



Uncertainty Aware Anomaly Detection to Predict Errant Beam Pulses and GradCAM Analysis

Kishansingh Rajput

Kishan@jlab.org

On behalf of Errant Beam Prediction Task
Force (Jefferson Lab and ORNL)

ORNL is managed by UT-Battelle, LLC for the US Department of Energy
JLab is managed by Jefferson Sc. Assoc., LLC for the US Department of Energy



U.S. DEPARTMENT OF
ENERGY

Outline

- Uncertainty Quantification in Deep Learning
- Errant Beam Prediction at SNS Accelerator
 - Introduction
 - Data Processing
 - Siamese Model
 - Uncertainty Aware Siamese Classifier
 - Equipment Fault Classification
 - Future Roadmap
- Uncertainty Aware Booster Surrogate for FNAL
 - Uncertainty Aware Deep Regression with single inference

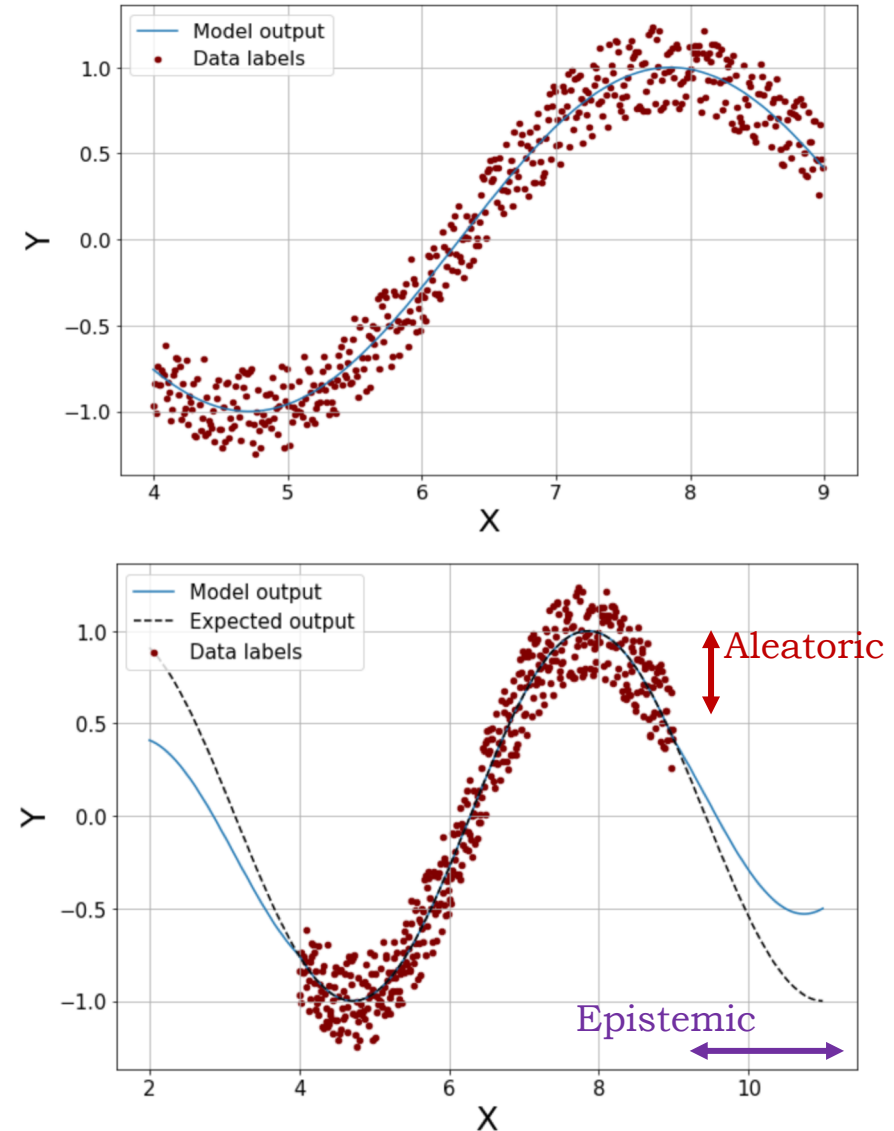
Uncertainty Quantification

- Deep Learning (DL) models are deterministic transformation functions from an input to the output
- DL models are very powerful and expressive
- It is important to know the confidence associated with each prediction from a DL models for decision making




