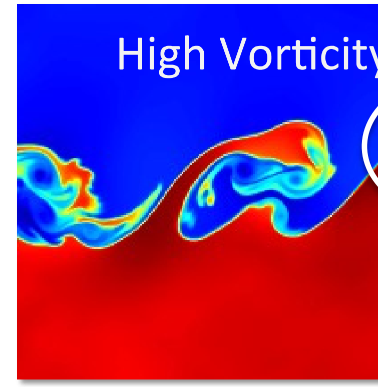

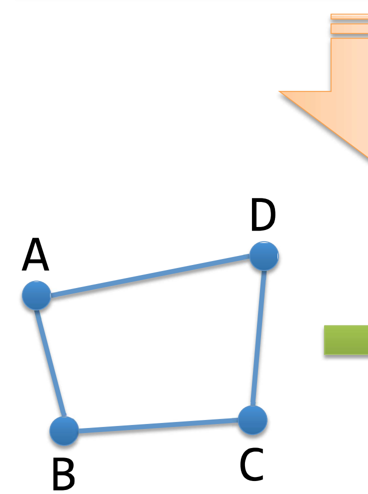

This problem is typically solved with hand-tuned relaxation strategi

Ming Jiang, Brian Gallagher, Keith Henderson, Alister Magui

Trained relaxation strategies can significantly reduce user bu

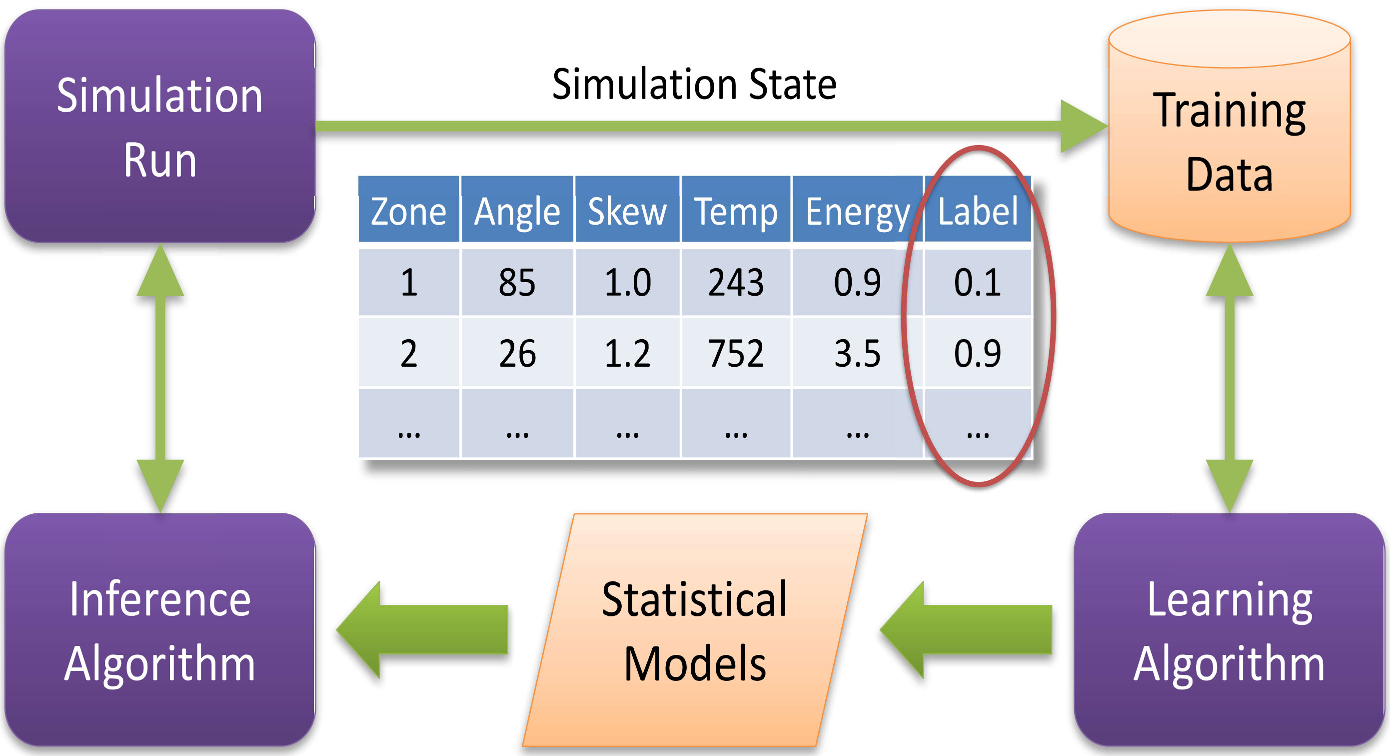

Zone	Angle	Skew	Temp	Energy	Label
1	85	1.0	243	0.9	0.1
2	26	1.2	752	3.5	0.9
…	…	…	…	…	…

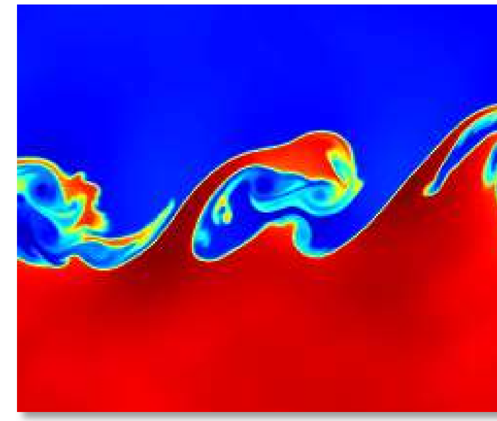

$$\underbrace{f(x_1, x_2, \ldots, x_n}_{}$$

Simulation state:
mesh + physics

Ming Jiang, Brian Gallagher, Keith Henderson, Alister Magui

Research project is moving into user community

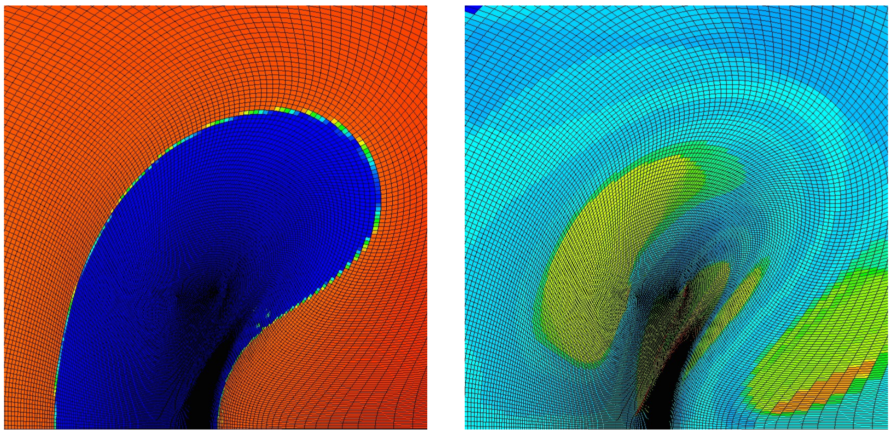

Bubble Shock

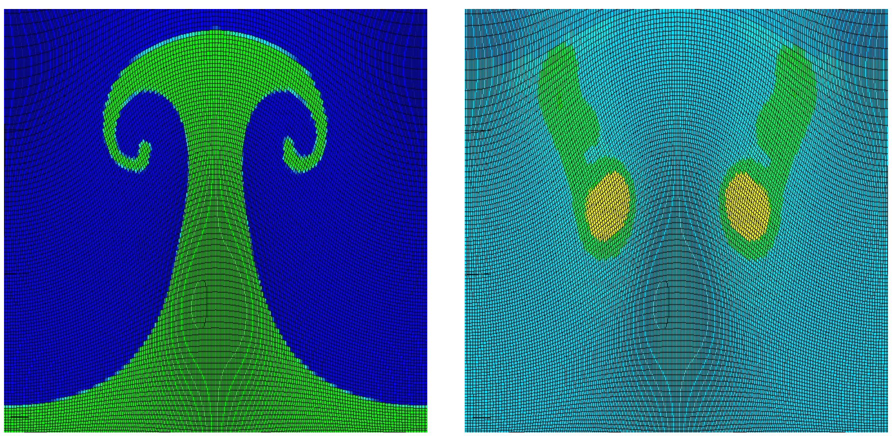

Shock Tube

Building on top of M. Jiang, B. Gallagher, J. Kallman, Supervised Learning Framework for Arbitrary Lagran Simulations," *IEEE International Conference on Mach Applications (ICMLA)*, pp. 977–982, 2016.

