Opportunities for Reinforcement Learning in Accelerator Automation

Jonathan Edelen, Joshua Einstein-Curtis, Christopher Hall, Morgan Henderson (RadiaSoft)

Auralee Edelen, Finn O'Shea (SLAC)

Stuart Calder, Christina Hoffman, Bhargavi Krishna (ORNL)

April 3rd 2025

RL4AA 2025: Hamburg Germany

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Accelerator R&D and Production award number DE-SC0023641

Automation of Sample Alignment in Neutron Beamlines

Optical Cameras → Machine Learning Model → Motor Controller → Detector Data

Fiducialized to sample and beam

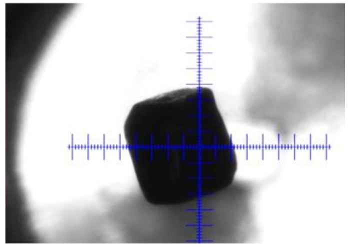

Machine Learning Model: Center of Mass of sample sent to external motor controllers

Motor Controller: Motor control software has pre-determined algorithm to move the sample based on user click

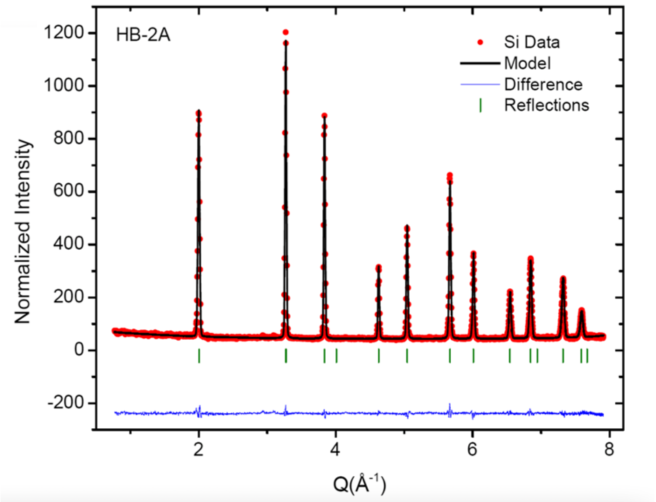

Neutron Camera → Beam Center Identification / Sample Position Identification; Sample Positioning Motors → Detector Data

Machine Learning + Classical Controls

Not Fiducialized to sample and beam

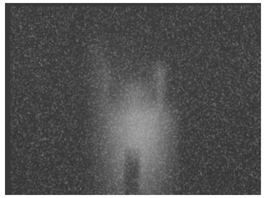

Detector Data: Provides verification of alignment status

TOPAZ Single Crystal Diffractometer

- Automation of sample alignment makes experiments more efficient and compensates for thermal drift during temperature scans
  - Utilize image contouring to identify sample location in the environment
  - Contour images using expert feedback
  - Train Convolutional Neural Network to compute the sample mask
  - From sample mask compute center of mass
- Implementation using Python based EPICS IOC
  - Optical image or neutron camera image (optical image shown in the top right)
  - U-Net computes sample center and serves up the mask to EPICS along with the sample offset from the beam
  - Motor software moves and rotates the sample for optimal alignment in the neutron beam

