Recent Machine Learning Applications at the ALS

Accelerator Technology & Applied Physics Division

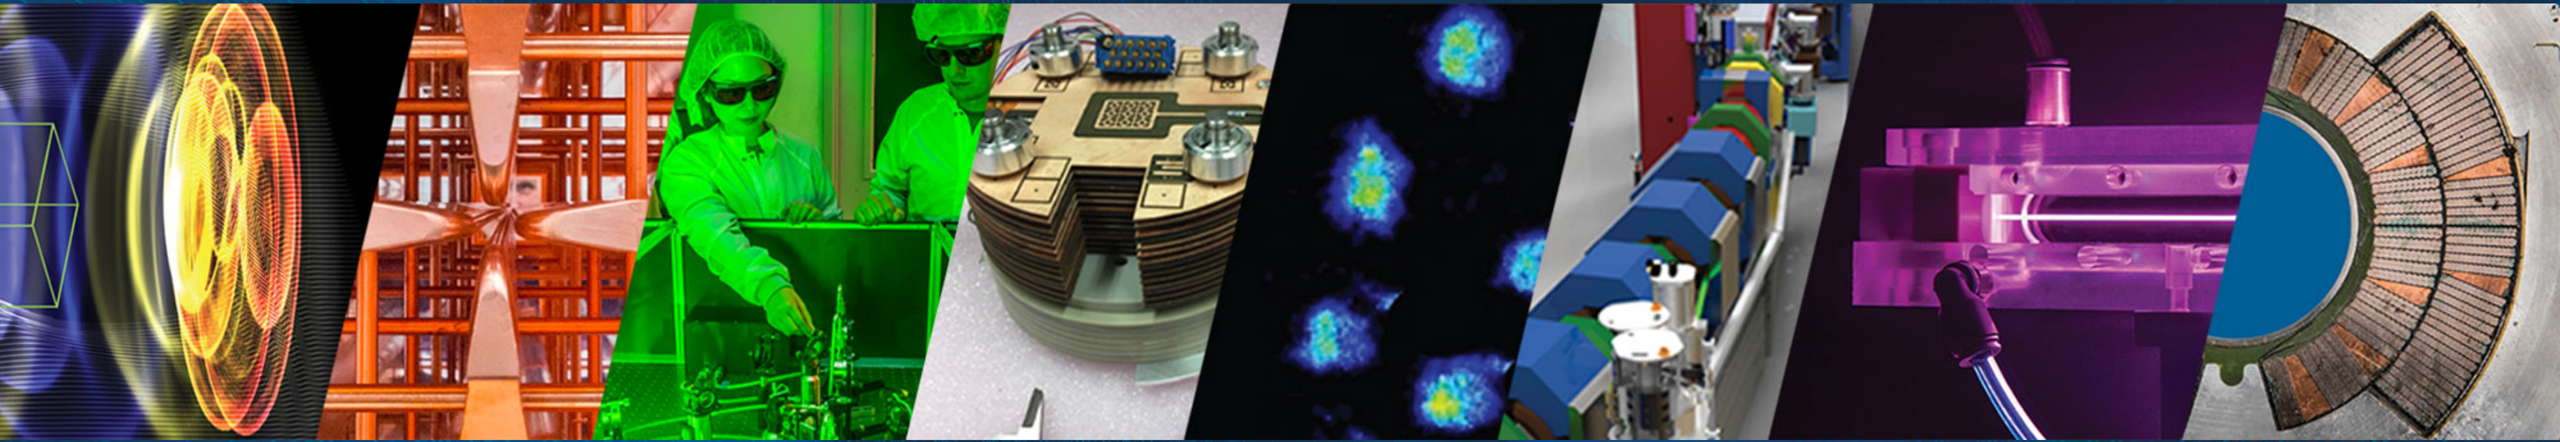

Simon C. Leemann • Aug 8, 2022

BERKELEY LAB

ACCELERATOR TECHNOLOGY & APPLIED PHYSICS DIVISION ATAP

U.S. DEPARTMENT OF ENERGY | Office of Science

Intro: Machine Learning (ML) at the ALS

- ALS ML efforts have so far been enabled by a 3-year grant funded jointly by **DOE BES** ADRP & ASCR

- Initial **ALS ML R&D effort:** use ML as powerful “new” tool to solve “old” accelerator problems:
  - **Accelerator operations:** automated tuning, replace feedback approaches, virtual diagnostics
  - **Accelerator development:** improve physics understanding, augment/extend lattice optimization, accelerate multi-objective optimization (e.g. MOGA)

- Two ALS examples today:
  - *Project #1:* ML stabilization demonstrated on operational accelerator published in **PRL**
  - *Project #2:* ML-enhanced optimization approach recently submitted to **PRAB**

What is the Advanced Light Source?

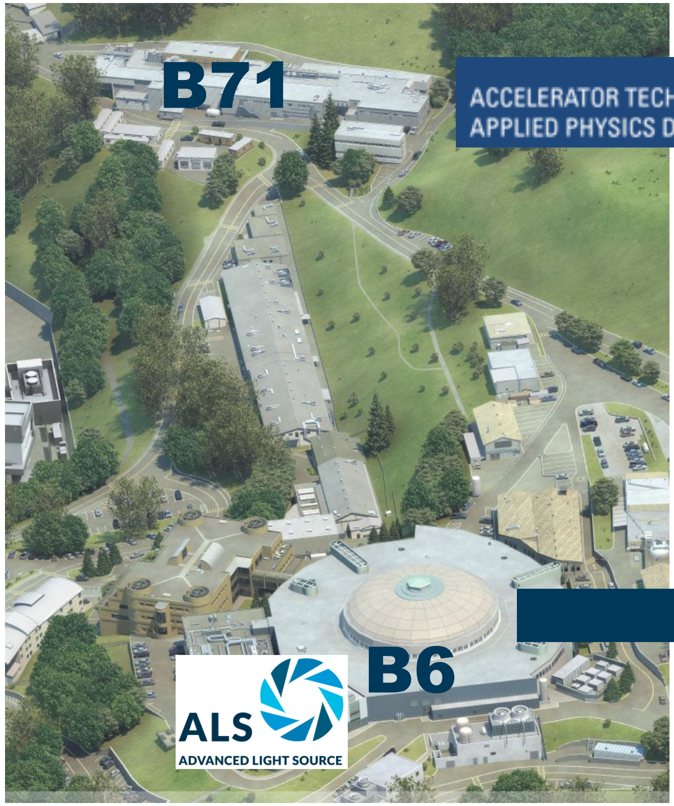

ACCELERATOR TECHNOLOGY & APPLIED PHYSICS DIVISION ATAP

184" cyclotron yoke (1940)

What is the Advanced Light Source?

