4TH GLOBAL RESEARCH PLATFORM WORKSHOP

Advanced Photon Source – Ongoing Upgrade, Data Deluge, Challenges and Opportunities

RAJKUMAR KETTIMUTHU

Computer Scientist and Group Leader
Data Science and Learning Division
Argonne National Laboratory

Nicholas Schwarz

Nino Miceli

Tekin Bicer


Argonne National Laboratory, IL


SLAC National Accelerator Laboratory, CA


SLAC National Accelerator Laboratory, CA

BASIC ENERGY SCIENCE

DOE LIGHT SOURCES


Brookhaven National Laboratory, NY


Lawrence Berkeley National Laboratory, CA

What is an X-ray Synchrotron?

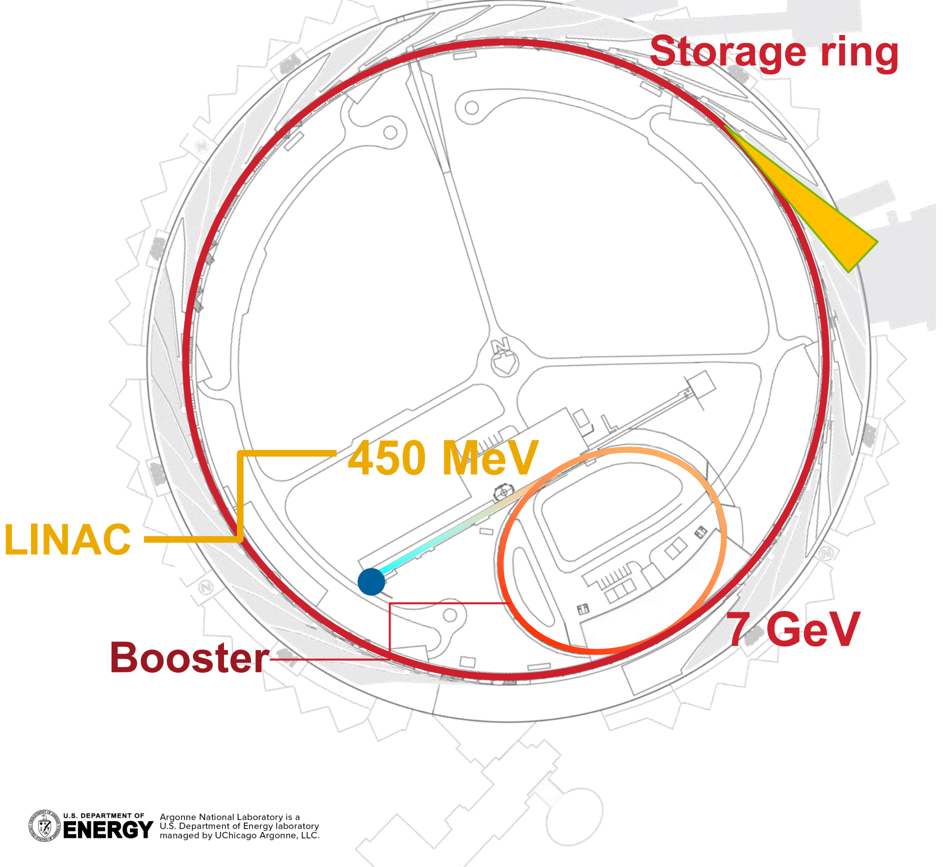

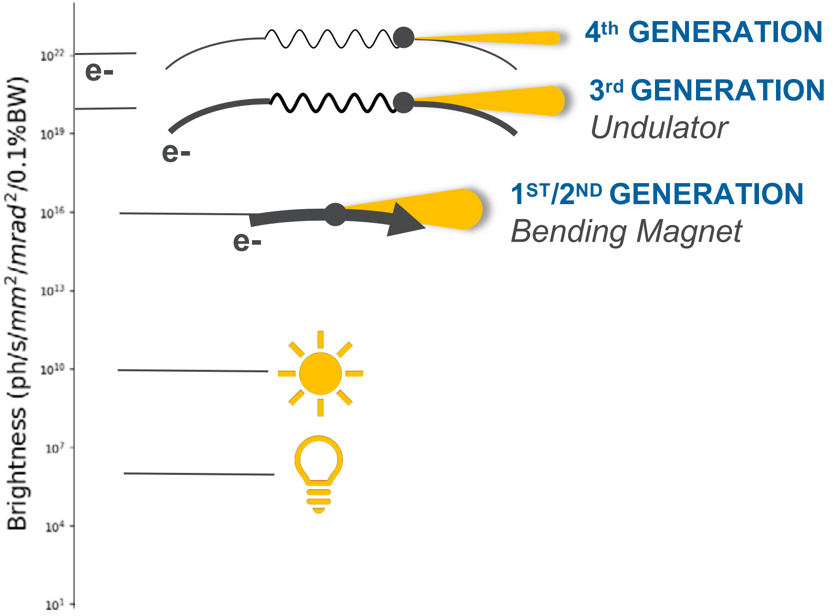

Advanced Photon Source

68
X-ray beamlines

6,000
Experiments
per year

2,000
Publications
per year

5,500
Unique users
in a typical year

Countless
Societal impacts

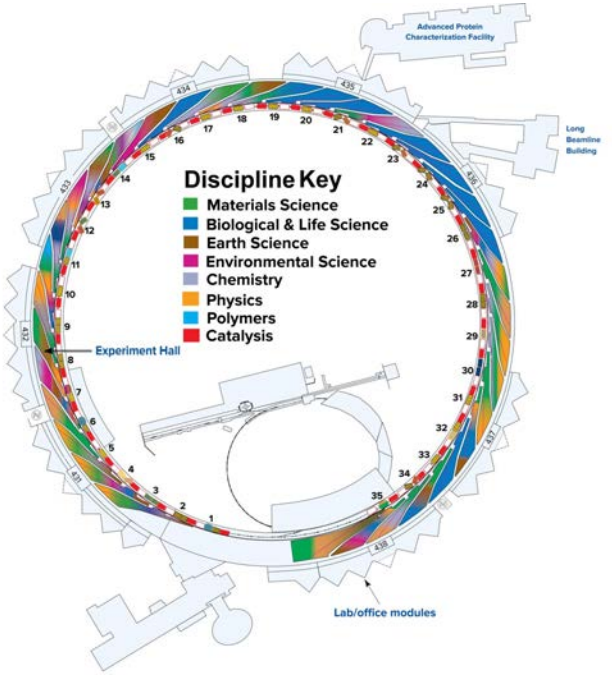

Building Longer-lasting Batteries

at the Advanced Photon Source

REAL-TIME BATTERY EXAMINATION DURING CHARGE/DISCHARGE
APS X-rays track decays and defects as they form.