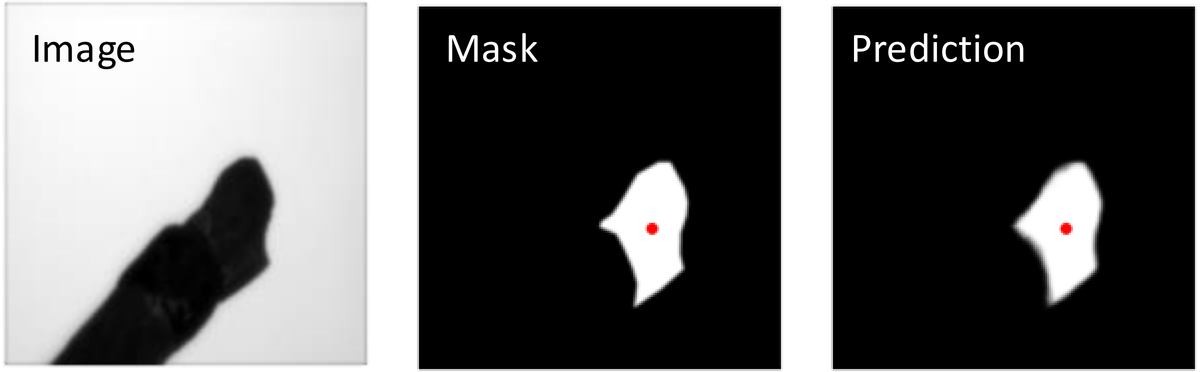

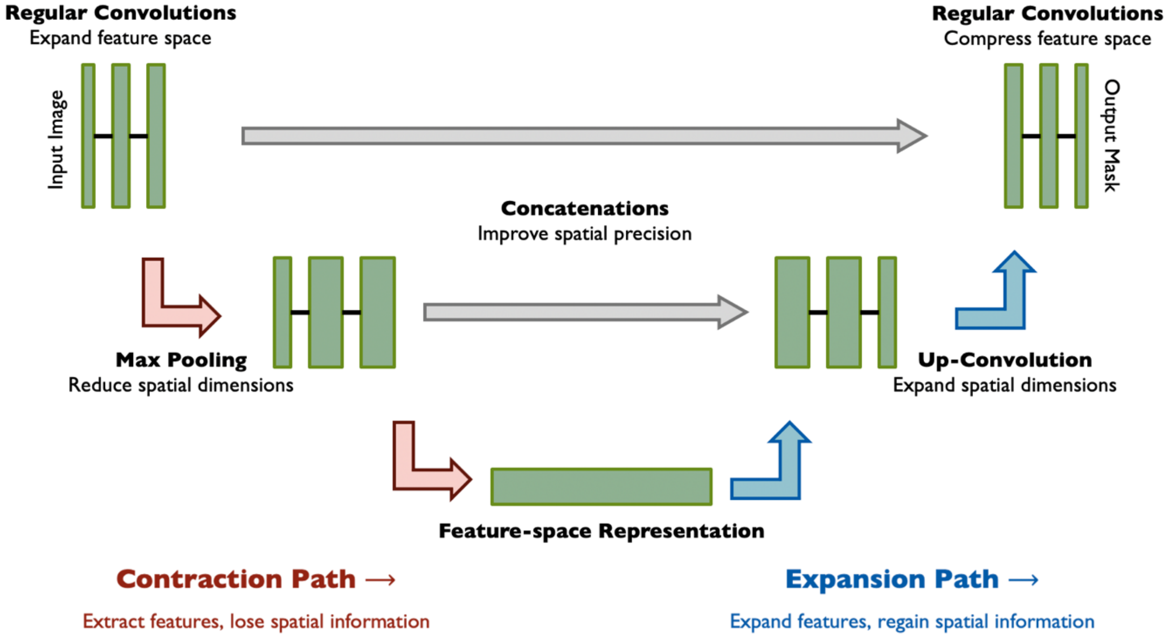

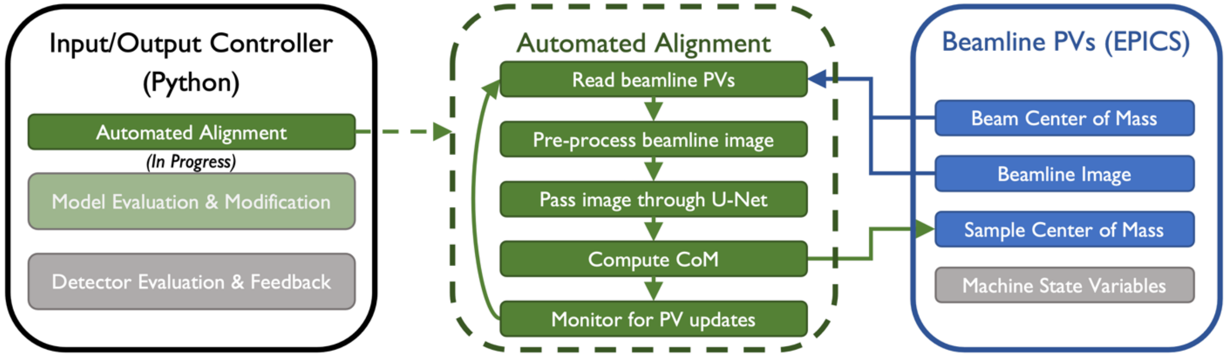

Uncertainty Quantification

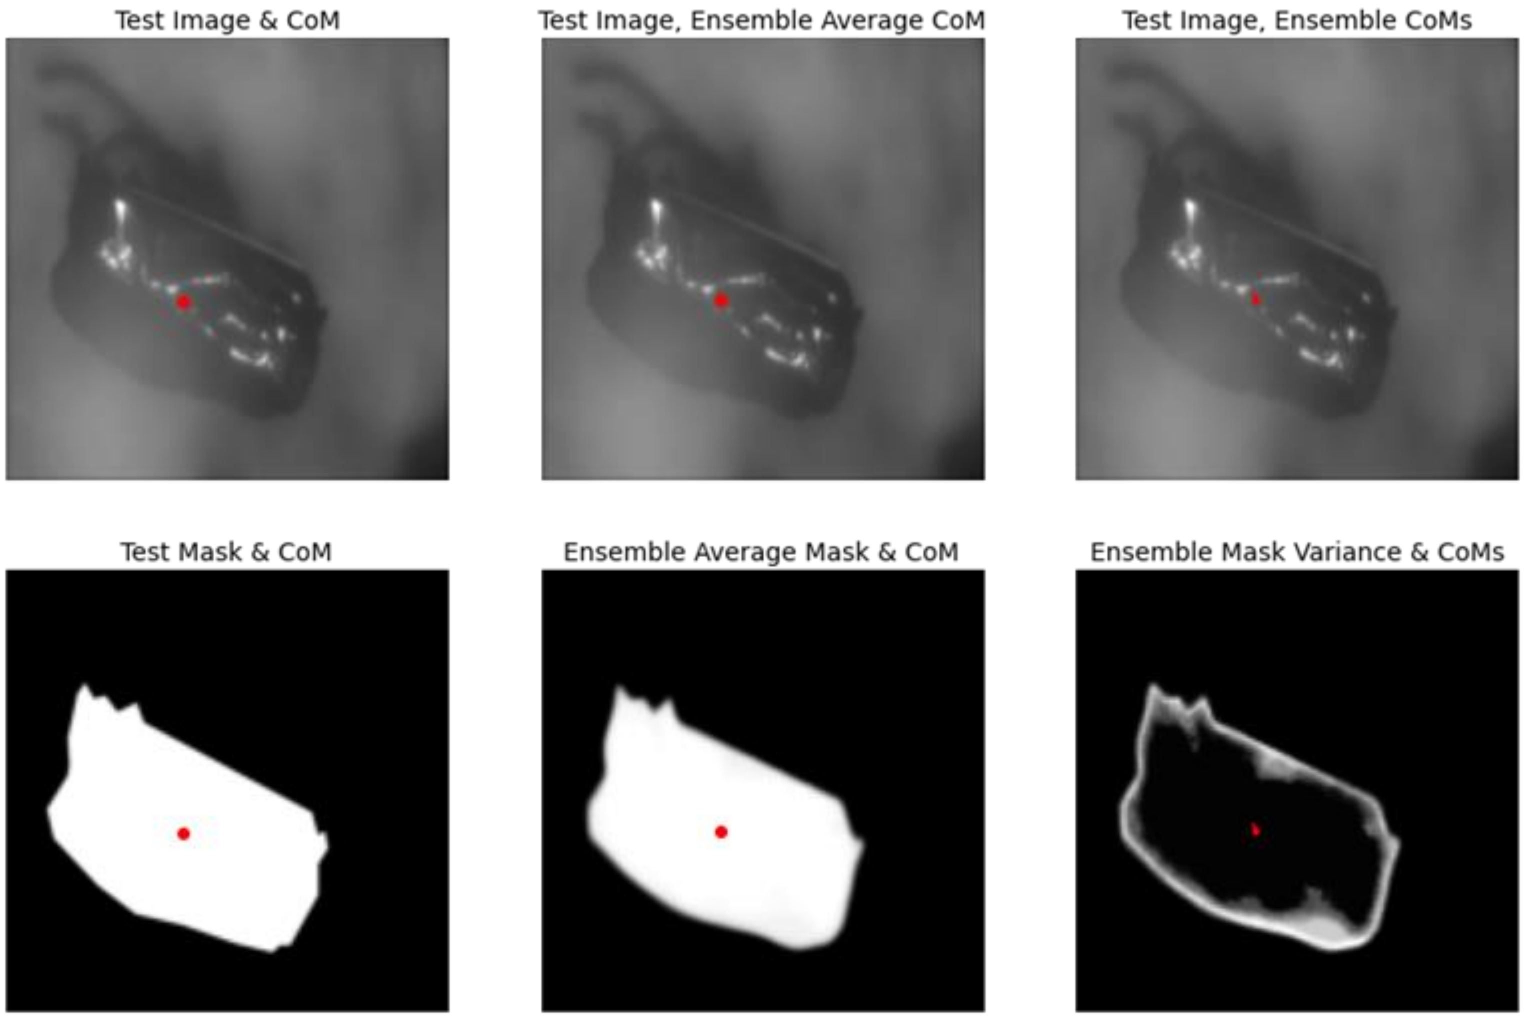

- During supervised training
  - Real error compared to human-defined masks
- During testing and operations
  - Ground-truth data not available
  - Employ statistics from ensemble predictions
  - Variance between many trained models
- Takeaways
  - Excellent CoM prediction
  - Negligible ensemble spread with low uncertainty
  - Mask variance is restricted to edges

HB-2A Powder Diffractometer

- Located at ORNL and receives neutrons from the Spallation Neutron Source
- Samples can be measured with high precision for volumetric sampling in reciprocal space (momentum measurements)
    - Samples are rotated to measure all aspects of the lattice
    - Temperature control from 5 K - 450 K
- Sample alignment
    - Neutron production time is limited
    - Some activities require constant realignment, such as temperature scans
    - User facilities especially face schedule constraints
    - Machine learning (ML) is a key automation tool
- Alignment protocols vary between beamlines
    - Opportunity to employ & test models
    - Broad applications for sample alignment