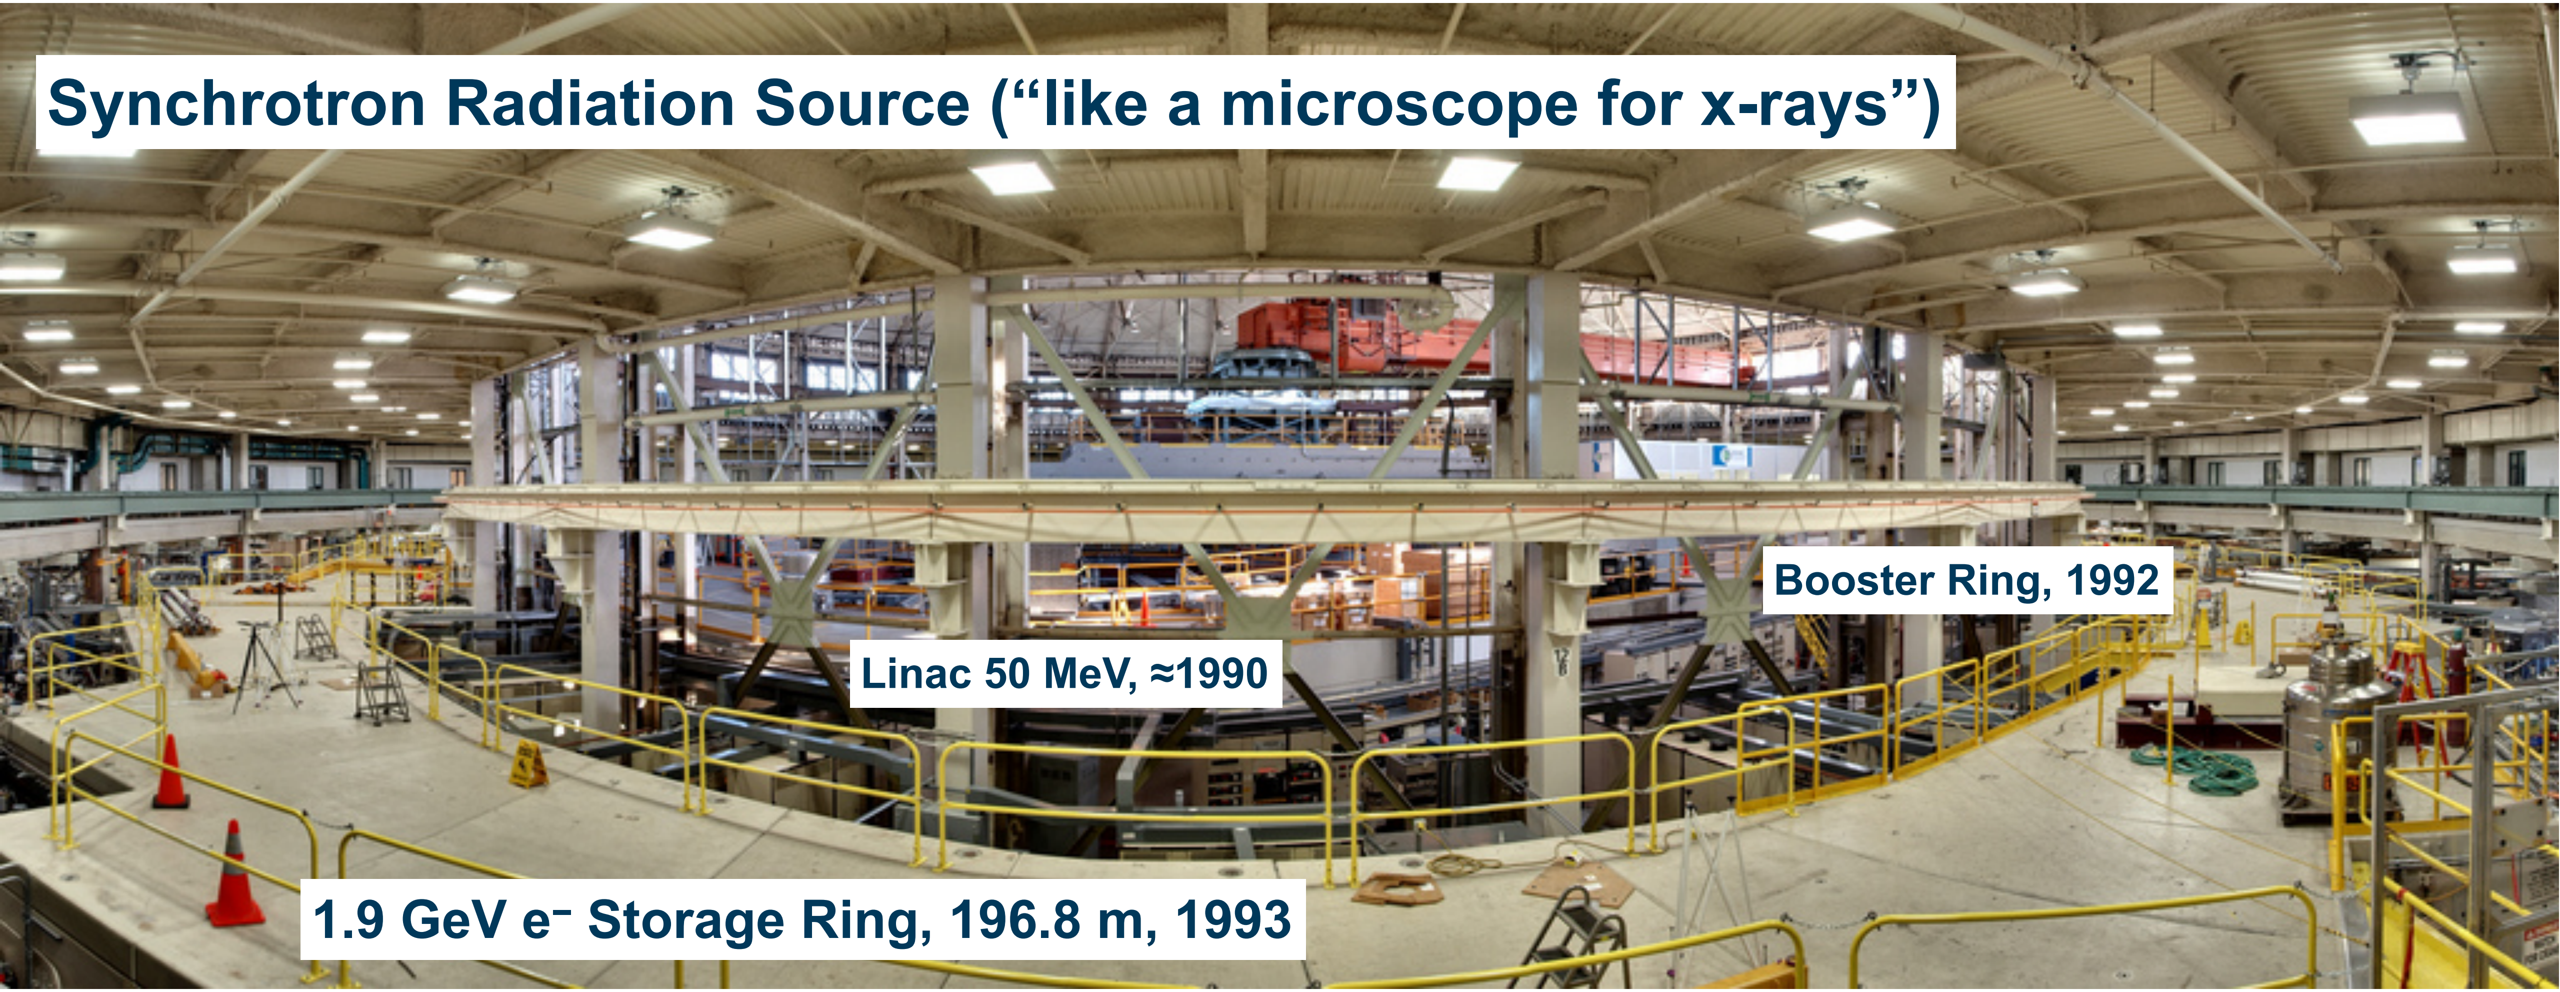

What is the Advanced Light Source?

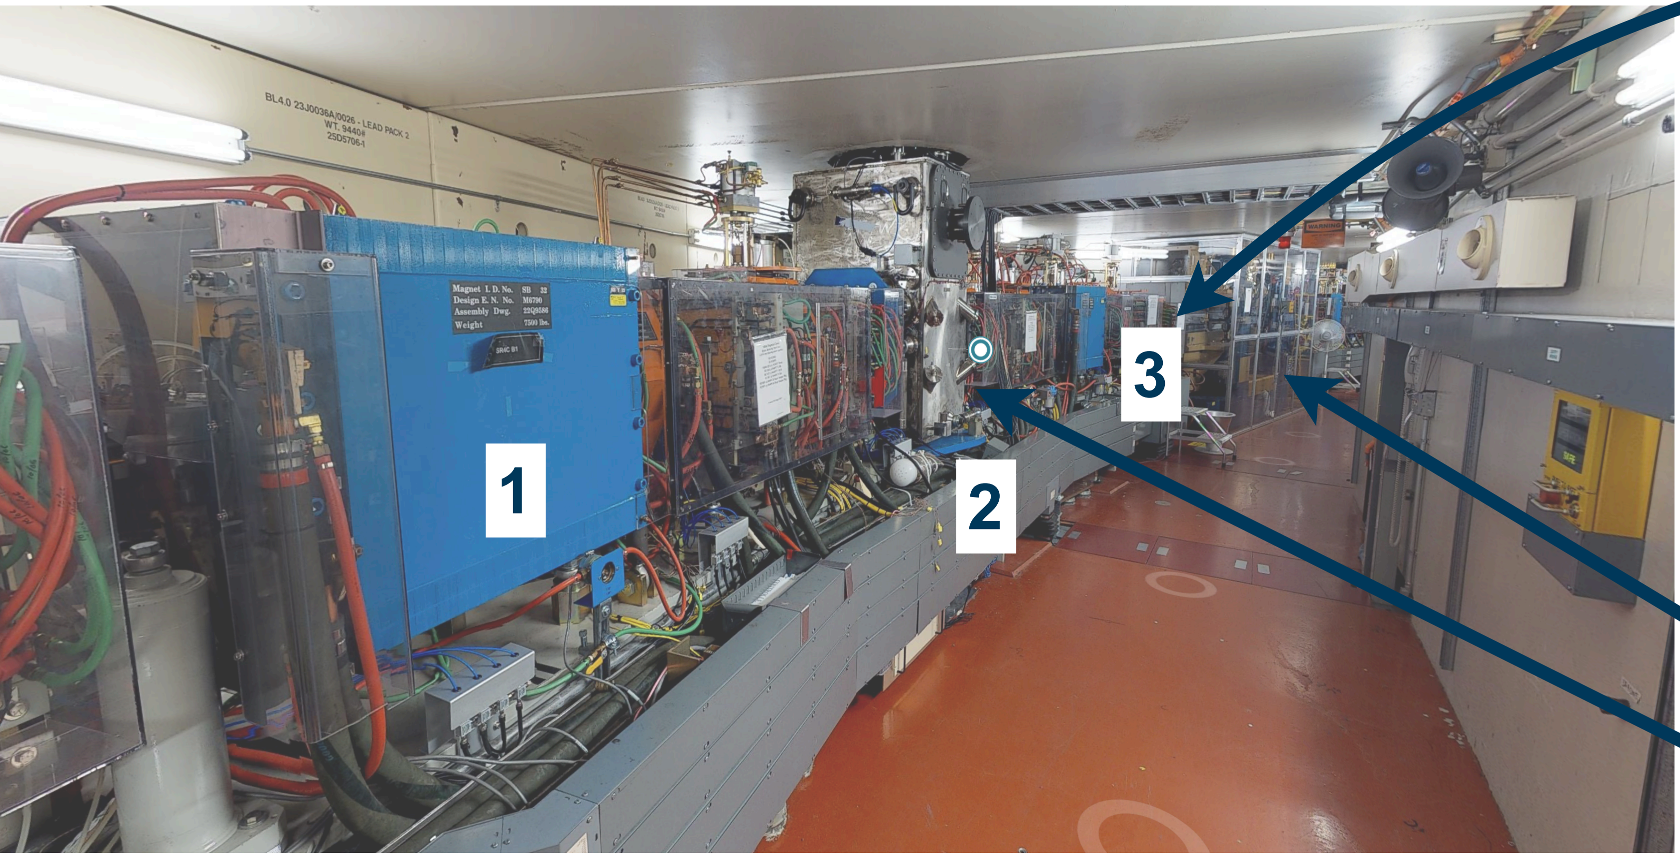

Triple-bend achromat
(12 sectors)