Uncertainty Types: Aleatoric vs Epistemic uncertainties

- Aleatoric → Data uncertainties
- Epistemic → Out of training distribution uncertainty (OOD)




Popular methods for UQ in DL




(a) MC Dropout

Use dropout during inference to create variability in the prediction which can be used to estimate the uncertainty. Requires offline calibration




(c) Ensemble

Create multiple copies of the same model architecture trained with different parameters initialization. Requires calibration after training




(d) Quantile Regression

Model is trained to predict quantiles for the regression problem



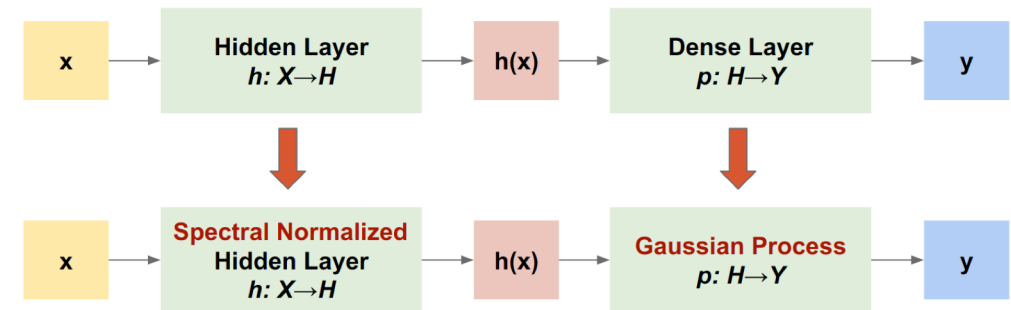

Popular methods for UQ in DL

- Unfortunately, majority UQ methods for DL do not account for OOD uncertainty
- This is critical in optimization or control problems
- For example, different methods yield vastly different uncertainty estimation
 - Deterministic (Prediction value)
 - MC Dropout
 - Deep Ensemble
 - Gaussian Processes (GP)



GP for UQ in DL

- GP transforms the input space into a higher dimensional space with the help of a kernel
- The inferences are based on the distance measure between different input samples
- This allows GP to intrinsically provide uncertainty estimates including OOD
- GP is limited in terms of Scaling and data reduction techniques are usually required for large data sets
- Recent study presented a way to introduce Gaussian Process approximation within a neural network
- This allows to use highly expressive deep networks and provide uncertainty estimation