BATTERY COMPONENTS
Assembled and tested in the APS electrochemistry lab.

X-ray image from within a battery

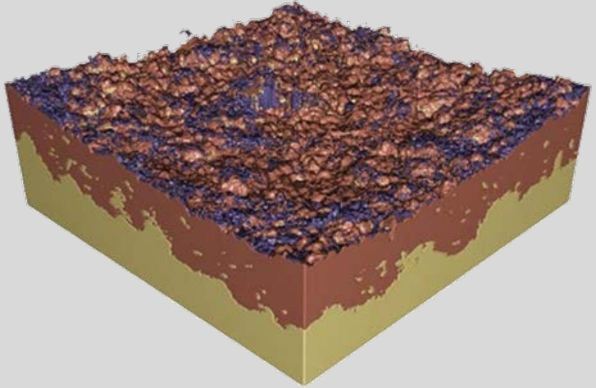

Better Catalysts for Cleaner Air

at the Advanced Photon Source

Microscopic catalysts can be examined by X-rays to improve their efficiency.

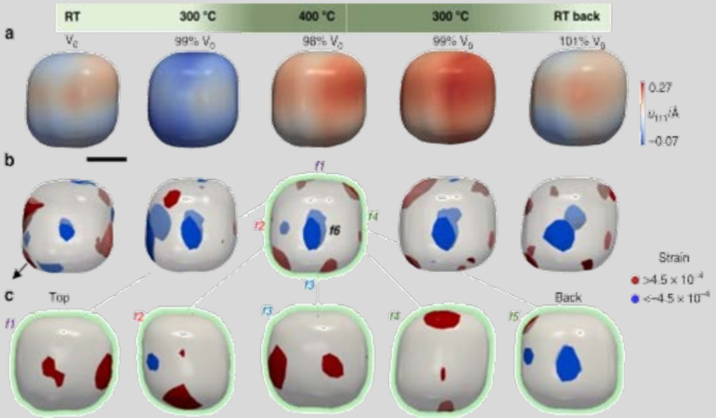

Better catalysts could lead to cleaner hydrogen production for the next generation of vehicles.