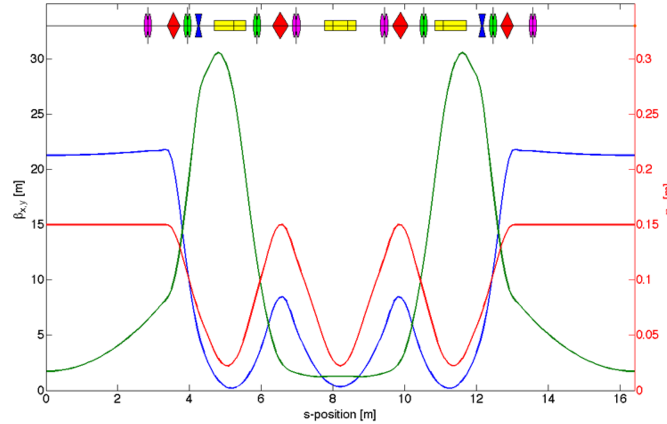

ALS TBA 2013, ε_0 = 2 nm rad

12 long straights
14 insertion devices (high brightness)

3 superbends (hard x-rays)

≈40 beamlines → IR, UV, soft & tender x-rays

What is the Advanced Light Source?

≈40 beamlines, ≈5000 hrs/y, ≈2000 users/y

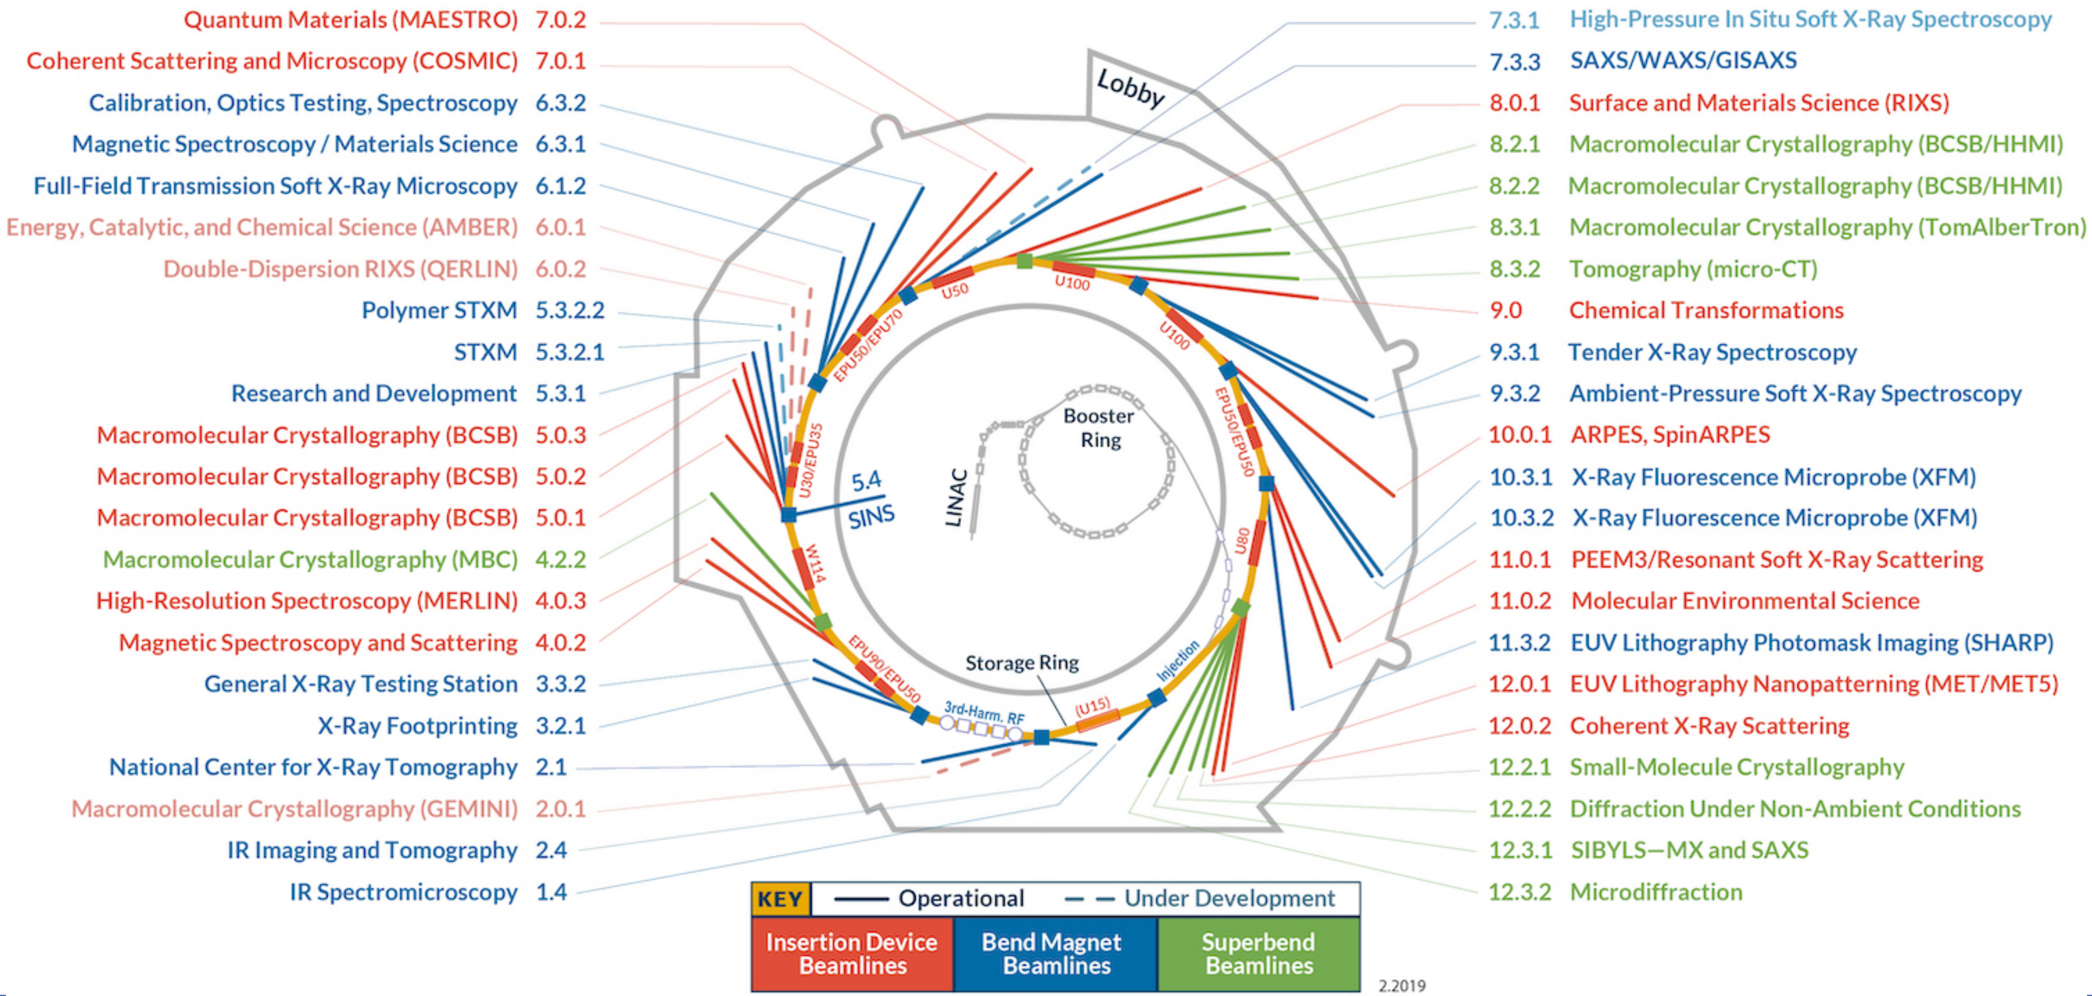

#1 ML for Acc Ops: Stabilizing Beam Size at ALS

- State-of-the-art light sources **achieve excellent stability** in terms of beam position/angle & current (orbit feedbacks, top-off injection)
- In spite of extensive correction efforts, **beam size** is still perturbed by **insertion device (ID)** config changes ➔ can affect experimental resolution
- Problem is nonlinear, complex, and non-stationary
- Previous solutions relied on approximations & required extensive **dedicated machine time** for frequent recalibration (feed-forward tables)
- Resulting level of performance has started to become a limitation at most demanding experiments & is expected to become a **serious issue** in next-generation light sources (diffraction-limited storage rings, eg. APS-U, ALS-U, …)

