Workshop on Artificial Intelligence Applied to Photon and Neutron Science

Machine Learning & Artificial Intelligence for Science at Argonne

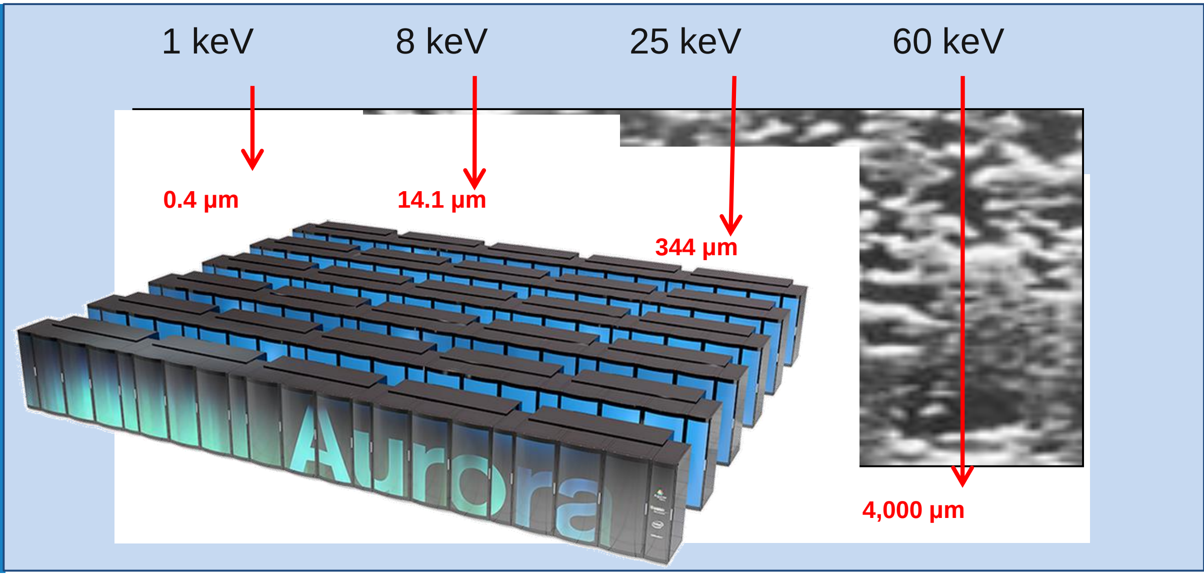

STEFAN WILD

Computational Mathematician, Laboratory for Applied Mathematics, Numerical Software, & Statistics
Deputy Director, Mathematics & Computer Science Division

Argonne National Laboratory

ESRF, Grenoble
November 13, 2019

Today:
1. ML/AI at Argonne
2. ML/AI at APS vignettes

AI for Science Requires New Infrastructure & Research

Applications	AI applications across science and engineering. Transformative approaches to simulation and experimental science.
Learning systems	AI software. Software infrastructure for managing data, models, workflows etc., and for delivering AI capabilities to 10,000s of scientists and engineers.
Foundations	Mathematics, algorithms; general AI, reinforcement learning, uncertainty quantification, explainability, optimization.
Hardware	Advanced hardware to support AI. Evaluation of new architectures and systems; exploration of neuromorphic and quantum as long term accelerators for AI.

Argonne NATIONAL LABORATORY

Data at Argonne: Scientific User Facilities

- ATLAS
- ARM
- CNM
- APS
- IVEM
- …

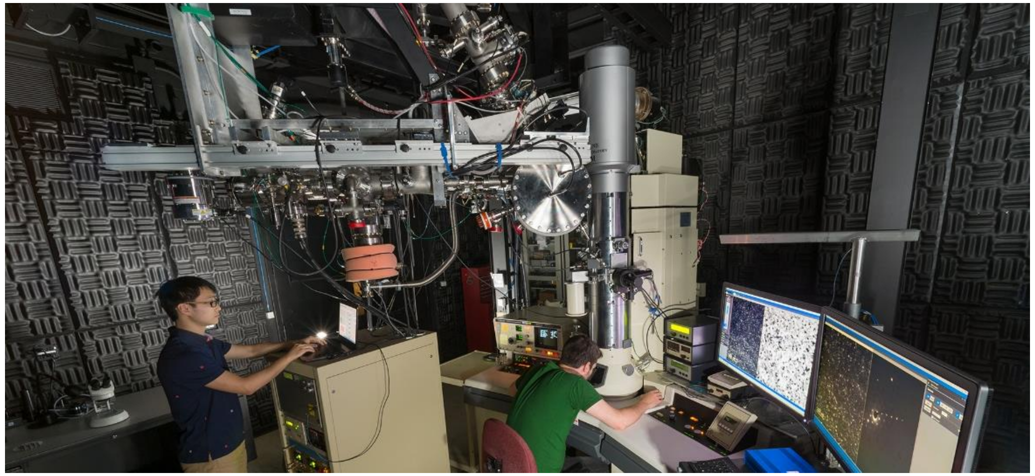

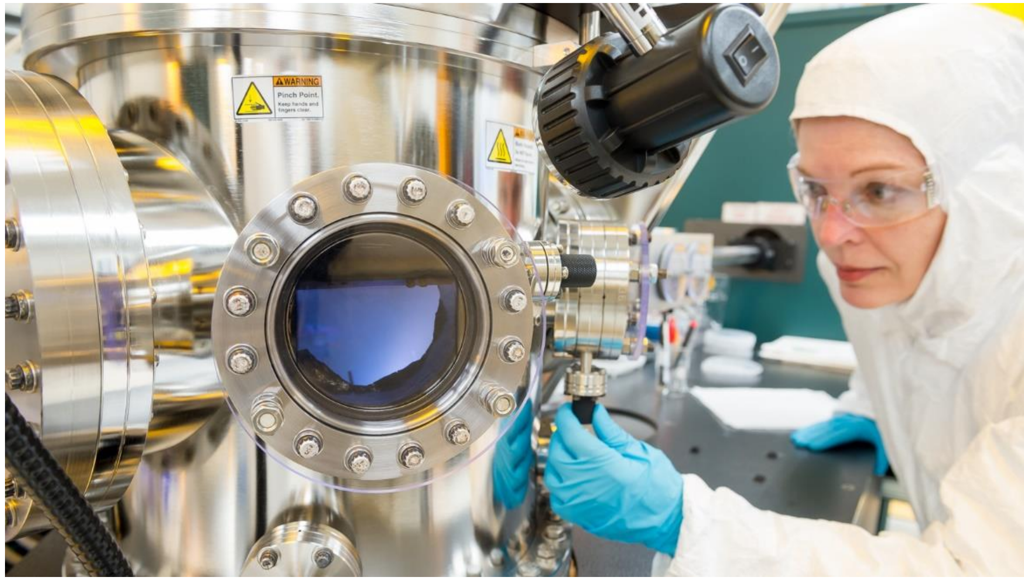

Intelligent Systems at Argonne: Energy, Manufacturing, Science, Synthesis

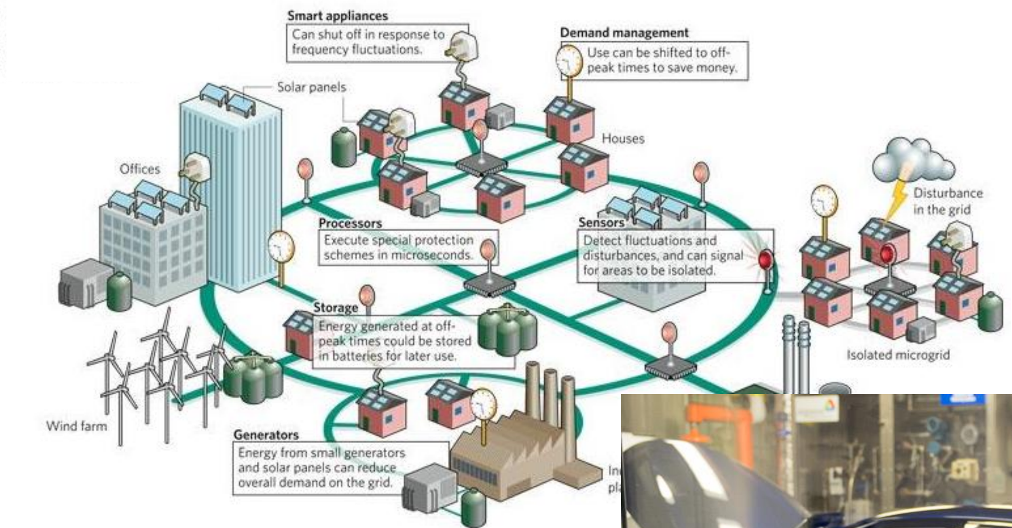

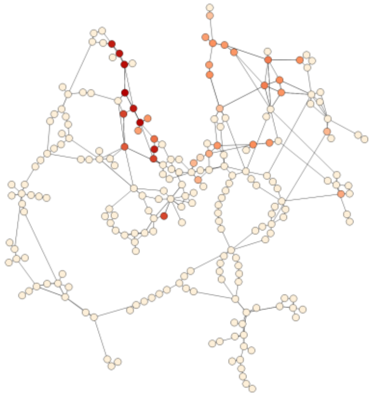

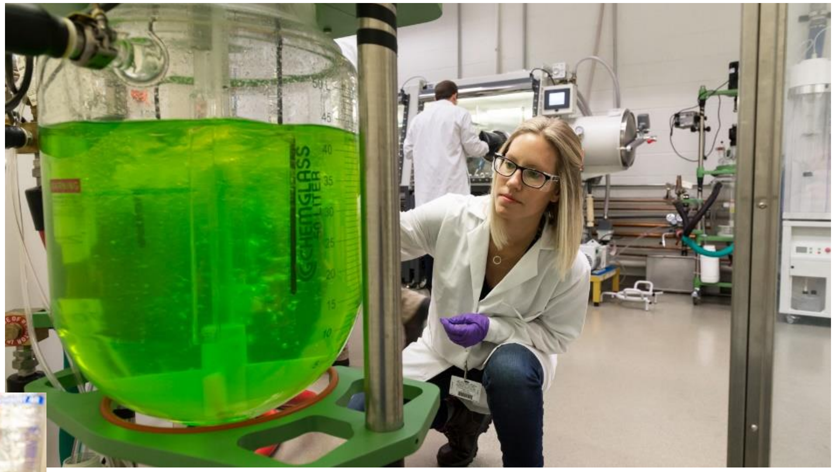

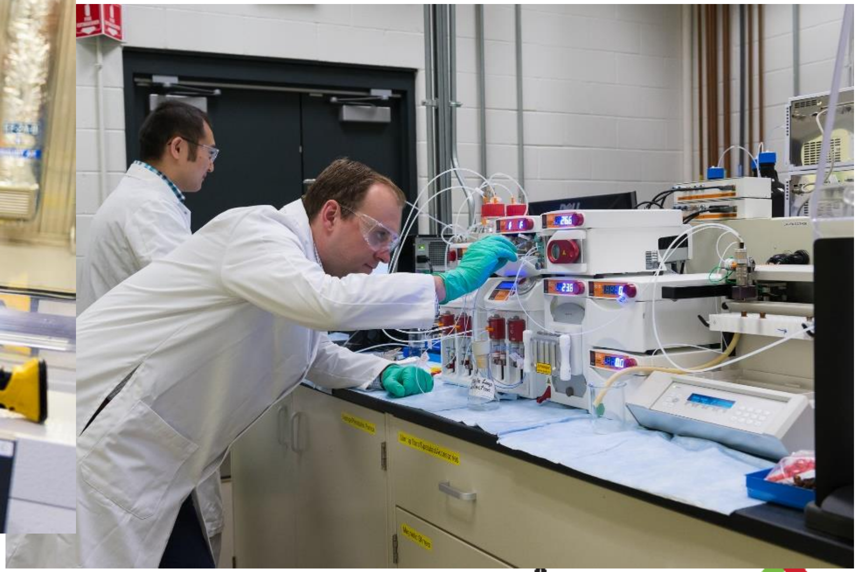

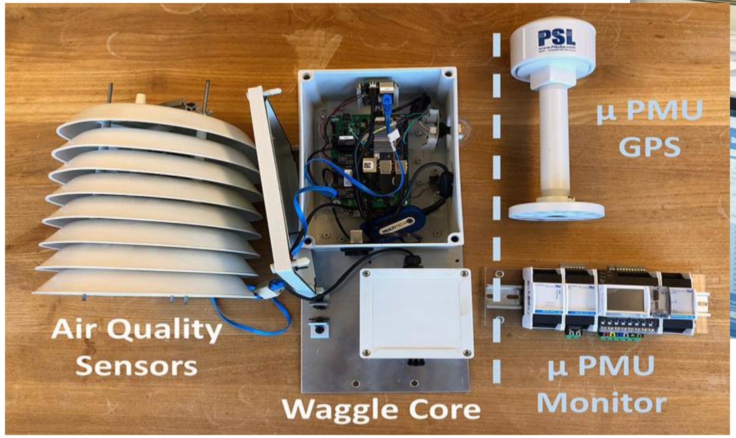

Aurora: US's First Exascale Supercomputer

Intel supercomputer to be delivered in 2021

Scaled up **to over 1000 PF**

Support for three "pillars" for applications in science, engineering, and security