- Initial results showed high accuracy using random
- Recent work improves the imbalance in training d generalization to noisy data using Convolutional N (CNNs)
- Working with user community to provide quantita results using the CNN for inference inline
  - Evaluate quantities of interest against experimental res
  - Develop a reward function for Reinforcement Learning
- Test case for proxy application on new hardware (

Inference in Action

We have an alpha capability to incorporate models into simulations

- Pre-trained models can be used directly
- Tools allow end users to retrain models using their simulations

Versioning and reproducibility are being addressed

Evaluating Reinforcement Learning and Graph Neural Networks

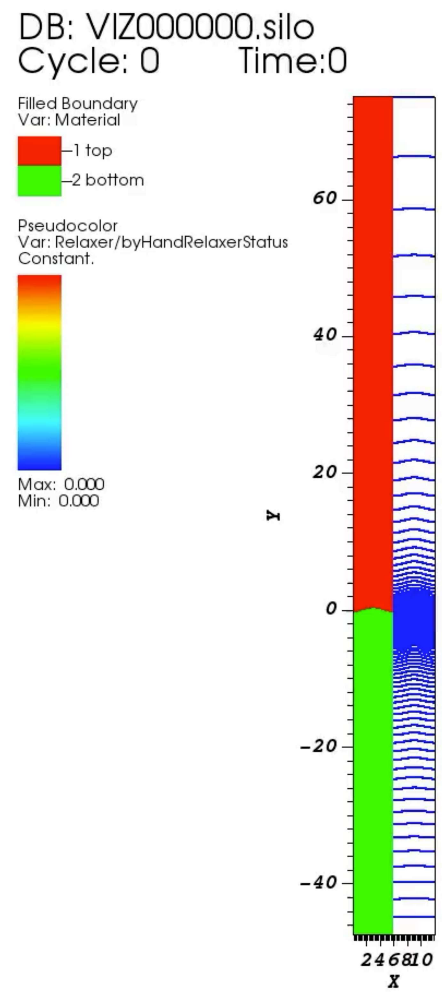

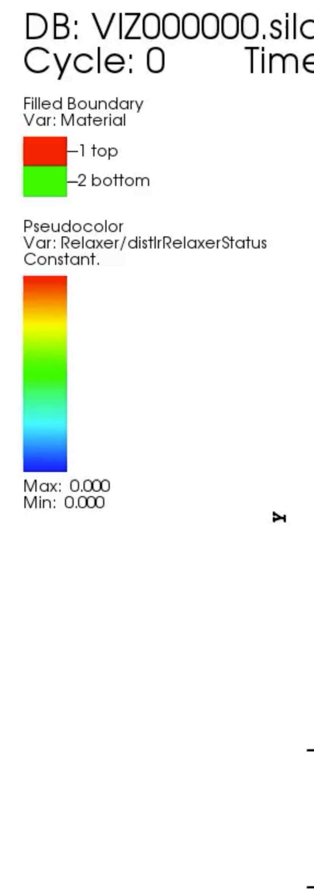

Kevin Hyunh and Alister Magu

Application: ML for Material Interface Reconstruction (M

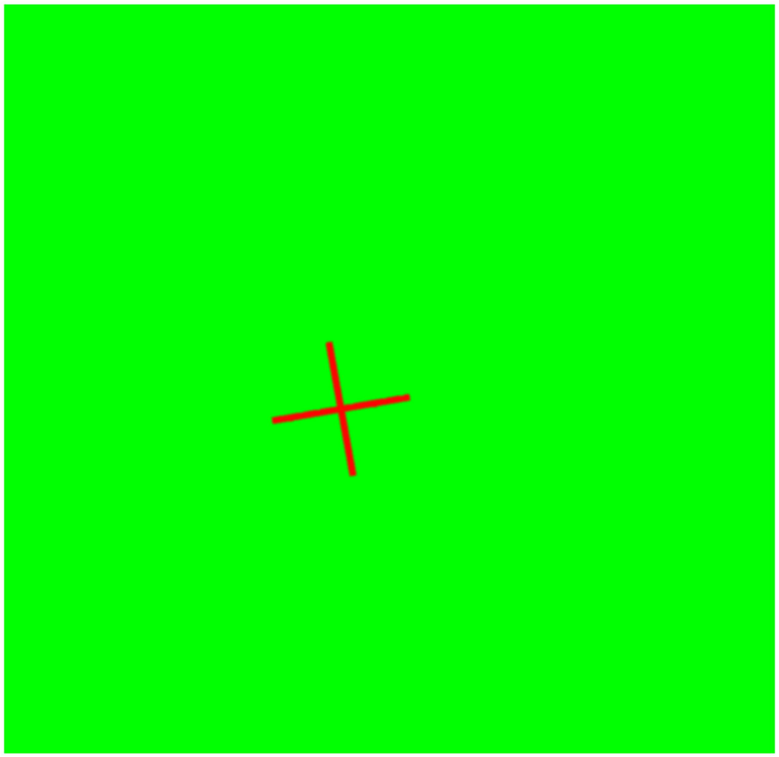

Actual Material
Boundaries

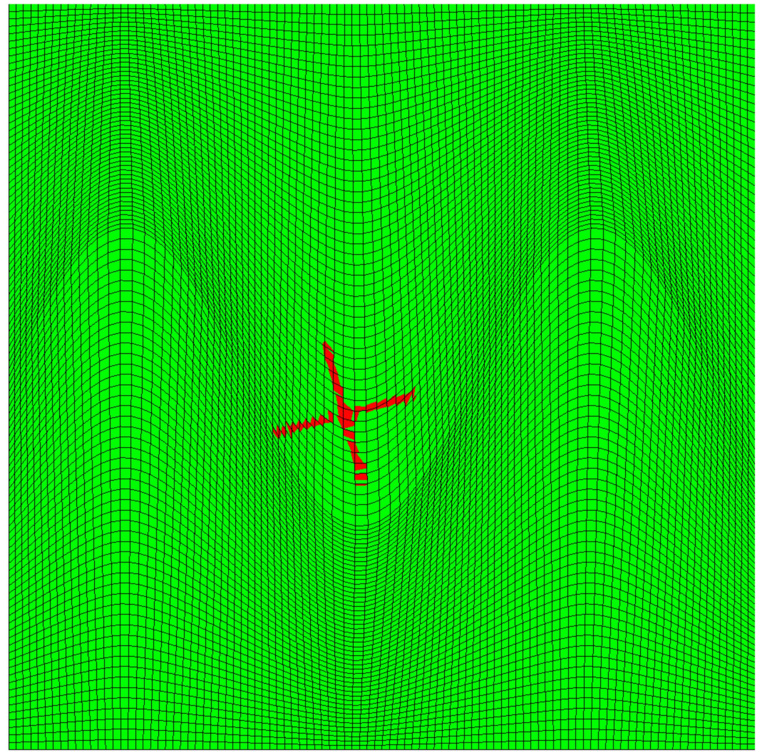

Current, High Res
Reconstruction

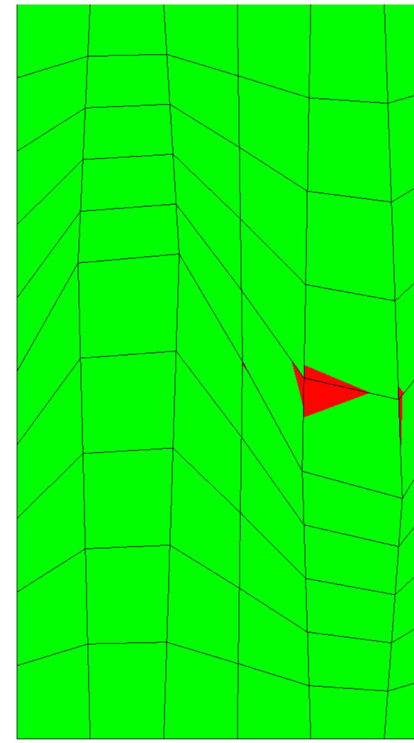

Current, Lo
Reconstru

Dan Fenn, Walt Nissen, Kenny Weiss

Training with actual geometry may avoid the common e[...] seen in heuristic solutions