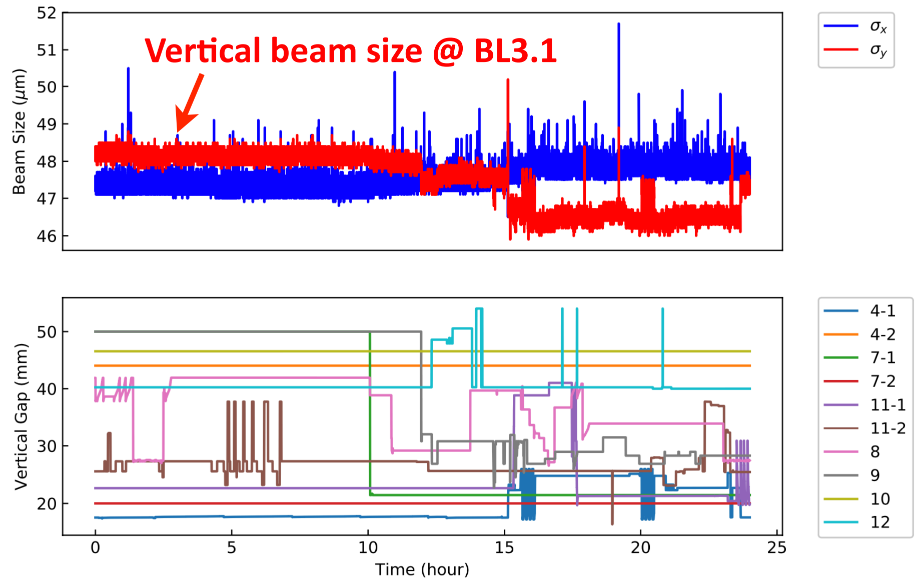

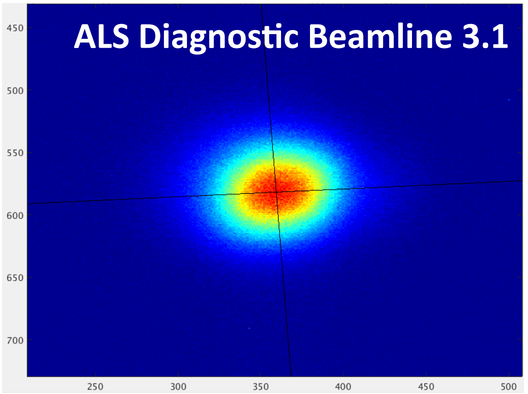

Developing a Solution Based on Machine Learning

- **Machine Learning** can exploit large amounts of data that are already collected during routine operations ➔ ***"training"***
- Once trained, **neural network (NN)** provides **predictions** for beam size changes that result from ID config changes & magnet corrections

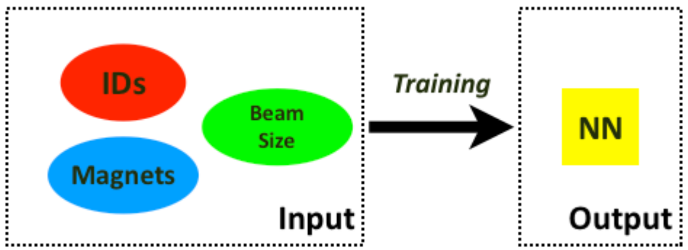

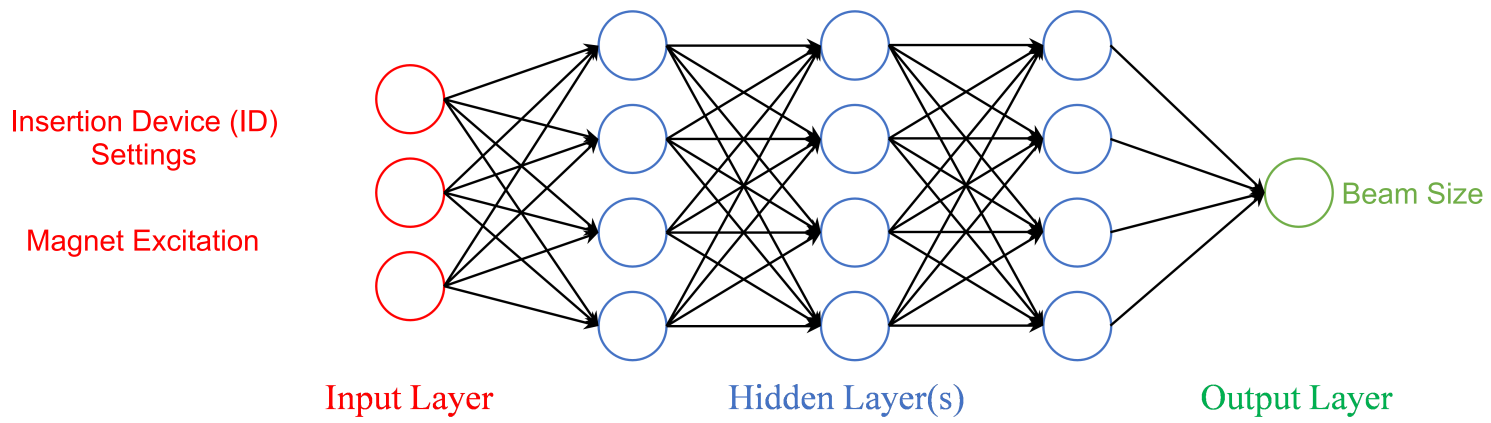

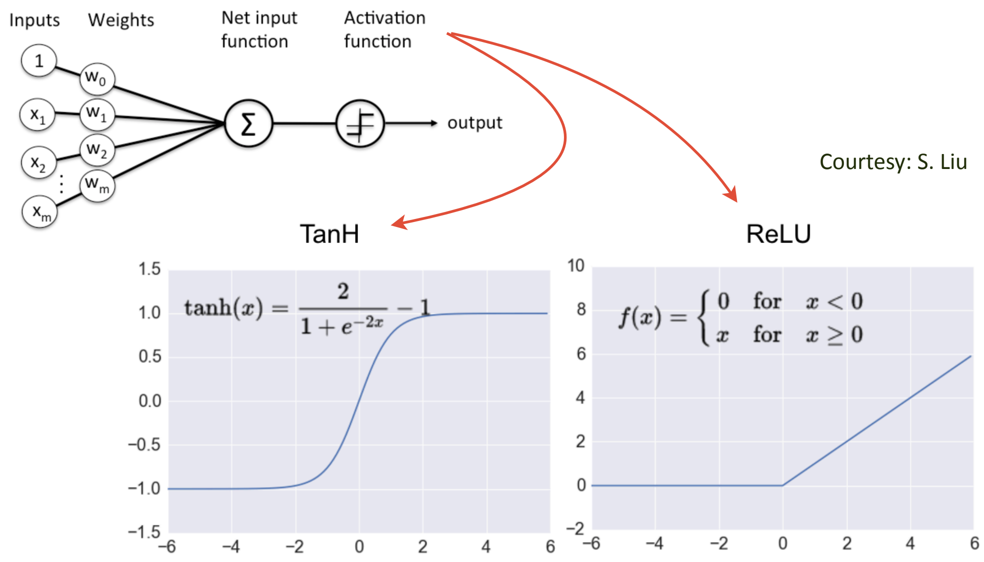

Developing a Solution Based on Machine Learning

- **Machine Learning** can exploit large amounts of data that are already collected during routine operations ➔ ***"training"***
- Once trained, **neural network (NN)** provides **predictions** for beam size changes that result from ID config changes & magnet corrections

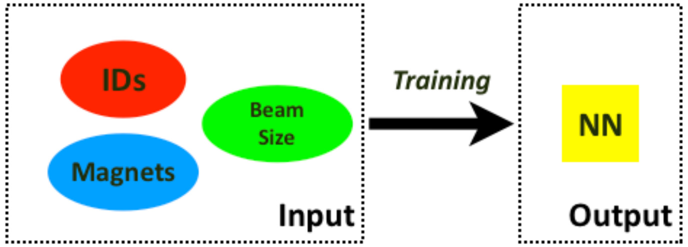

- These predictions can serve as a **dynamic lookup** ➔ which magnetic correction required to compensate for changes resulting from currently applied ID config?

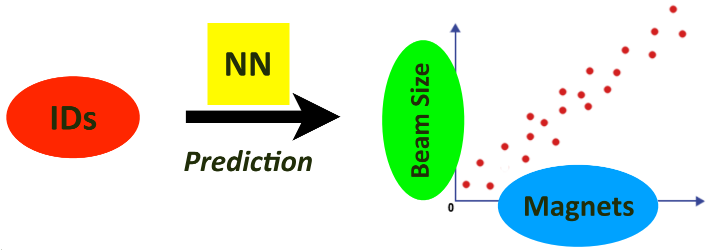

- If such a lookup is incorporated into the accelerator control system as a **feed forward (FF)**, we can stabilize the storage ring over prolonged periods of time & **online retraining** can mitigate drift

PRL **123**, 194801 (2019)