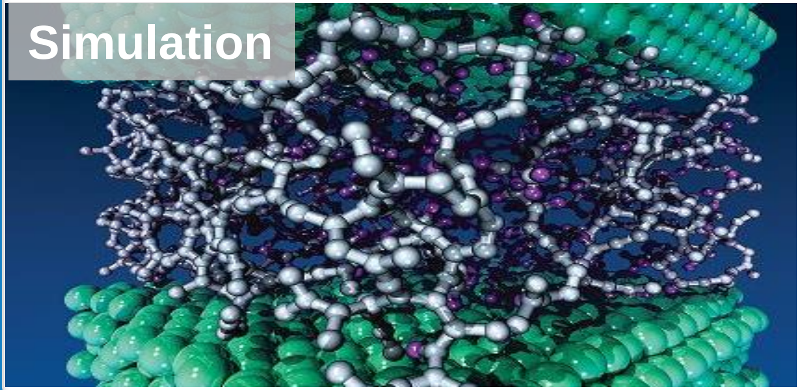

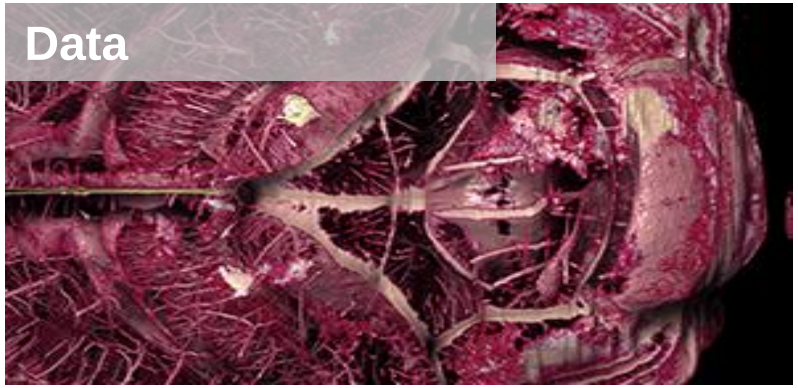

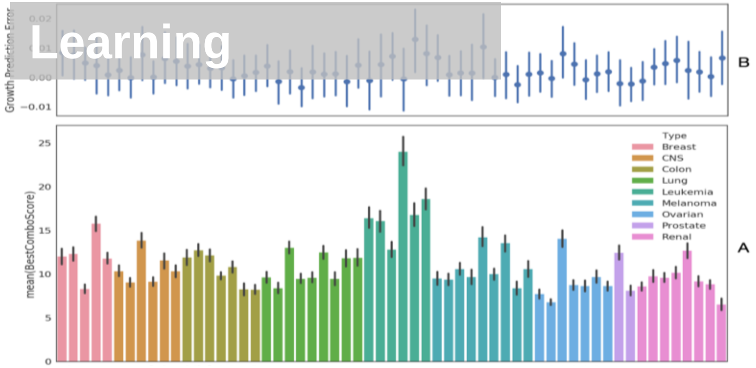

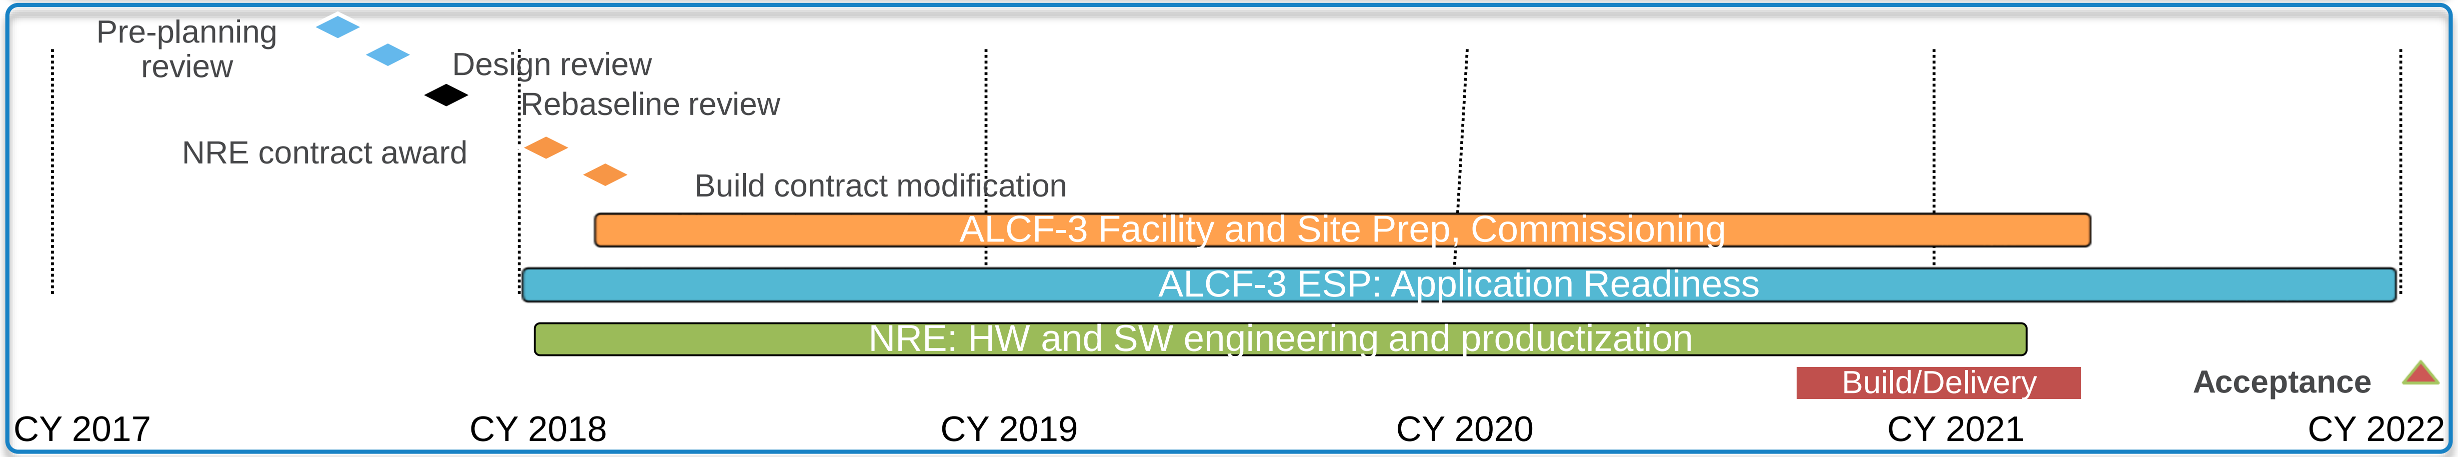

Argonne NATIONAL LABORATORY

Advancing the Foundations of AI for Science

Mathematical, statistical, & information-theoretic building blocks

Some exemplar grand challenges to address open problems

- How should scientific and engineering domain knowledge and governing principles from time-tested observations about natural phenomena be exploited in an AI era?

  Enable understanding

- What are the limits of AI techniques? What assumptions and circumstances can lead to establishing assurance of AI predictions and decisions for science & engineering?

  Increase trust

- Which AI techniques can best address different sampling scenarios and enable efficient AI on various computing and sensing environments?