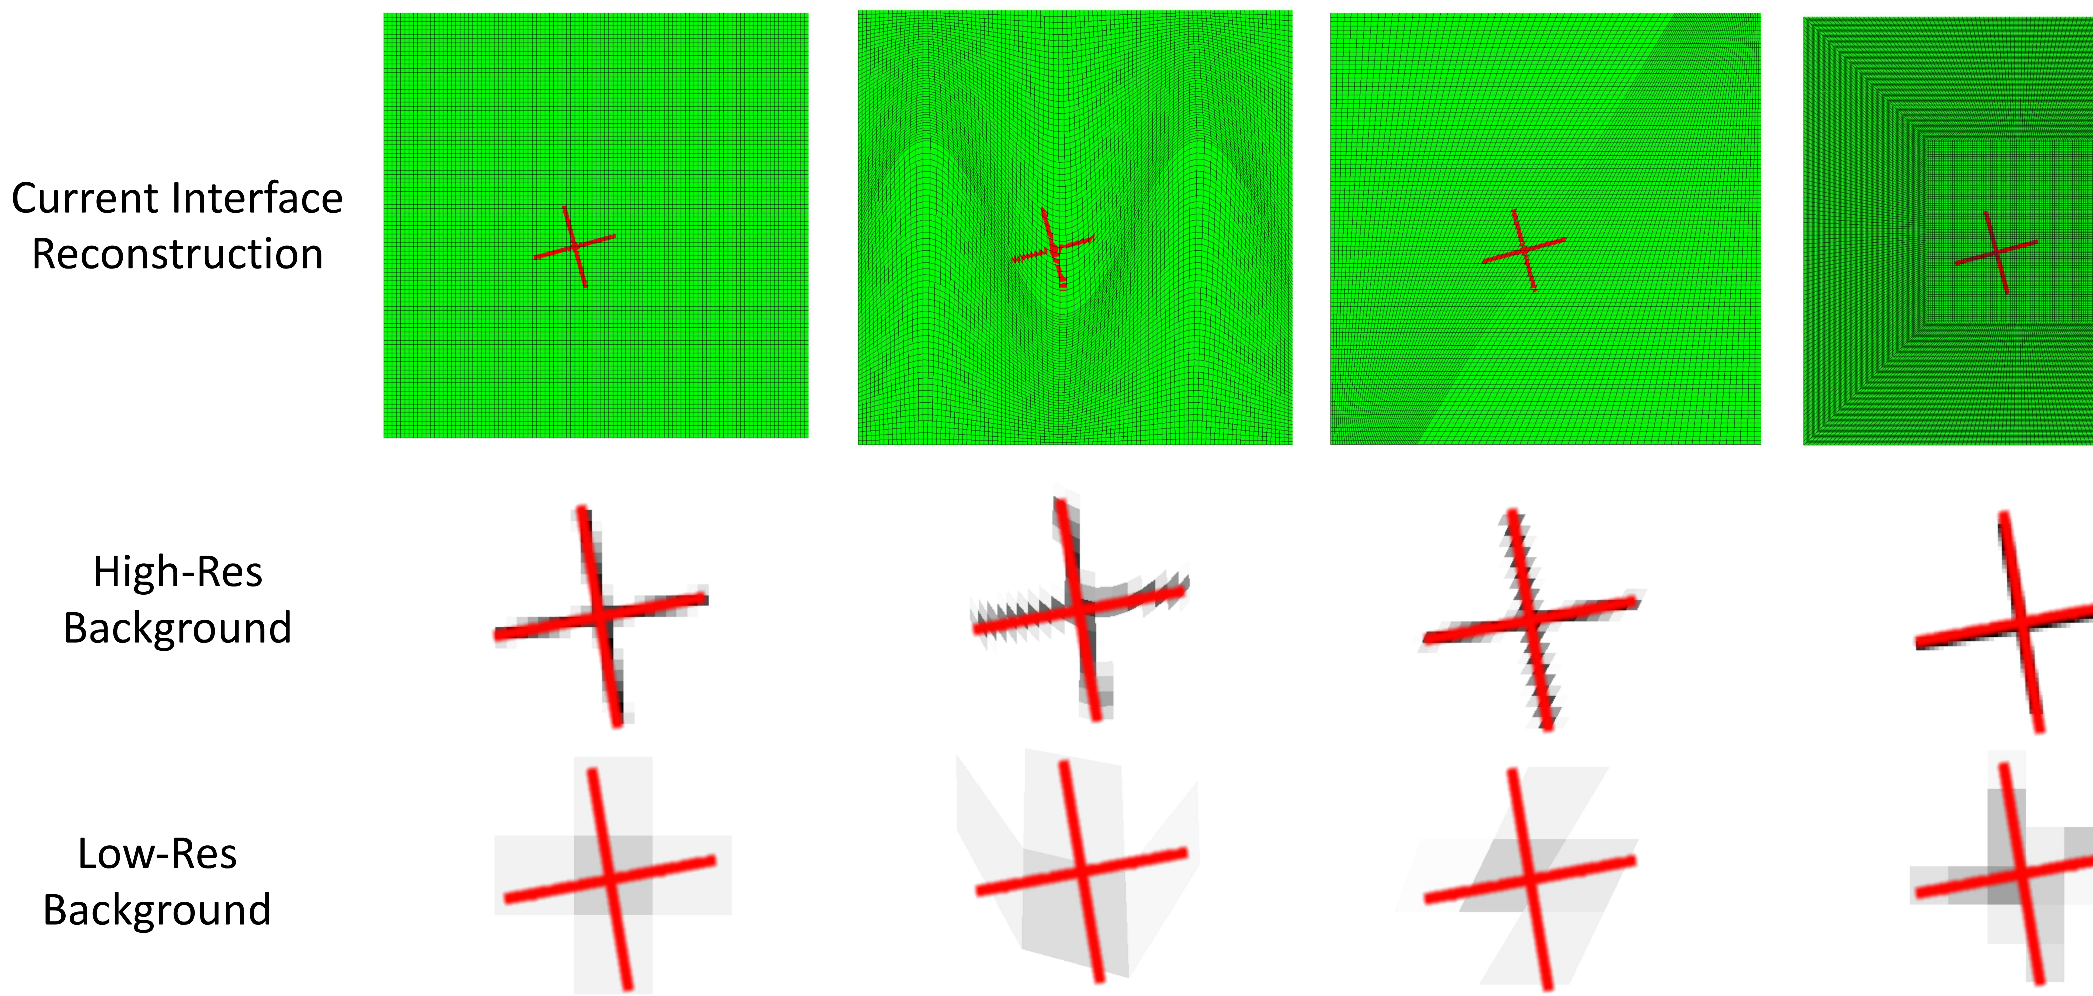

We can use this methodology to train on many shapes

Varying:
- Position
- Size
- Rotation
- Background vs. Foreground

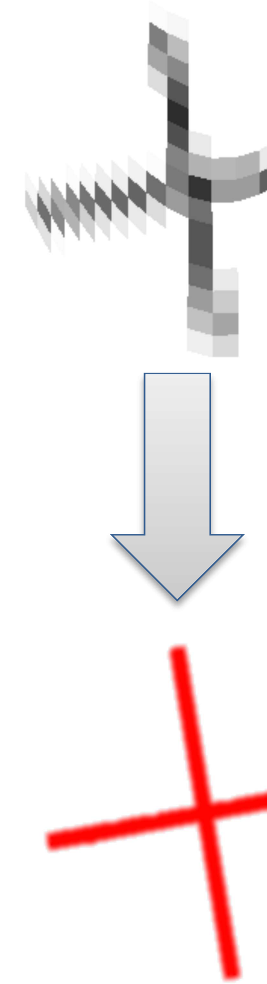

Dan Fenn, Walt Nissen, Kenny Weiss

Initial results are very promising!