Results: NN-based FF Off vs. On During User Ops

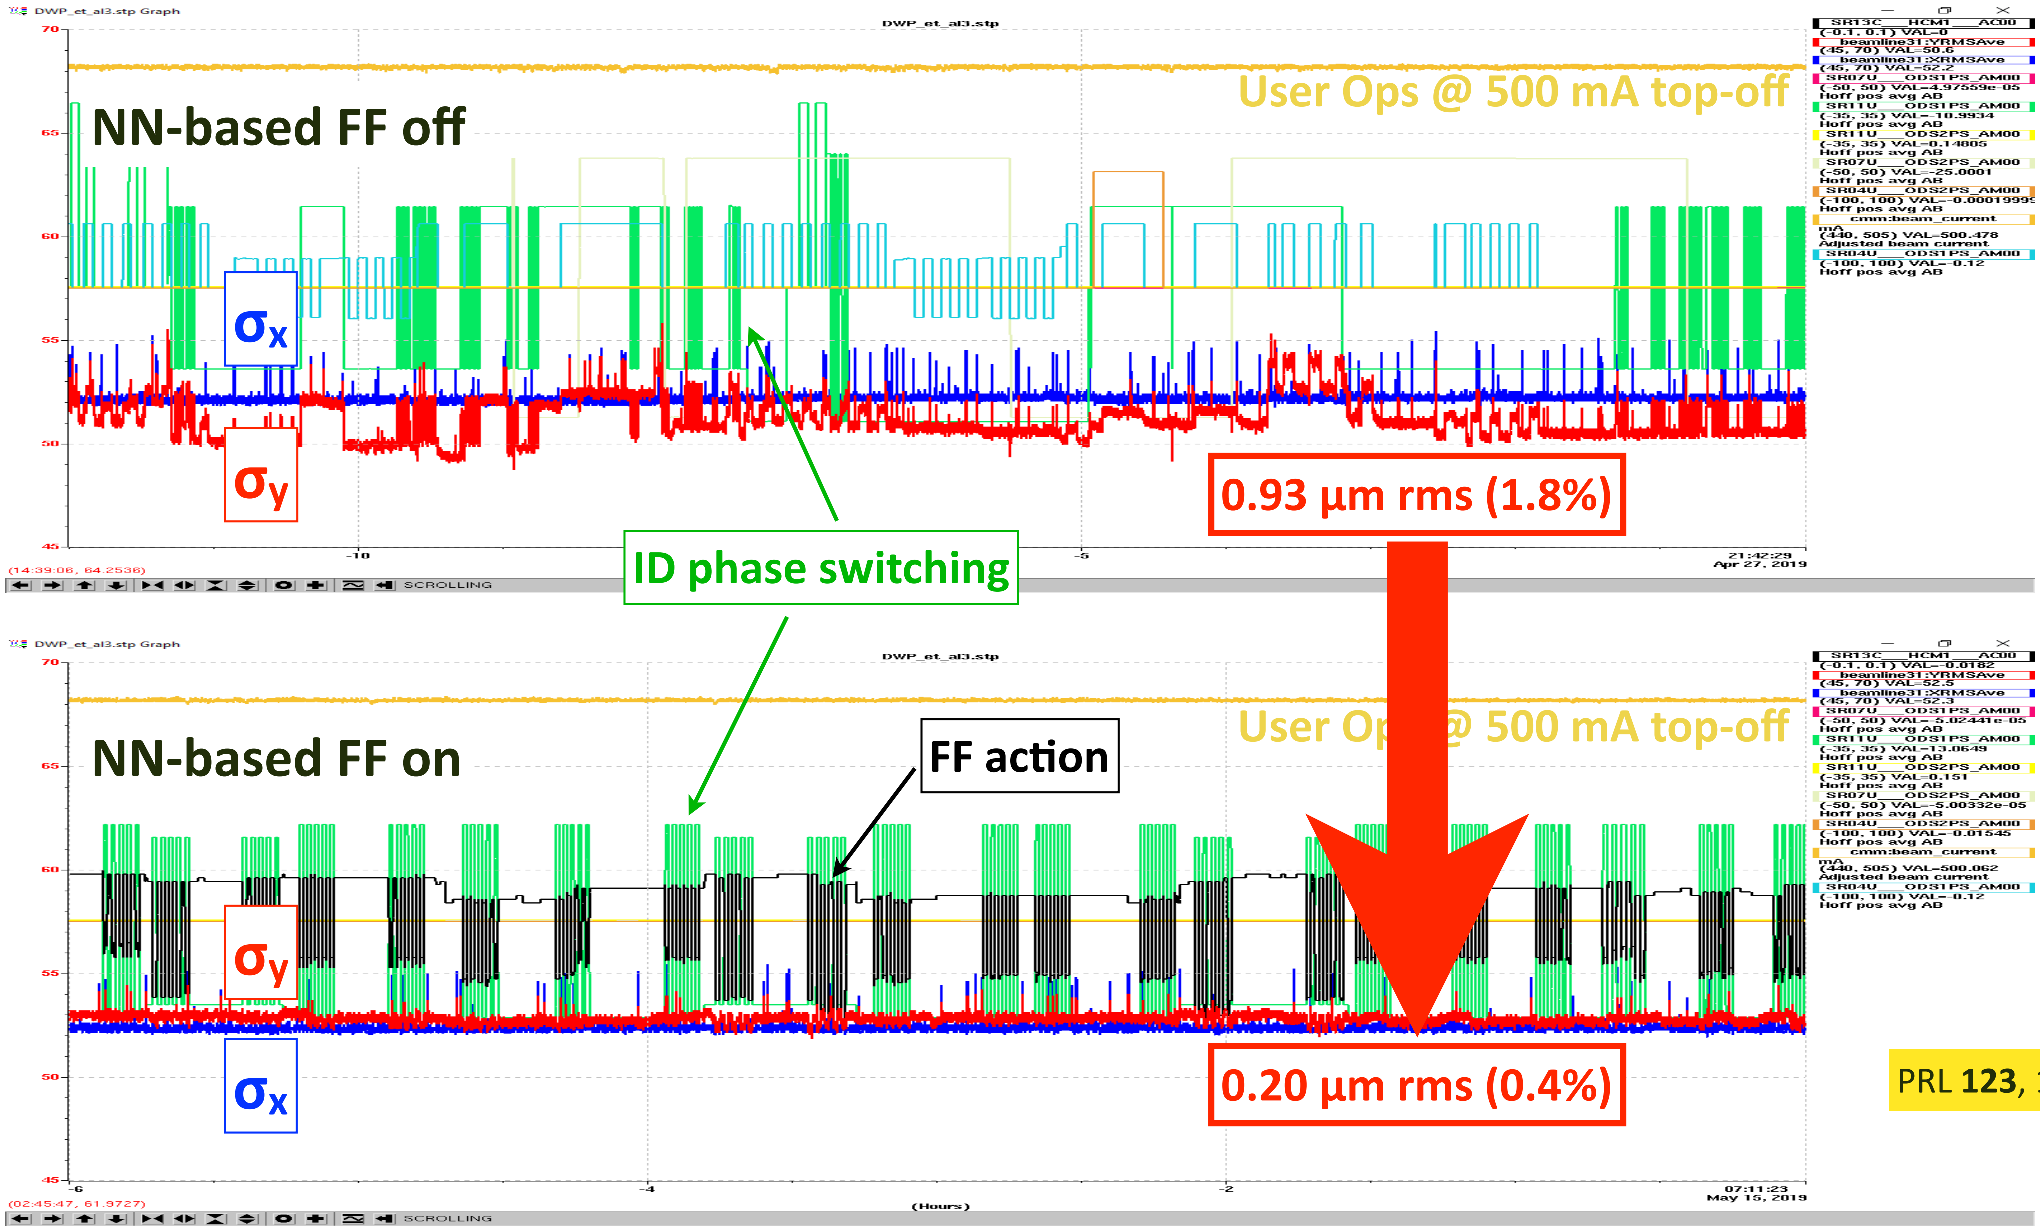

PRL **123**, 194801 (2019)

Stabilization Confirmed at Experiment (ALS BL 5.3.2.2)