  Broaden applicability

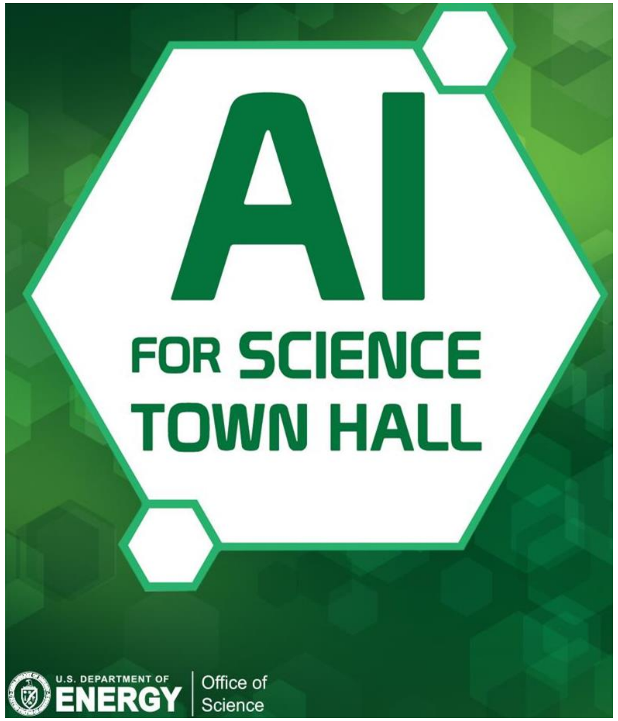

Expect report end of 2019

Emerging Examples of Leveraging Domain Knowledge

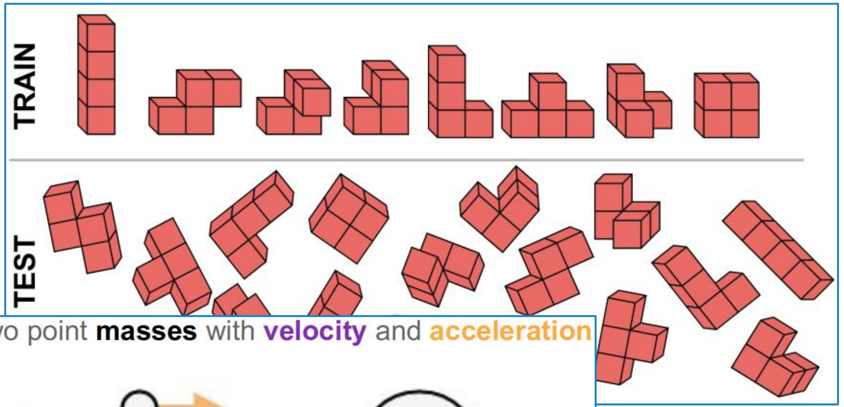

Smidt et al 2018

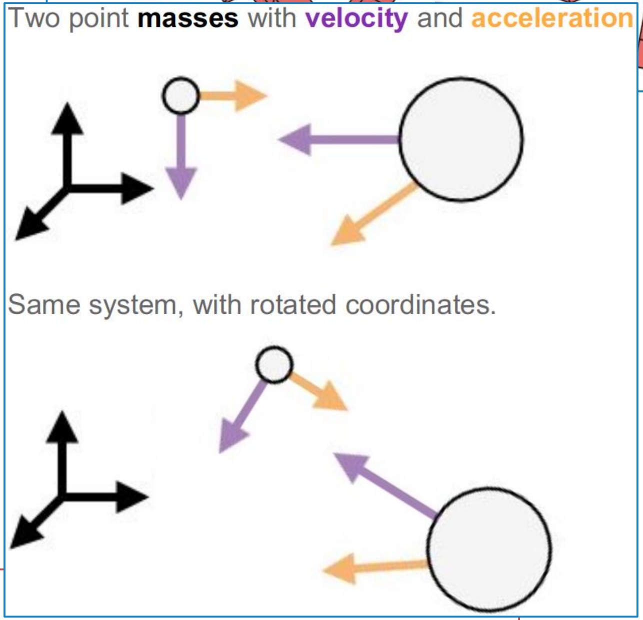

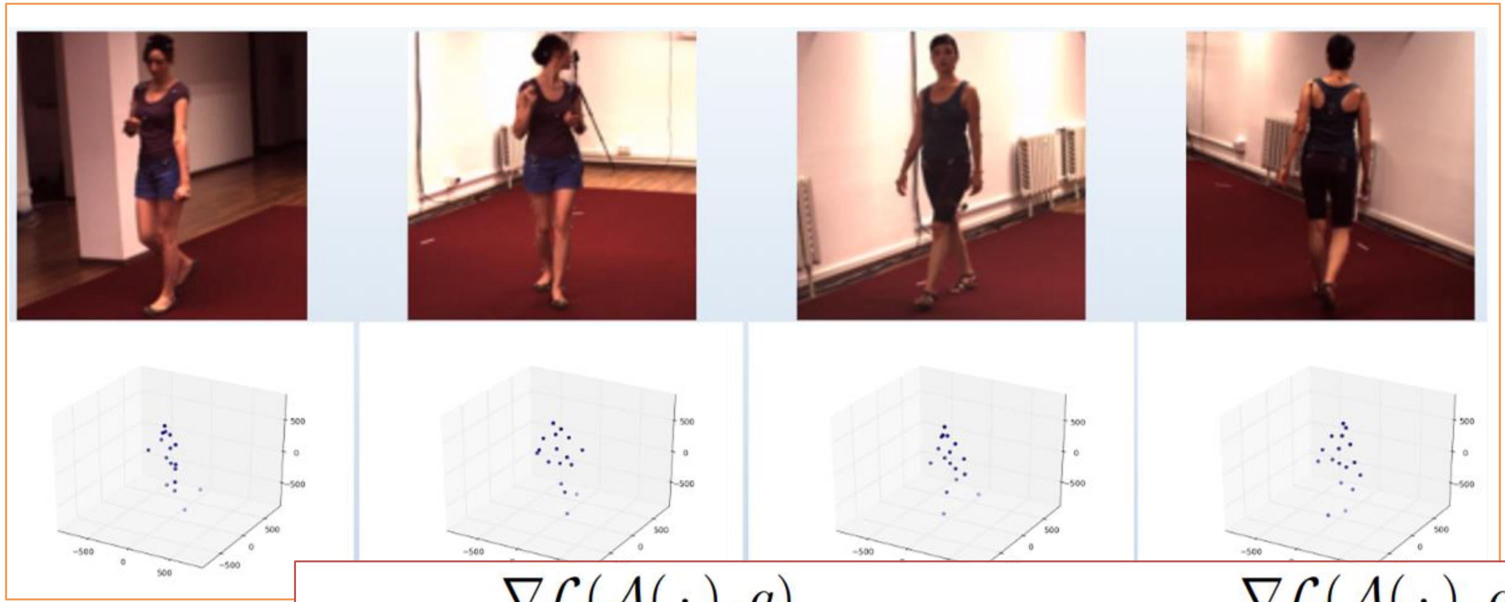

Marquez-Neila et al 2019

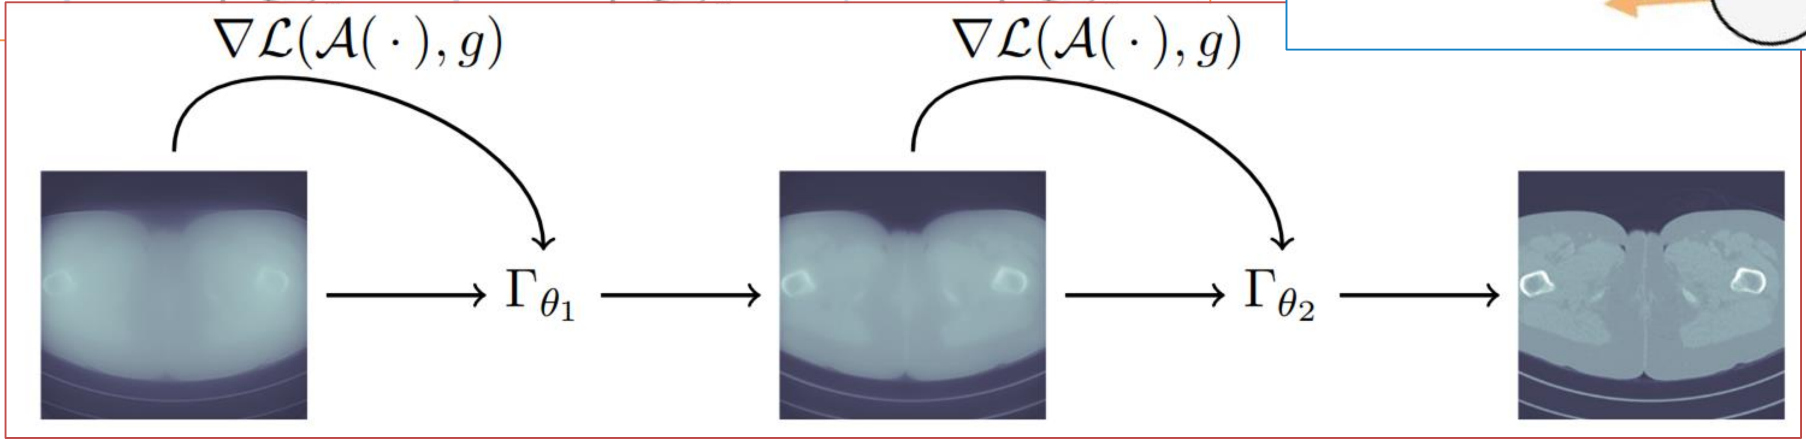

Arridge et al 2019

Moore’s Law for X-Ray Source Brightness

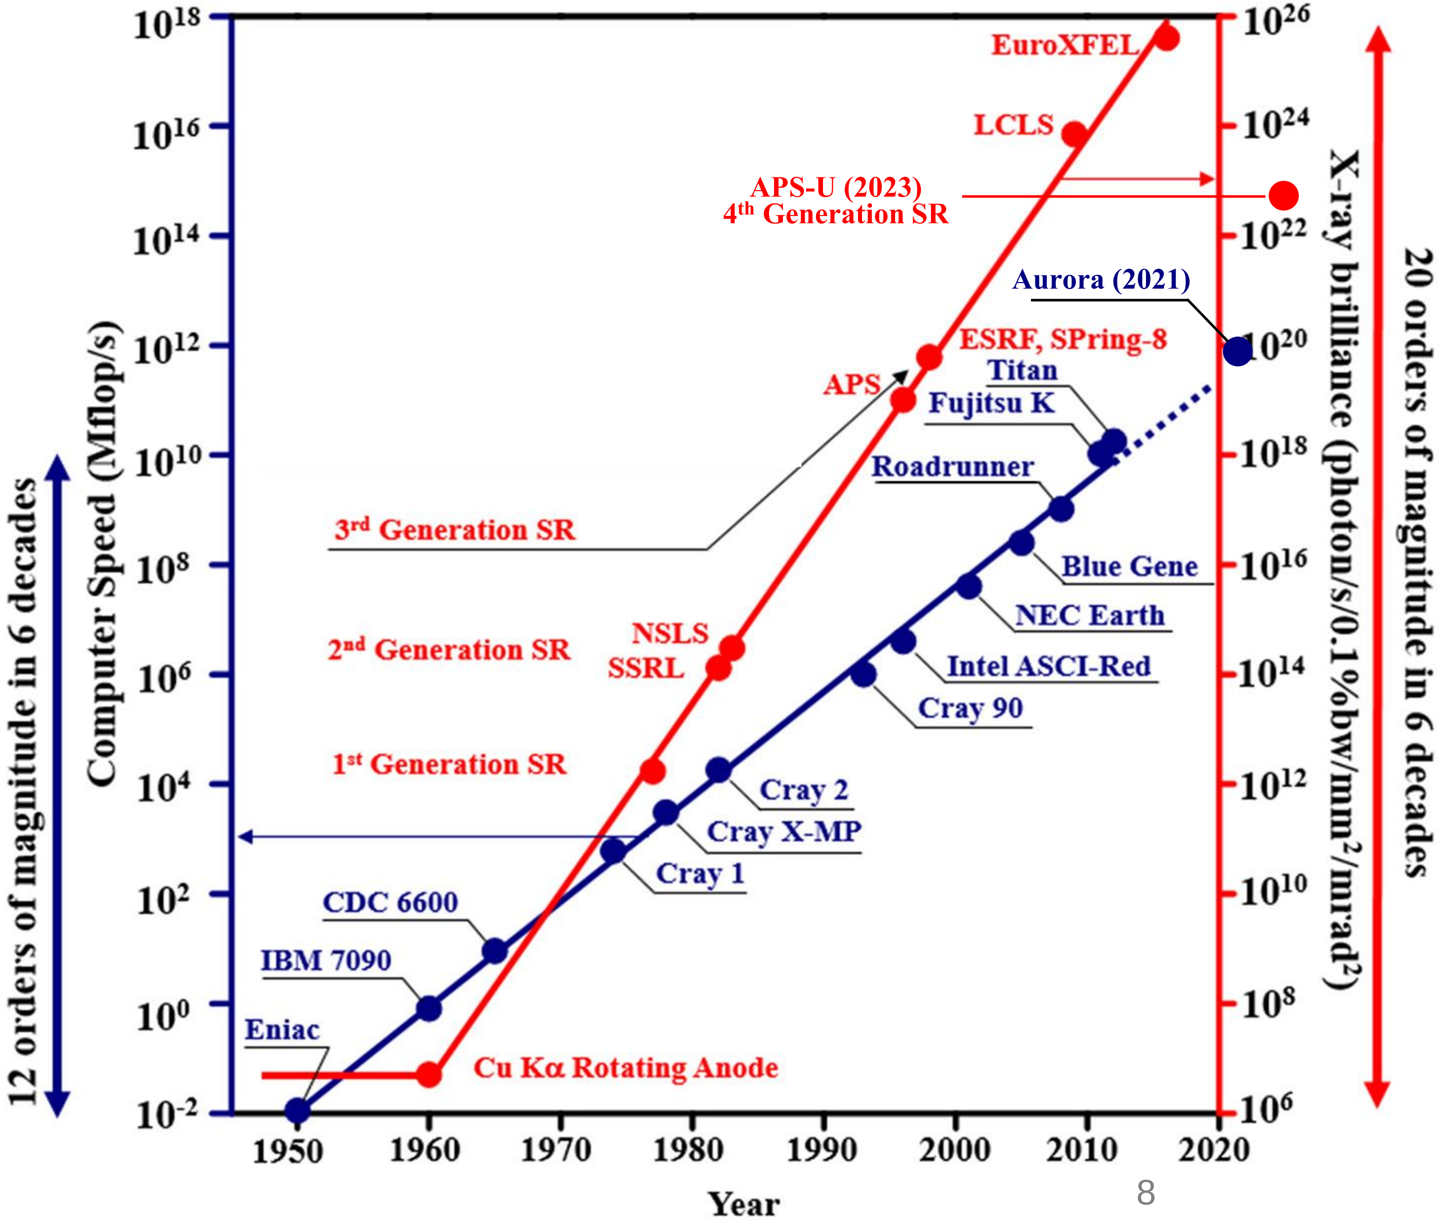

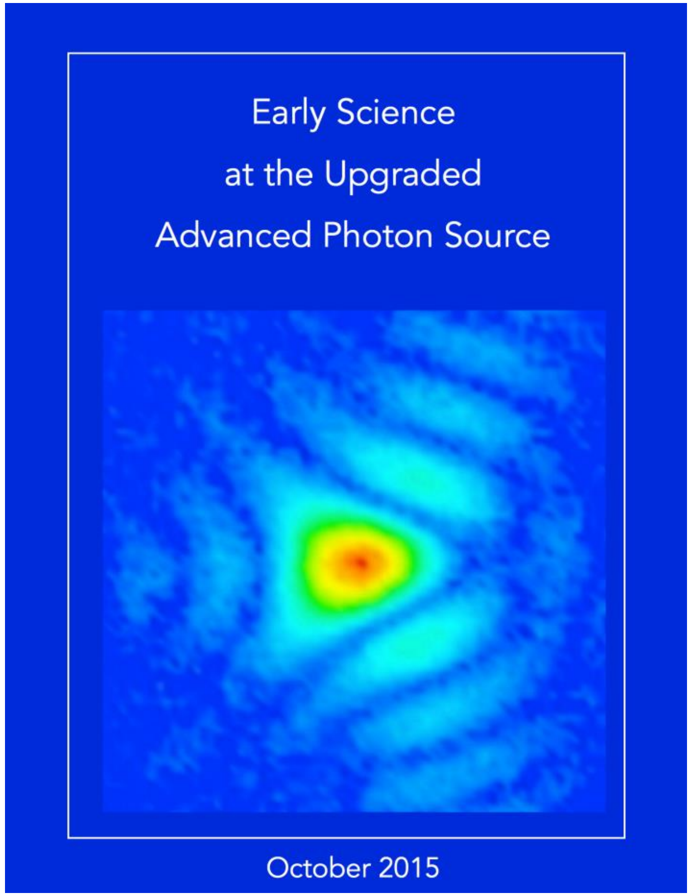

AI + X-Rays Motifs Emerging at Argonne

Ideal convergence: Intelligence in data, sensing, & computing

- **Scientific & measurement motifs include**
  - Detecting rare events
    large volumes, nanoscale resolution
  - Capturing dynamic processes
    the ultra fast to the ultra slow
  - Enabling multidimensional inquiry
    exploring spaces of higher dimension and size

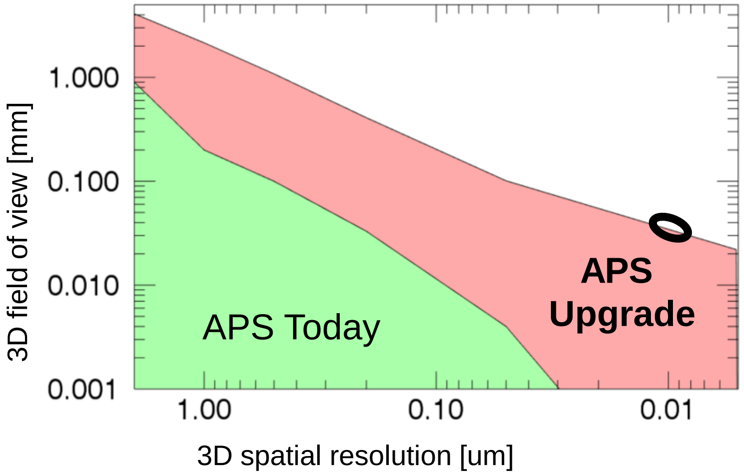

- **AI & data science motifs include**
  - Offline learning and experimental design:
    HPC & ML as offline scale-bridging, prediction, and experiment-generating engine
  - Analysis and reconstruction of massive, multimodal data volumes
  - Online/real-time feature detection, dimension reduction, reinforcement learning, UQ, and active learning to drive an experiment

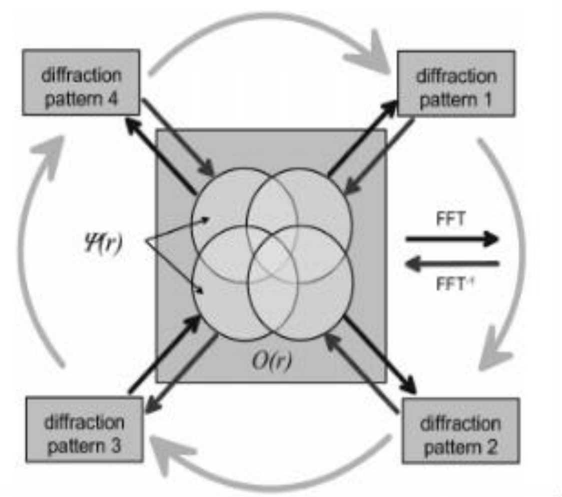

Neural Network-Enabled Diffractive Imaging

4D coherent diffraction imaging (CDI)

- CDI excels at 4D *operando* measurement
- Coherent imaging a primary driver for APS-U
  - 100X coherent flux:
    - How to effectively use these photons?
    - How to analyze 100X data increase?

Today: 512x512x512 arrays	APS-U: 5120x5120x5120 arrays
• ~30 Gb (for phasing)	• ~30 Tb (memory for phasing!)
• ~7 nm resolution	• ~ 7 A resolution

Current CDI limitations

- Real-time experimental feedback not possible
  - 1000s of iterations; multiple restarts
- Strong phase objects (e.g., materials with defects) do not reconstruct
- Novice users struggle to tune convergence parameters for successful object reconstruction

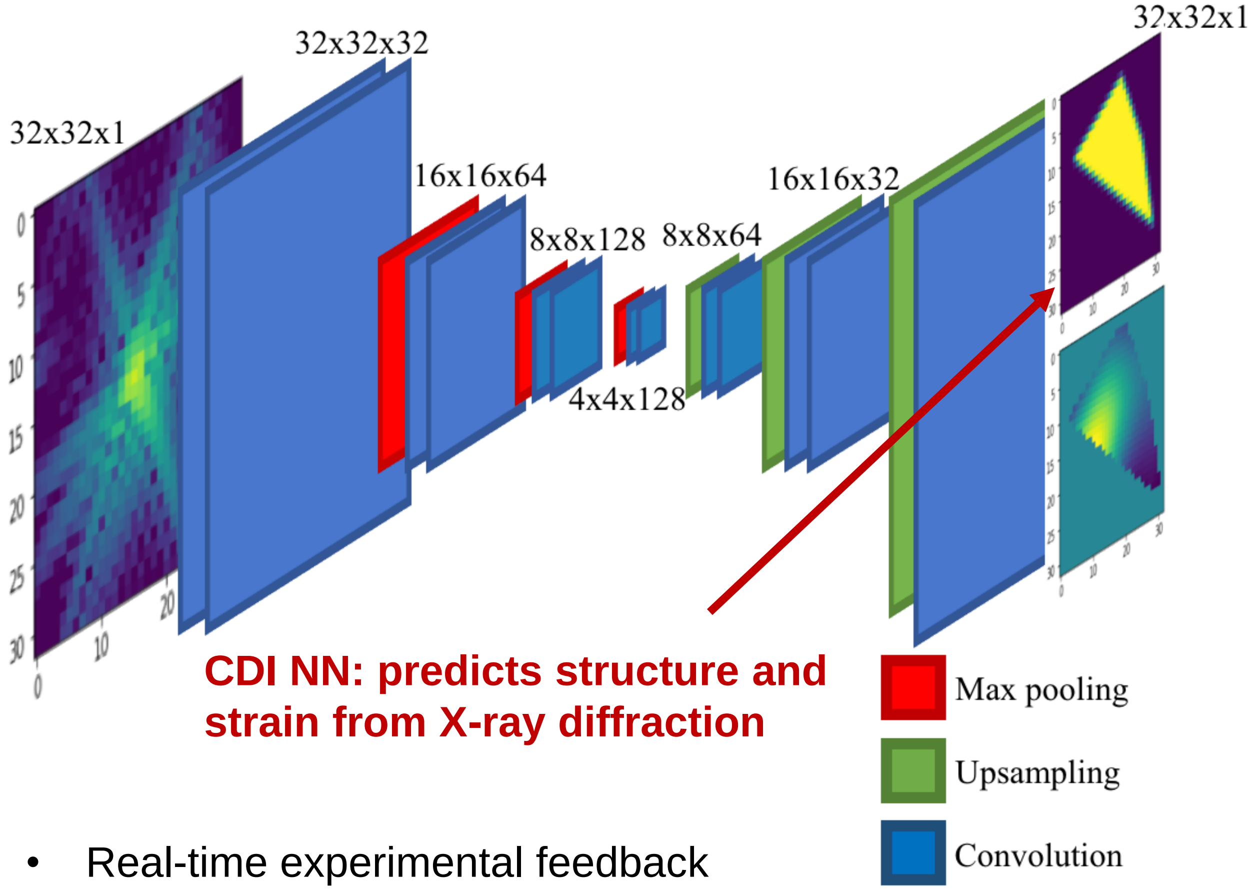

CDI NN: predicts structure and strain from X-ray diffraction

- Real-time experimental feedback
  - CDI NN is 500X faster
- Excellent predictions on strong phase objects
  - Can now image defective materials
- Once trained, no parameters for users to tune