Scene coordinate system

3D

Microelectronics

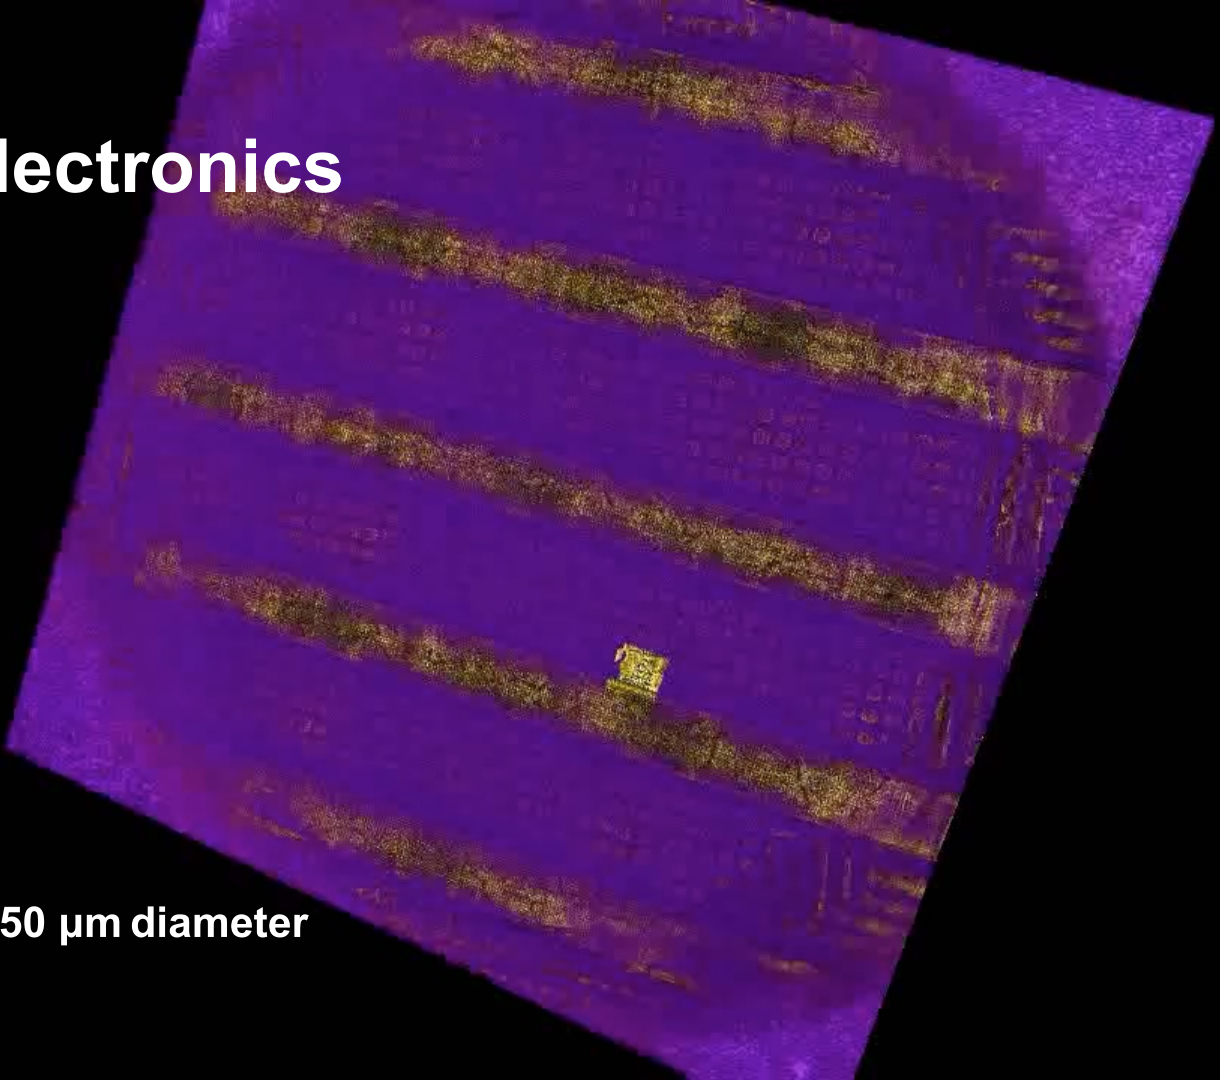

Scan FOV: 50 µm diameter

PROTEIN CRYSTALLOGRAPHY

APS: A leader in structural biology

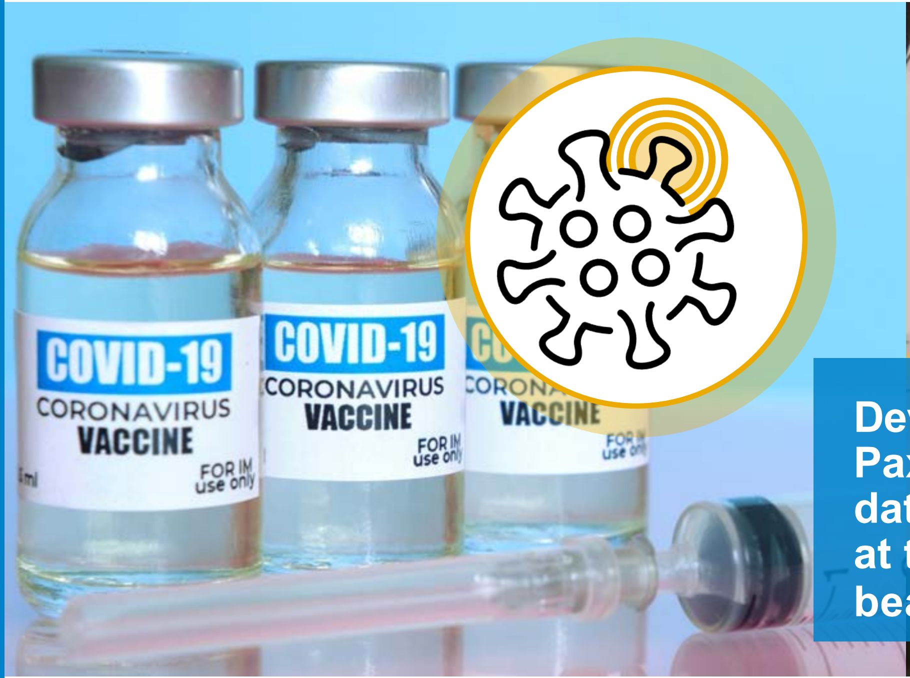

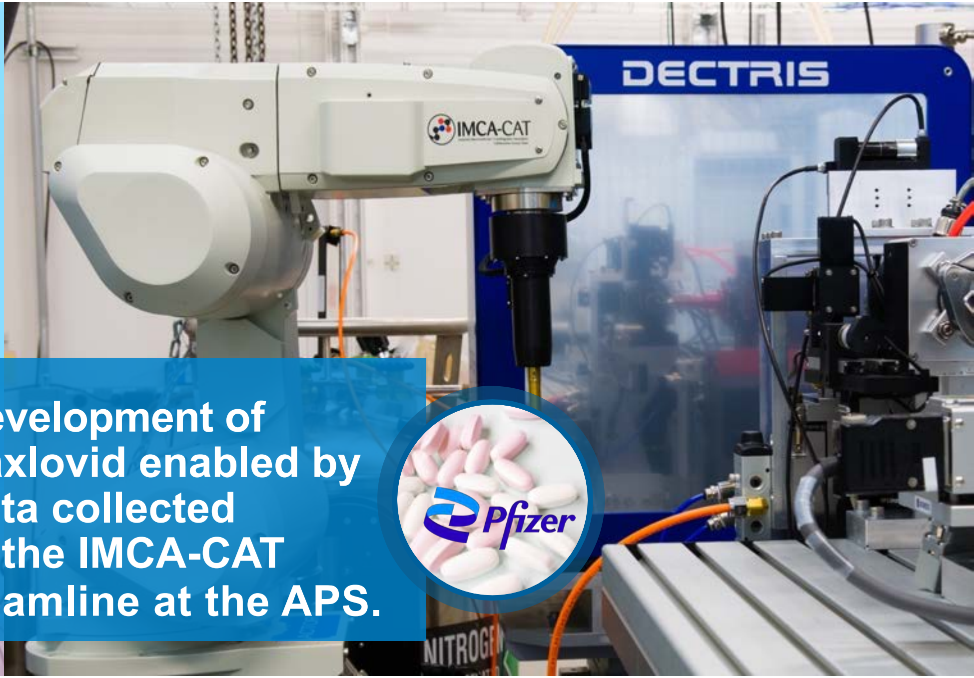

Development of Paxlovid enabled by data collected at the IMCA-CAT beamline at the APS.