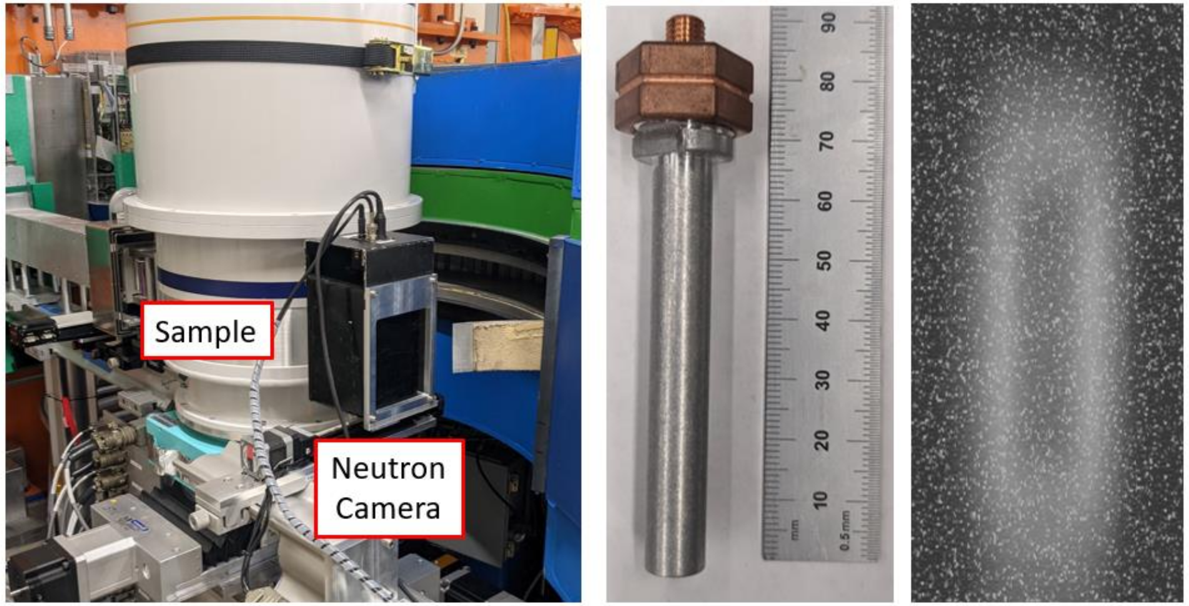

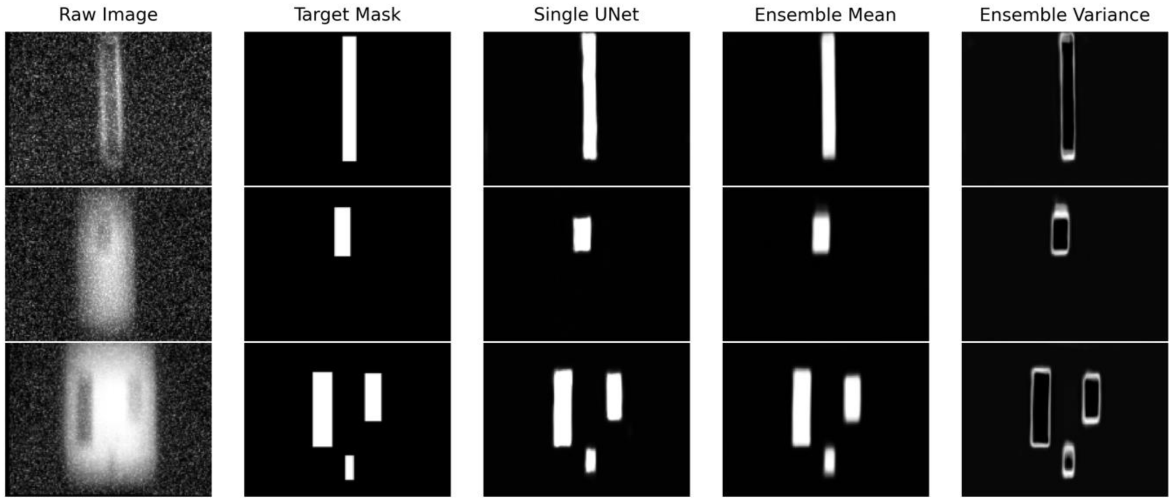

RL for Automation of Sample Alignment

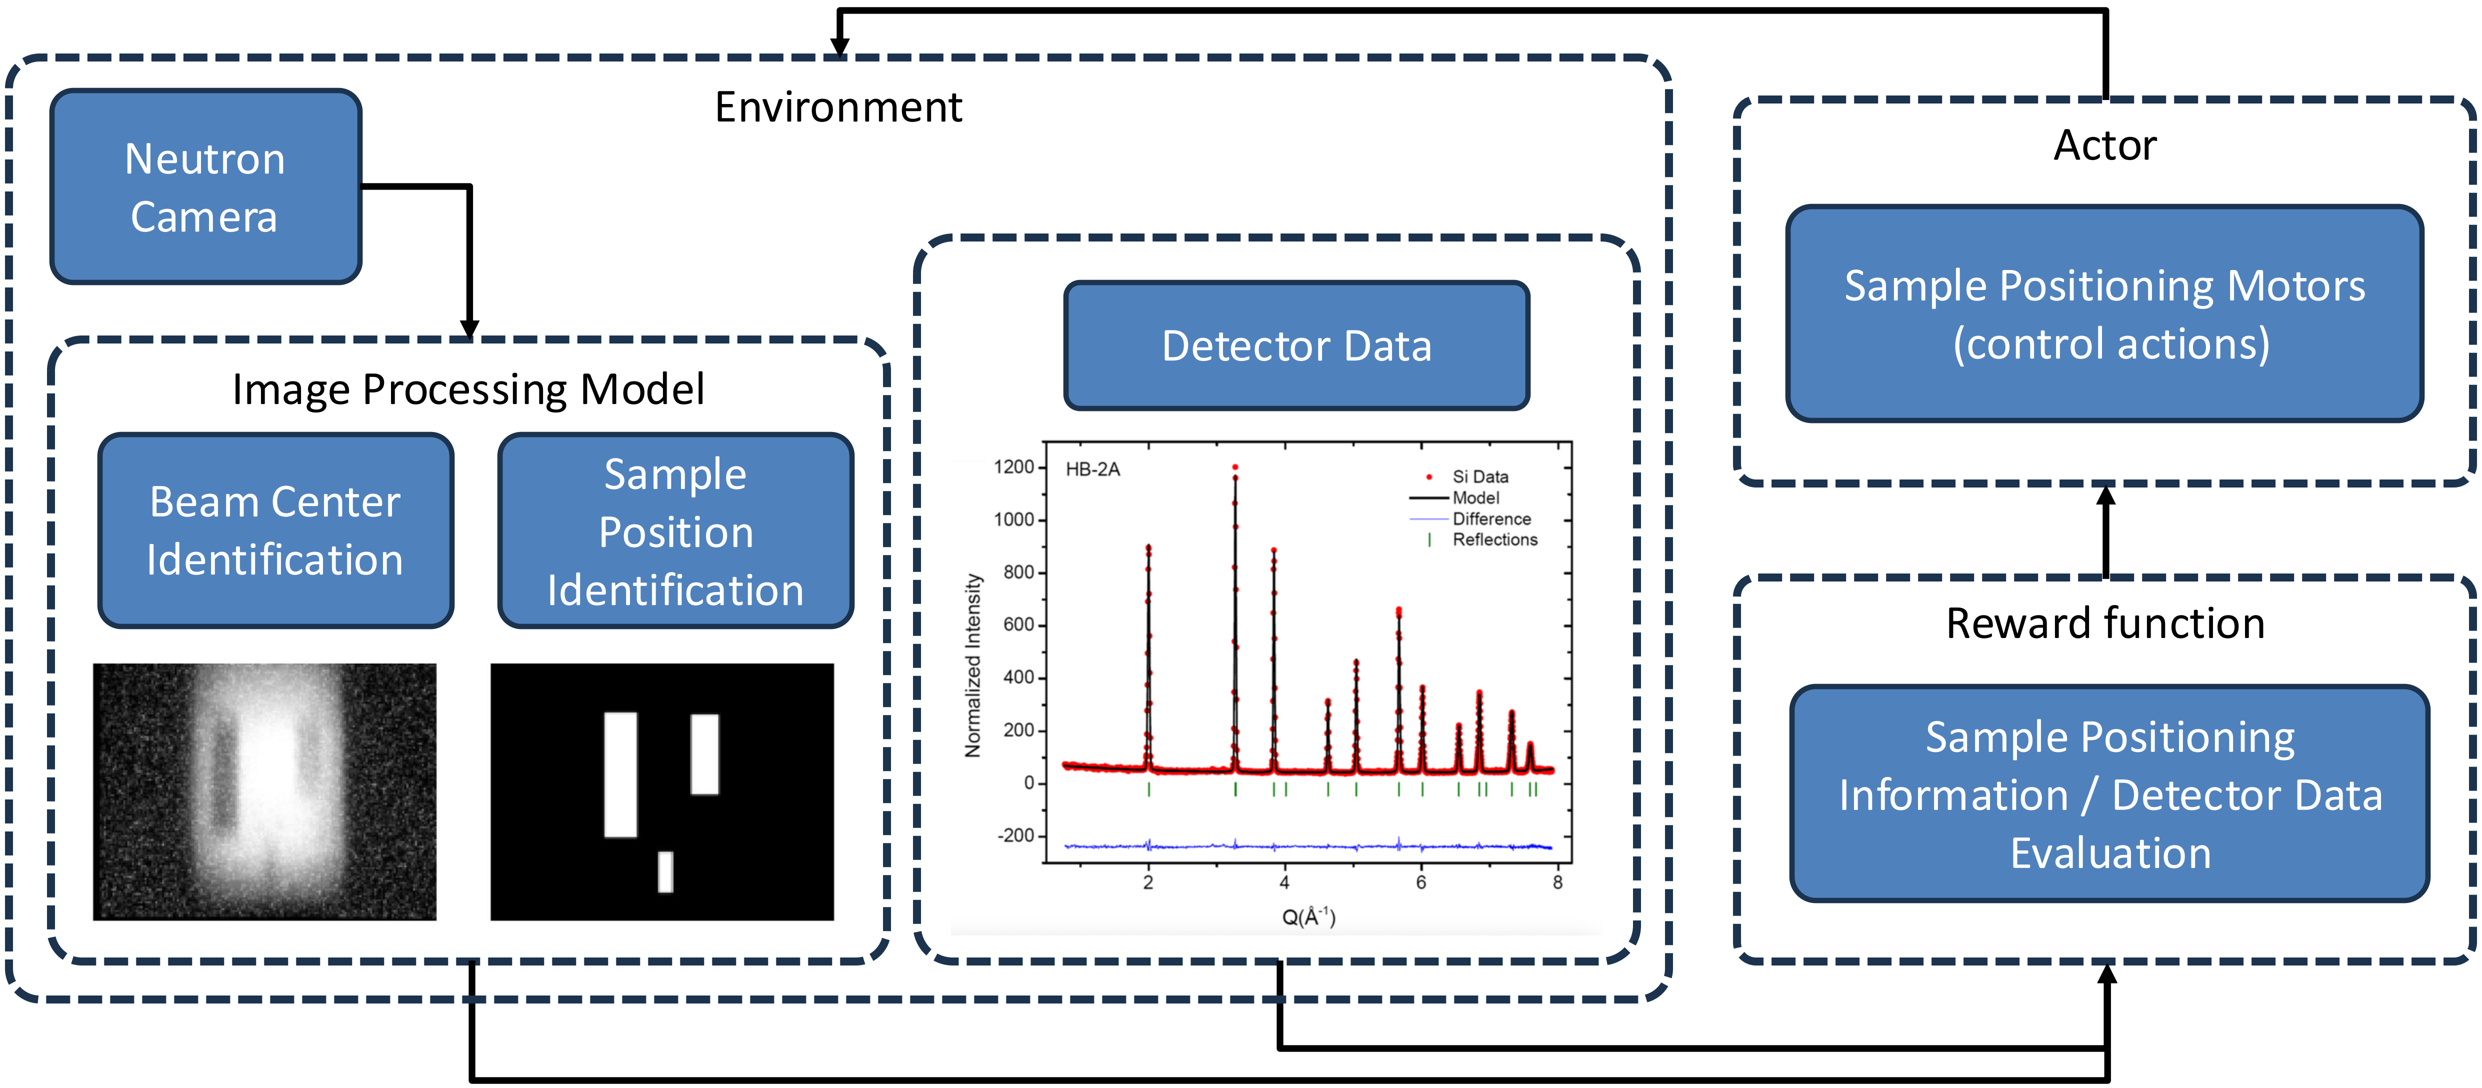

What is FLASH Radiotherapy?

- Ultrafast delivery of radiation
  - dose rates that are several orders of magnitude greater than those used in conventional radiotherapy
  - 40 Gy/s (FLASH) *vs* 0.5–5 Gy/min (conventional)

- Preclinical data suggesting that FLASH could achieve better disease control with fewer side effects