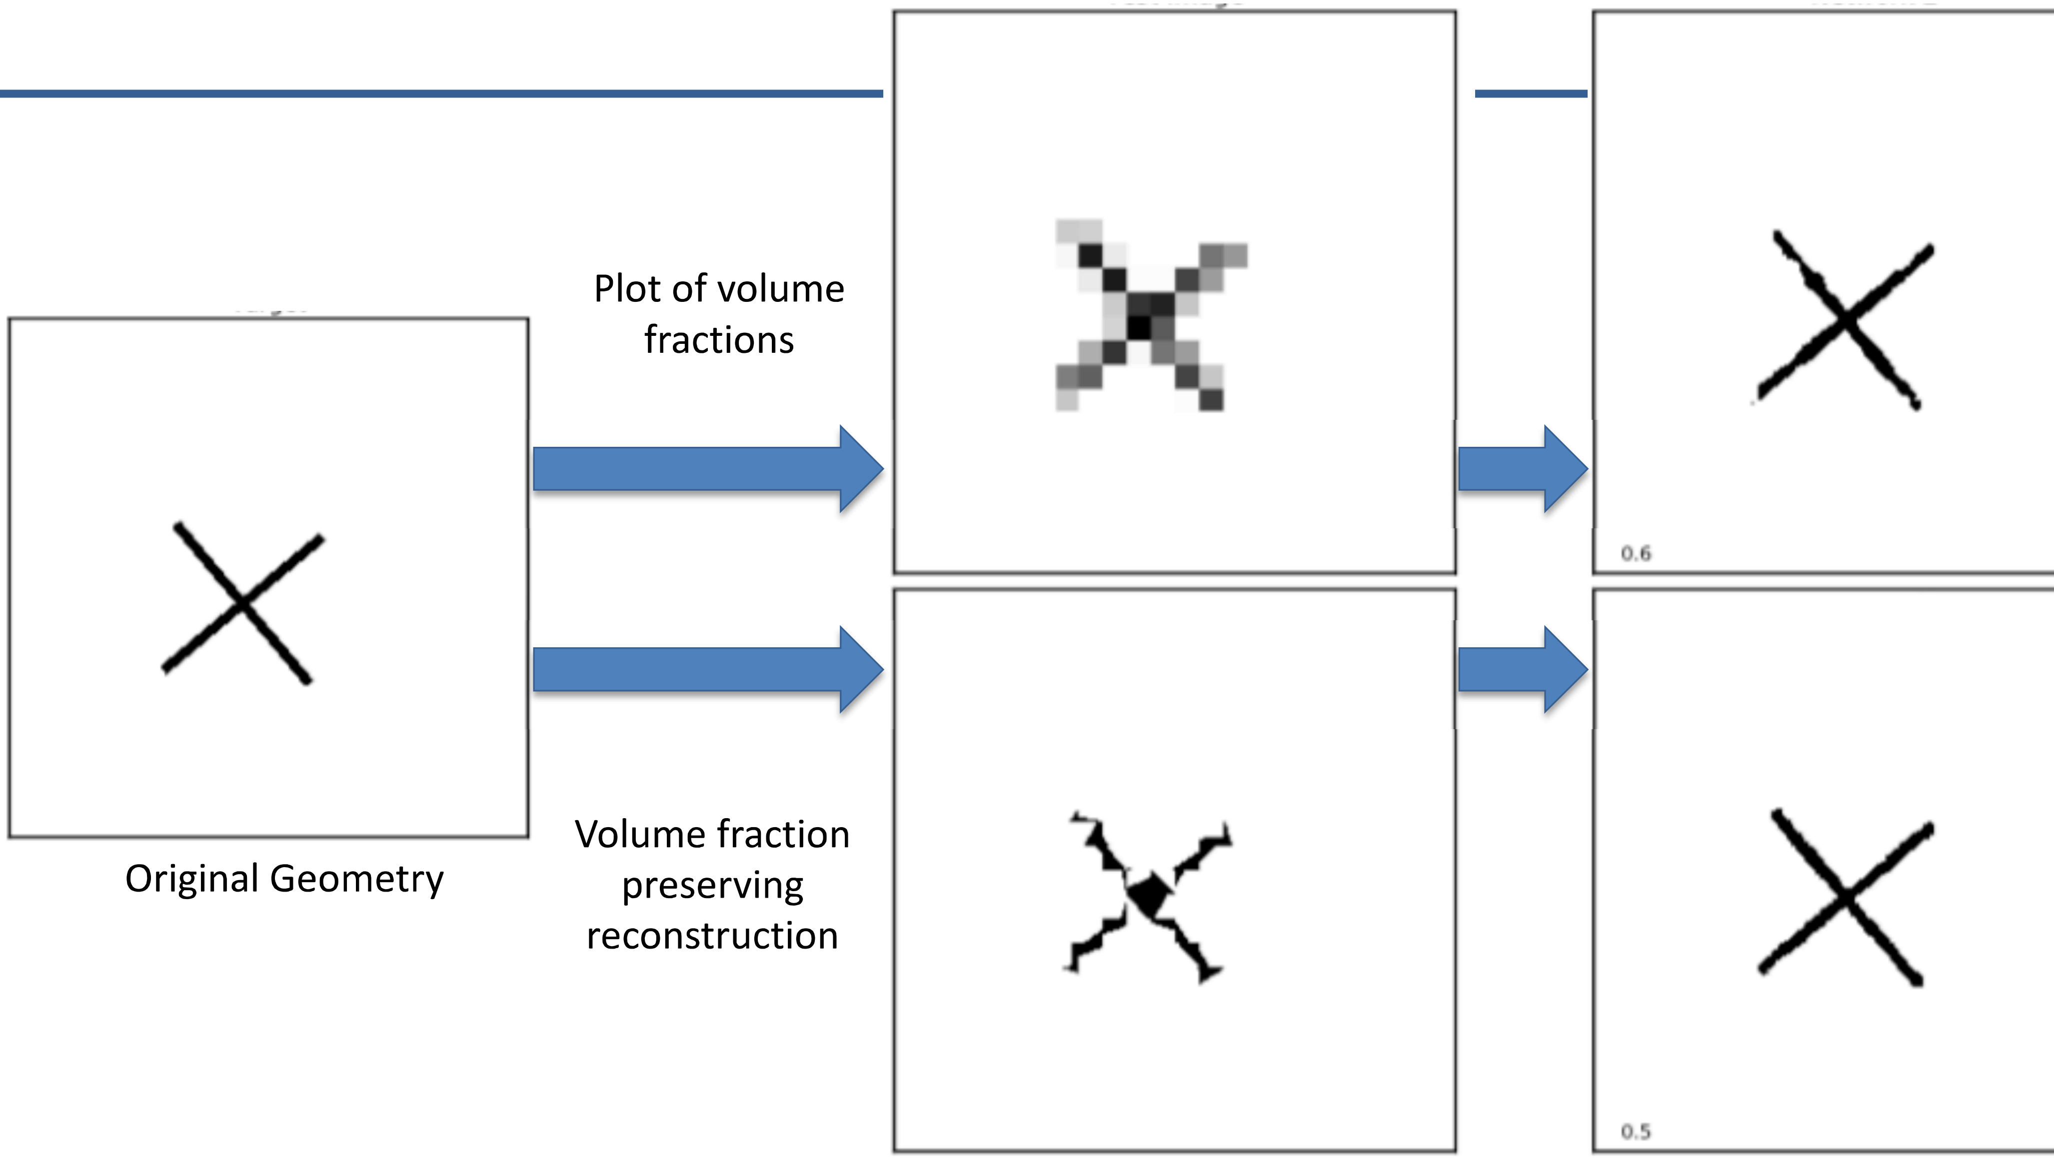

Dan Fenn, Walt Nissen, Kenny Weiss

Material Interface Reconstruction – Next Steps

- Incorporate volume fraction information into training as a loss/reward fur
  - Modifying threshold to meet volume fractions was unsuccessful (i.e., too noisy)
  - Applying a weight to focus on interfaces is showing improvements
- Incorporate active learning techniques to handle new types of geometries
- Evaluate how the algorithm will work in-situ, accounting for conserved ph quantities and parallelization schemes
- Investigate reconstruction for multiple materials

Application: ML for Fast Surrogate Modeling

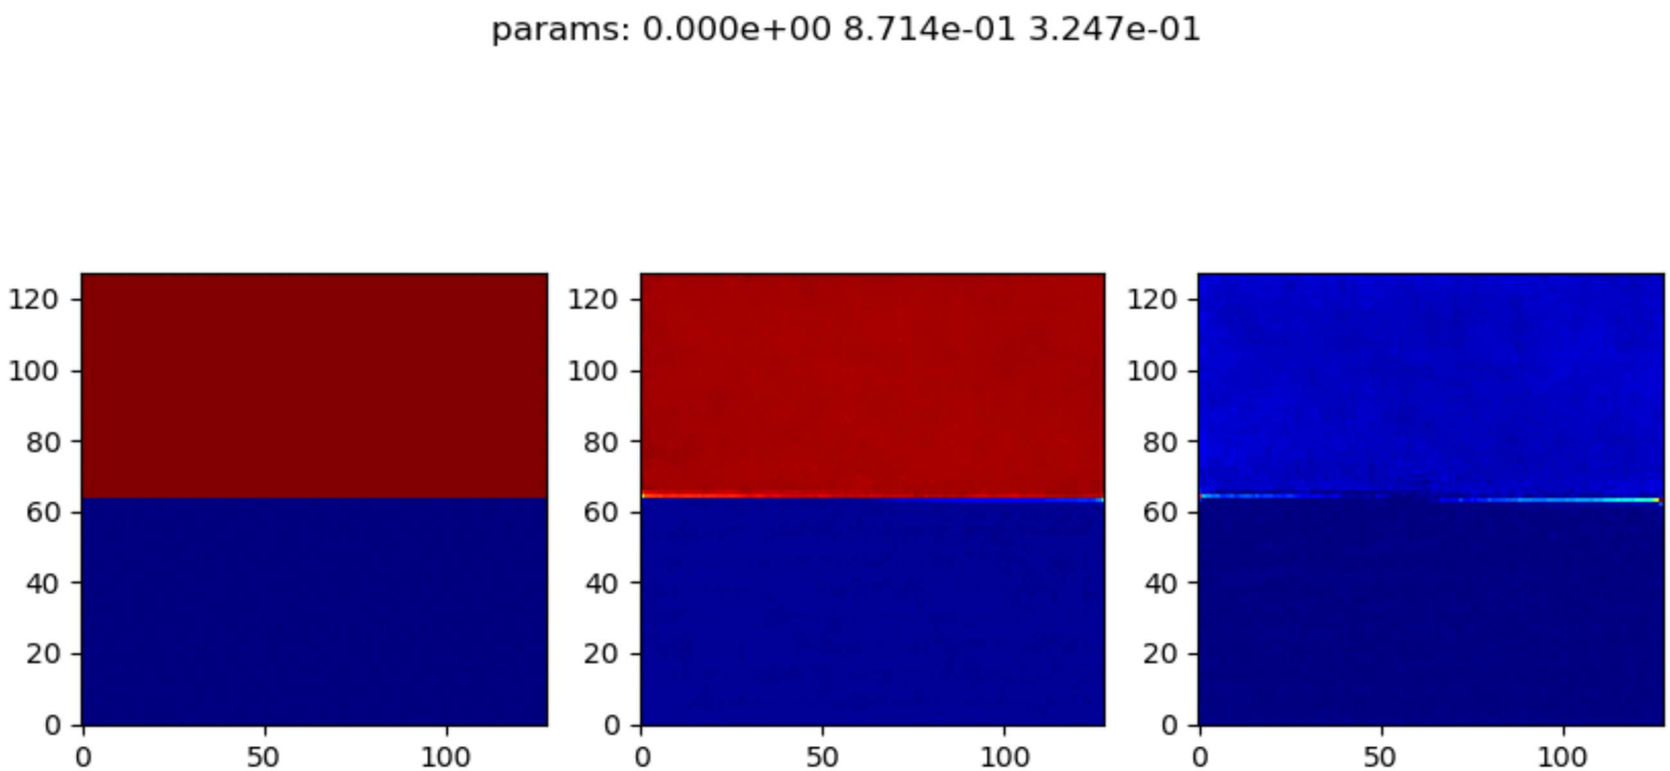

2D Hydro Simulation

Neural Network Surrogate

Difference

100s of hyd
used to trai

Inference is
the full simu

Maximum n
across cycle
maximum e
higher.

This method
optimize pa
simulation.

Application: ML to improve multi-scale coupling

Ab-initio	Atoms	Long-time	Microstructure	Dislocation	Crystal	Continuur
Inter-atomic forces, EOS, excited states	Defects and interfaces, nucleation	Defects and defect structures	Meso-scale multi-phase, multi-grain evolution	Meso-scale strength	Meso-scale material response	Macro-scale material respo
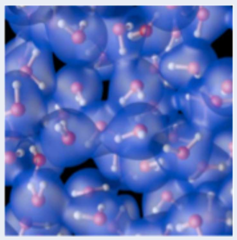	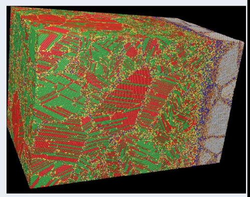	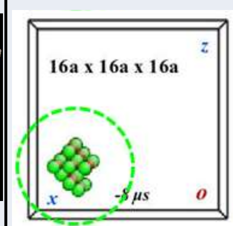	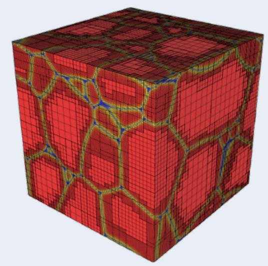	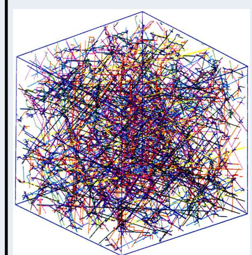	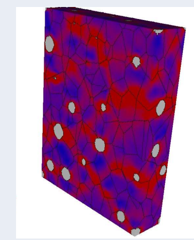	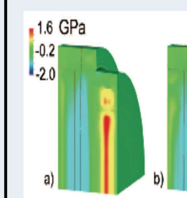

Can ML replace interpolation schemes used within continu
when querying opacity or equation of state models, ***reduci***
footprints while maintaining accuracy?