Advanced Photon Source Upgrade

APS-U: The Ultimate 3D Microscope

A next-generation synchrotron light source for science and industry

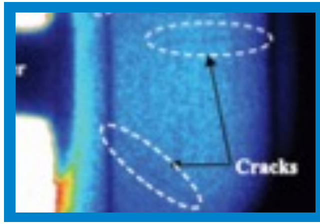

HIGH ENERGY

Penetrating bulk materials and operating systems

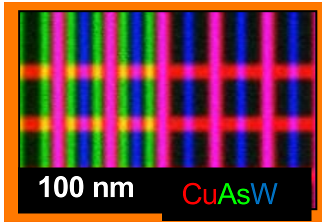

BRIGHTNESS

Providing time-resolved, macroscopic fields of view with nm-scale resolution

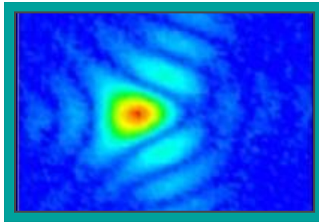

COHERENCE

Enabling highest spatial resolution even in non-periodic materials

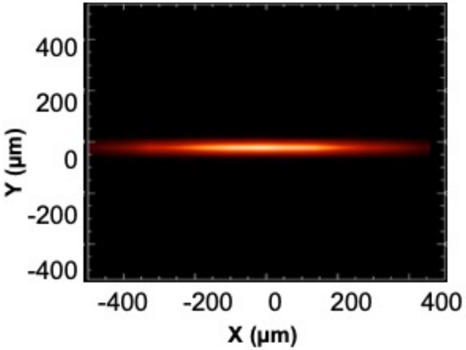

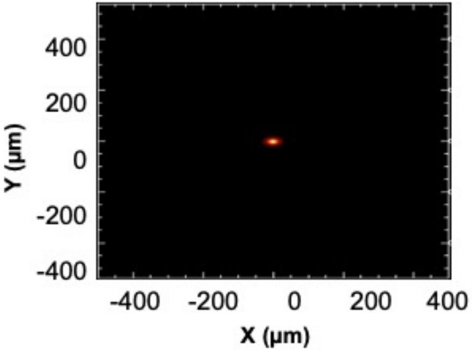

4th Generation Synchrotron Projects

- 22 [illegible] 4[illegible] [illegible]
- A[illegible] 2024

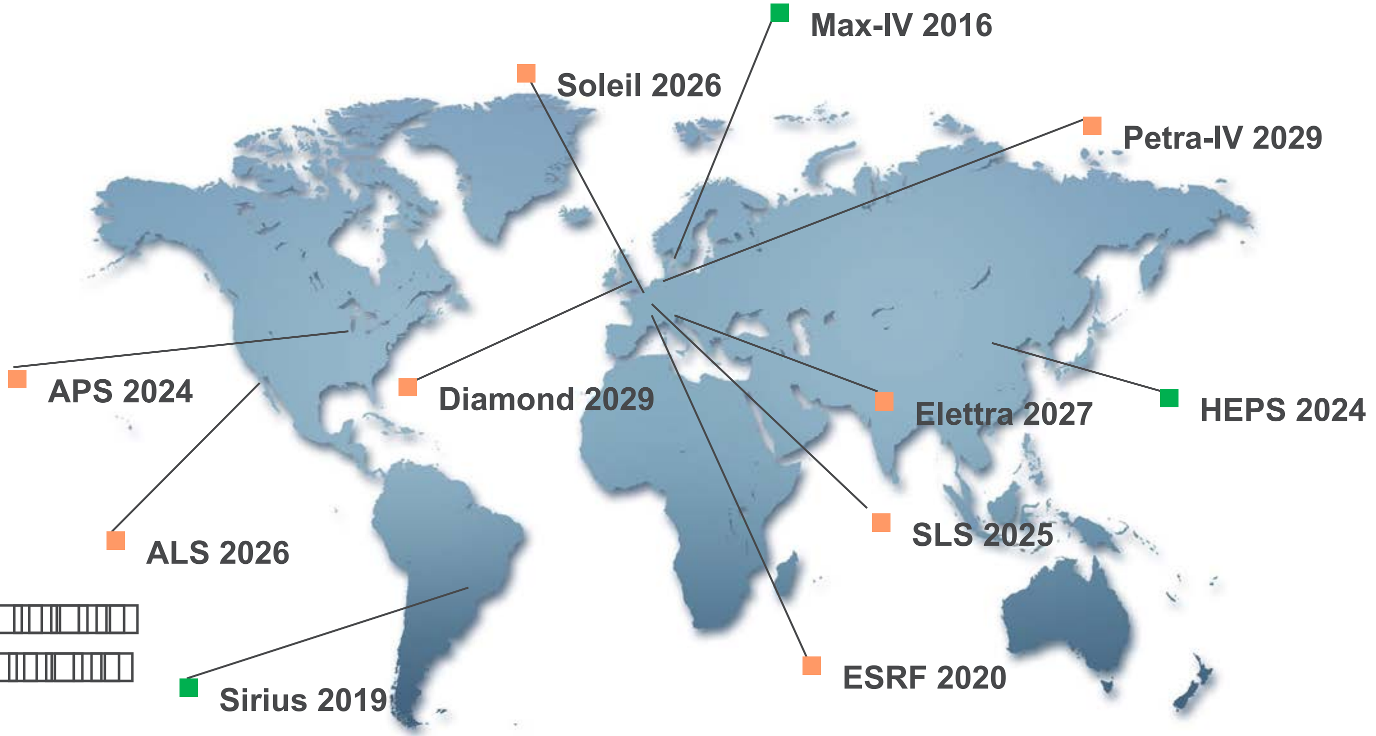

- B[illegible]
- G[illegible]