Cherukara, Nashed, Harder. *Scientific Reports,* 2018

AI CDI: Deep Learning + Simulation + AD

Workflow

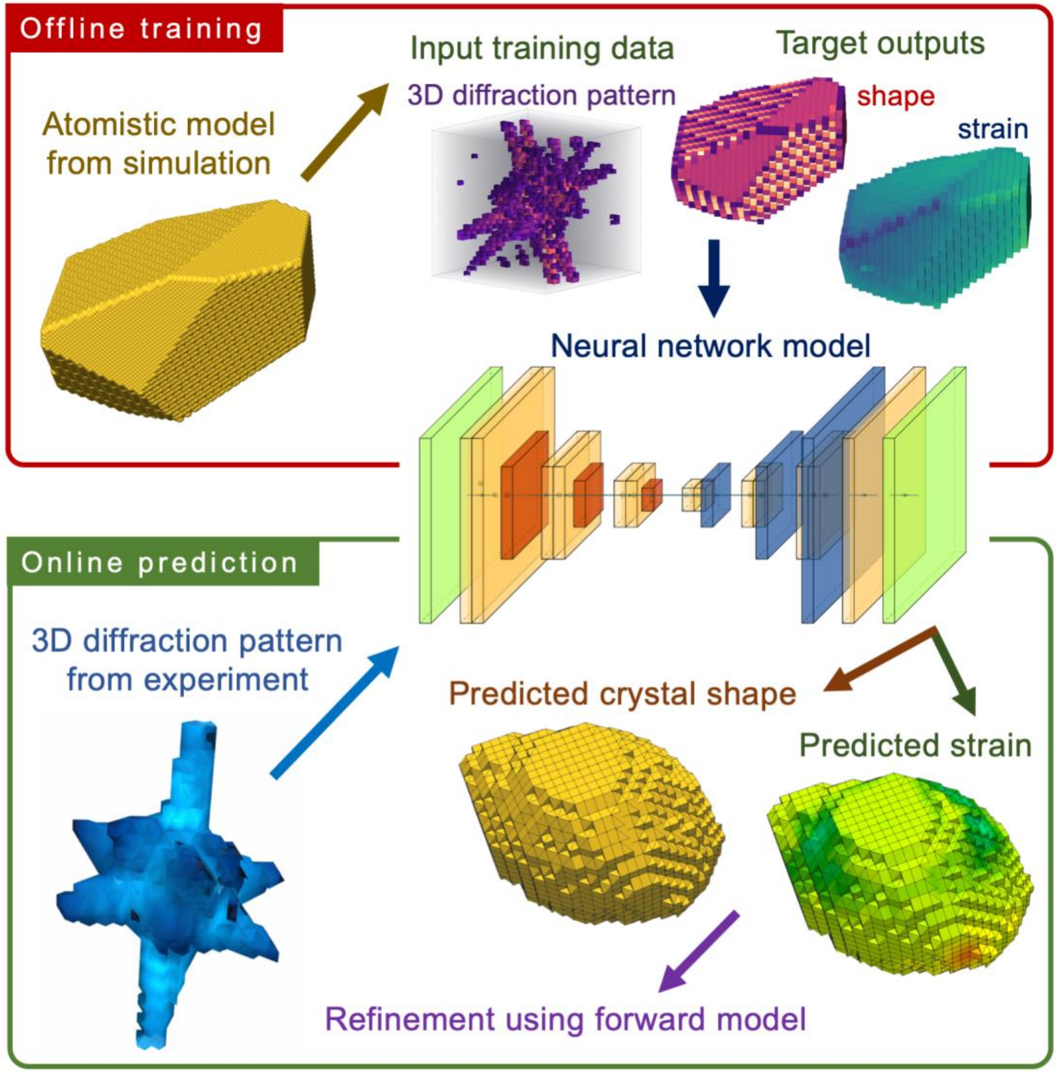

Performance on experimental data

Phase retrieved | NN prediction | AD refined

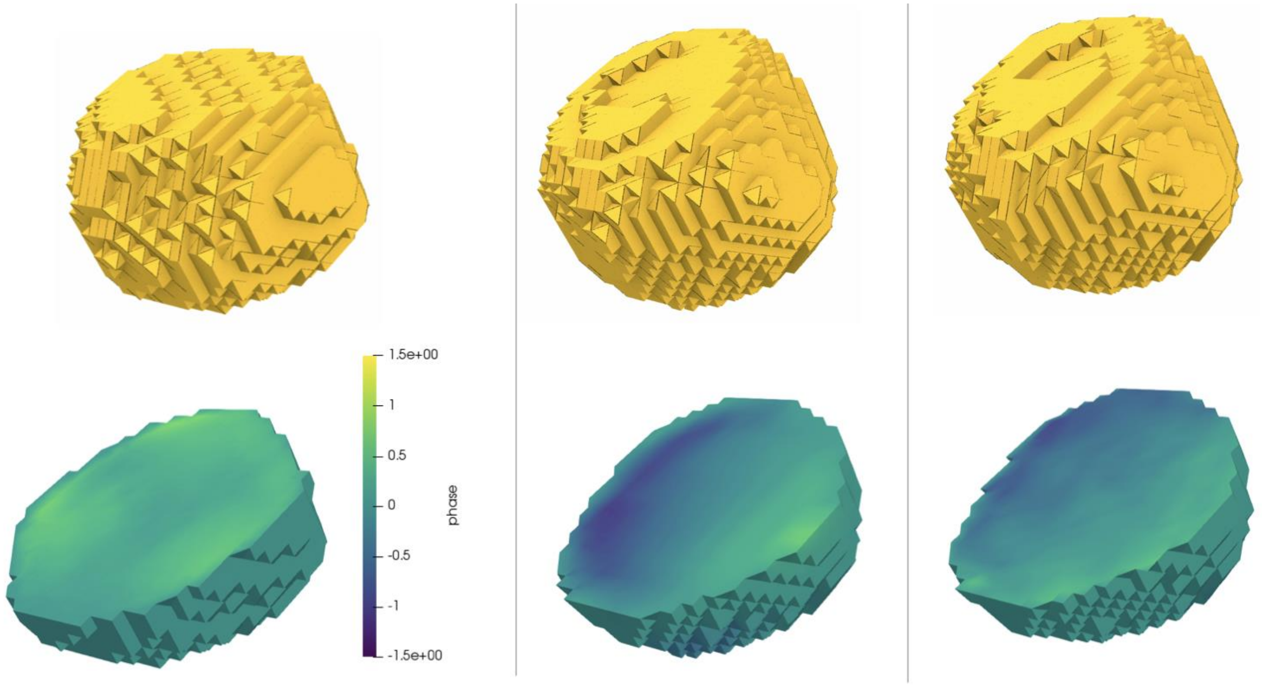

Minutes | Miliseconds | Seconds

Reconstruction time

H Chan, Y Nashed, S Sankaranarayanan, R Harder, M Cherukara

ENHANCED DIFFRACTIVE IMAGING

High-Resolution Ptychography on Large-Scale Samples

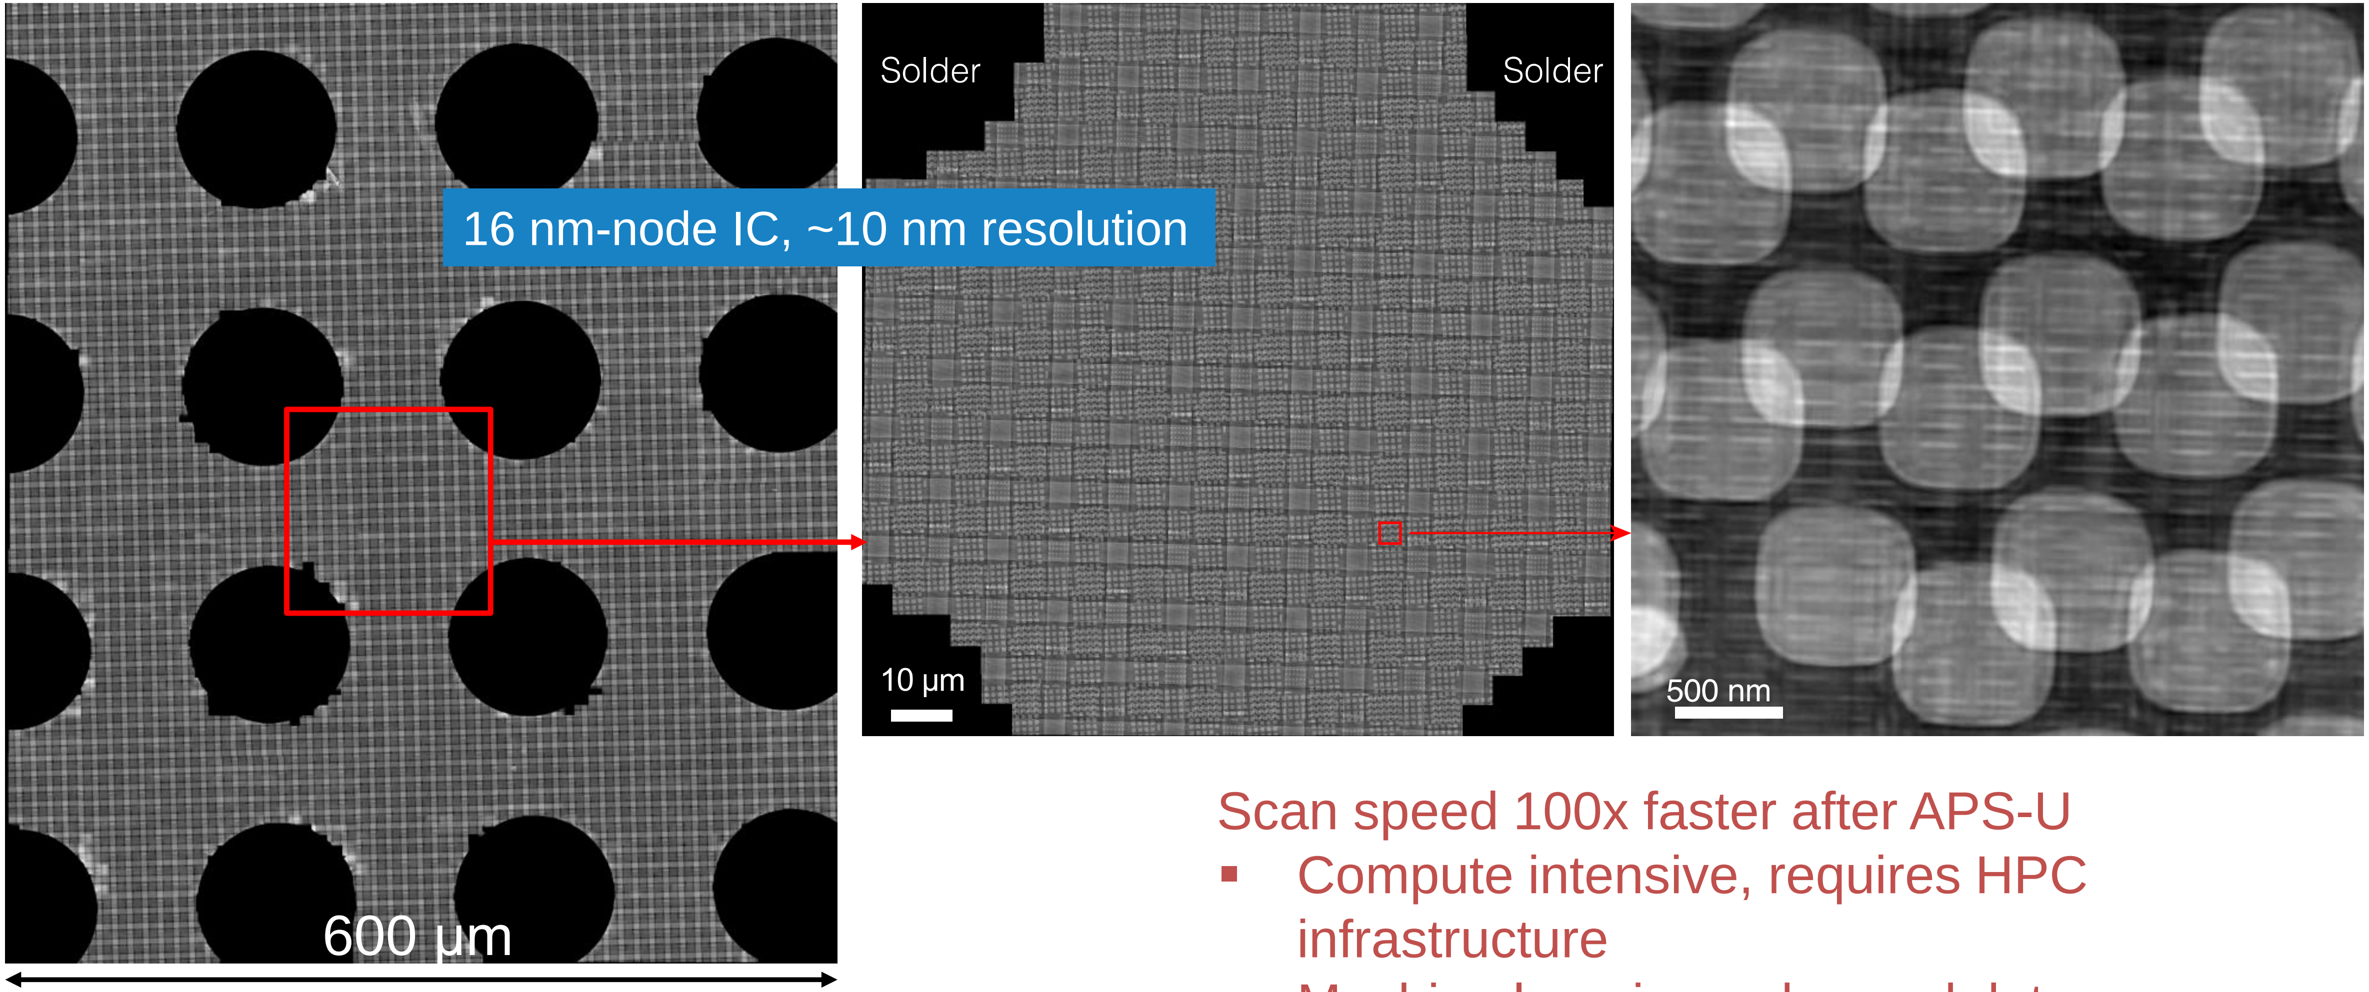

- 0.5 mm^2 area, 4 scan days, 42 TB raw data
- 10 GPUs, 30 days of reconstruction

Scan speed 100x faster after APS-U

- Compute intensive, requires HPC infrastructure
- Machine learning-enhanced data acquisition and processing

Courtesy Junjing Deng

Argonne NATIONAL LABORATORY

Tools for Beyond-Depth-of-Focus Imaging

Exploit better understanding of the forward problem

- Looming problem: depth of focus goes like

  DOF≃5(transverse resolution)^2/λ

  – Ex.- 5(5 nm)^2 at 10 keV gives DOF=4 μm
- Entering a regime where we no longer obtain pure projection images from *thick specimens*
  – x-ray imaging is otherwise perfect for
- Must account for beam propagation within the specimen, and how that changes illumination of downstream planes
  – violation of the Born approximation

