This work was supported by the Department of Energy, Laboratory Directed Research and Development program at SLAC National Accelerator Laboratory, under contract DE-AC02-76SF00515

SLAC NATIONAL ACCELERATOR LABORATORY

Virtual Diagnostic For Longitudinal Phase Space Prediction at FACET-II

J. Park [1], C. Emma [1], A. Marinelli [1]

[1] SLAC National Accelerator Laboratory

Longitudinal Phase Space (LPS) Control at FACET-II

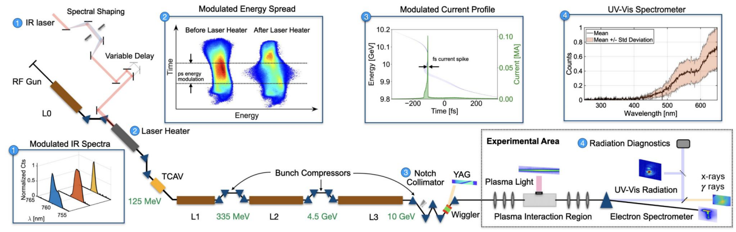

Experimental Generation of Extreme Electron Beams for Advanced Accelerator Applications, C. Emma et al, Phys. Rev. Lett. (2025)

At SLAC's Facility for Advanced Accelerator Experimental Tests II (FACET-II), precise control of the electron beam's Longitudinal Phase Space (LPS) is critical for driving advanced particle acceleration techniques. LPS represents the two-dimensional distribution of electrons in relative longitudinal position or time (ζ) versus energy (E).

High-gradient PWFA requires careful shaping of the drive bunch's longitudinal current profile to achieve **high peak currents** and maximize the transformer ratio. It also demands precise tuning of the **witness bunch's energy chirp** to compensate for wakefield slopes and preserve low slice energy spread.

However, LPS diagnostic tools, such as **Transverse Deflection Cavity** (TCAV), is destructive in nature, so cannot be deployed simultaneously with the plasma experiment to diagnose the characteristics of the beam that is used in interaction. This research aims to reconstruct shot-to-shot LPS from **Beam Synchronous Acquisition machine process variables** (BSA PV).

In FACET-II, the LPS is controlled with **four major knobs**: photocathode UV laser timing, IR laser heater timing at the end of the injector, the acceleration phase of L1 which accelerates the beam up to 335 MeV, and the acceleration phase of L2 which accelerates the beam up to 4.5 GeV.

Performance Benchmark

Example LPS and current profile output from VTCAV:

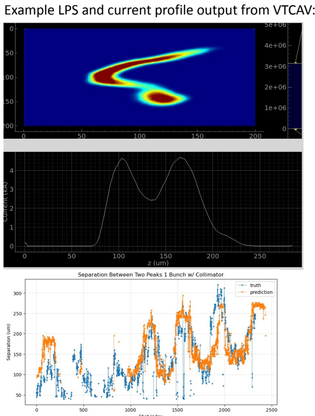

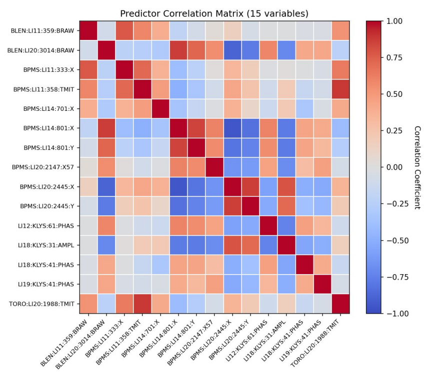

True vs Prediction Benchmark using Bunch Separation:

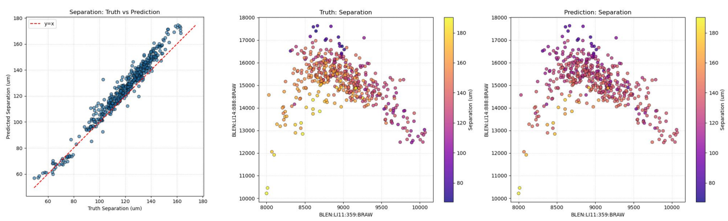

Conclusions

The virtual diagnostics can **predict the LPS from the live machine state, with some caveats**: its fidelity drops when operating outside of the training dataset. Because such anomalous machine states are also detrimental to experimental programs, a complementary **anomaly detection** algorithm, which is derived directly from the inference model, is integrated to detect when prediction confidence is low, so operators can intervene to turn the machine back to the initial state.