APS Upgrade: The Ultimate 3D Microscope

A next-generation synchrotron light source for science and industry

$815 M project to update and renew the facility

Re-uses $1.5 B in existing infrastructure

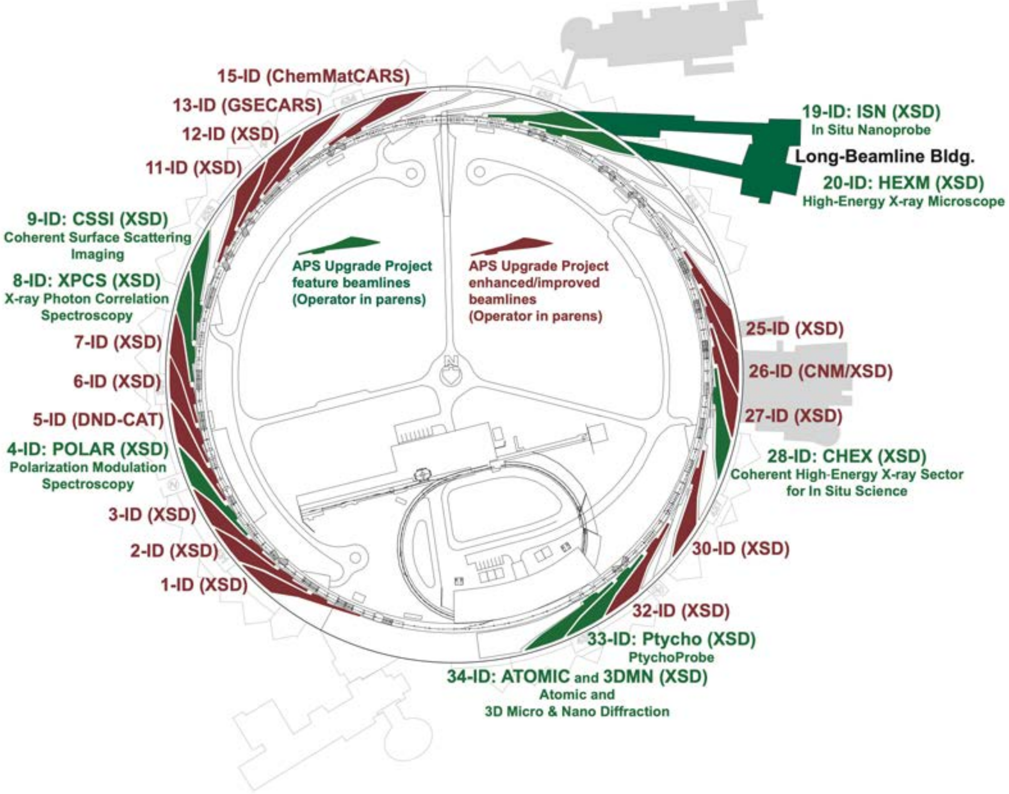

- Completely new storage ring, 42 pm emittance @ 6 GeV, 200 mA
- New and updated insertion devices
- Combined result in brightness increases of up to 500x
- 9 new feature beamlines (green)
- 15 beamline enhancements (red)