Spectral Neural Gaussian Process*

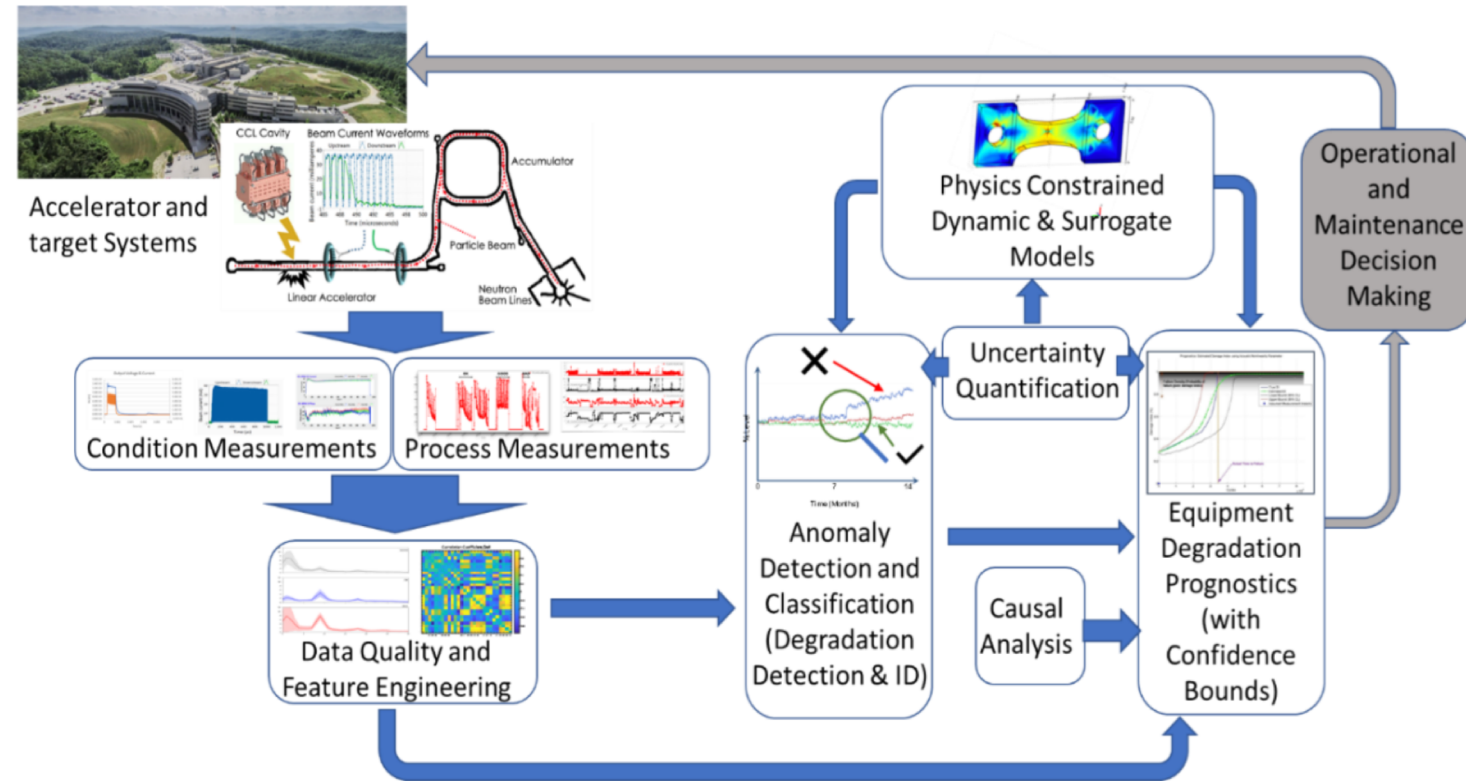
* <https://arxiv.org/abs/2006.10108>

Outline

- Uncertainty Quantification in Deep Learning
- Errant Beam Prediction at SNS Accelerator
 - Introduction
 - Data Processing
 - Siamese Model
 - Uncertainty Aware Siamese Classifier
 - Equipment Fault Classification
 - Future Roadmap
- Uncertainty Aware Booster Surrogate for FNAL
 - Uncertainty Aware Deep Regression with single inference

