Longitudinal Phase Space Tomography via Machine Learning at the Los Alamos Neutron Science Center

Christopher Leon[1*], Alexander Scheinker[1] and Petr M. Anisimov[1]

[1*]Los Alamos National Laboratory, Los Alamos, NM, 87545, USA.

*Corresponding author(s). E-mail(s): cleon@lanl.gov;

Supplementary Note: Full-Dataset Test Set Examples

Supplementary Figures S1 and S2 show examples of model predictions on the test set.

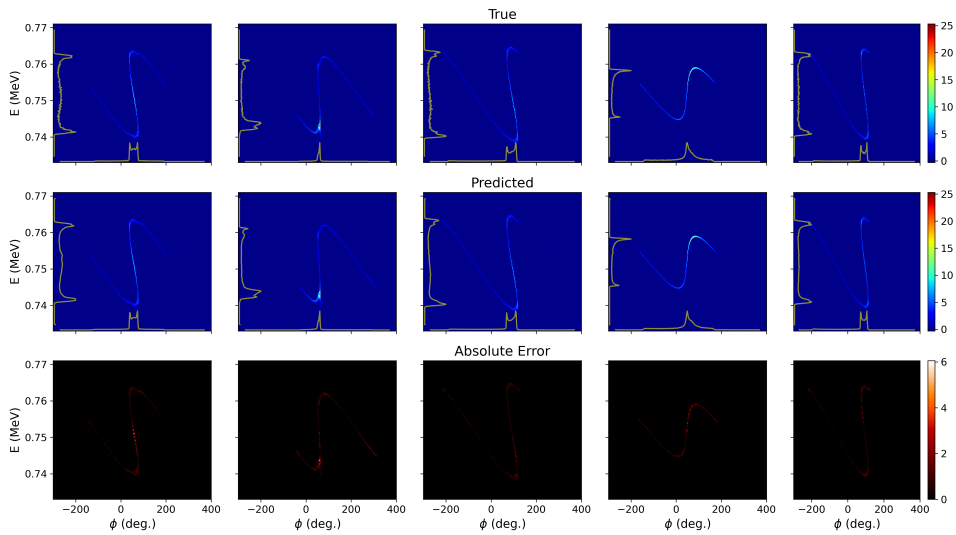

Figure S1 U-Net-Single-Full on Test Data – comparison of true LPS vs. prediction for some samples in the test set. The 1D projection curves for E and ϕ are shown in yellow.

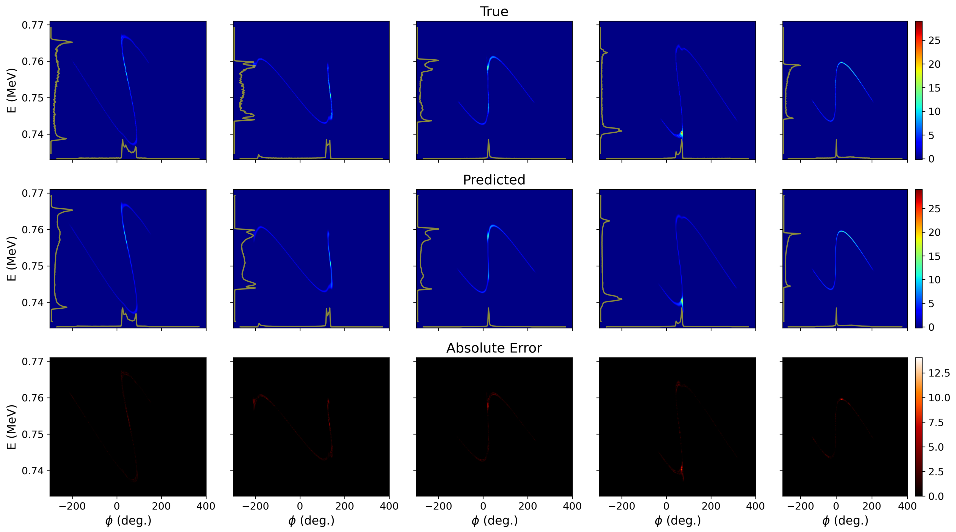

Figure S2 U-Net-Comp-Full on Test Data – comparison of true LPS vs. prediction for some samples in the test set. The 1D projection curves for E and ϕ are shown in yellow.

Supplementary Note: Upsampling – Model Exploration

Two models for the upsampling U-Net were investigated, both depicted in Supplementary Fig. S3. The first has a non-trainable residual connection that performs bilinear interpolation. The input is also passed through a subnetwork that acts as a perturbative corrector residual connection. The second architecture starts with the bilinear interpolation as an initial layer, whose output is then passed through a U-Net. The number of filters in each model was chosen so that both models had approximately 4.3 M parameters.

We required both upsampling models to significantly outperform simpler baseline algorithms. As baselines, we compared them with bilinear and bicubic interpolation. The test-set error is shown in Supplementary Fig. S4. Both models significantly outperform the baseline cases, while the U-Net without the residual upsampler does slightly better than the U-Net with the residual upsampler.