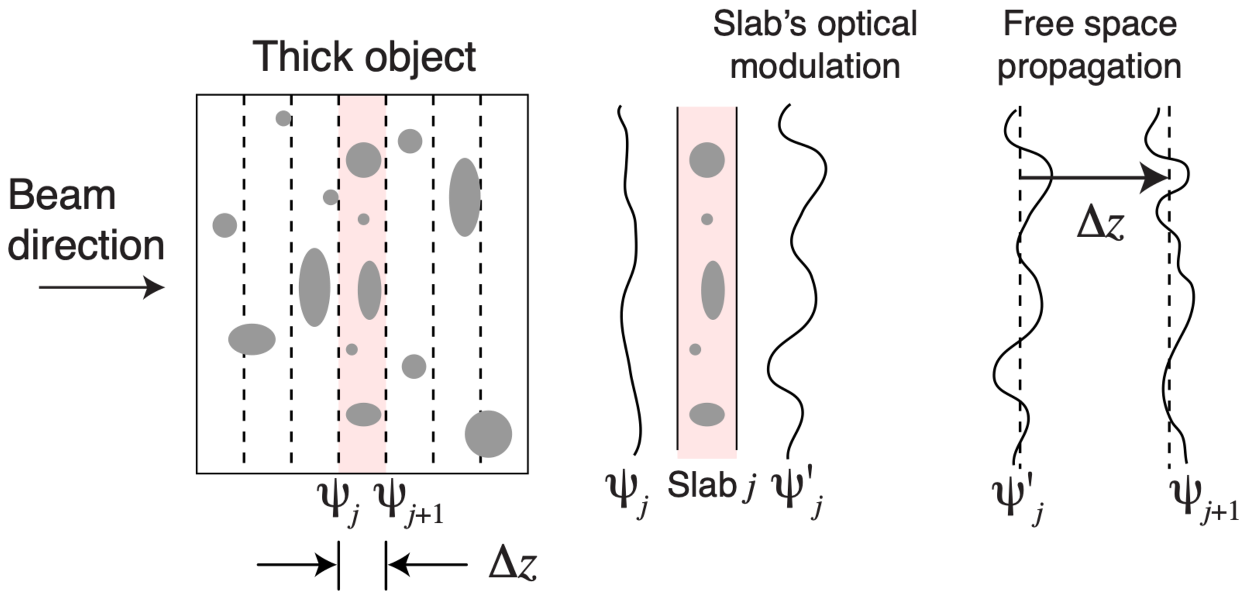

Multislice method: Cowley & Moodie, *Proc PhySocLondon B* (1957) *Acta Crystallographica* (1957)

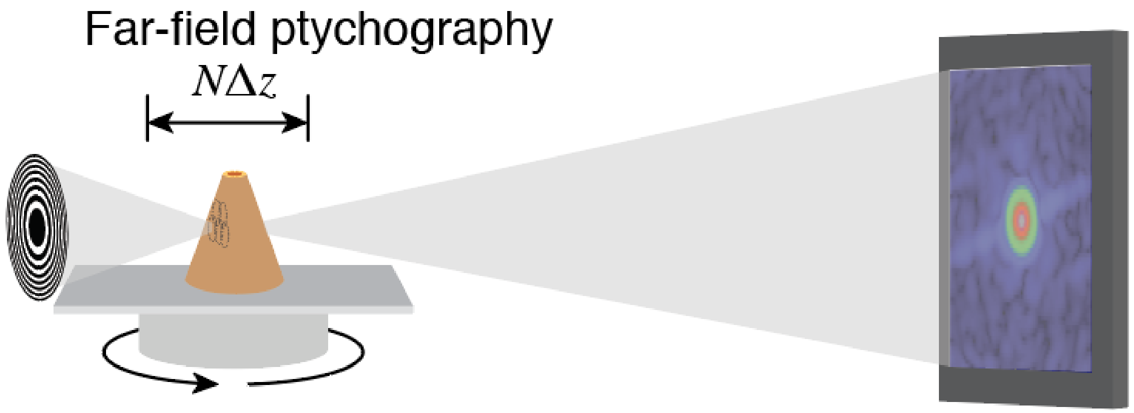

M Du, Y Nashed, S Kandel, D Gursoy, C Jacobsen

Tools for Beyond-DOF Imaging

- Gradient-based optimization approach[1]
- Automatic differentiation (AD) lets one easily change noise models, and add multiple nonlinear regularizers
- AD has been demonstrated for ptychography[2] and several other coherent diffraction imaging methods[3]
- Now being used for beyond-DOF imaging[4] (right)

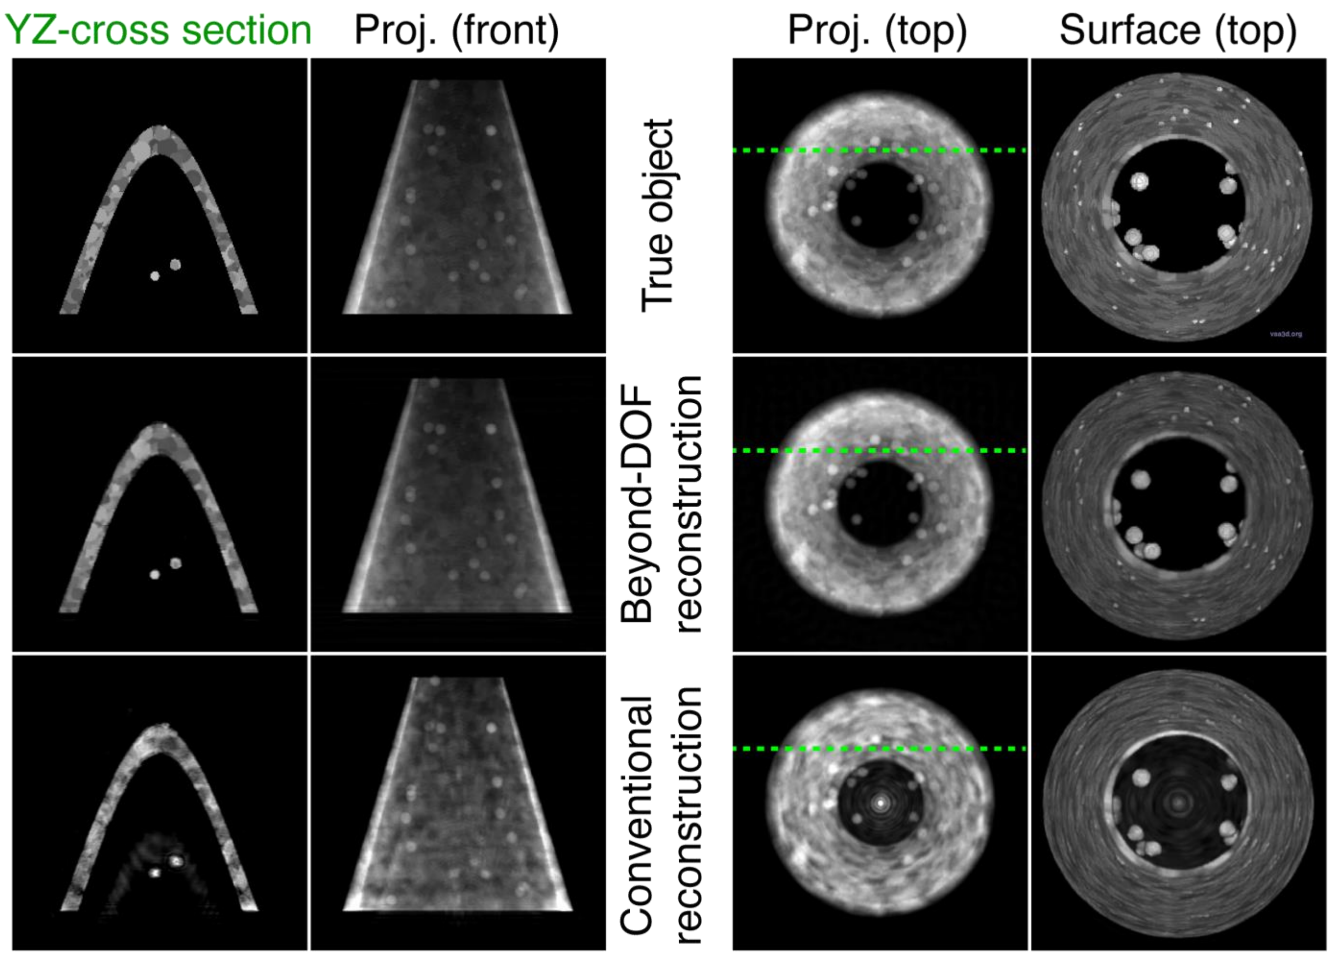

1. Gilles, Nashed, Du, Jacobsen, Wild, *Optica* **5**, 2018
2. Nashed, Peterka, Deng, Jacobsen, *Proc Comp Sci* **108C**, 2017
3. Kandel, Maddali, Allain, Hruszkewycz, Jacobsen, Nashed, *Opt Exp* **27**, 2019
4. Du, Nashed, Kandel, Gursoy, and Jacobsen, Preprint 2019

Optimization-Based Sensing Error Recovery

Automatically correcting experimental errors in tomography

Problem in tomographic imaging:
- Sample/center-of-rotation (CoR) drift

Approach
- Embed CoR drift parameters in the nonlinear optimization problem of sample recovery
- Iteratively reconstruct sample and drift (implicitly and explicitly)