  - Improved safety and efficacy (confirmed by clinical trials)

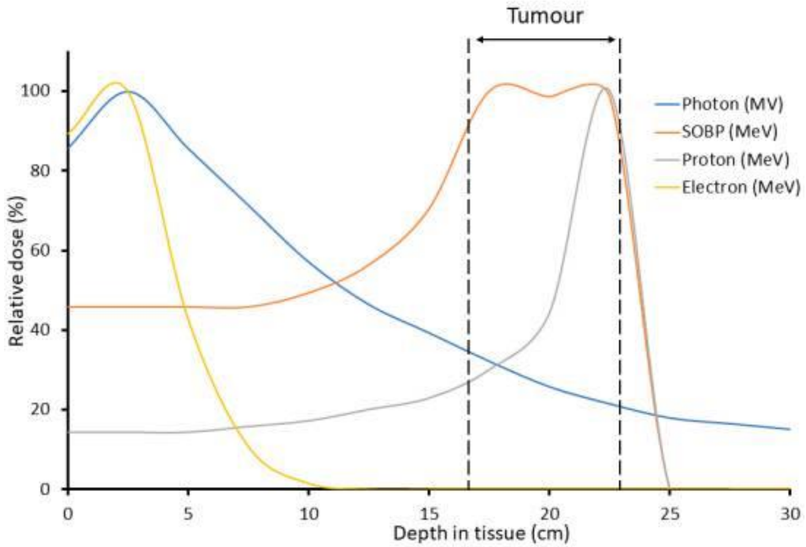

Hughes JR, Parsons JL. FLASH Radiotherapy: Current Knowledge and Future Insights Using Proton-Beam Therapy. Int J Mol Sci. 2020 Sep 5;21(18):6492. doi: 10.3390/ijms21186492. PMID: 32899466; PMCID: PMC7556020.

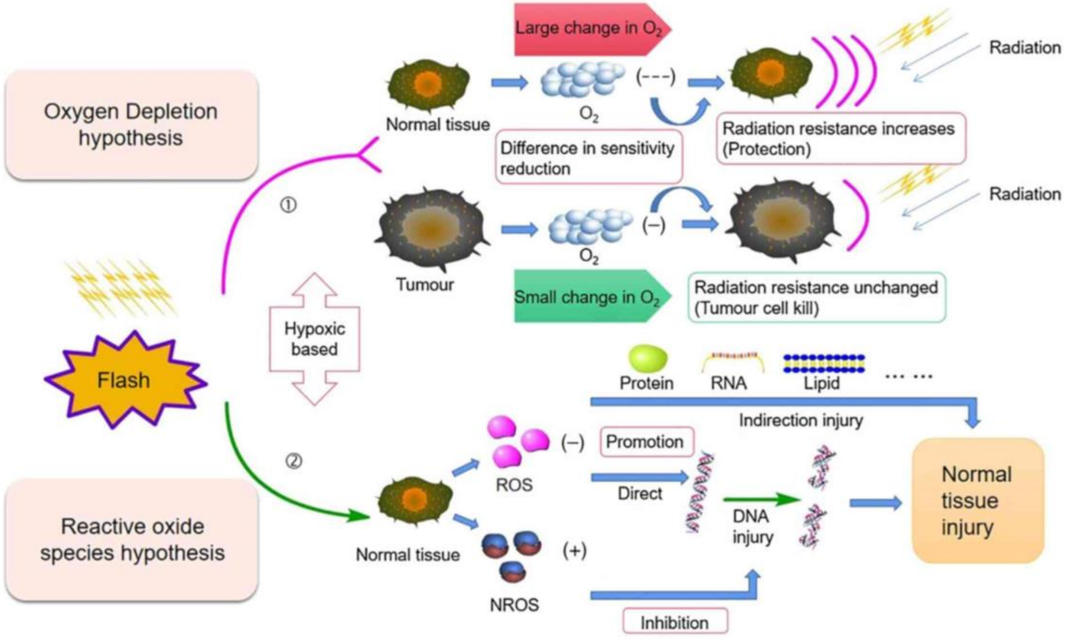

- Oxygen consumption hypothesis: High-dose transient irradiation reduces the presence of oxygen, and this effect is greater on normal cells, resulting in stronger radiation resistance.