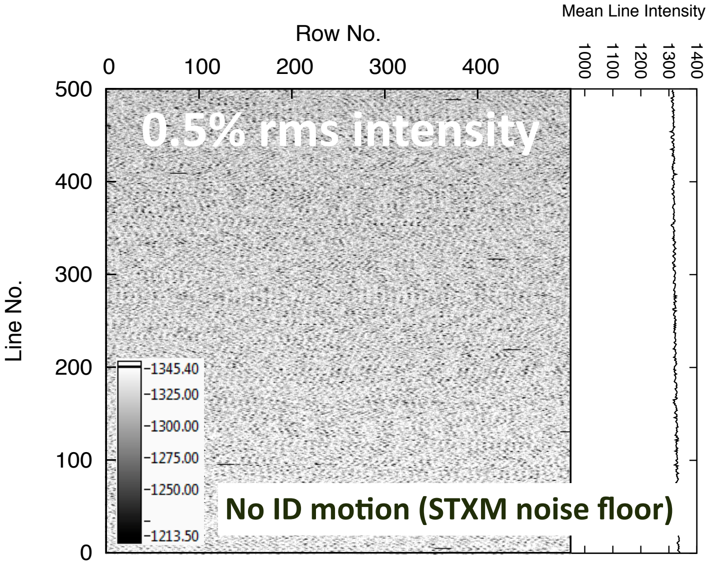

ID Motion

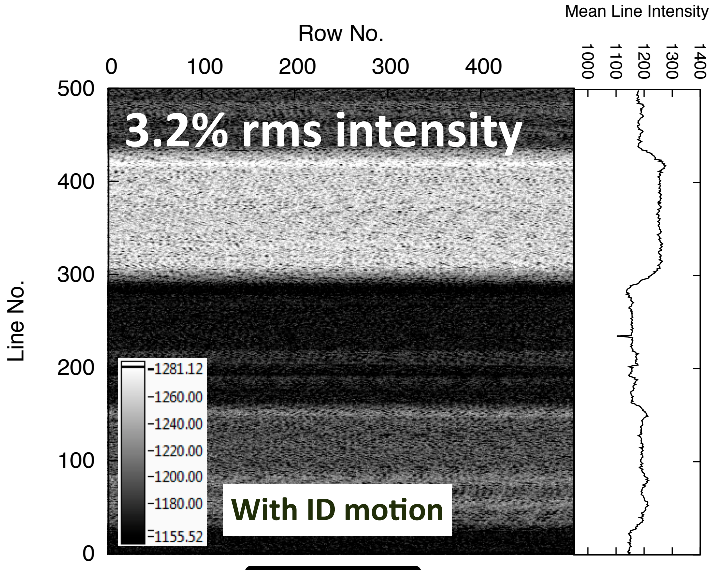

NN-based FF on

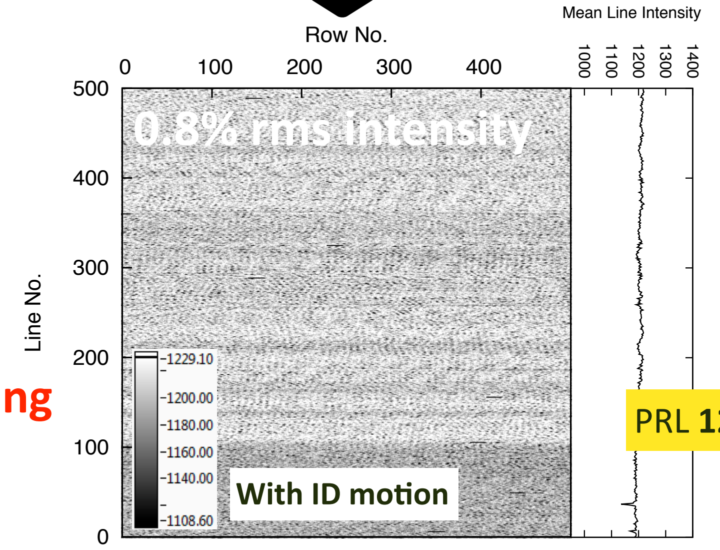

Online Retraining

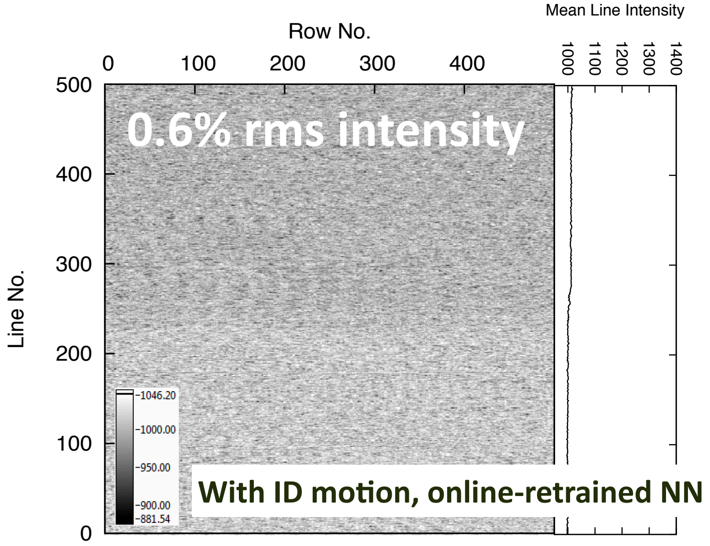

Noise reduced to almost floor level

PRL **123**, 194801 (2019)

#2 ML for Acc Dev: Improving Multi-Objective Optimization

- **4th-generation storage rings (4GSRs)** leverage multi-bend achromat (MBA) lattices to deliver ultra-high brightness & large coherent fraction
- But **MBA lattices** are very challenging: dense & exploit very strong focusing ➔ drives large chromatic terms & higher-order corrections
- Solutions not only highly nonlinear but involve many degrees of freedom (DoF) ➔ **demanding optimization**:
  - tough objectives, many often in direct competition
  - large number of parameters, many boundary constraints

➡ **Multi-objective genetic algorithms (MOGA)** are highly successful at such optimization & have become tool of choice among community

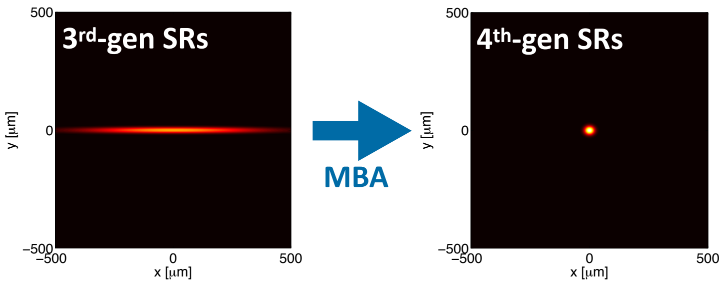

Courtesy: Dave Robin

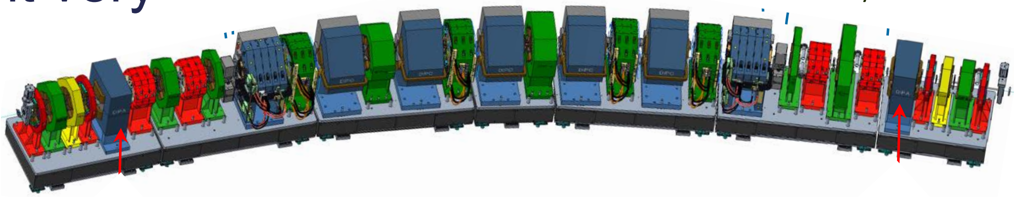

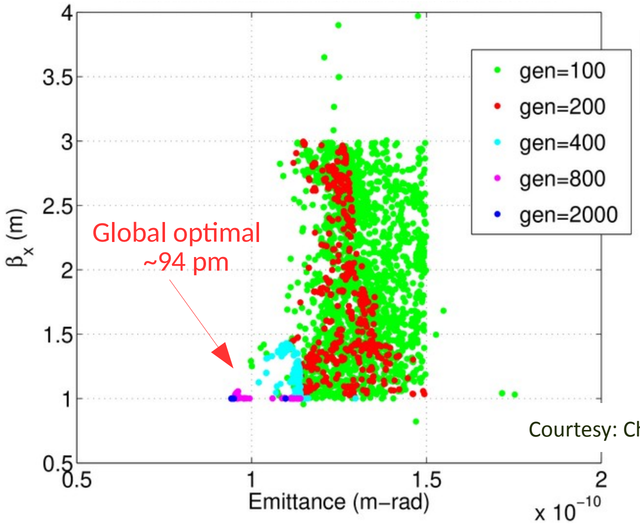

Courtesy: Changchun Sun

Improving MOGA: ML to the Rescue

- But MOGA's **stochastic nature** is inherent weakness ➔ need to evaluate vast number of lattice candidates, most ultimately rejected
- Do *not* want to artificially limit DoF, search ranges, or make many initial assumptions about attractive solutions ➔ so what can we do?
- ML can be employed to render **deep neural networks (DNNs)** ➔ **surrogate models** used in place of computationally expensive evaluation

➡ Evaluation of lattice candidates becomes almost **instantaneous**

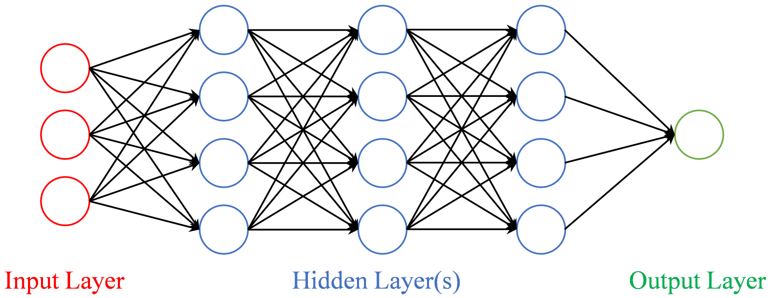

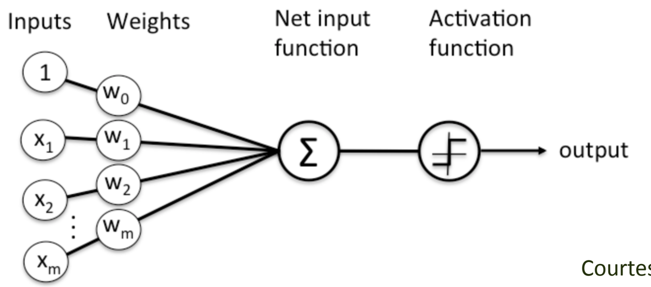

Courtesy: S. Liu

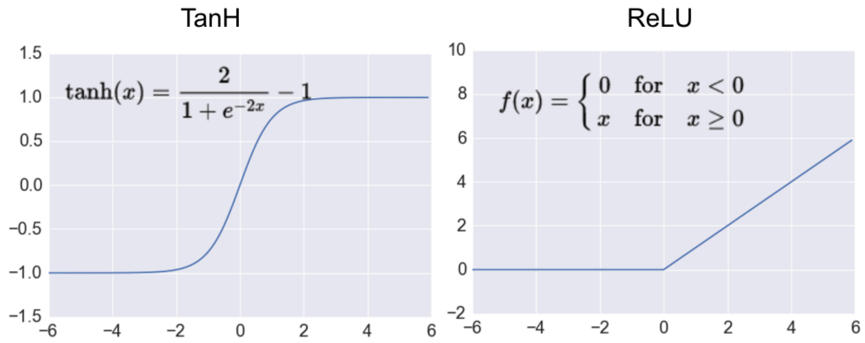

ALS-U Optimization as a Test Case for ML

- **ALS-U storage ring** calls for challenging 9BA lattice to achieve ≈75 pm rad (round beam) @ 2 GeV in <200 m
  ➔ dense, strong focusing, **very strained optics**