x500

$815M

$1.5B

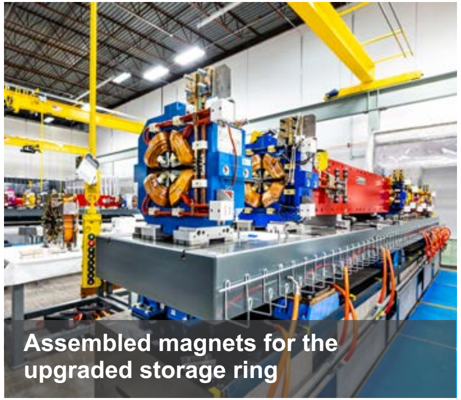

Assembled magnets for the upgraded storage ring

First new beamline instrument up and running

New front end systems to deliver X-ray beams to experiments

Long Beamline Building, which will house two feature beamlines

APS-U – High Brightness Storage Ring Lattice

APS Today

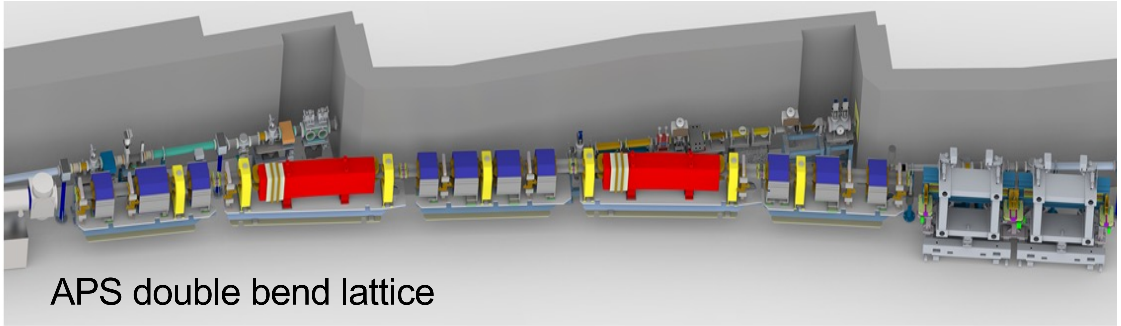

APS Upgrade

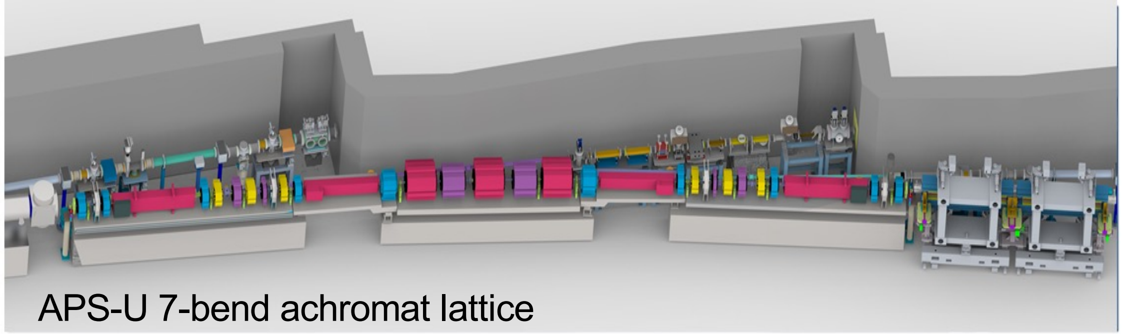

APS Today

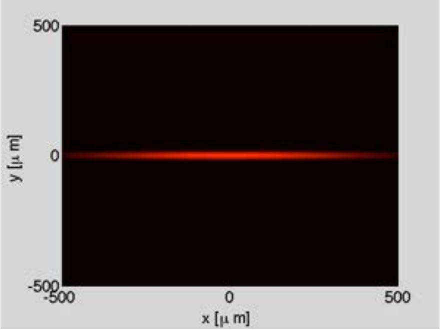

ε_o = 3100 pm.rad

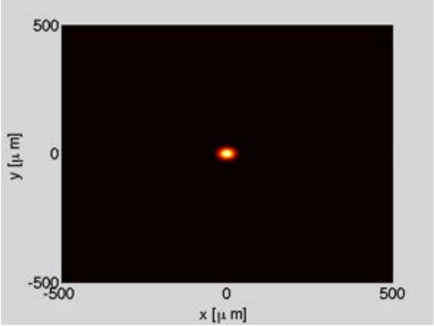

ε_o = 42 pm.rad

Data Deluge, Challenges and Opportunities

A data deluge

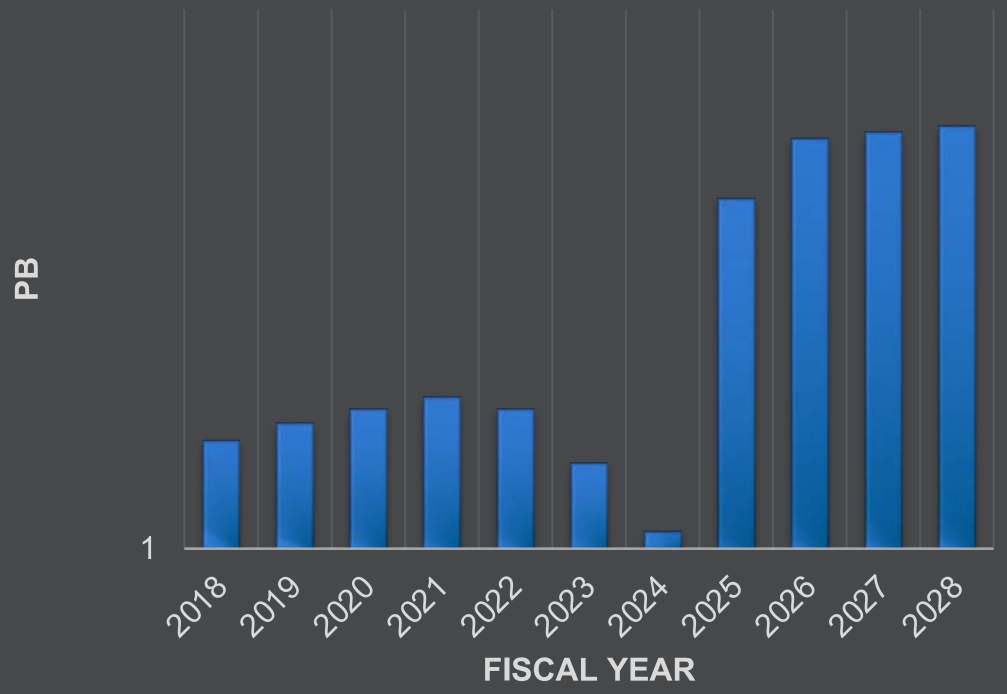

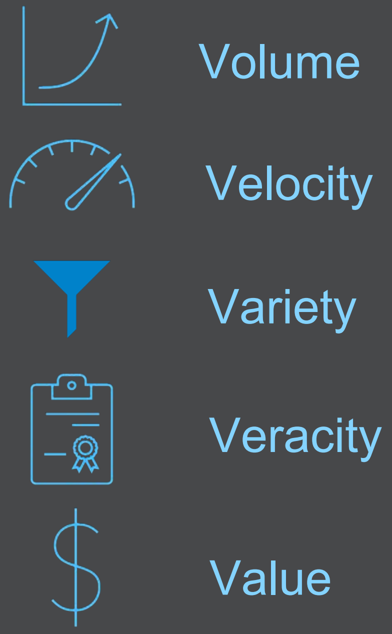

Advanced Computing and Data Management is Crucial to Address Drastic Increases in Data

Increased source brightness (orders-of-magnitude brighter)
- Due to facility upgrades and accelerator improvements

New and more complex experiments
- Multi-modal experiments that combine data from multiple samples, techniques, and facilities
- *In situ* and *in operando* experiments require real-time feedback and autonomous control

Detector advances (orders-of-magnitude faster)
- Increased dynamic range
- Faster readout rates
- Larger pixel arrays

→

- Analyze and reconstruct massive multi-modal data volumes
- Identify and classify features and patterns
- Merge simulation and experiment data to drive experiments and new results
- Execute experiments dynamically using real-time reduction and AI/ML

Computing Resources

Multi-tiered approach spanning local and remote resources

Local compute resources

- Perform pre-analysis/data reduction (including compression and running ML models) to a form that allows quality control and experiment steering
- This may include, for example, a GPU workstation at a beamline, or the APS computing cluster

High-end compute resources

- Large data processing tasks, ML training, post-processing, and data refinement
- The APS has facility allocations at the Argonne Leadership Computing Facility
- The Argonne Leadership Computing Facility now provides a resource allocation queue/policy that suites APS job size and frequency profiles

Advanced Photon Source (APS)

Orthros – General purpose distributed-memory compute cluster
~27 TFLOP/s CPU cores
Sayre – Single node GPU system for Bragg CDI reconstructions
~111 TFLOP/s
5 x Ti 2080 | 2 x P100 | 1 x Ti 1080 | 1 x Quadro RTX 8000 GPUs
Axinite – Single node GPU system for CSSI and XPCS data processing
~155 TFLOP/s
4 x A6000 GPUs
Monas – 4 node GPU cluster for ptychography reconstructions
~430 TFLOP/s
8 x Ti 2080 GPUs per node

Argonne Laboratory Computing Resource Center (LCRC)

BeBop
~1,750 TFLOP/s
43,344 Intel Broadwell cores | 65,536 Intel Phi cores
Swing
~925 TFLOP/s
48 NVIDIA A100s | 768 AMD EPYC cores
Blues
~198 TFLOP/s
6,000 compute cores

Argonne Leadership Computing Facility (ALCF)

Theta & Theta GPU
Theta: 281,088 Intel Phi cores (~11.3 PFLOP/s)
Theta GPU: 192 NVIDIA A100s

Polaris
~44 PFLOP/s (~4 PFLOPS/s for exploring use by experimental facilities)

Aurora
Anticipated 2023
Intel CPUs / GPUs
> 1 EXAFLOP/s

Argonne Leadership Computing Facility (ALCF)

Coupling APS instruments with ALCF supercomputers to accelerate scientific discovery

Polaris Supercomputer
44 petaflop/s peak performance

Aurora Supercomputer (online in 2023)
2 exaflop/s peak performance

- APS jobs can launch on-demand within seconds, preempting other running jobs
- Deploying a >1 terabit/s network between the APS and the ALCF