- Reactive Oxygen Species (ROS) levels that causes DNA, RNA, protein and lipid injury, and an increase in the protective non-reactive oxygen species (NROS) levels that inhibits DNA injury.

PHASER: A solution for FLASH-RT

- PHASER (pluridirectional high-energy agile scanning electronic radiotherapy)
  - 16 klystrinos power combined to drive a given linac with 5.3 MW of peak power
  - Switching between LINACs occurs at 300ns
- Understanding the accelerator
  - Modeling the power combining is challenging (compensating for phase and amplitude jitter in the klystrinos)
  - Different LINACs need to operate at different energies
  - Beam steering using magnets
- Patient treatment
  - Rapid computation of optimal dose
  - Compensation for breathing

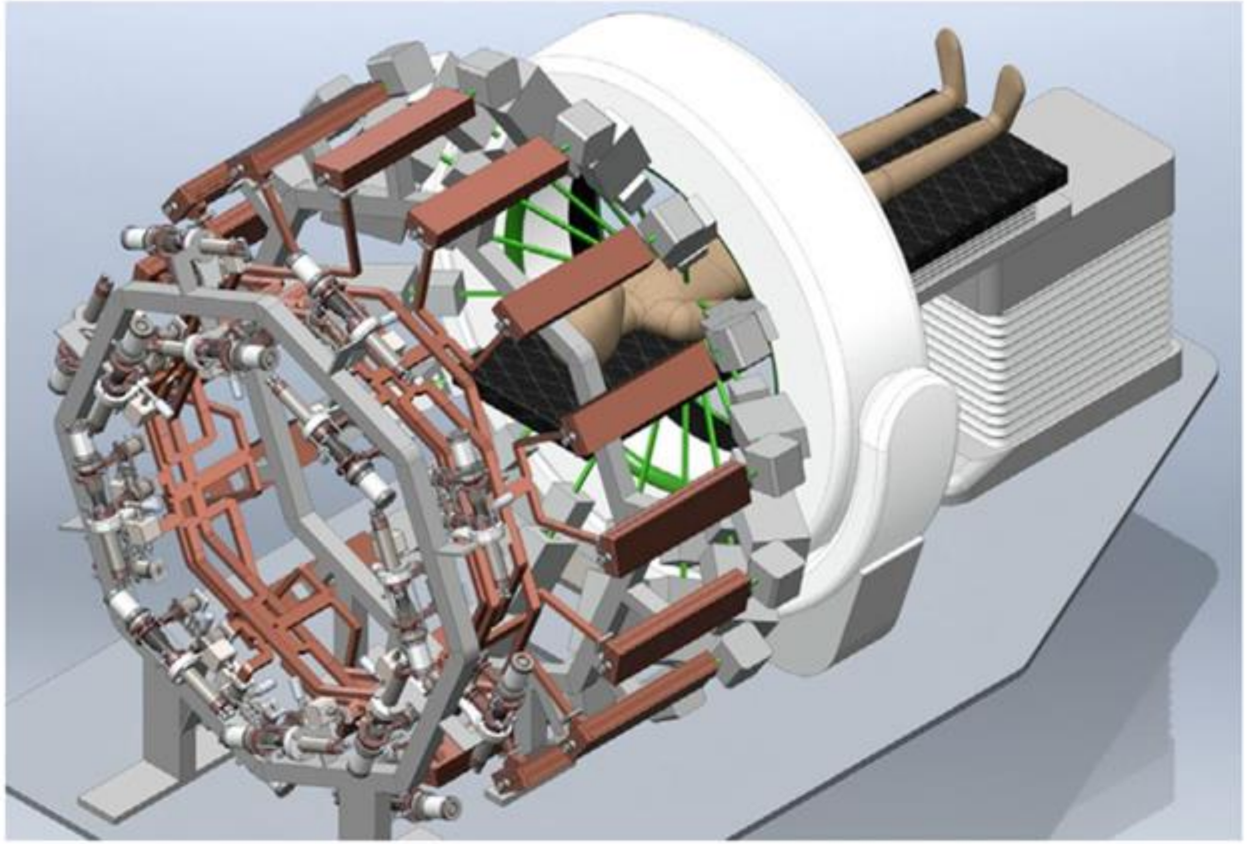

A Toy Model for A PHASER-like system

- Water phantom simulated in GEANT-4
  - Modeling 1-D energy loss / deposition
- Sixteen different x-ray sources with the ability to tune energy and steering to optimize the dose delivery profile
  - Toy model assumes a single target plane (2D)
  - Energy deposited in a water phantom
  - Energy range of the x-rays is 1-20 MeV
  - Can adjust commensurate with PHASER parameters
- Compute the energy loss as a function of position inside a water phantom
  - Scan energy then use interpolating function to generate a continuous control knob for the RL model
  - Dose computed for 10^6 x-rays – realistic beams would deliver ~10k times this dose