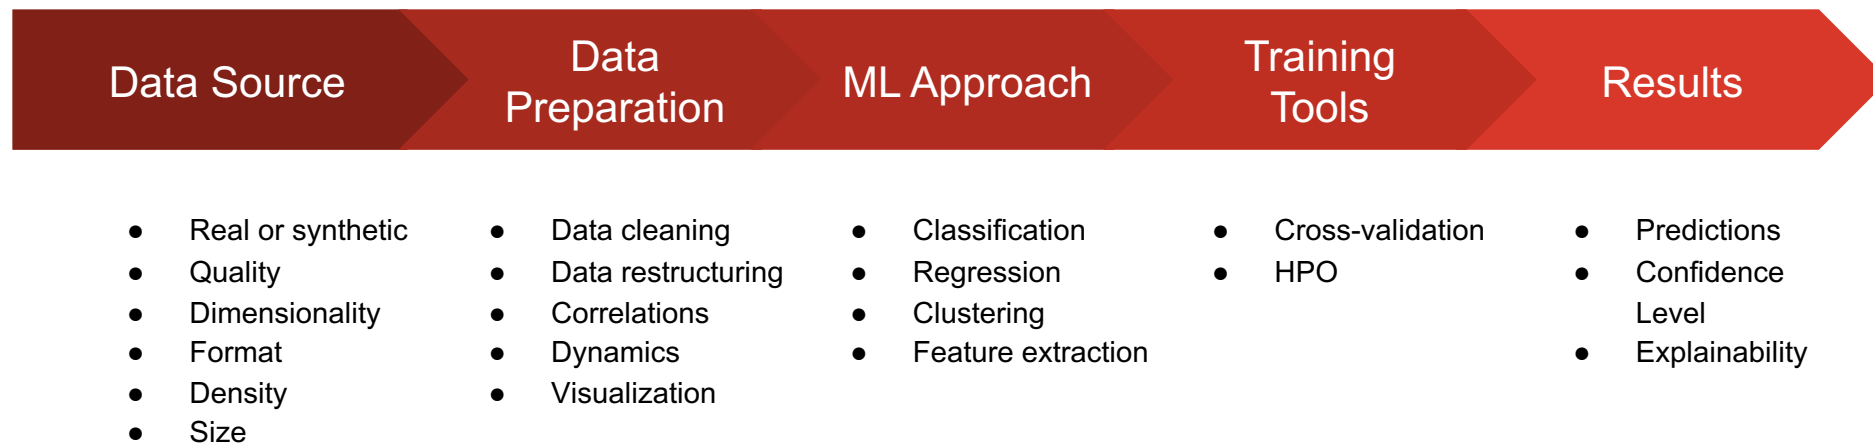
Errant Beam Prediction for SNS Accelerator

- Spallation Neutron Source (SNS) accelerator at ORNL delivers 1.4 MW of a 1 GeV pulsed beam at 60 Hz
- Ongoing work to **predict errant beam pulses** as well as equipment degradation and prognostics
- Continuous data collection is done by **Differential Current Monitor (DCM)**, Beam Position Monitor (BPM) etc.
- Errant beam prediction on one pulse before it occurs to potentially avoid it



Errant Beam Prediction for SNS Accelerator

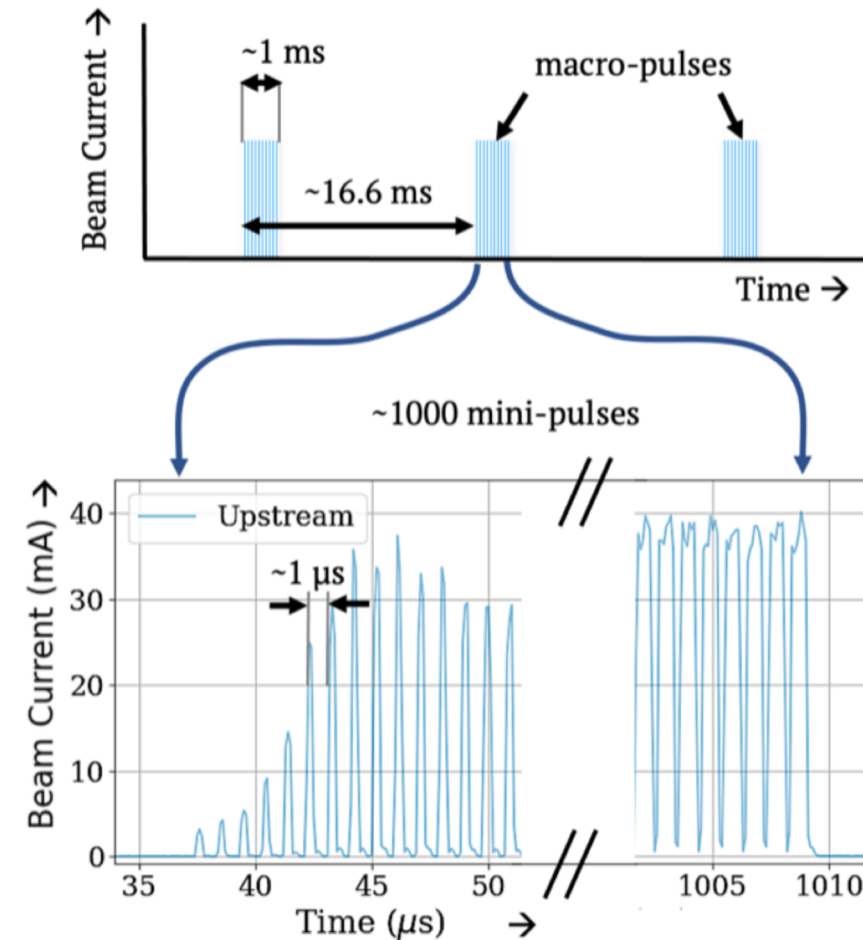
- **Goal:** Predict an upcoming machine trip before it occurs to potentially allow the crew to change settings to avoid it
- **How:** We use pulses leading to a trip (tagged "Before") and identify features that indicate an upcoming failure
- **Data science pipeline used:**