To give a qualitative sense of upsampling performance, Supplementary Fig. S5 shows an example of upsampling from an element in the test set. The first row shows the true distribution in the test set. The second row shows upsampling by the simple algorithms of bilinear and bicubic interpolation. The third row shows predictions from the two upsampling models.

The models were then combined with the initial U-Net in the composite model and fine-tuned. The results of the models' performance on Full-Dataset-Test can be seen in Supplementary Fig. S6. The U-Net without the residual performs significantly better.

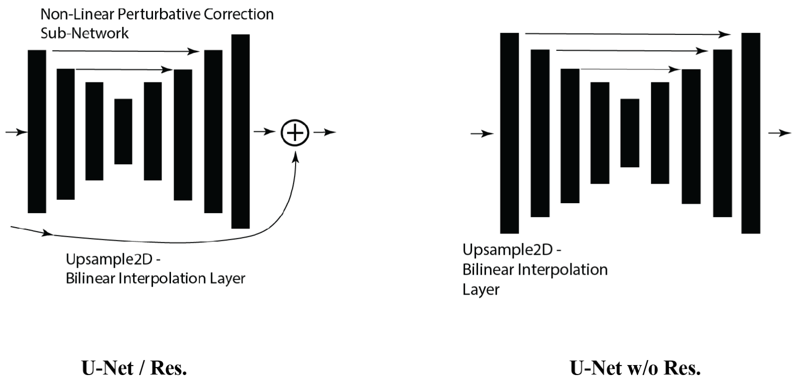

Figure S3 The two architectures explored for the upsampling task in the composite model.

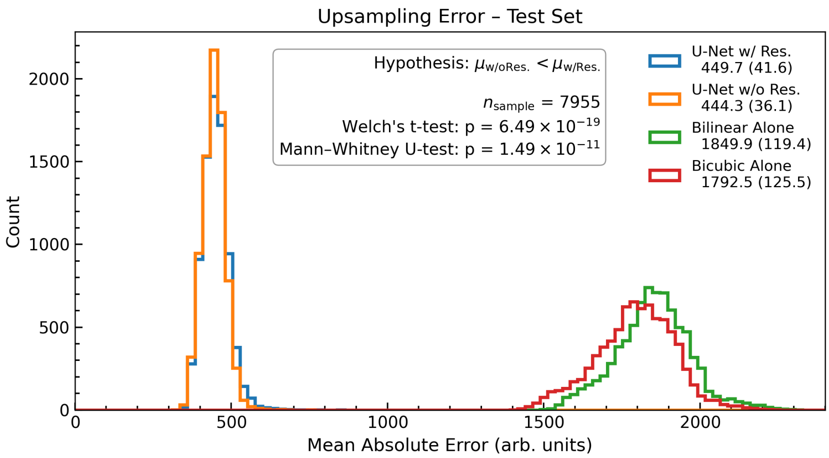

Figure S4 Error histogram for the test set. The mean (standard deviation) of the model MAE are shown in the legend.

Supplementary Note: ML LPS – Removing Background Noise Via Cutoff

The ML predictions exhibit a low-amplitude background away from the main beam distribution. Supplementary Fig. S7 shows the effect of applying a cutoff to the ML predictions obtained from the filtered experimental 2DPS data. After the cutoff is applied, the predicted LPS distributions more closely resemble the calibrated simulation LPS and are more physically plausible. The cutoff value was selected empirically to suppress the low-amplitude background while preserving the main features of the predicted LPS distribution.

We tested the sensitivity of the ensemble central values to the choice of LPS cutoff, as depicted in Supplementary Fig. S8. The energy-related metrics showed very weak cutoff dependence because the relevant energy range is small. In particular, the mean

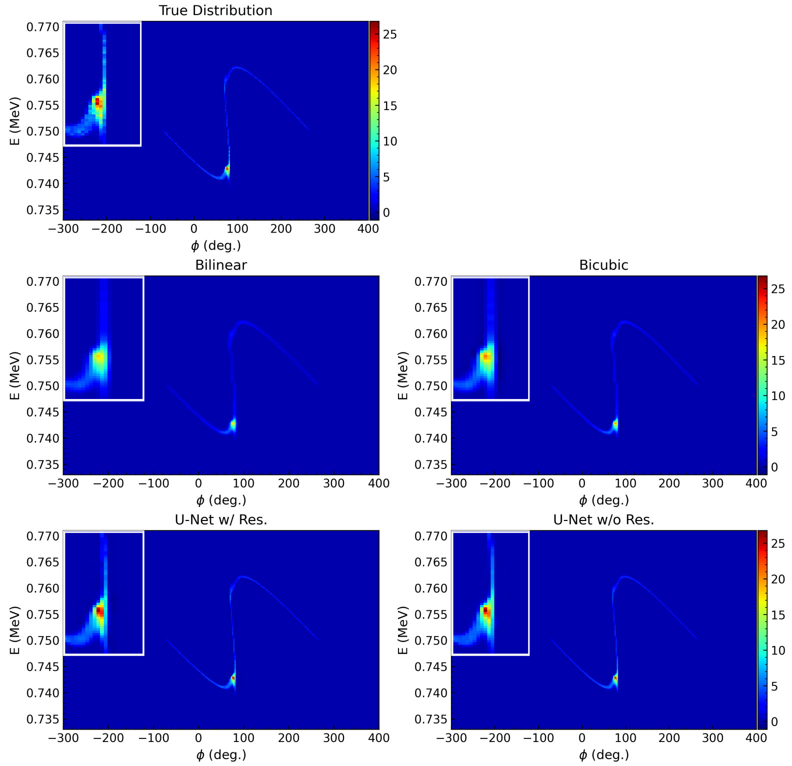

Figure S5 Test Set – example of upsampling. The insets on the upper left of the subplots are ×3.0 zooms on the beam core.