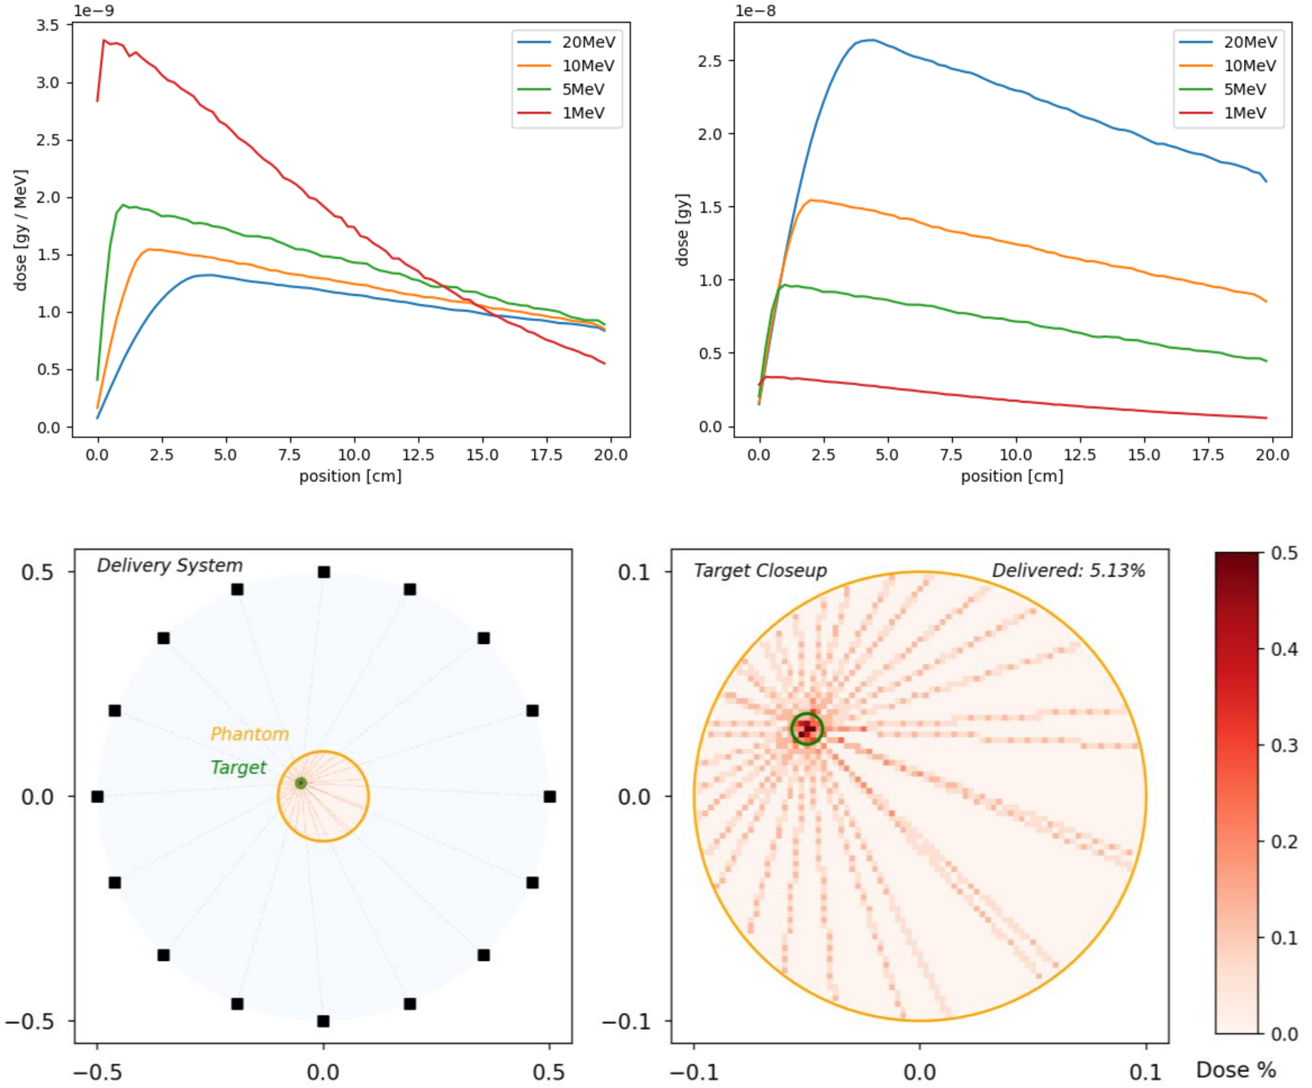

RL for Dose Optimization

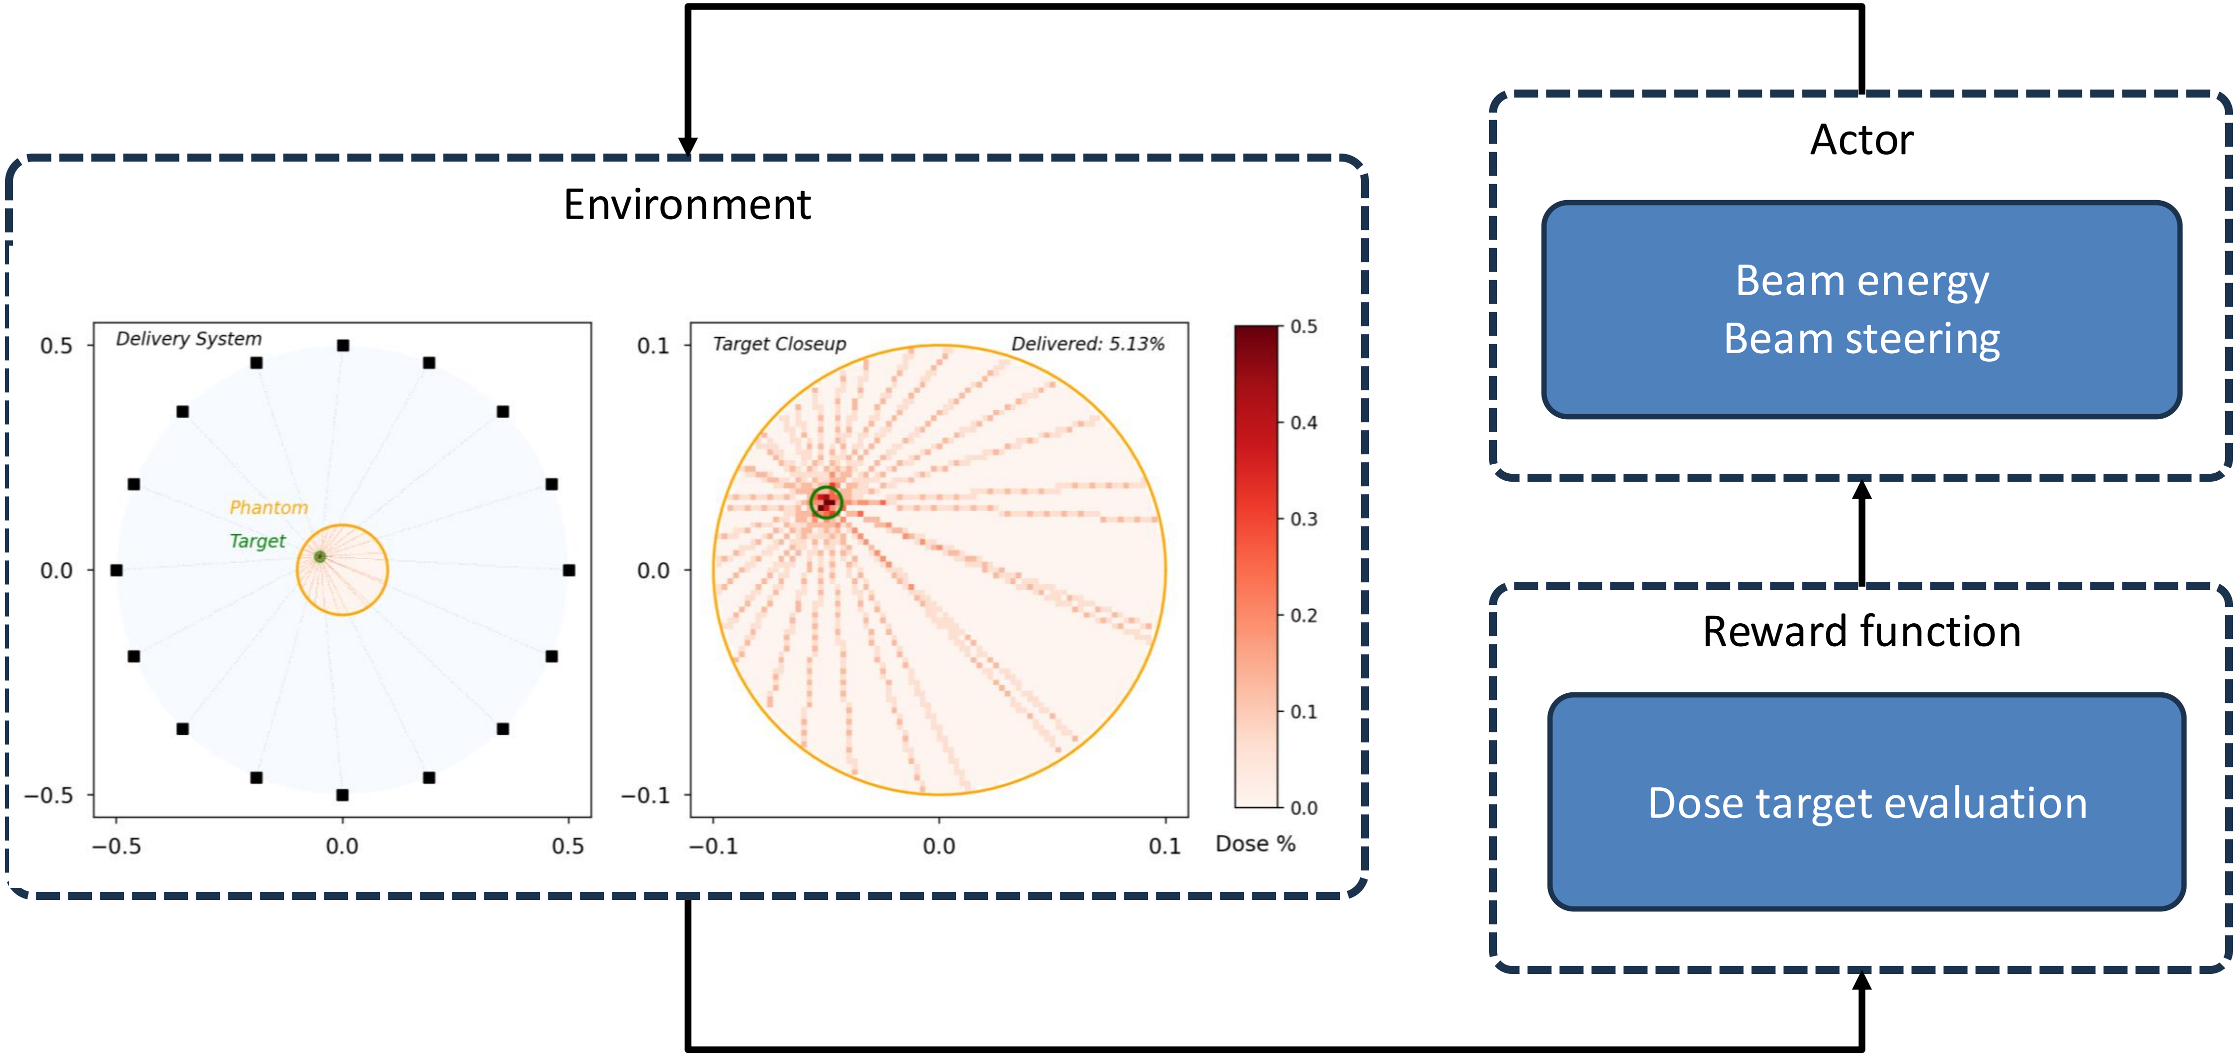

A Toy Model for A PHASER-like system

- Target defined as a circular region with a fixed center point
- Right – reward function, actor loss, and critic loss as a function of training iteration
- Bottom – dose delivery results
  - Random errors in the steering (left)
  - Correct beam steering (middle)