Network Architecture

Updates underway to support APS-U Era data and computing requirements

- 3-tier network infrastructure: facility, sector, hutch
- Supervisory Control and Data Acquisition (SCADA) architecture to better support controls, data, and regular network traffic
- Installed a new fiber plant for all APS beamlines; 768 pairs of new single mode fiber from the APS computer room to beamline networks
- Installed new core network switches capable of 100 Gbps links
- Procuring new sector and hutch switches for APS sectors capable of 100 Gbps links
- Recently upgraded the APS <-> ALCF network connection to 200 Gbps; upgrade to a terabit/s network in the future
- Adding wireless access points inside hutches and installing CAT 6A 10 Gbps copper cable at beamlines

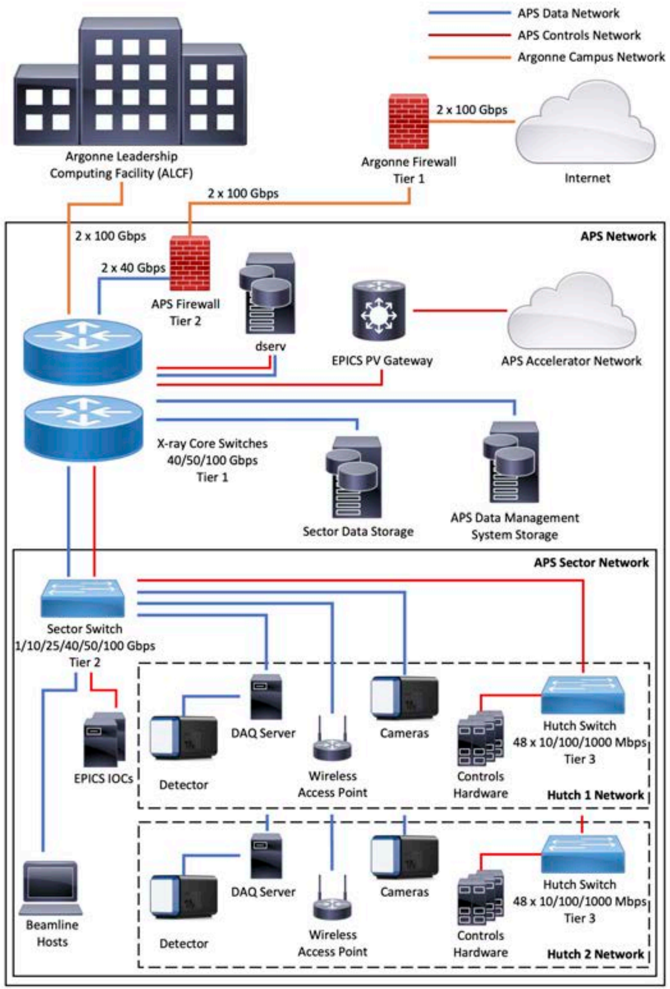

Globus is the *Glue* Connecting the APS to Advanced Computing Resources

The APS is leveraging Globus as a computational and data fabric to enable advanced computing and data management

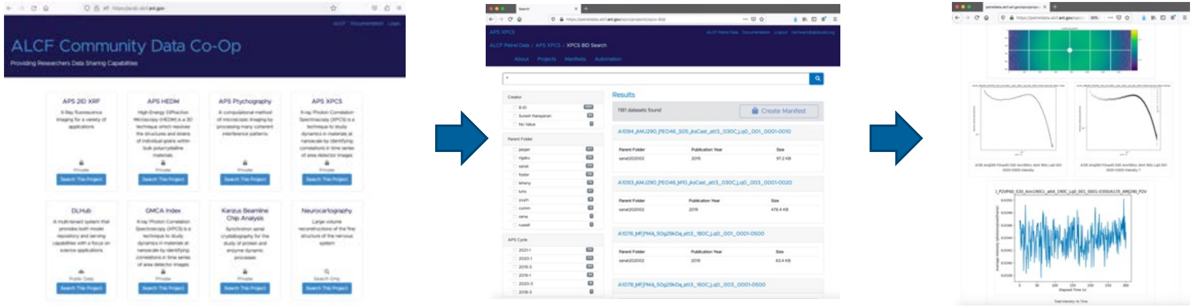

- Globus Automation and Compute Services are enabling technique and instrument specific data processing workflows
- Technique / instrument tailored web portals enable viewing and searching data, and re-run data processing

The Polaris supercomputer and Globus enable on-demand data analysis at the APS

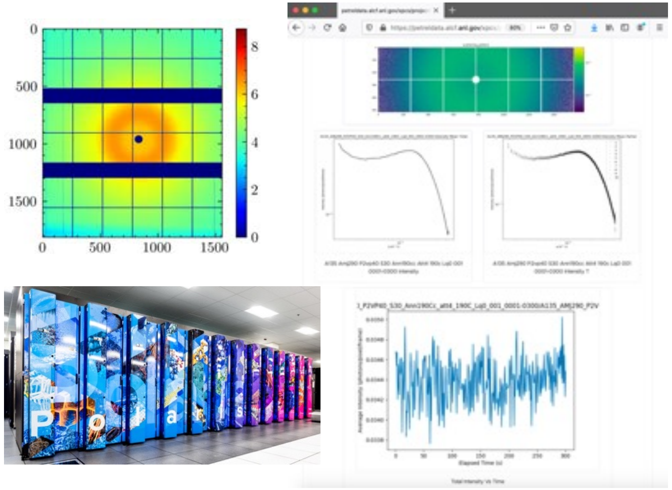

Speckle data from the APS 8-ID-I beamline (top left) is automaticity transferred to the Polaris supercomputer (bottom left) where it is processed on-demand and displayed in a Globus web portal (right).