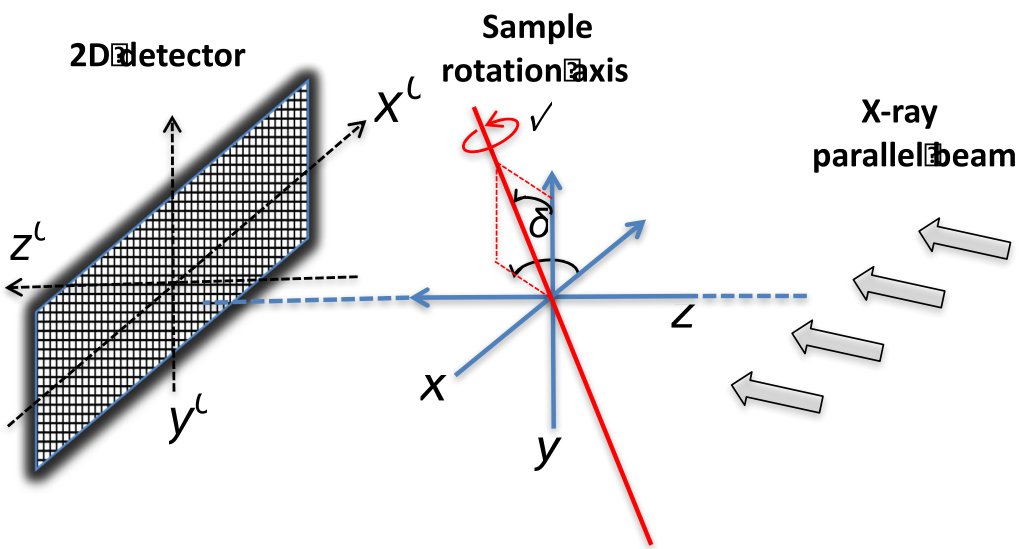

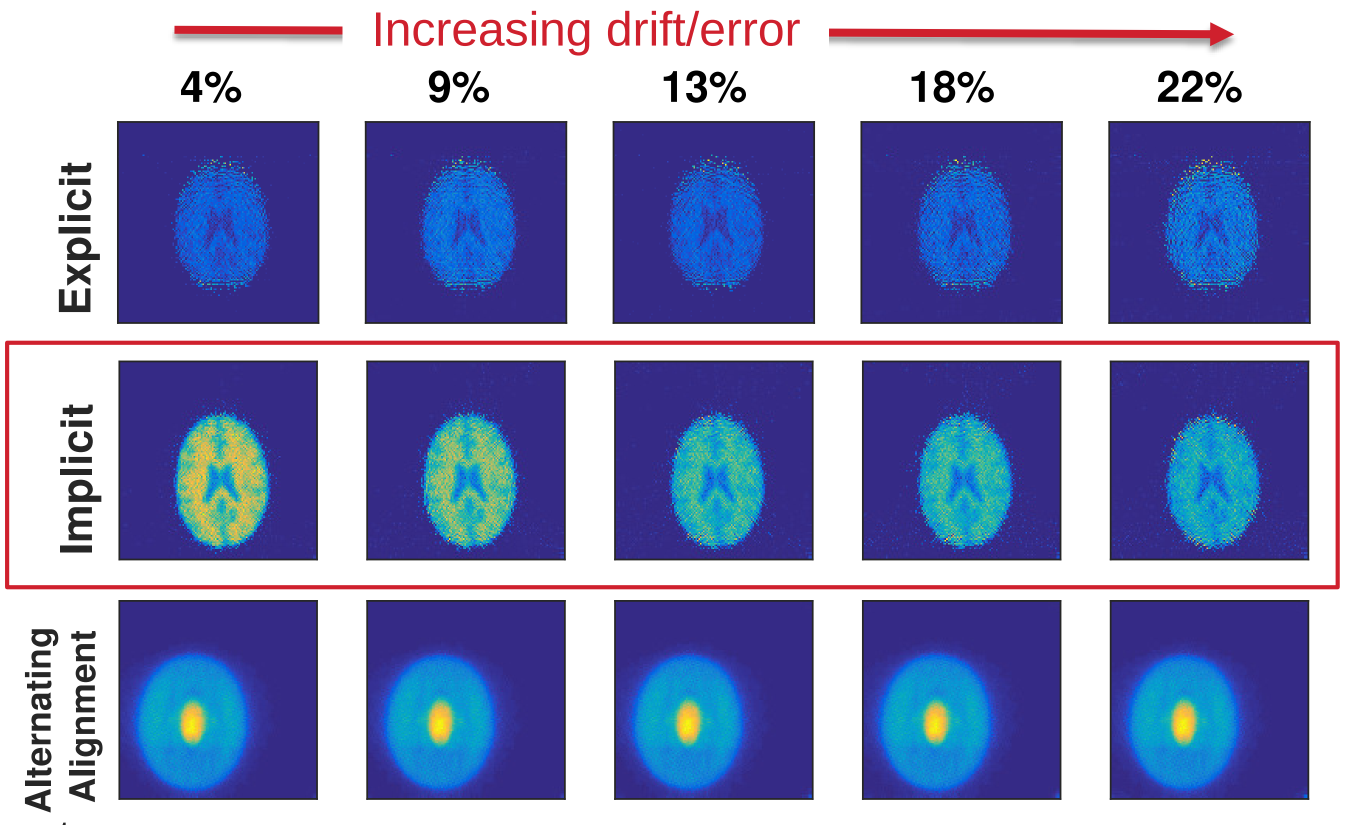

Austin, et al. *SIAM J. Scientific Computing*, 2019
Di, et al. *IEEE ICIP*, 2019

SEMI-SUPERVISED LEARNING

DL-Based Automatic Segmentation

Nano tomography volume segmentation

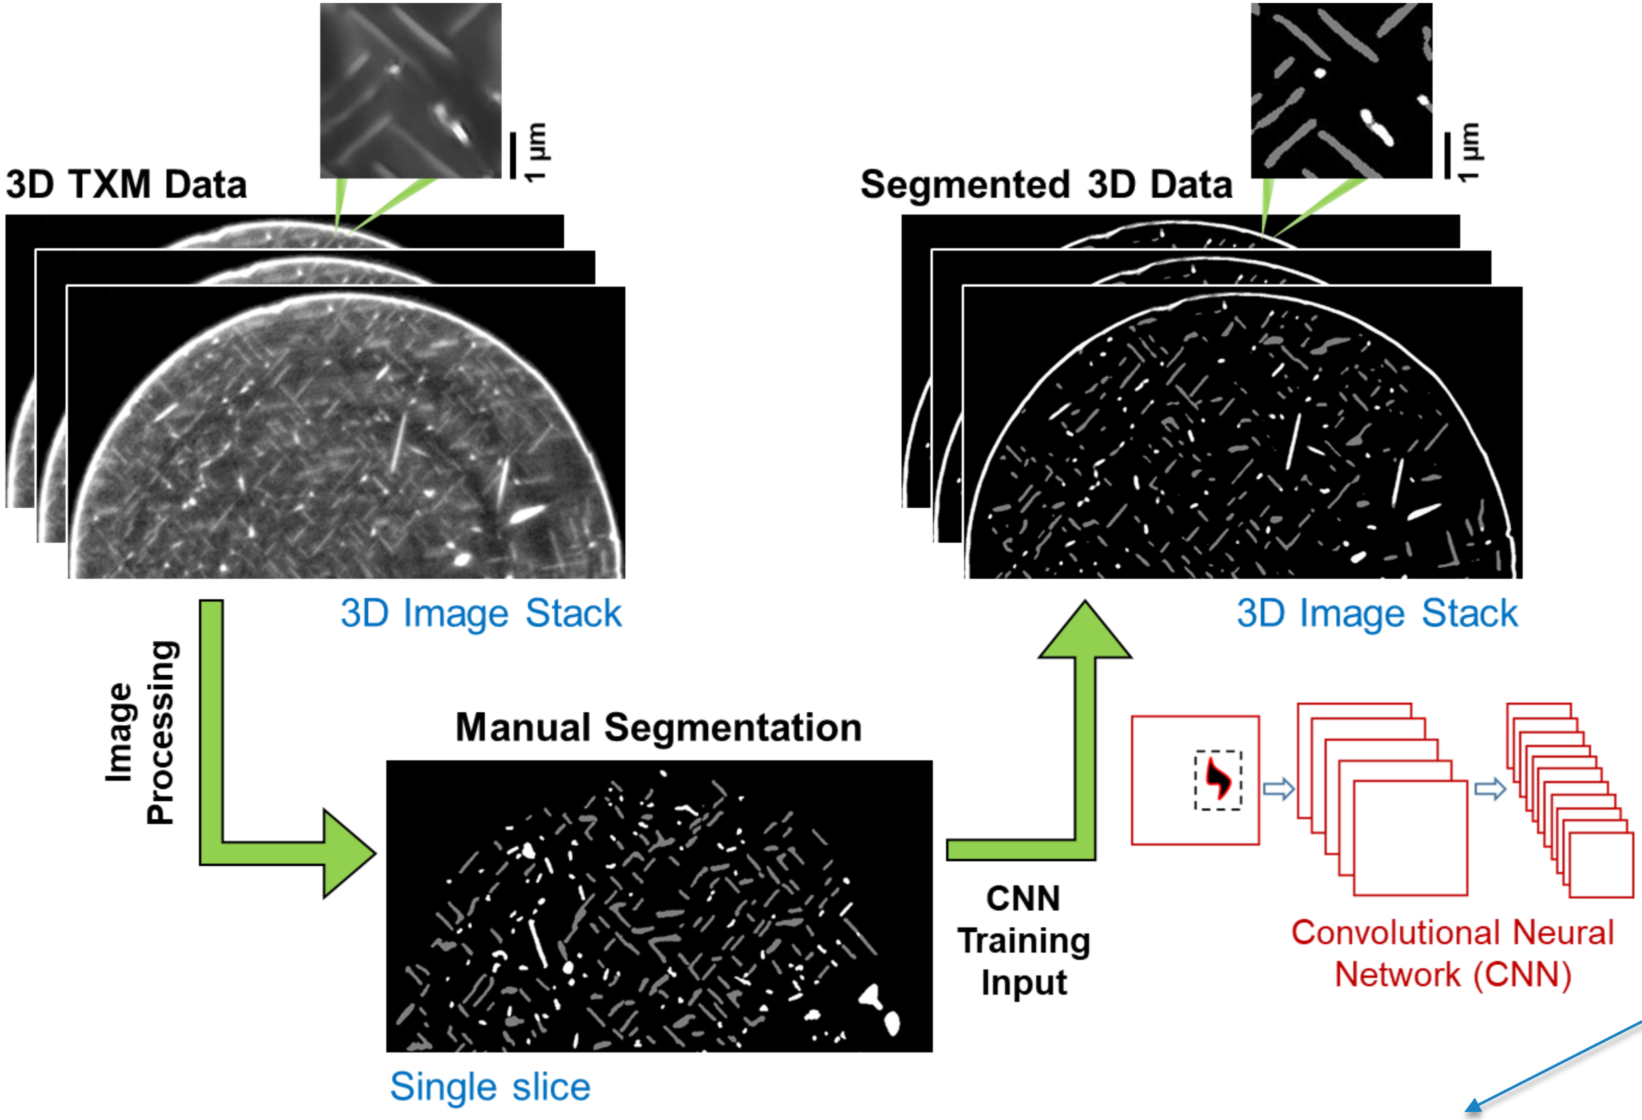

- Patch-based training on one slice to learn the unsupervised segmentation
- Deploy model from rapid segmentation on remaining slices

Xiaogang Yang (now DESY) talking Thursday @ 10:15

C.S. Kaira, X Yang, *et al.*, Mater. Charac., (2018)

Argonne NATIONAL LABORATORY

Phase Segmentation in Atom-Probe Tomography

Deep-learning-based edge detection

Transfer learning for segmentation:

- Transfer knowledge from common images (abundant labels) to segment data obtained from APT (scarce labels) into different phases
- Efficiently segments data phases & extracts interfacial properties without need for expensive interface labeling
- Segmentation done on 2D slices along 3 orthogonal directions; combined into 3D edge map of the interface delineating the 2 phases
- Uses holistically-nested edge detection (HED), based on CNNs with side outputs & deep supervision to significantly improve edge detection
- Demonstrated approach is qualitatively & quantitatively as accurate as proprietary solutions

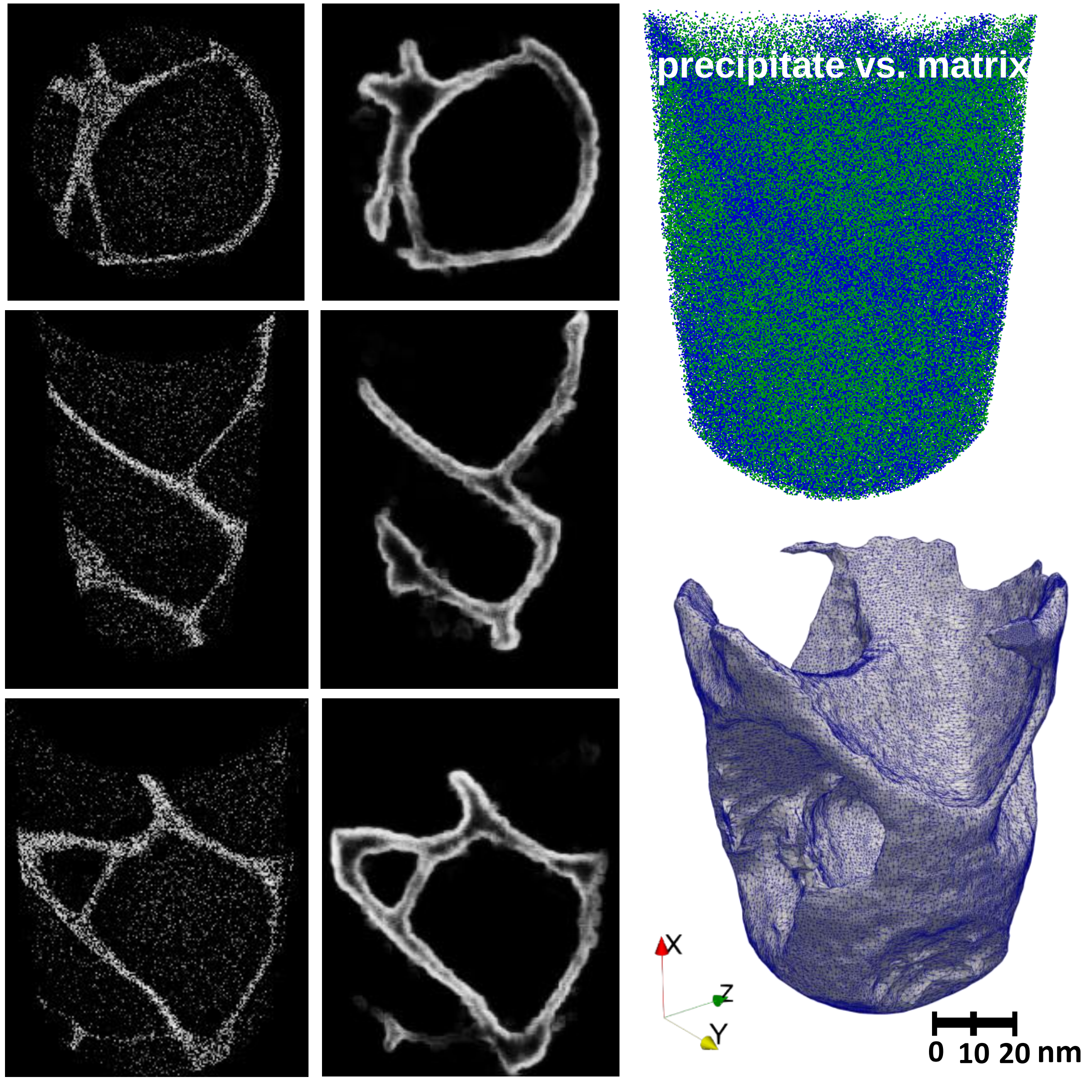

2D & 3D edge maps learned from APT data of a Cobalt-based superalloy

Madireddy, Chung, Loeffler, Sankaranarayanan, Seidman, Heinonen, Balaprakash. Segmentation in 3D Atom Probe Tomography using Deep Learning based Edge Detection. *Preprint, 2019*

UNSUPERVISED LEARNING

Multi-Element-Based Automatic Identification of Aerosol Particles

- *Multi-element* morphology-based learning method for particle identification
- Per-particle-based extraction of elemental concentration

In X-ray fluorescence imaging:

- Significantly reduces analysis time & *human bias* in comparing different particle samples
- Enables detailed, statistical characterization of observed atmospheric particles
- Provides critical constraints for earth system models representation of particle-related processes, including
  - iron chemistry in dust particles
  - morphological effects on aerosol optical properties
  - particle ice nucleation efficiency

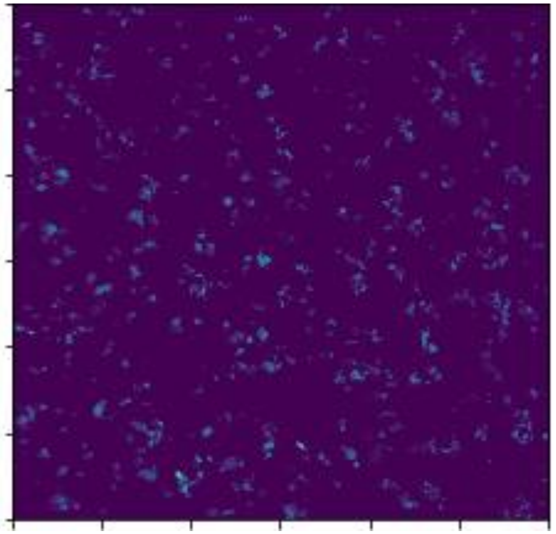

Original: total fluorescence yield from all existing elements

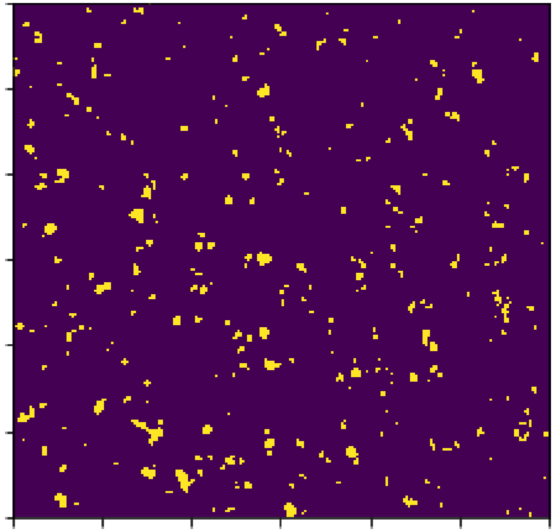

Morphology-based identification of particles

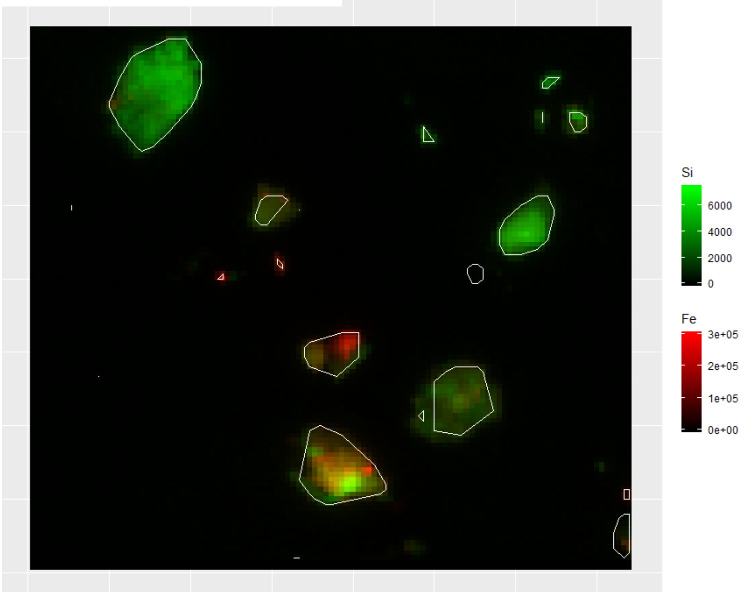

Figure 2: identified particle with different color representing different elemental concentration

Online Analysis of Streaming Experimental Data and Autonomous Steering

ADVANCED AUTOMATION

Online analysis

- Rethink the process of writing experimental data to disk --- stream directly to compute resources for online analysis
- Provides continuous feedback to beamlines while experimental data is captured
- Provides valuable information for timely decisions and real-time experimental steering