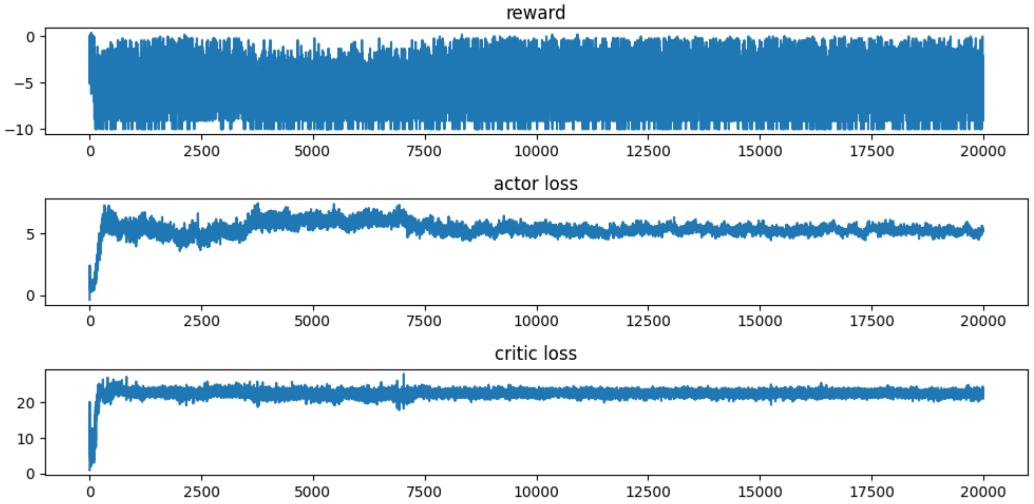

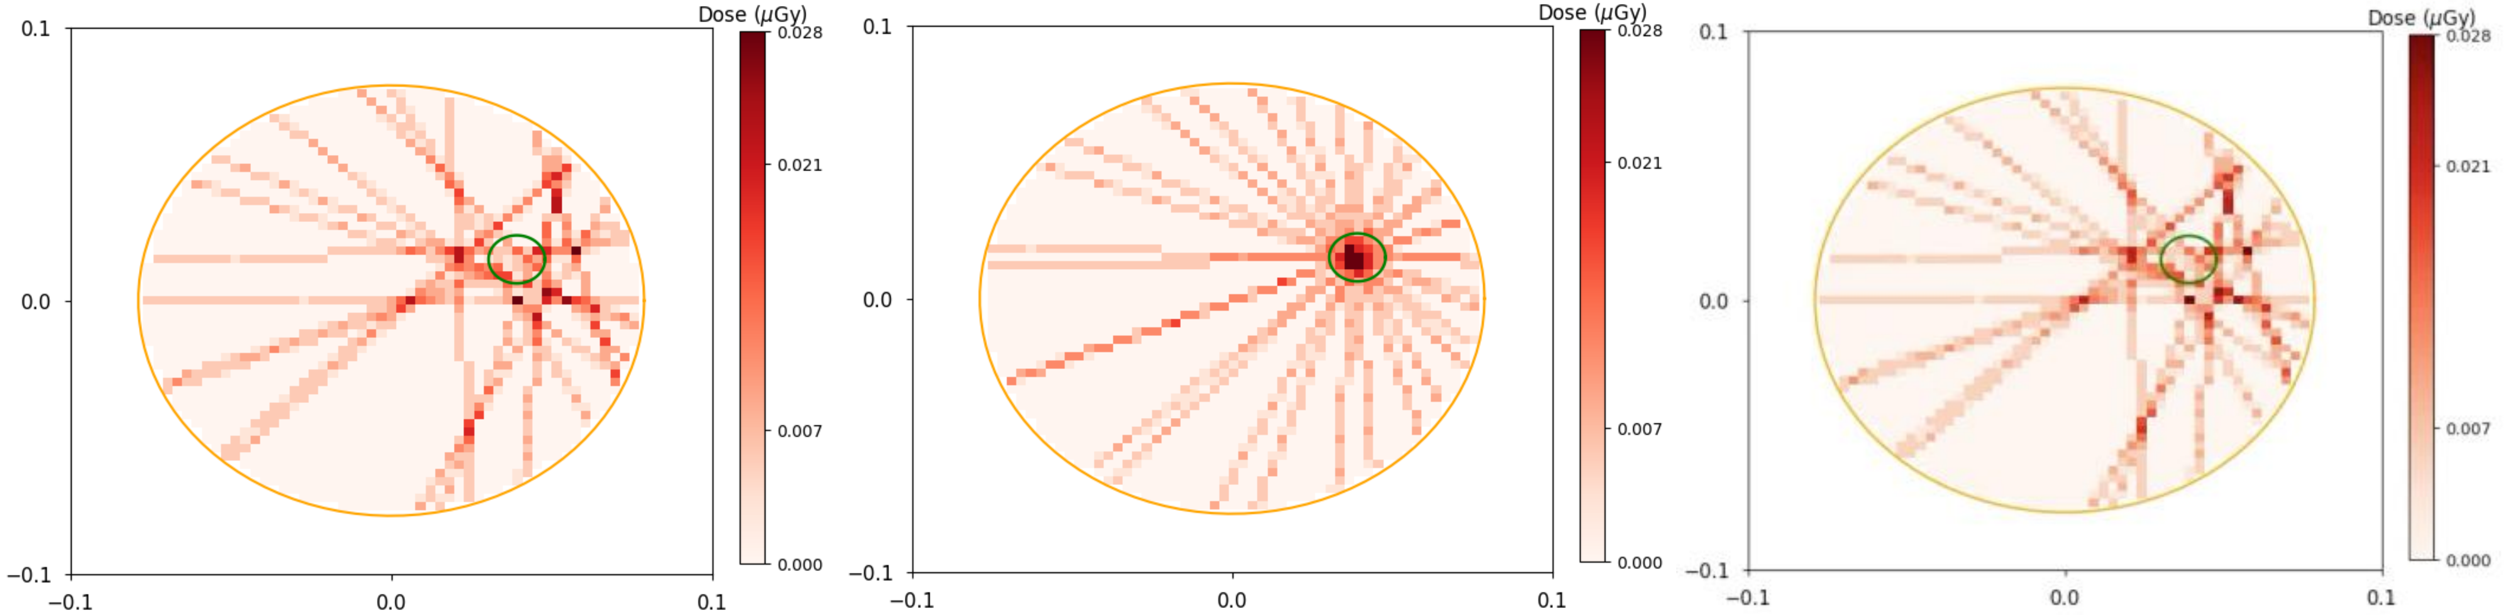

Applying RL to Industrial Accelerators

- Objective: Rapidly switch between machine states for different sample imaging requirements
  - Control parameters: RF power / frequency (energy), electron beam pulse format and current, solenoids