Team: Miaoqi Chu, Hannah Parraga, Sinisa Veseli, John Hammonds, Qingteng Zhang, Eric Dufresne, Suresh Narayanan, Ryan Chard, Nickolaus Saint, Rafael Vescovi, Ben Blaiszik, William Allcock

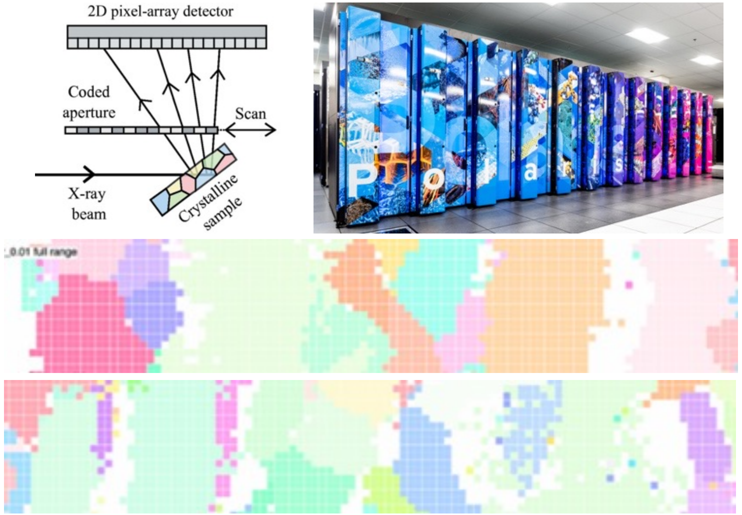

Data from the new coded aperture at APS 34-ID-E (top left) is automaticity transferred to the Polaris supercomputer (top right) where it is reconstructed on-demand (bottom)

Team: Michael Prince, Ryan Chard, Bill Allcock, Gürsoy, Doğa, Barbara Frosik, Hannah Parraga, Dina Sheyfer, Jonathan Tishler

AI/ML Enabled Science at the APS

Data Reduction

PtychoNN: Machine learning ptychography reconstruction

- 100s of times faster and requires up to 5 times less data than conventional iterative approaches

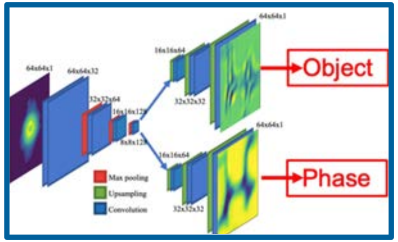

BraggNN: Machine learning method for determining Bragg peak locations from far-field high-energy diffraction microscopy data

- >200 times faster than conventional pseudo-Voigt profiling approach

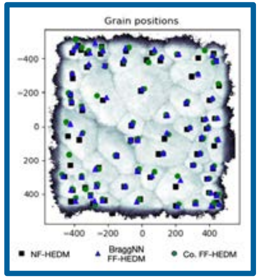

Cherukara, M., Zhou, T., Nashed, Y., Enfedaque, P., Hexemer, A., Harder, R.J., Holt, M. V., "AI-enabled high-resolution scanning coherent diffraction imaging," Applied Physics Letters 117, 044103 (2020).

Liu, Z., Sharma, H., Park, J. S., Kenesei, P., Miceli, A., Almer, J., Kettimuthu, R., Foster, I., "BraggNN: fast X-ray Bragg peak analysis using deep learning," IUCrJ 9, 104-113 (2022).

Experiment Steering

Smart Data Acquisition: Machine learning optimizes acquisition scanning path in real-time

- Motor movement is reduced by 80%

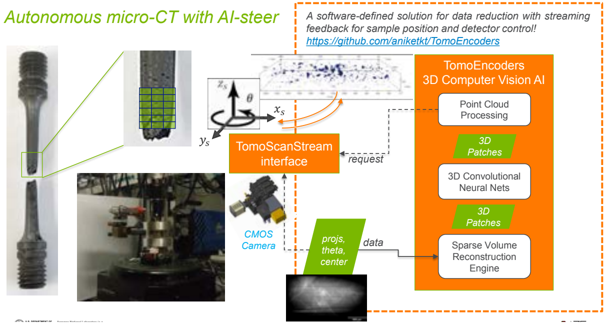

Contacts: Saugat Kandel, Tao Zhou, CD Phatak, et al.

Zhang, Y., Godaliyadda, G. M., Ferrier, N., Gulsoy, E. B., Bouman, C. A., & Phatak, C., "SLADS-Net: supervised learning approach for dynamic sampling using deep neural networks," Electronic Imaging, 2018(15), 131-1.

Knowledge Extraction

TomoGAN: Generative adversarial network improves the quality of tomographic reconstructions

- Uses up to 1/16th less dose or projections

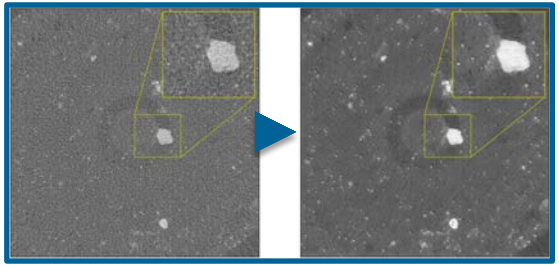

Generating interatomic potentials: Unsupervised machine learning generated interatomic potentials for a refractory oxide

- Diffraction measurements initialize an active-learner that iteratively improves an ML model

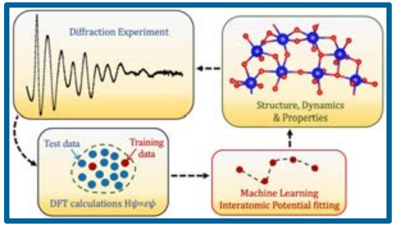

Liu, Z., Bicer, T., Kettimuthu, R., Gursoy, D., De Carlo, F. and Foster, I., "TomoGAN: low-dose synchrotron x-ray tomography with generative adversarial networks: discussion," JOSA A, 37(3), pp.422-434 (2020).

Sivaraman, G., Gallington, L., Krishnamoorthy, A. N., Stan, M., Csányi, G., Vázquez-Mayagoitia, Á., Benmore, C. J., "Experimentally driven automated machine-learned interatomic potential for a refractory oxide," Physical Review Letters, 126(15), 156002 (2021).

SUMMARY

- Big data problem very real at APS, existing and next generation light sources
- Hundreds of Pb per year will be generated at APS after the upgrade
- Complexity, multi-modality, operando science is becoming the norm
- Exploit full computing continuum, including learning from real time data
- Delivering rapid analysis at scale is critical and will provide competitive advantage
- Many other opportunities, for example:
  - Accelerator control and fault detection with edge devices
  - Coupling simulations, advances in surrogate models with experimental science


Argonne
NATIONAL LABORATORY