Rob Blake, Ben Yee, and Mike Hohensee

Evaluation of ML for Opacity Interpolation

Current method

Before running continuum code:

- Perform atomic physics calculations to obtain detailed data
- Store data in a 3-D table

During continuum code:

- Table lookup & linear interpolation

Fast, but inaccurate and memory intensive

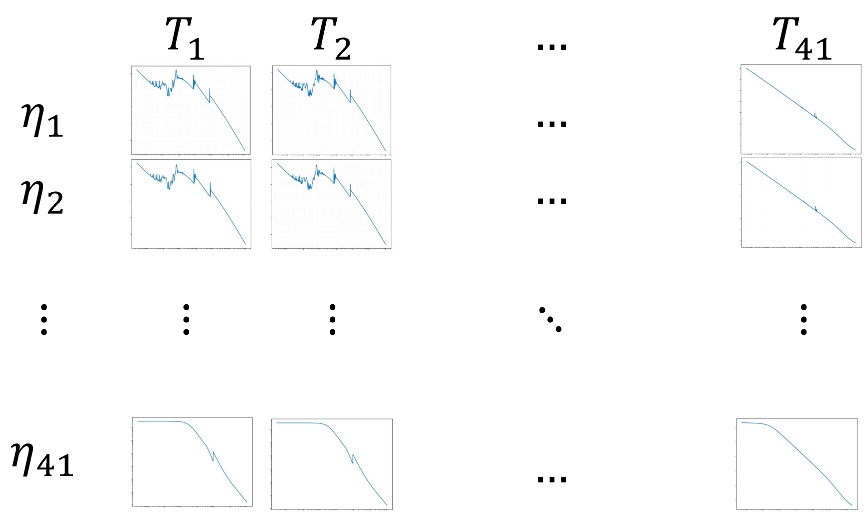

Proposed method with machin

Before running continuum code:

- Perform atomic physics calculations t detailed data
- Regression problem: Use data to trai

During continuum code:

- Apply inference on neural net

FLOPs, accuracy, memo

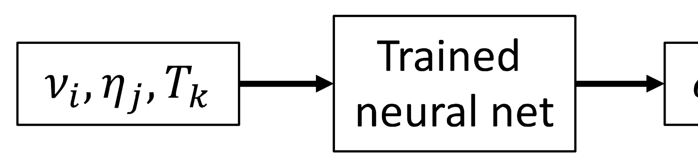

Rob Blake, Ben Yee, and Mike Hohensee

Networks trained on a subset of the domain

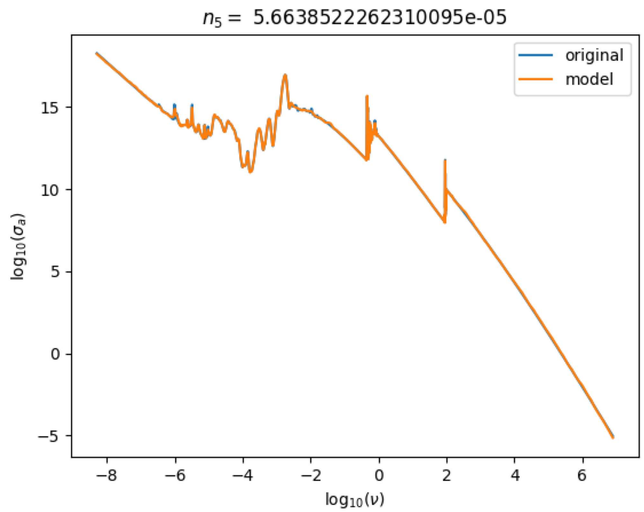

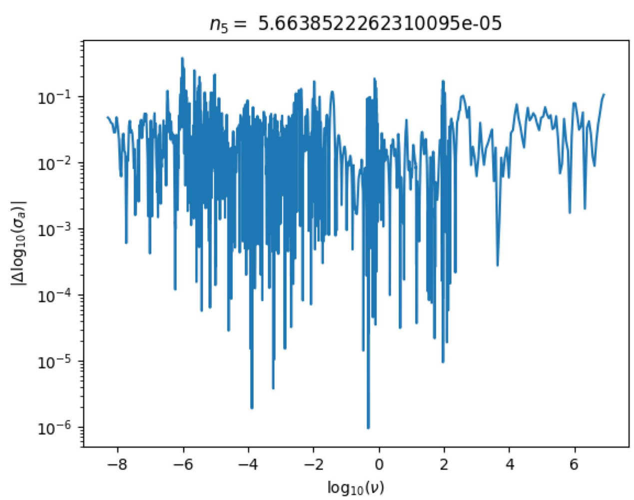

- Initial Evaluation:
  - Network trained on 2D slice of iron, varying density
  - Specific density slices omitted for network validatio
- Results:
  - Current network has accuracy comparable to existin
  - Current network has improved accuracy between da
  - Network consumes less memory, ~100x savings.
- Unfortunately:
  - The highest error is where is matters most (spiky da
  - Table data is highly curated
- Investigating ML to solve other problems relat