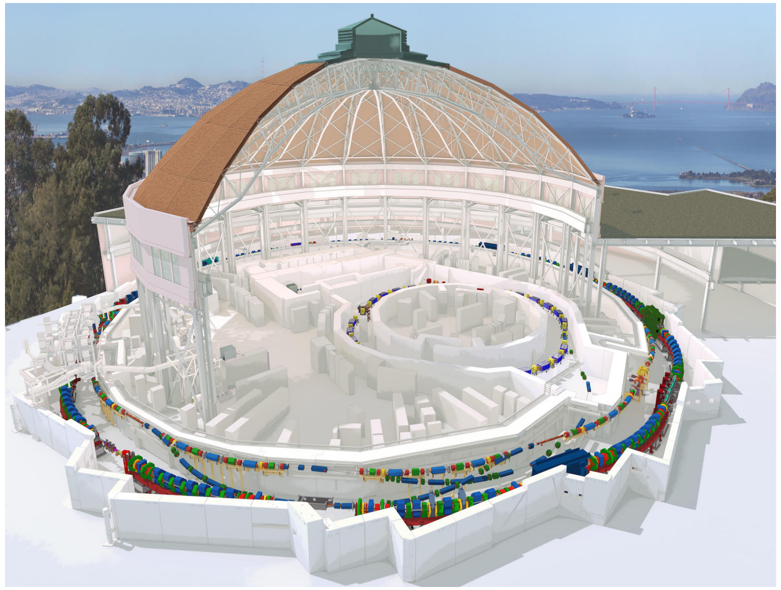

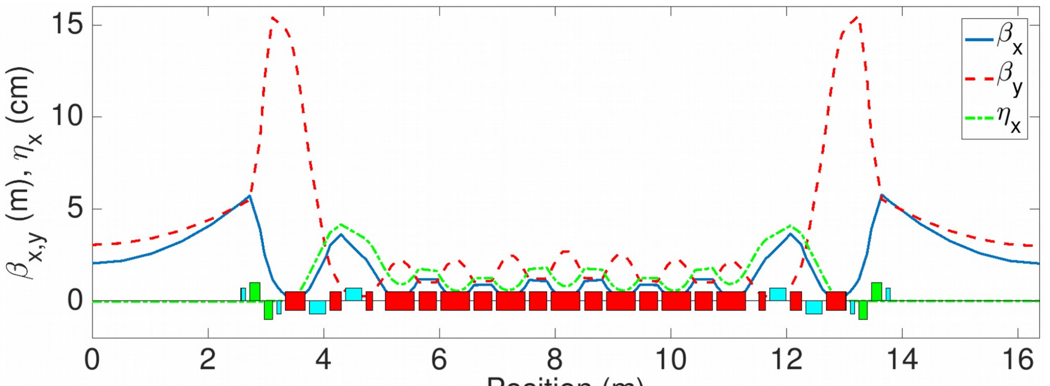

Courtesy: Changchun Sun

ALS-U Optimization as a Test Case for ML

- **ALS-U storage ring** calls for challenging 9BA lattice to achieve ≈75 pm rad (round beam) @ 2 GeV in <200 m ➔ dense, strong focusing, **very strained optics**
- Initial optimization: 9 quadrupoles, 4 sextupoles ➔ **11 free knobs** (later: include reverse bending & superbends)
  - Roughly a dozen magnet/lattice **constraints** on top of pre-determined quadrupole ranges
  - **Objectives**: ε_0, MA, and on-momentum DA (modeled as integrated diffusion rate)

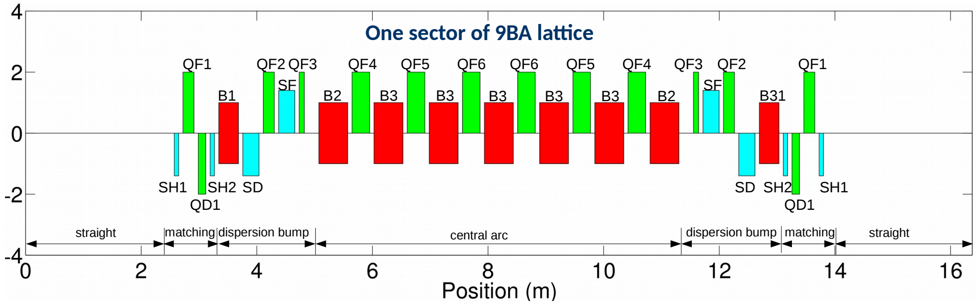

Natural emittance	$\varepsilon_0 < 155\,\mathrm{pm\,rad}$
Maximum beta	$\beta_{x,y} < 30\,\mathrm{m}$
Maximum dispersion	$\eta_x < 15\,\mathrm{cm}$
Fractional tunes	$0.1 < \nu_{x,y} < 0.4$
Dispersion at center of straight	$\lvert\eta_x^*\rvert < 1\,\mathrm{mm}$
Beta at center of straight	$1\,\mathrm{m} < \beta_{x,y}^* < 5\,\mathrm{m}$
Beta in central arc bends (B3)	$\beta_{x,y}^{\mathrm{B3}} < 4\,\mathrm{m}$
Fractional tune difference	$\lvert\nu_x - \nu_y\rvert < 0.01$
Chromatic sextupole strength (SF, SD)	$b_2 < 900\,\mathrm{m}^{-3}$

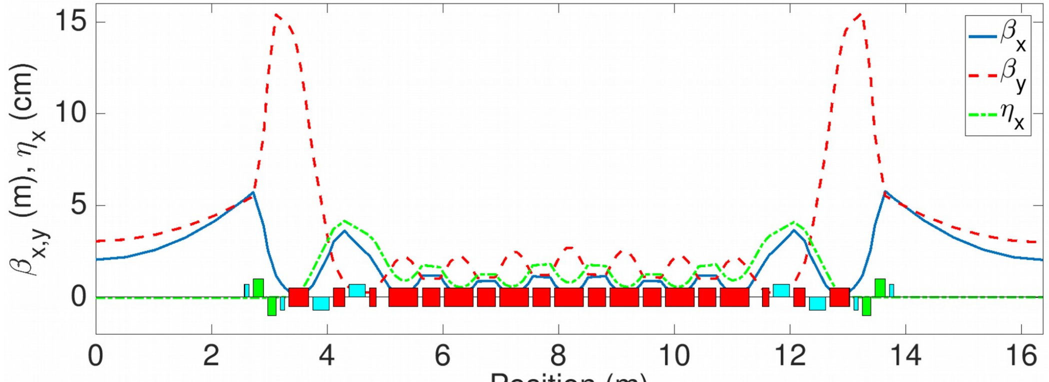

Courtesy: Changchun Sun

ALS-U Optimization as a Test Case for ML

- **ALS-U storage ring** calls for challenging 9BA lattice to achieve ≈75 pm rad (round beam) @ 2 GeV in <200 m ➔ dense, strong focusing, **very strained optics**
- Initial optimization: 9 quadrupoles, 4 sextupoles ➔ **11 free knobs** (later: include reverse bending & superbends)
  - Roughly a dozen magnet/lattice **constraints** on top of pre-determined quadrupole ranges
  - **Objectives**: ε_0, MA, and on-momentum DA (modeled as integrated diffusion rate)
- Ultimately, a highly **staged MOGA approach** resulted in
  - ±1 mm DA (compatible with on-axis swap-out & AR)
  - ≈1 hr overall lifetime (including x4 boost from 3HCs)
- …but required ***months* of CPU time** on large clusters

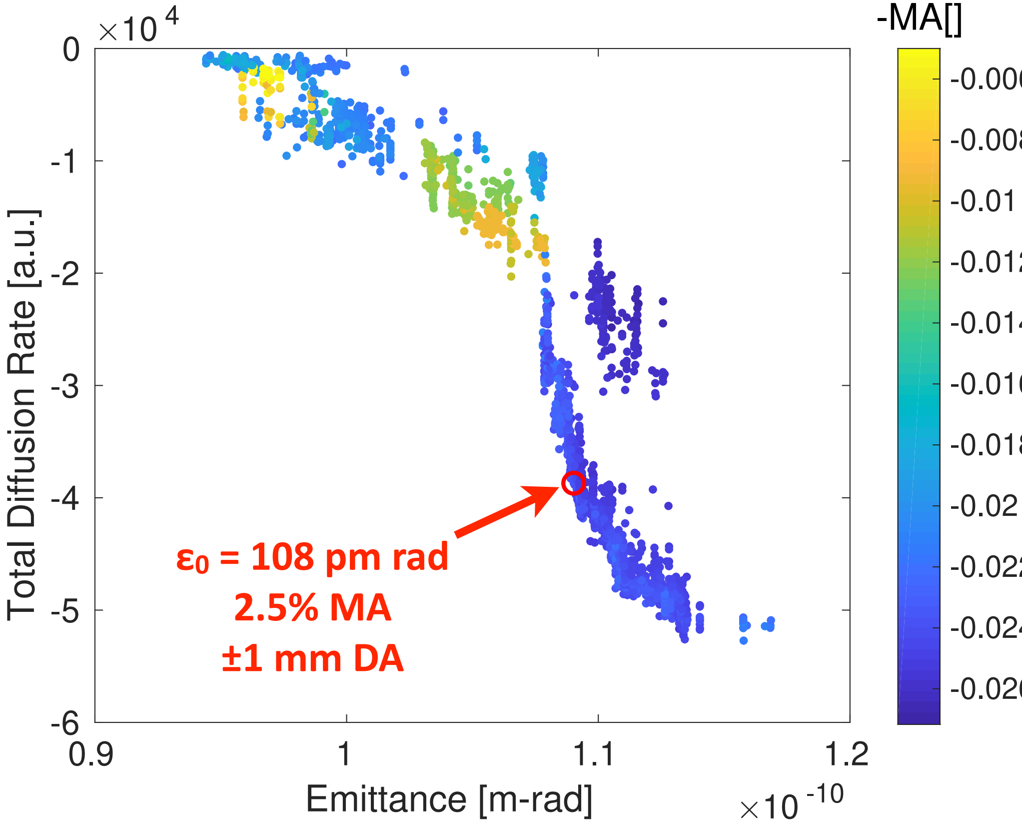

Courtesy: Changchun Sun

ML for Full Linear & Nonlinear ALS-U Optimization

- **Training data** for 11D problem cannot be acquired through systematic sampling of input space
- Instead: use first few generations of conventional **MOGA** as training data for DNNs
- **Two 8-layer DNNs** used in MOGA instead of calls to Tracy for DA and MA (via many-turn tracking)
  - Traditional MOGA requires about **640 gen** (5000 children/gen) ➔ ≈8 days on 1000-core cluster
  - Training 2 DNNs to get DA/MA predictions ≈1% rms requires data from about **10 gen**
- Training DNNs takes just ≈30 min on desktop CPU
- DNN provides quasi-instantaneous **lookup** (16 ms) vs. conventional DA/MA tracking (88 sec)