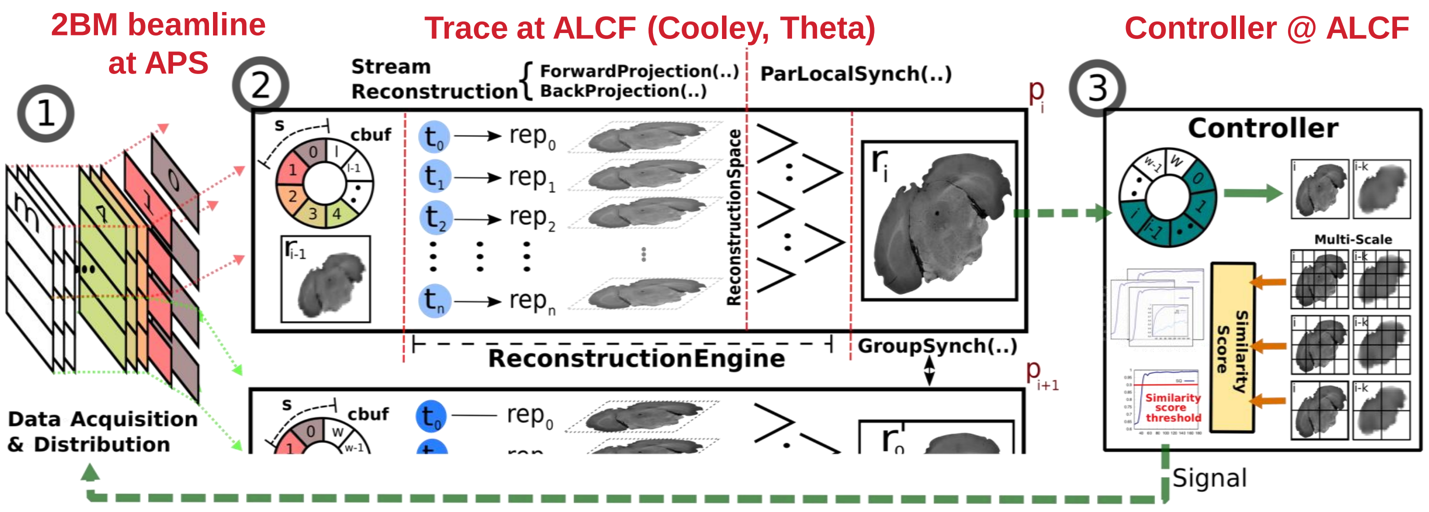

Link APS & ALCF for real-time analysis

- Trace scales to 10k+ cores and addresses high volume, high velocity data streams
- Evaluated performance at 2BM (microCT)
- Achieved >200 projections/s data consumption rate using 1200 cores of ALCF/Cooley
- Trace runtime system efficiency demonstrated on up to 32K cores at ALCF/Mira
- Feedback from multiple stages of analysis workflow can be observed

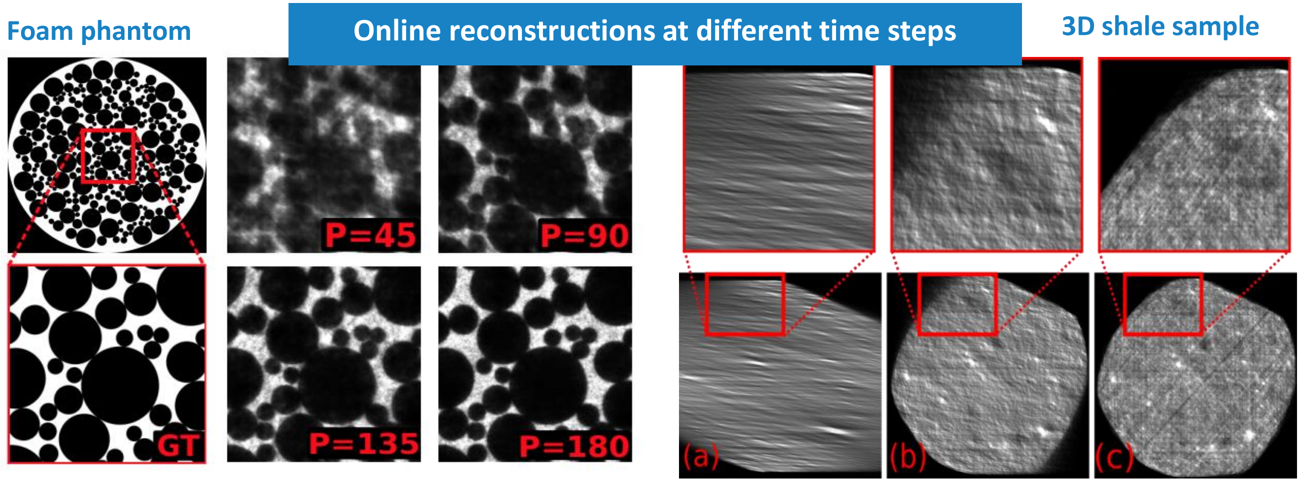

Bicer et al., Supercomputing 2019 Technology Challenge: Real-Time Analysis of Streaming Synchrotron Data

Bicer et al., *APS Science*, 2018

Closing the Loop: Chemical Separations by Design

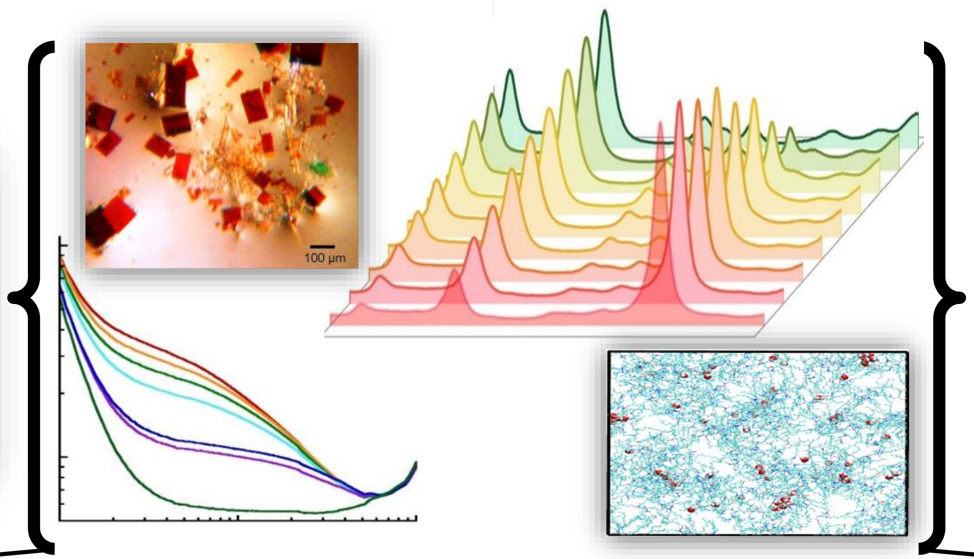

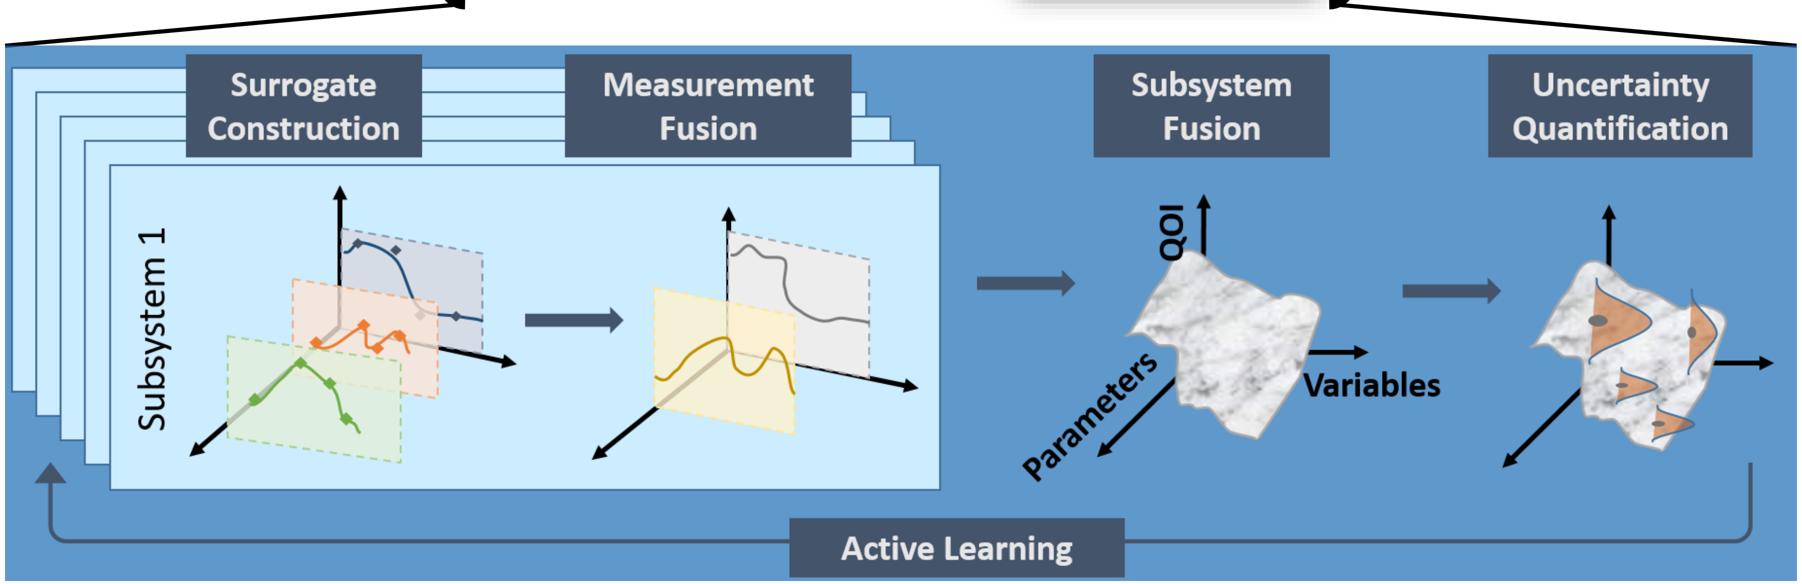

L. Soderholm et al

Learning for Learning: Automating Neural Architecture Search for Deep Learning

DeepHyper: A scalable AutoML package

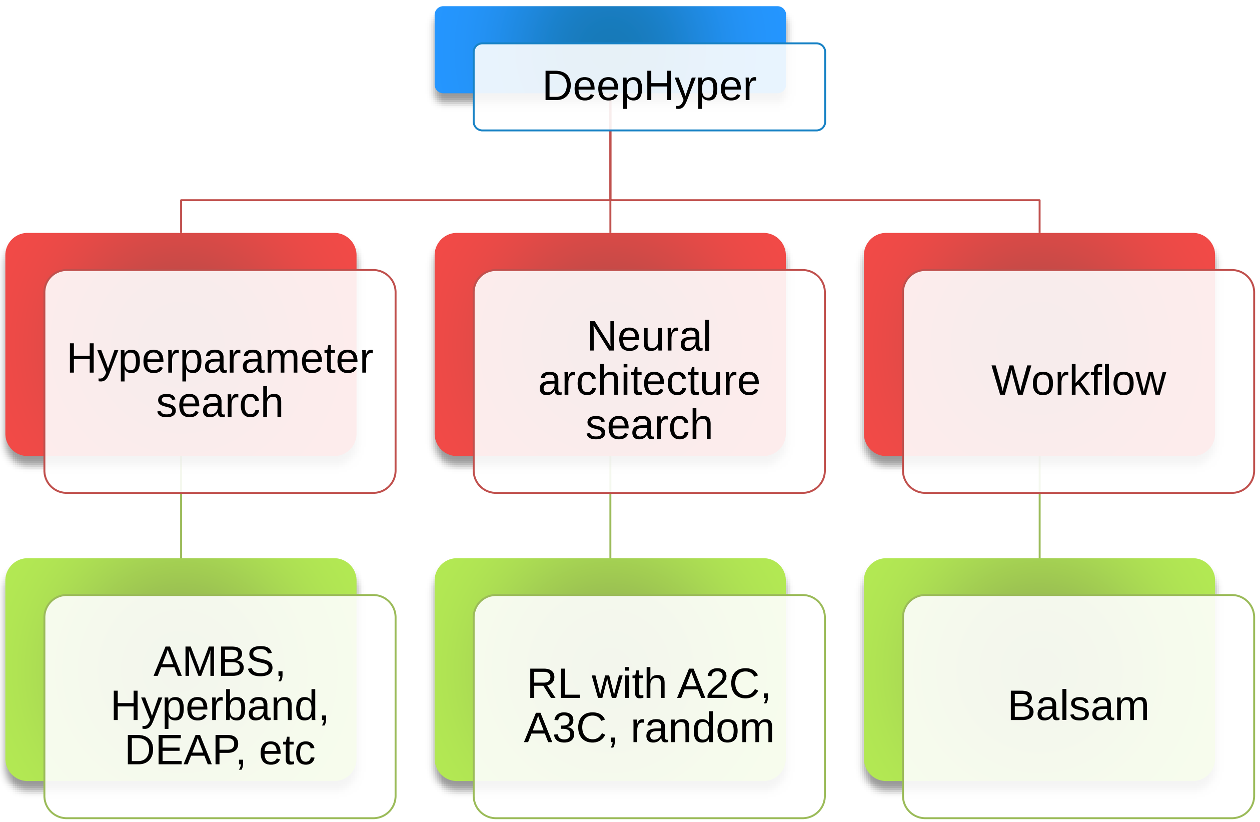

https://github.com/deephyper

On CANDLE benchmark problems

	Trainable Parameters	Training Time (s)	R^2 or *ACC*
Combo			
manually designed	13,772,001	705.26	0.926
A3C-best	**1,883,301**	**283.00**	**0.93**
Uno			
manually designed	19,274,001	164.94	0.649
A3C-best	**1,670,401**	**63.53**	**0.729**
NT3			
manually designed	96,777,878	247.63	0.986
A3C-best	**120,968**	**16.65**	**0.989**

Balaprakash, Egele, Salim, Wild, Vishwanath, Xia, Brettin, Stevens; Scalable reinforcement-learning-based neural architecture search for cancer deep learning research. SC ’19

Flutter-Shutter Computational Detector System for Imaging Dynamics

Detectors that can encode motion for high-speed applications

- Computational detector in which a shutter is used to control which photons reach the detector during each exposure
- Demonstrated improved time-resolutions for transmission x-ray computed tomography through simulations
- Superior performance to conventional post-acquisition image deblurring and applicable to fly-scan instruments or for high-speed radiography

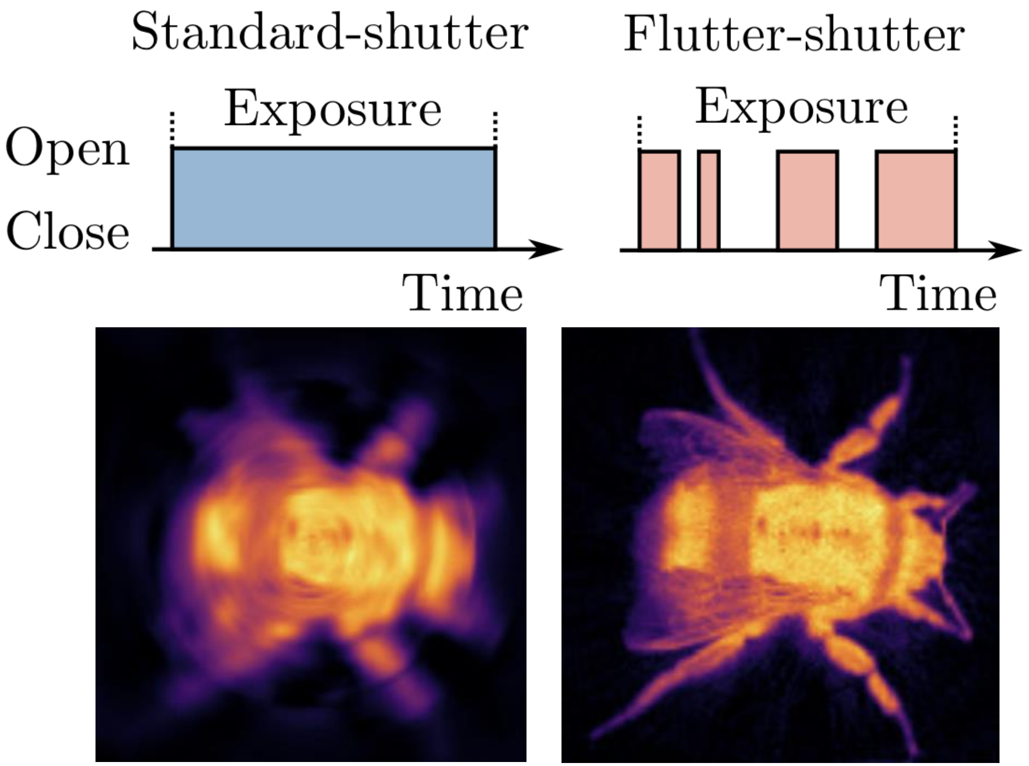

Tomographic reconstruction of a bumble bee without (left) and with (right) the known shutter code pattern incorporated in the computational process to compensate for rotational blur resulting from a 10X faster imaging rate