  - Measurables: beam transmission, beam size/energy
- Establish simulation model in SPIFFE that is representativ of the system
  - See poster by Finn O'Shea

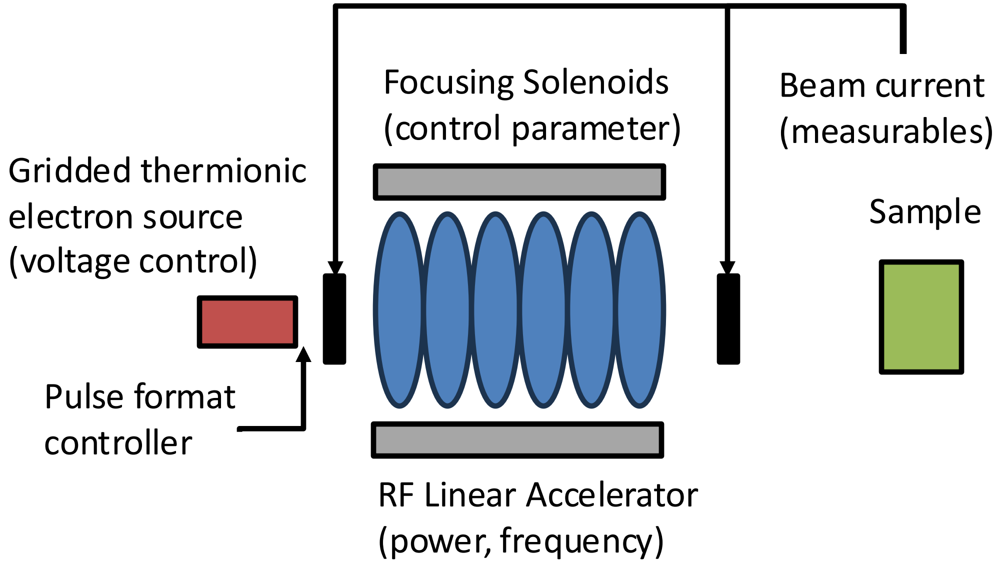

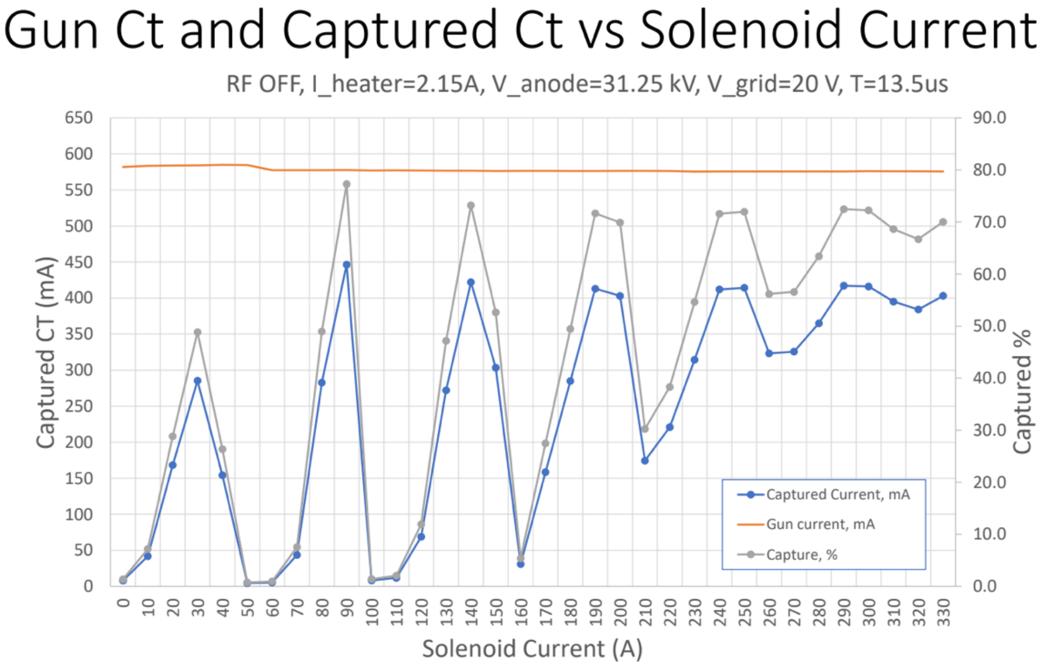

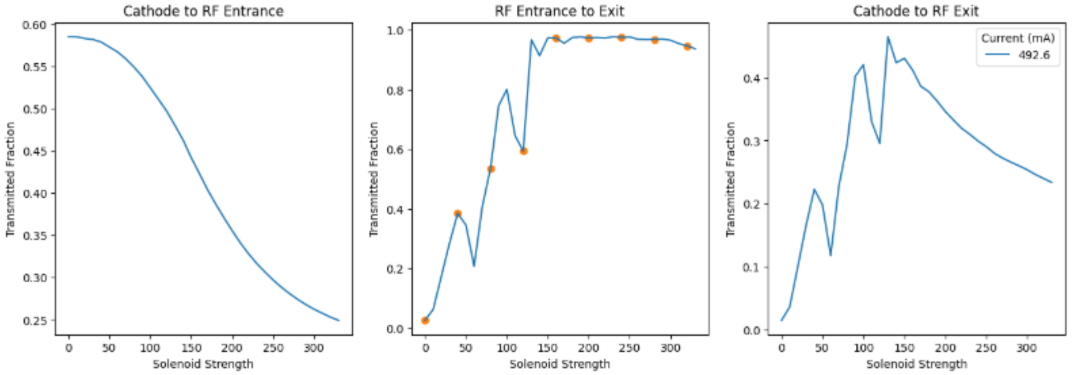

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.