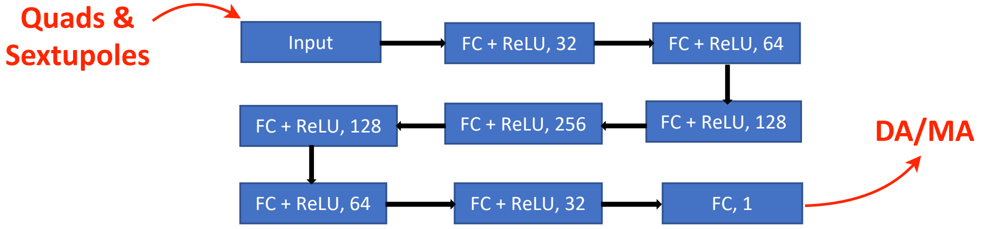

Fully-connected (FC) NN, using ReLU as activation function, # = node depth

Courtesy: Yuping Lu

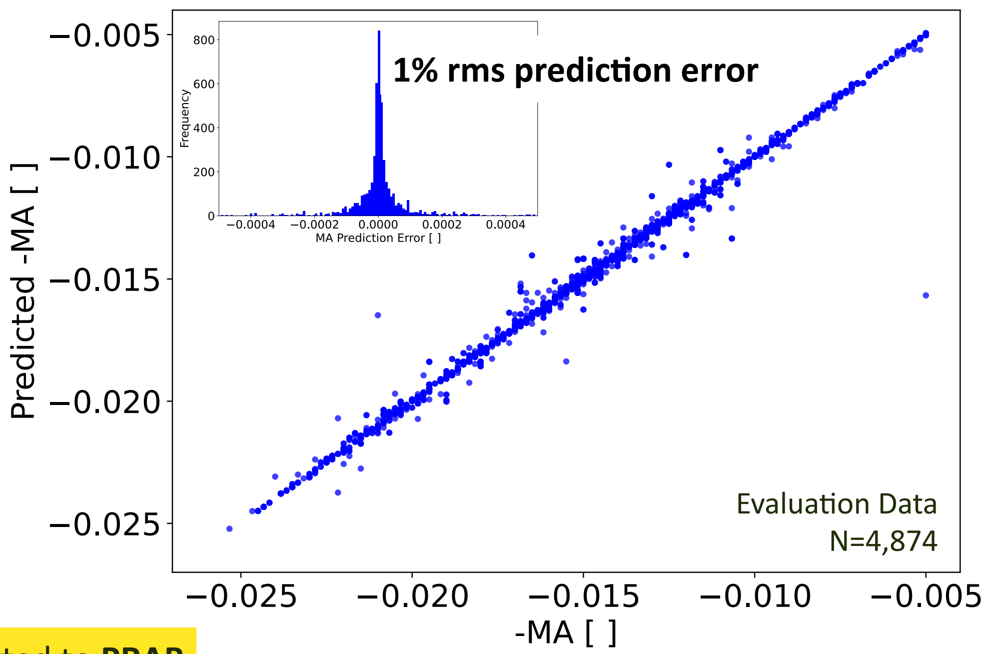

Submitted to **PRAB**

Results: ML-MOGA Successful & 40× Faster

- Initial **ML predictions** are not 100% accurate (training based on early data)
- First ML-MOGA solutions show disagreement compared to **tracking validation** ➔ but can **retrain DNNs** with data from validation step
- Iterate cycles of *validation–retraining–ML-MOGA* using **model-independent distance metrics** to determine **convergence**

➡ ML-MOGA very quickly converges (6-8 iterations) towards true **Pareto-optimal front** ➔ overall **speedup ≈ 40×** (incl. training effort)

- Once fully converged, **ML-MOGA inputs & objectives match** those of traditional MOGA to within "noise floor"
- **Flexible**: can be adapted to other lattice optimization problems as long as can provide reasonably accurate DNNs
- Potential to **fully automate** entire optimization campaign & **optimize in parallel** from the start for many error lattices is highly attractive ➔ derive truly ***global*** optimum

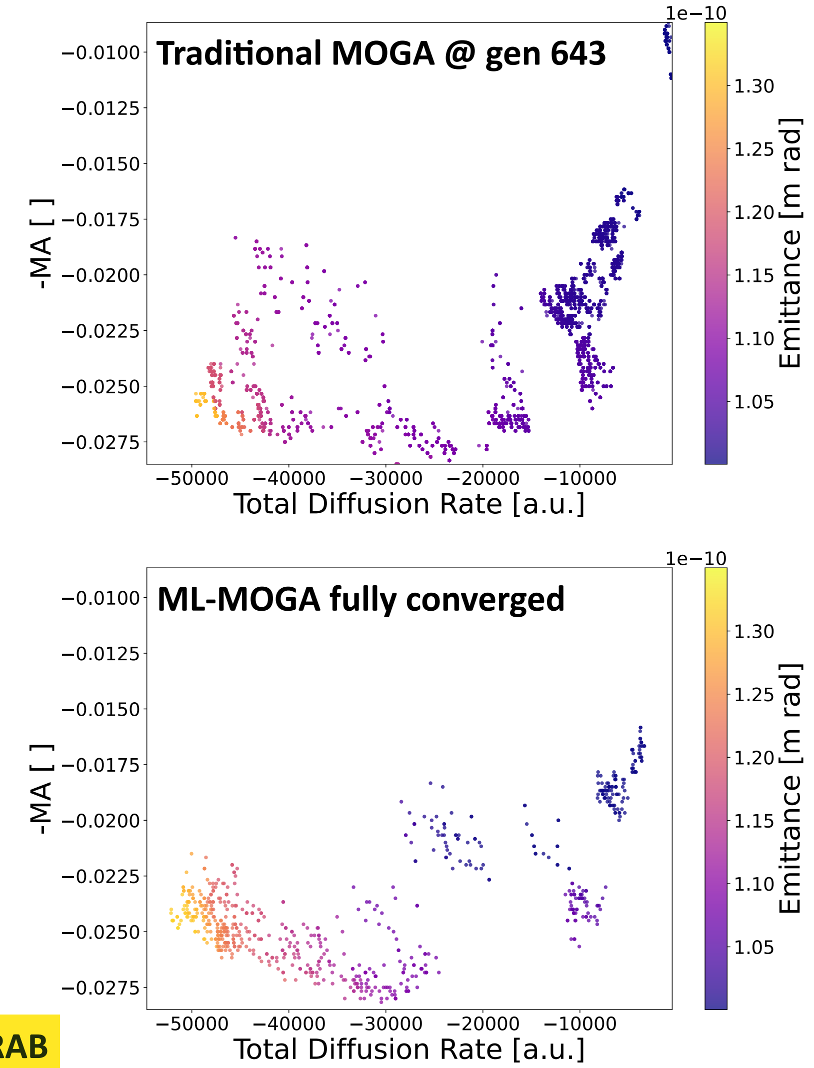

Submitted to **PRAB**

Outlook & Opportunities for Collaboration

- ML applied to accelerators shows tremendous potential to enable
  - new **more aggressive designs**, but also
  - **exploit full performance** of existing & soon to be commissioned rings

➡ These are highly relevant issues in both present (ALS) & future **4th-gen. storage rings** (ALS-U)

- ATAP & LBL can **foster great collaboration** on **ML for accelerators**
  - CRD (*the* experts on ML) ← Daniela Ushizima
  - AMP (*in-house* experts on modeling) ← Rémi Lehe (ATAP POC), Gregg Penn, Axel Huebl, Chad Mitchell, Ji Qiang

  ➡ https://atap.lbl.gov/machine-learning-artificial-intelligence-and-particle-accelerators/
  - BACI (eg. virtual diagnostics & adaptive control) ← Daniele Filippetto, Dan Wang, Du Qiang
  - Magnets & ENG (eg. magnetic field mapping & fiducialization, diagnostics, image analysis) ← Laura Garcia Fajardo, Maxim Martchevsky, Al Baskys
  - IBT (accelerator optimization using ML) ← Qing Ji, Arun Persaud
  - ALS & CAMERA (beamline instrumentation → *"Digital Twin"*) ← Antoine Wojdyla, Alex Hexemer

➡ https://ml4sci.lbl.gov/projects

Thank You

Questions?

Acknowledgments: Shuai Liu, Nathan Melton, Yuping Lu, Hiroshi Nishimura, Changchun Sun, Matthew Marcus, David Shapiro, Alex Hexemer, Dani Ushizima, Mike Ehrlichman, Gregg Penn, Thorsten Hellert, Erik Wallen, Warren Byrne, Fernando Sannibale, Marco Venturini, Andreas Scholl, Rob Ryne, DOE Office of Science Contract No. DEAC02-05CH11231

BERKELEY LAB

ACCELERATOR TECHNOLOGY & APPLIED PHYSICS DIVISION ATAP

U.S. DEPARTMENT OF ENERGY | Office of Science