Ching, Aslan, Nikitin, Gursoy. "Time-coded apertures for x-ray imaging" *Optics Letters,* 2019

Additive Manufacturing Research at APS

Address the critical issues in metal additive manufacturing

In situ/operando synchrotron x-ray experiments ➡

- Structure defects
- Failure mechanisms
- High-fidelity models
- Reliability and Repeatability

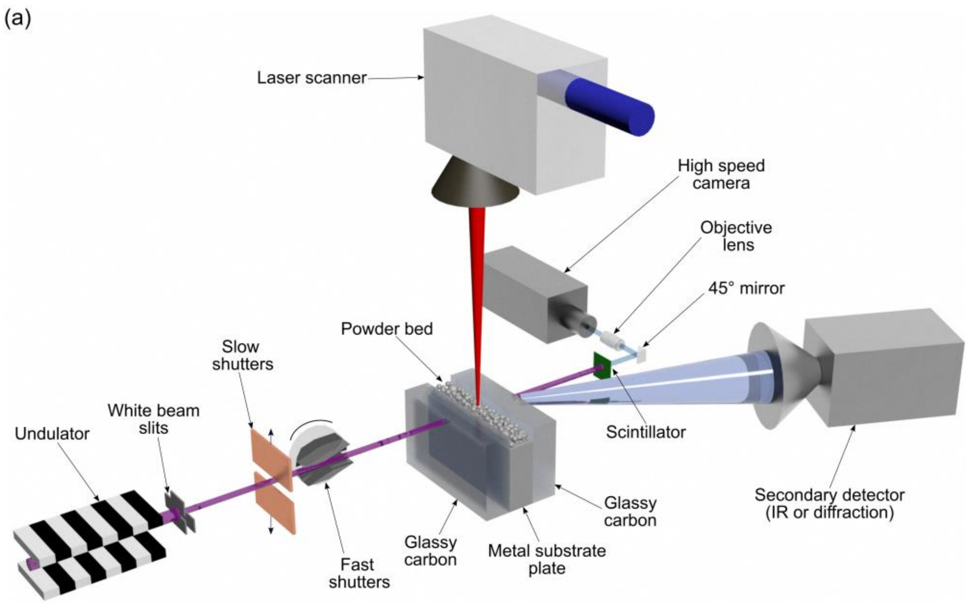

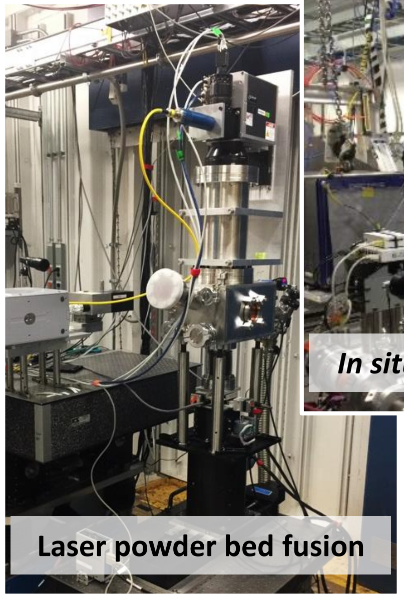
Laser powder bed fusion


In situ diffraction experiment

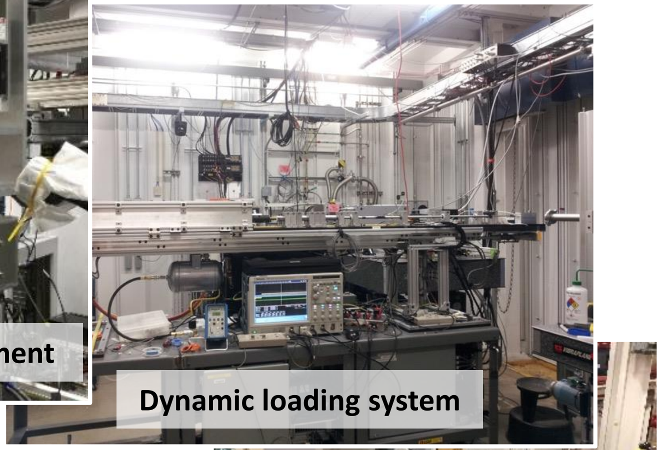
Dynamic loading system


Binder jetting (ExOne machine)


Piezo-nozzle direct energy deposition (with NU)

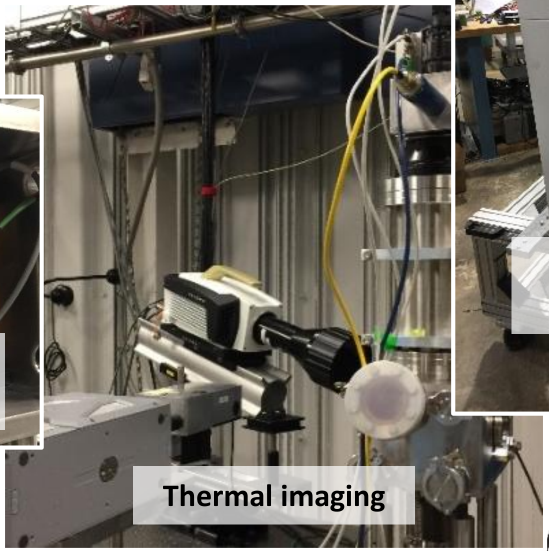
Thermal imaging

DESIGN OPTIMIZATION & CONTROL

Additive Manufacturing Research at APS

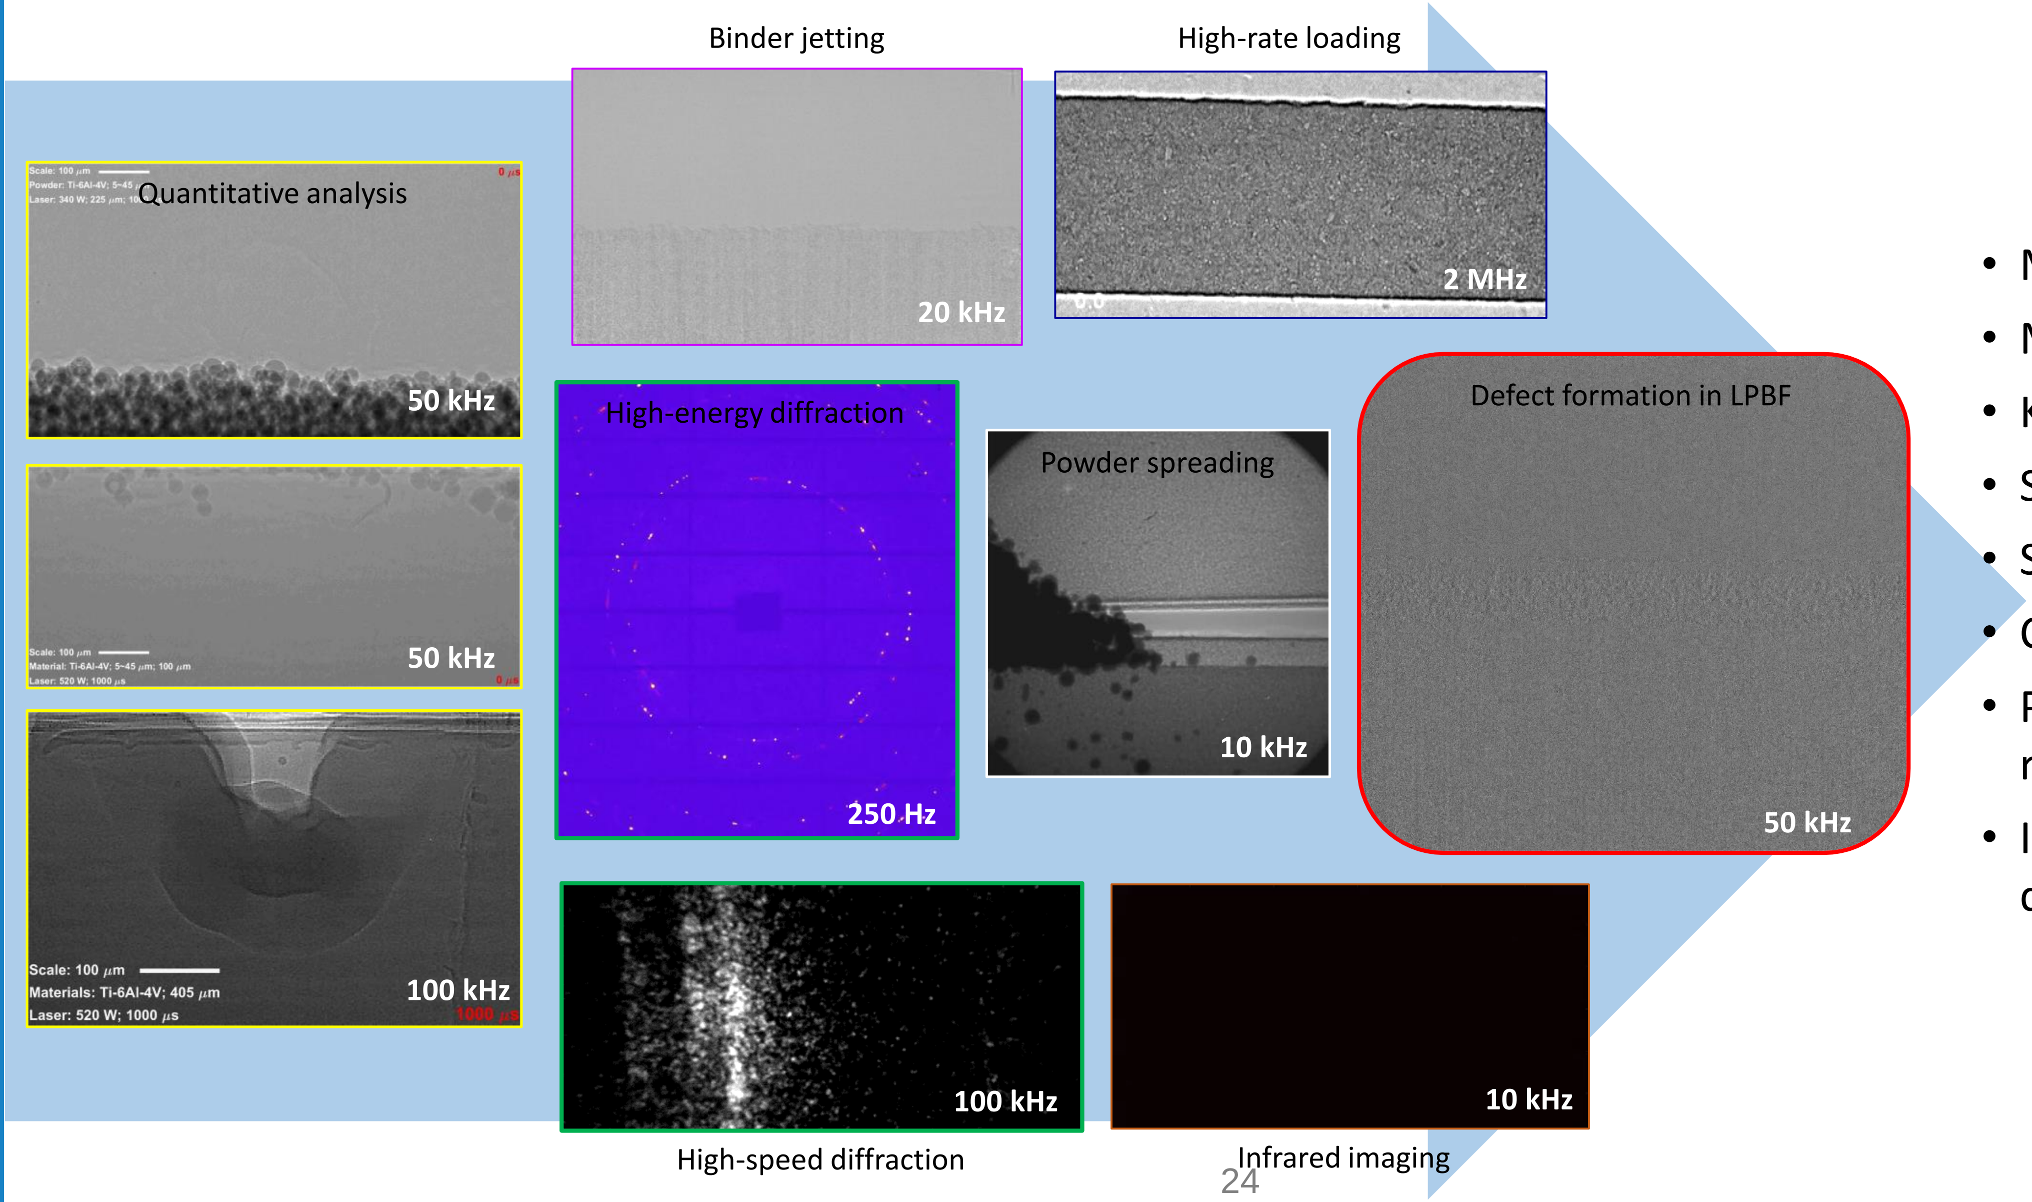

- Melt pool morphology
- Melt flow velocity
- Keyhole dynamics
- Spattering velocity
- Solidification rate
- Cooling rate
- Phase transformation rate
- Internal temperature distribution

The Future's So Bright…

Science enabled by intelligence in computing, data, & X-rays

Abundant opportunities for AI

- Process and analyze huge, complex, multimodal data volume & find rare events
- Accelerate discovery through optimal design of experiment (sample, acquisition, …)
- Quantify correlations & uncertainties in high dimensions
- Real-time control & experimental steering

Requires multidisciplinary teams & advances on several simultaneous fronts

Argonne & APS Hiring ML+Imaging Postdocs

http://bit.ly/2X9hU6P

http://bit.ly/338oH2f

Thanks!

wild@anl.gov

Argonne NATIONAL LABORATORY