ML Architecture

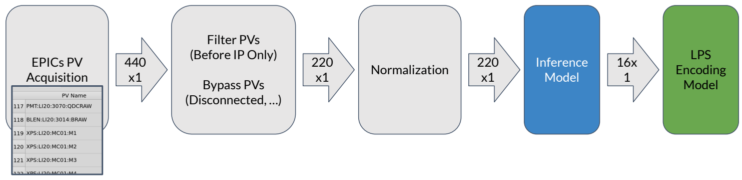

PV acquisition
Each shot begins with a beam-synchronous read of the relevant scalars. BSA-eligible signals are pulled through a synchronous buffer that guarantees the readings share a common pulse ID, while slower non-BSA process variables (klystron settings, control-room readbacks) are fetched directly and merged into the same vector.

Regression from PVs to latent parameters
The standardized, importance-weighted vector is fed into the trained regressor, which maps the machine state onto a compact set of latent coordinates. The regressor output is then passed back through the target-space scaler to recover the latents in their native, physical units before they are handed off to the image model.

LPS Encoding model reconstruction
The recovered latents are pushed through the image model's decoder, which turns them into a two-dimensional LPS distribution on the diagnostic screen coordinates and normalizes the result by the measured toroid charge so the intensity carries physical units of current per longitudinal bin.

Anomaly Detection

Only a **small subset of the roughly 1,000 PVs** recorded at FACET-II directly affects the LPS. During operation, frequent machine reconfigurations make tracking individual PV changes difficult and create cognitive overload for operators. In practice, the LPS is primarily governed by the chirp and compression chain, including the **L1 and L2 RF phases as well as the BC11 and BC14 chicane orbits**.

Unexpected behavior or hardware failures within this compression chain, which happens daily, severely disrupt downstream user experiments. At the same time, these unmodeled states push the machine outside the training domain of the VTCAV model, causing **prediction accuracy to drop, as shown below**. Because anomalous states degrade both beam quality and model reliability, a real-time anomaly metric is necessary.

To address this, the VTCAV diagnostic continuously calculates the topological distance between the live machine state and the historical training dataset. By evaluating this distance across multiple algorithms, the system quantifies prediction confidence and provides operators with immediate anomaly alerts.

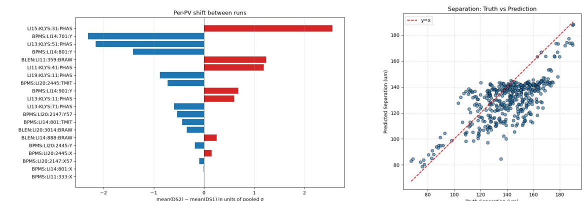

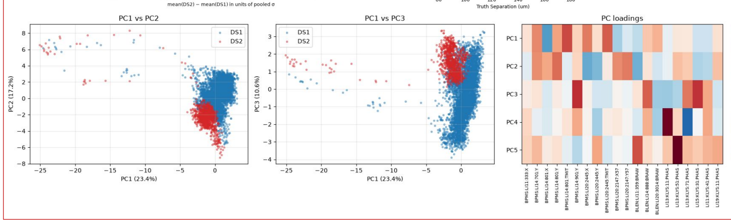

Application / Future Direction

- Photon Science
  - Plasma Attosecond X-ray source
  - Attosecond XFEL in LCLS
- Agentic AI Autonomous Operation
- Digital Twin
  - Run real-time BMAD simulations for comparison
- Inference Model Improvement
  - Linear Forest Regressor
  - Orbit Parametrization