Rob Blake, Ben Yee, and Mike Hohensee

Application: ML to speed up radiographic analysis

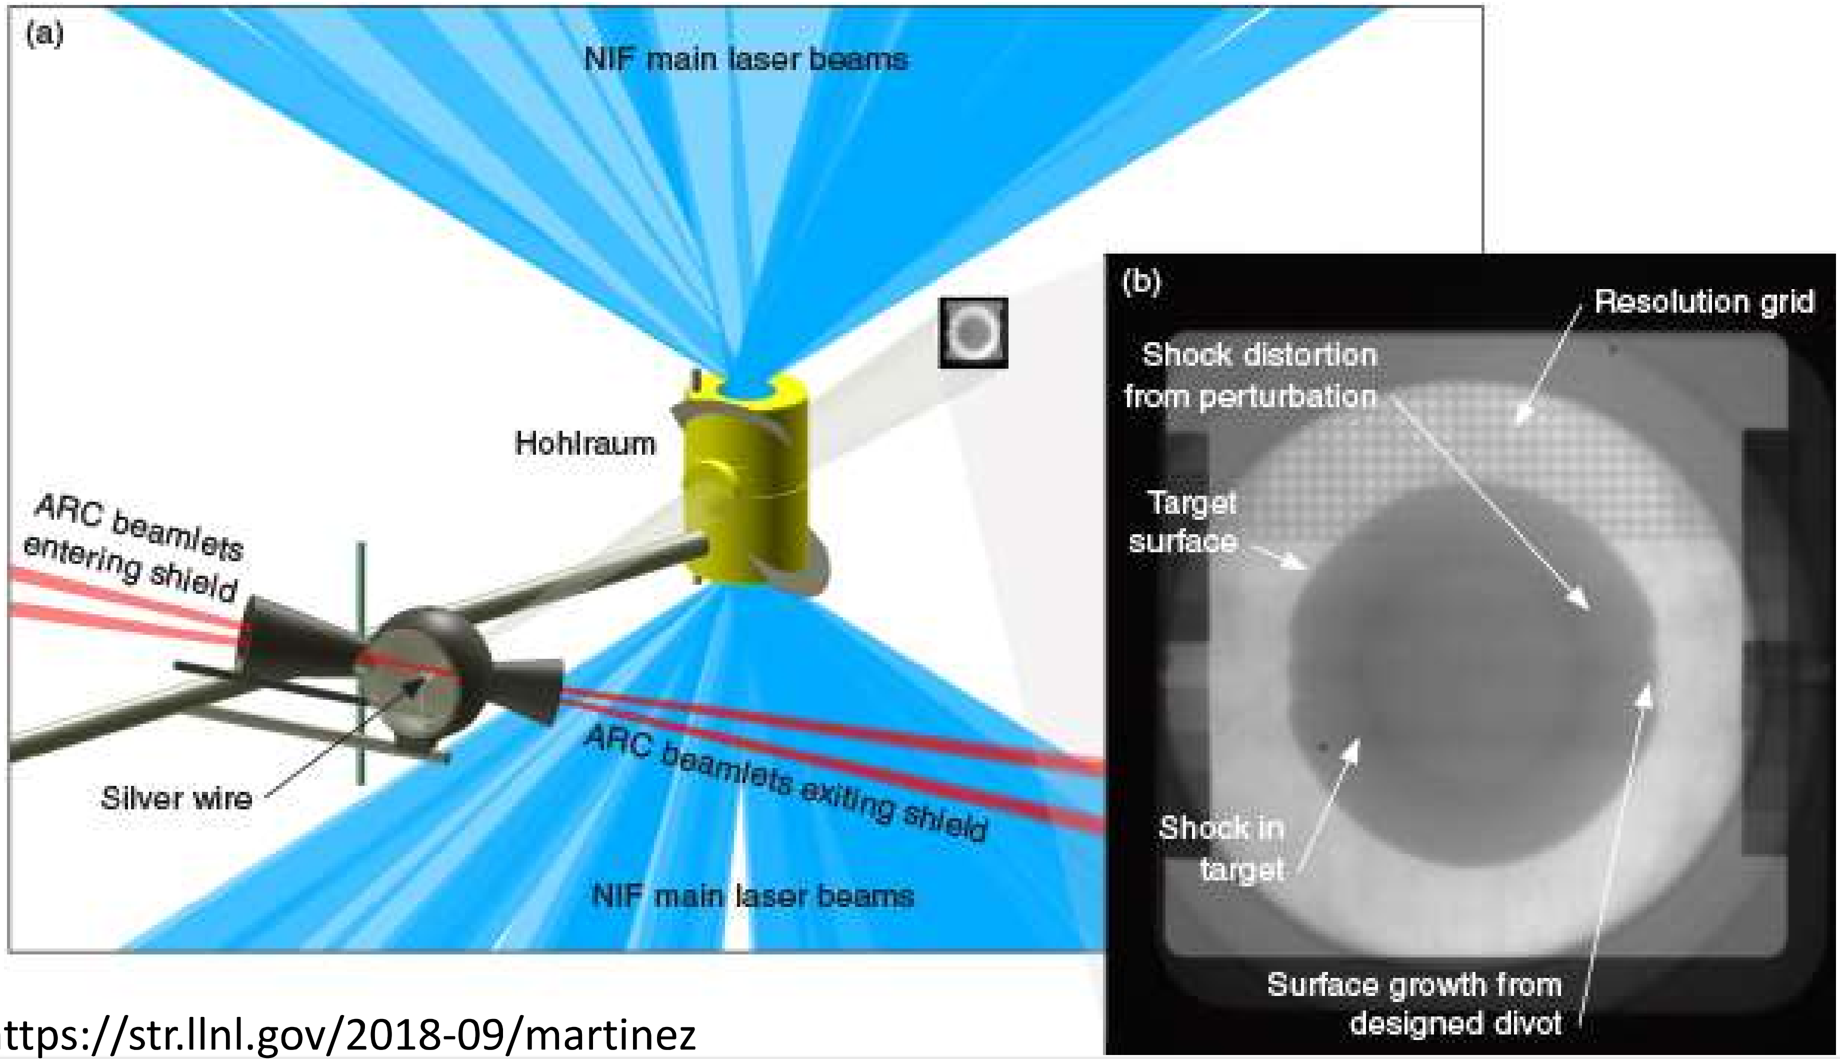

https://str.llnl.gov/2018-09/martinez

Medical imaging using Machine Learning may be transferable to our needs

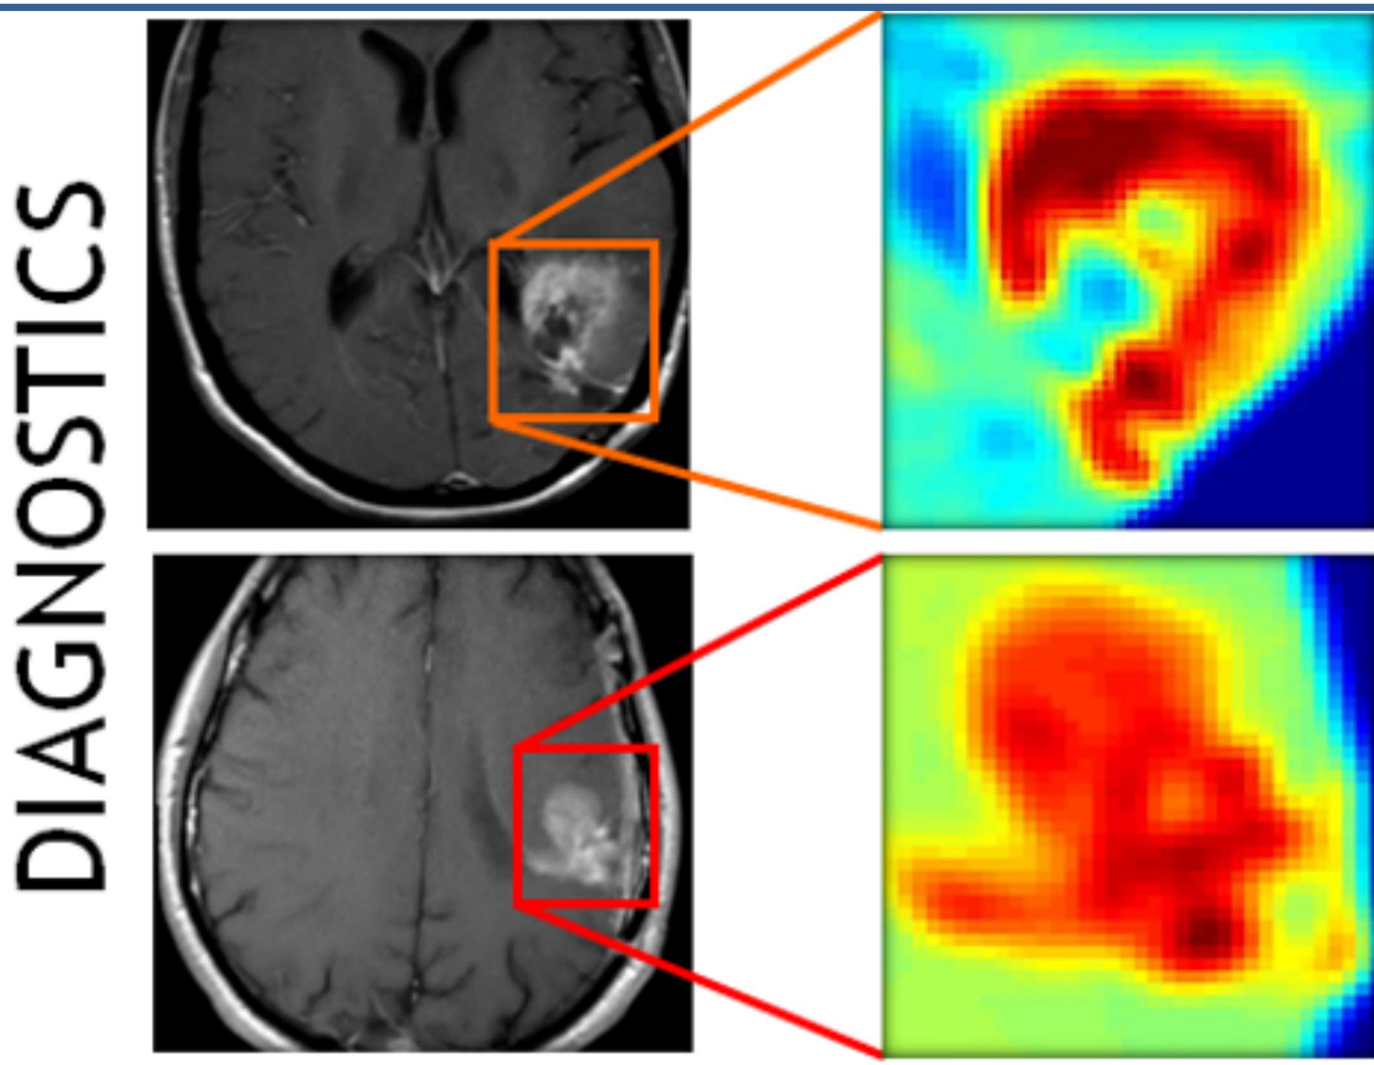

https://algorithmia.com/blo
machine-learning-for-health

https://github.com/mateusz
segmentation-pytorch

Creating a training set of labeled images using simulatio

- Approach: use simulation to generate synthetic radiographs and image masl
  - Start with clean radiographs and then introduce distortions normally found in experime using different channels to identify different features.
  - End goal is to detect features in experimental radiographs, while limiting manual labelir

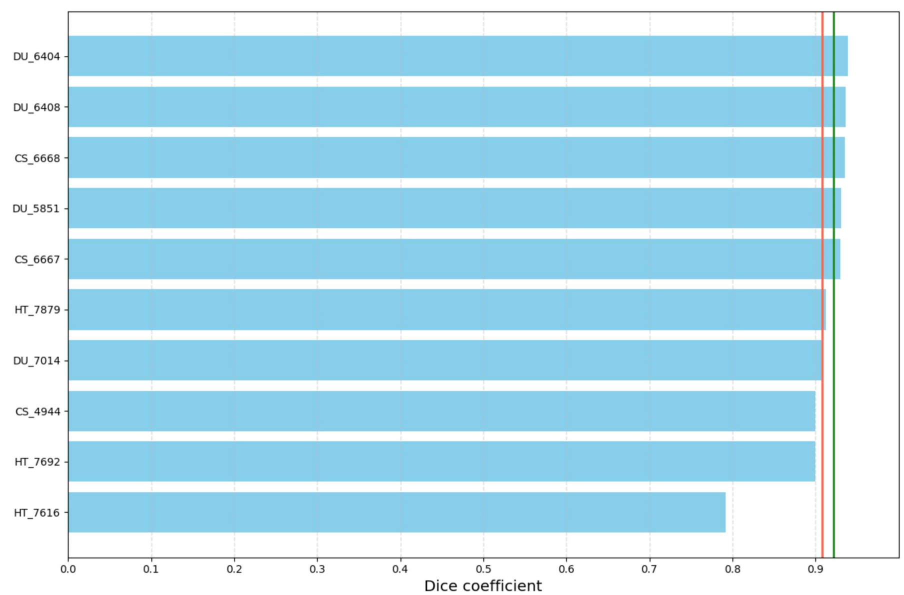

Brain Tumor Dataset

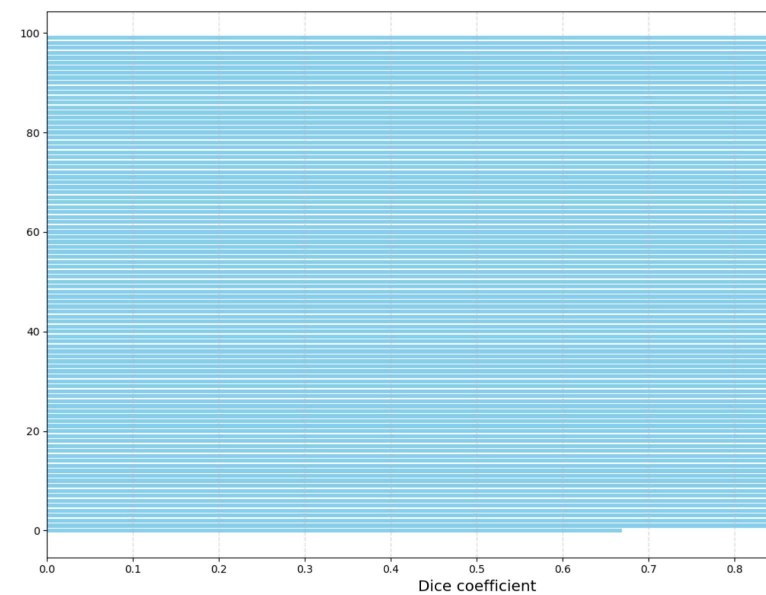

Radiograph Dataset

Cross-validation Dice score (2*overlap/total pixels) for 100 clean radiographic test images looks

Morry Aufderheide, Kevin Chen, and Hardeep Sul

Data Infrastructure Nee

Data Infrastructure is a fundamental need for ML

- ML (esp. DL) needs **a lot** of data, with verified **labels and provenance**
- Traditional databases are not common in the HPC environment

Kosh (Sanskrit for Treasury) is being developed to solve these problems

- Multi-modal data sources seamlessly searchable and accessible by authent users
  - Plan to incorporated sampling algorithms (spatial and temporal)
  - Datasets can have multiple files associated with them and multiple file formats
- Data can be distributed across organization/lab/compute centers.
- Users can query data to get only what they want for training.