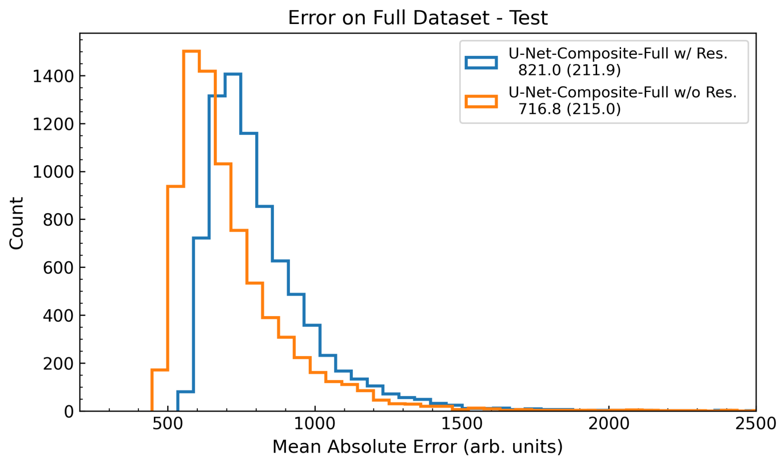

Figure S6 Error histogram on Full-Dataset test set for two composite models. The mean (standard deviation) of the model MAE are shown in the legend.

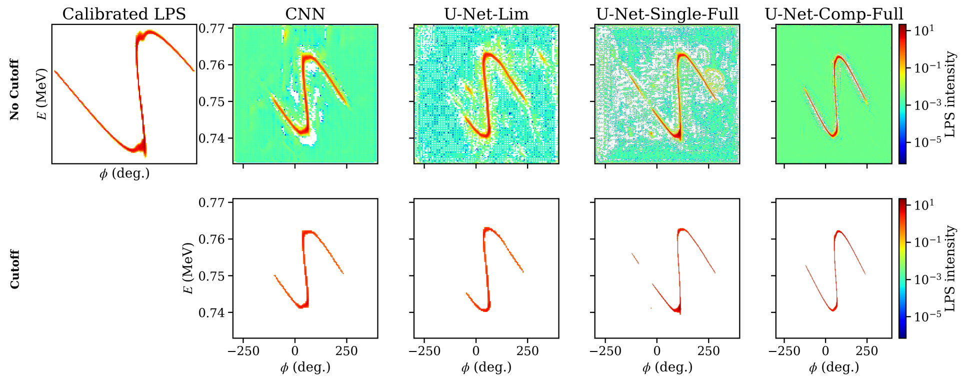

Figure S7 Top row: calibrated LPS and ML-predicted LPS distributions before applying the cutoff. Bottom row: corresponding ML-predicted LPS distributions after applying the cutoff. The color scale is logarithmic; white regions indicate zero-valued bins.

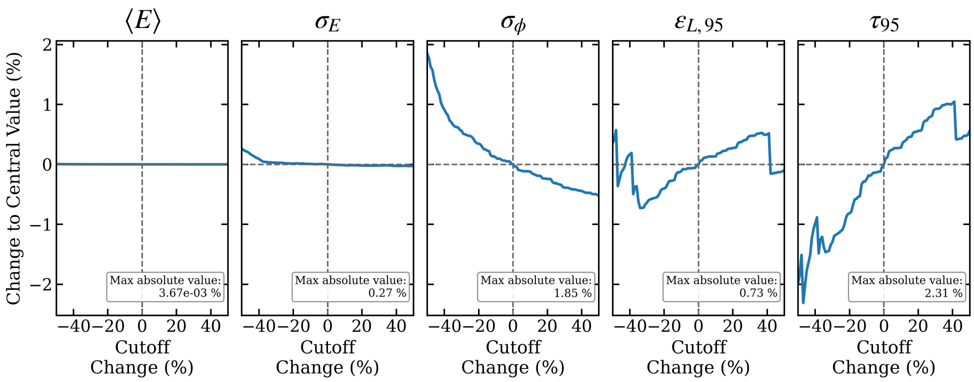

Figure S8 Dependence of the ensemble metric central values on the LPS cutoff.

energy was largely unaffected because the low-amplitude background is approximately uniform, with its contributions mostly canceling during averaging. For the standard deviations, the spreads decrease with an increasing cutoff as expected. Nonetheless, for all metrics, the maximum absolute change induced by varying the cutoff was smaller than the variation among the models in the ensemble.