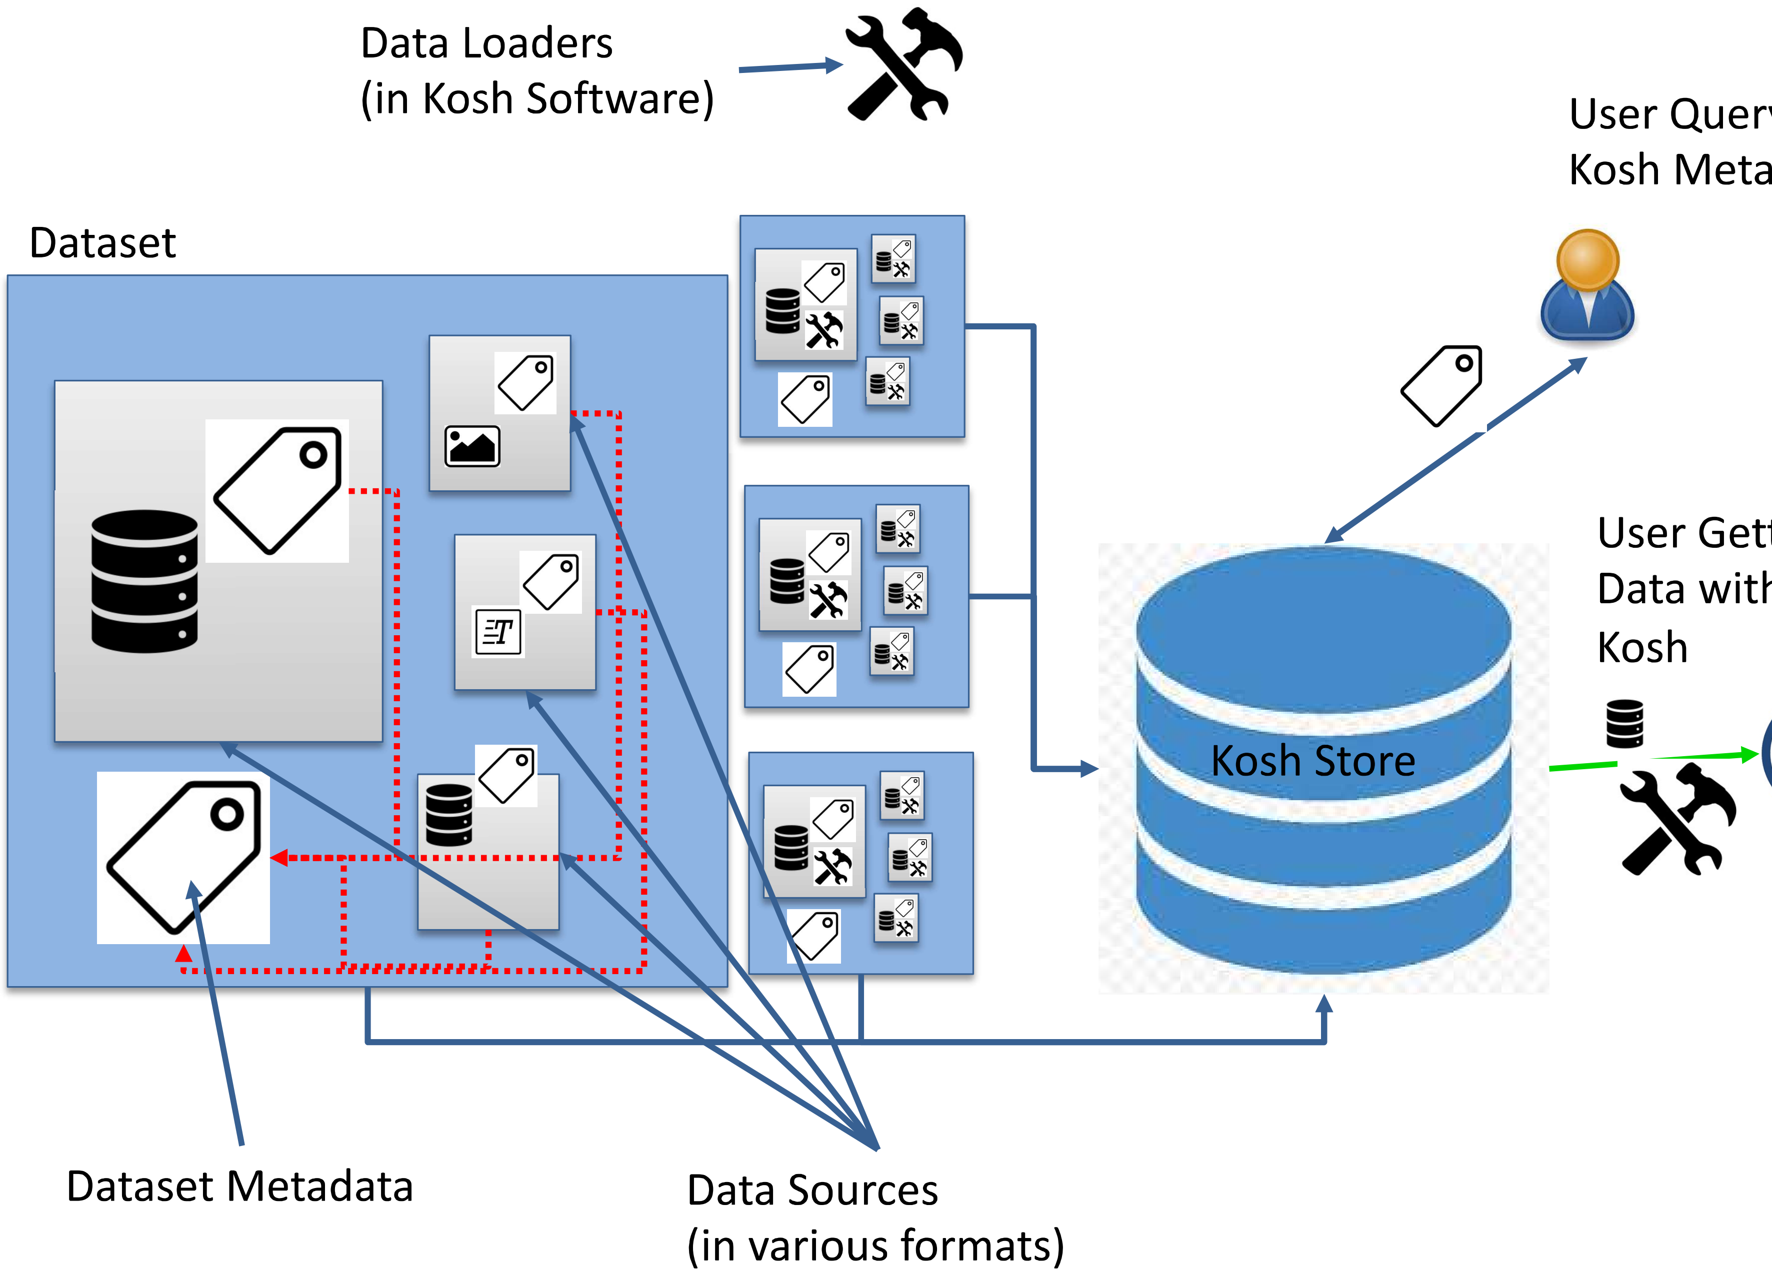
Data Loaders
(in Kosh Software)
Dataset
Kosh Store
Dataset Metadata
Data Sources
(in various formats)

Exploration of new, M
hardware

Specialized hardware is also emerging in the HPC space Cerebras CS-1 is being integrated into Lassen

LLNL is strategically looking at AI test applications across Scientific Computing programs

- High-precision scientific **simulation**
- Frequent ML **training**
- Potentially very high-frequency **inference**

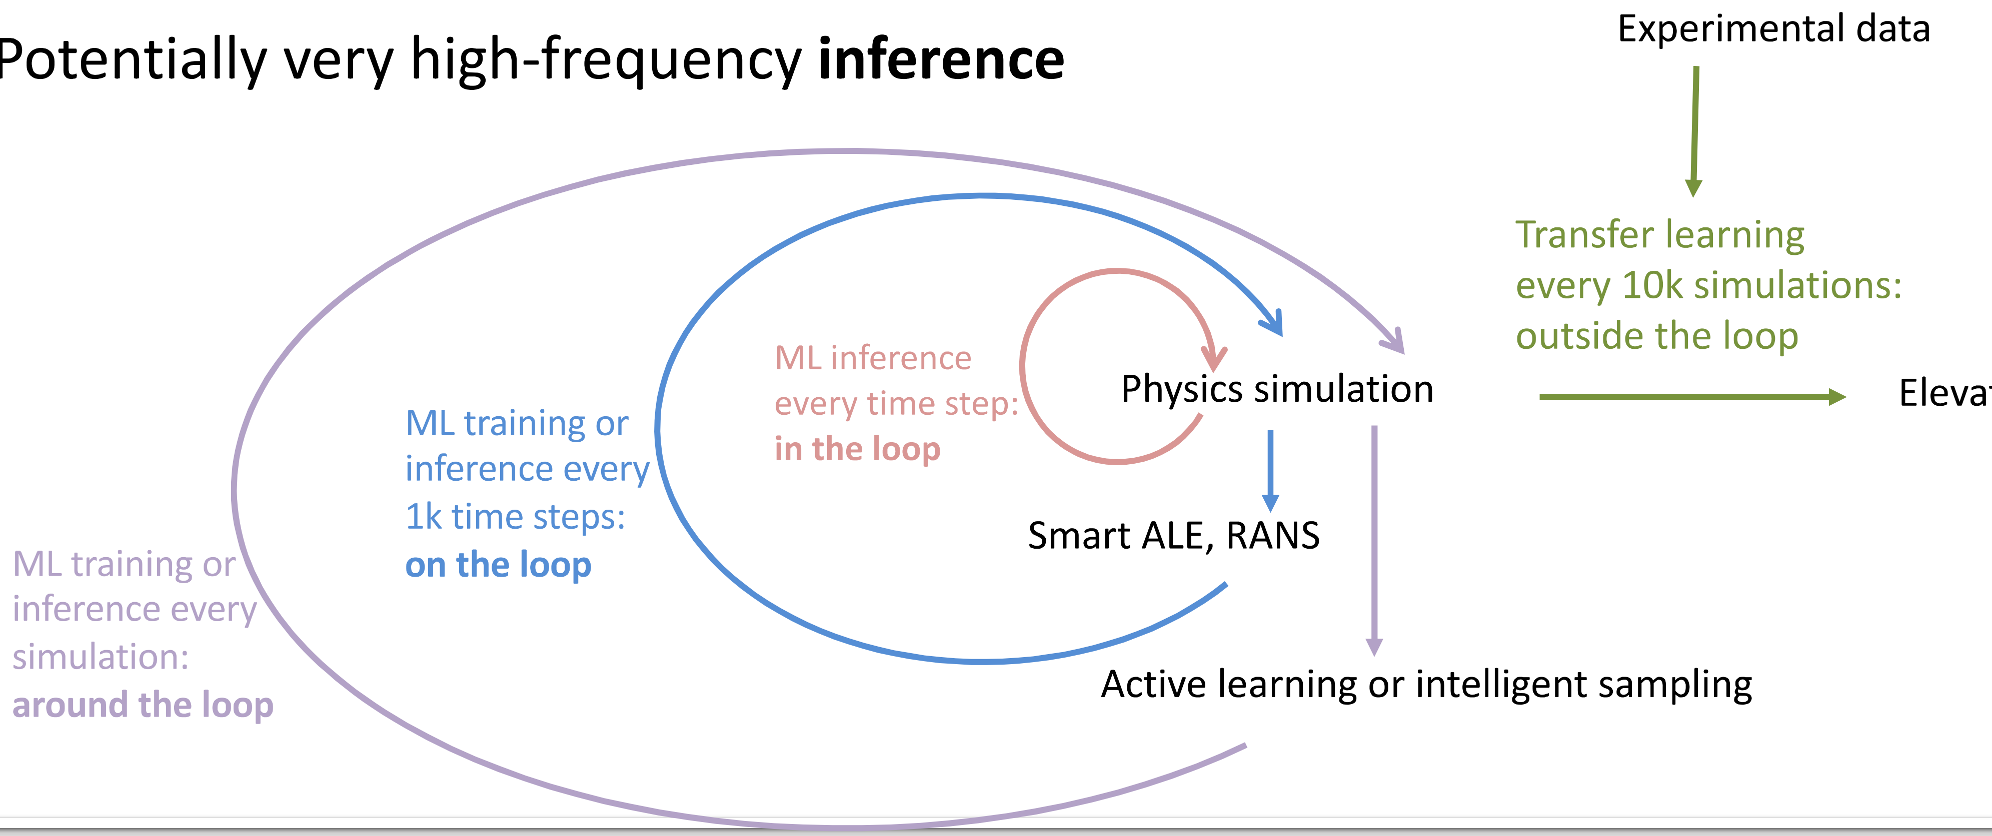

We are developing a proxy application to understand memory and bandwidth issues with accelerators for ALE

- High-precision scientific **simulation**
- Frequent ML **training**
- Potentially very high-frequency **inference**

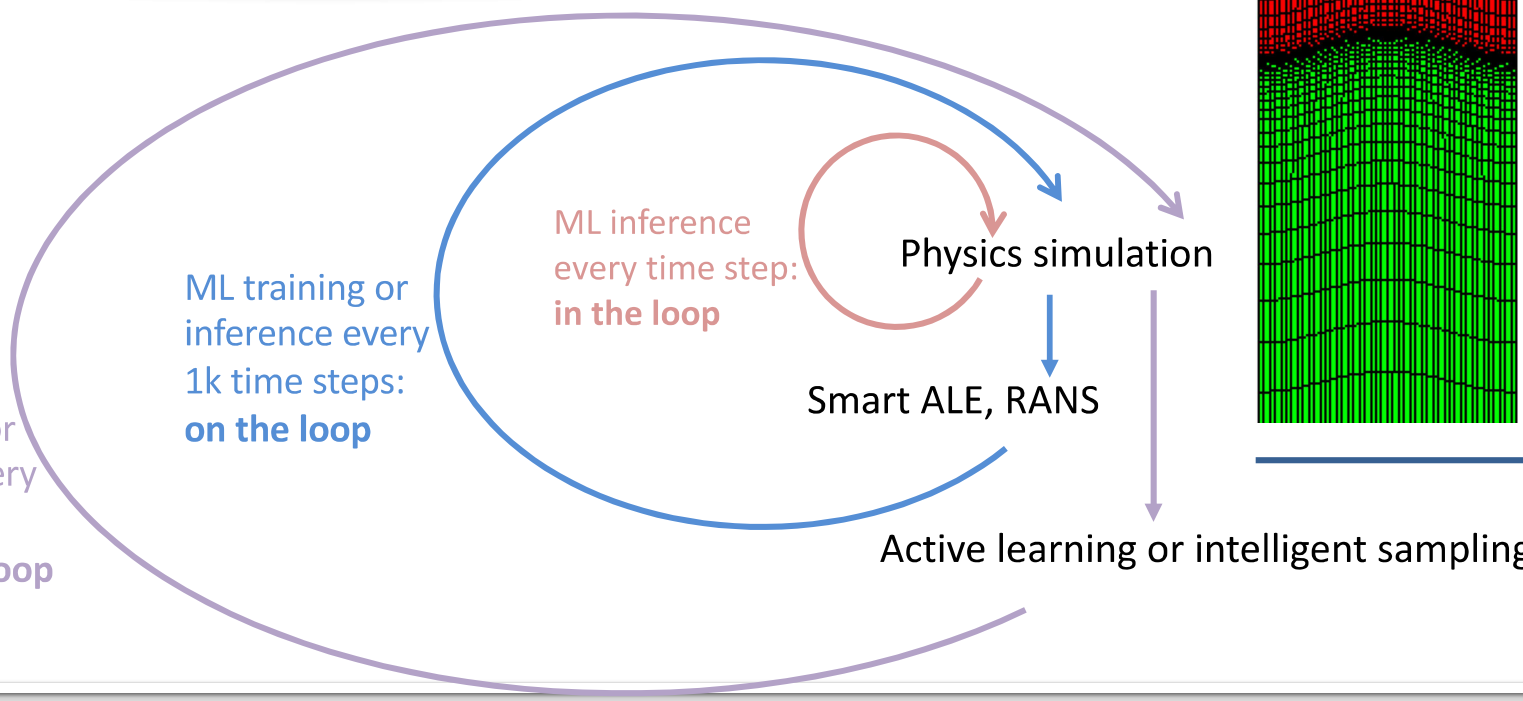

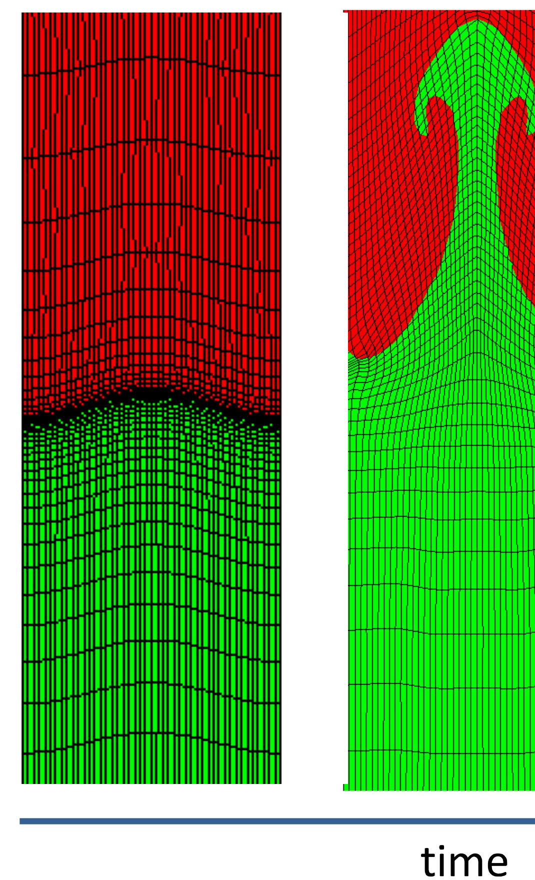

Data Science potential spans Scientific Computing sp

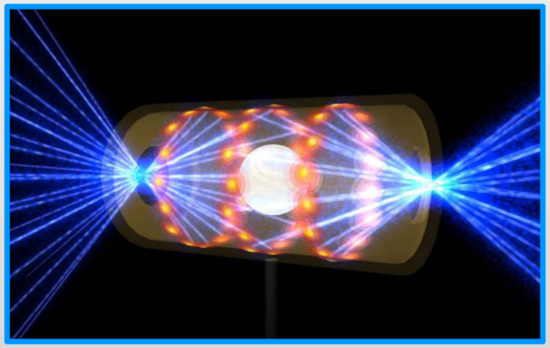

Enhanced Design Workflow

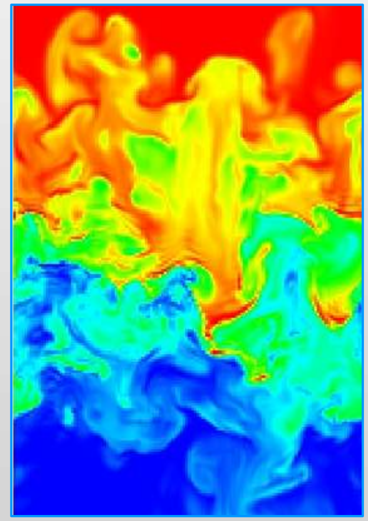

Enhanced Modeling

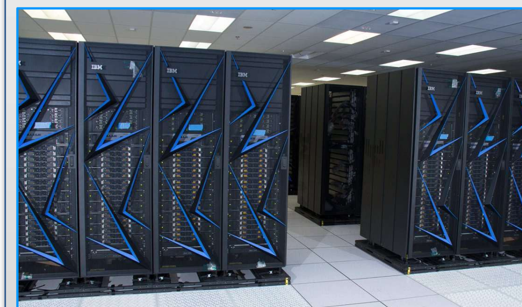

Improved System Performance

Physics Constrained Predictions

Interpretable Predictions

Community Engagement

Many Thanks!

Morry Aufderheide
Rob Blake
Kevin Chen
Sean Copeland
Charles Doutriaux
Dan Fenn
Brian Gallagher
Becky Haluska
Keith Henderson
Kevin Huynh
Ming Jiang

Josh Kallman
Ian Karlin
Alister Maguire
Walt Nissen
Brian Spears
Tom Stitt
Hardeep Sullan
Brian Van Essen
Ping Wang
Kenny Weiss
Kris Zieb

Lawrence Livermore National Laboratory