Data Collection and Preparation

How was the data collected and labeled?

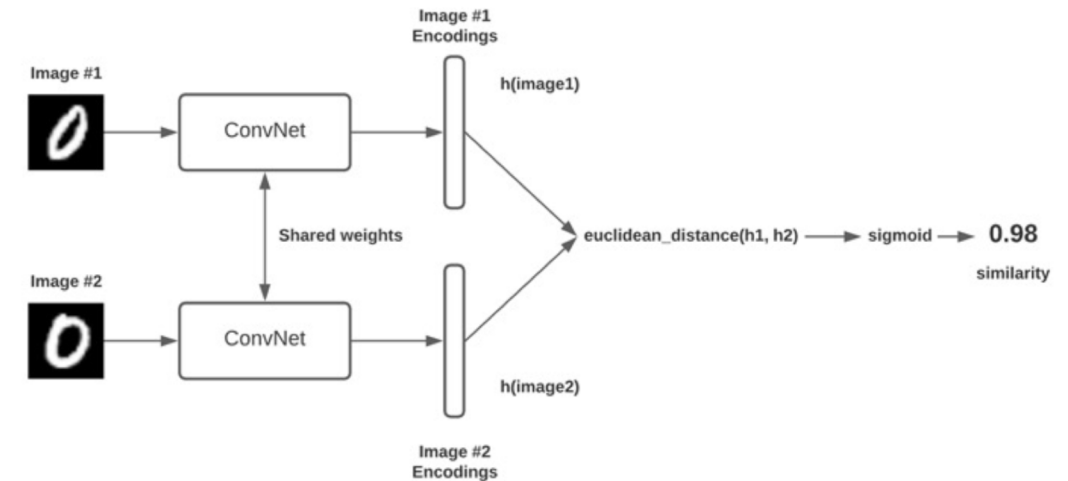
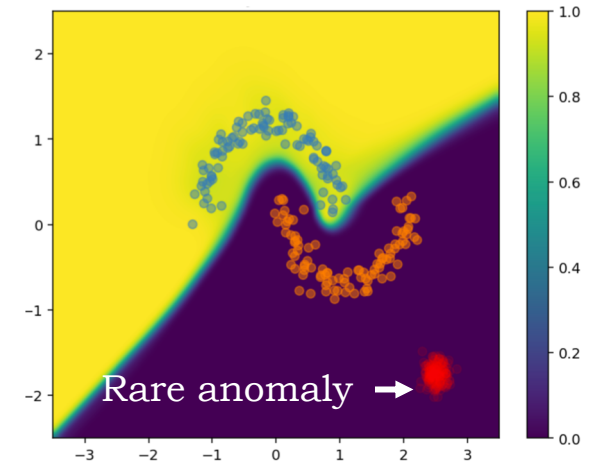
- DCM creates a series of pulses (“macro-pulses”) with each macro-pulse composed of ~1k mini-pulses
 - An errant-beam data file is composed of 25 “good” macro-pulses followed by the errant beam pulse
 - A “normal” data file has no errant beam pulse
- We used the **macro-pulse before the errant beam pulse (and labeled it as anomaly)** and **macro-pulses from the normal file (and labeled them as normal)** for our studies
- Our hypothesis: there is a sign about upcoming anomaly in macro-pulses even before it happens
- We also need to forecast the fault within a short time window to be actionable
- Samples were divided into 3 orthogonal dataset:
 - Train (64%)/Test(20%)/Validation(16%)



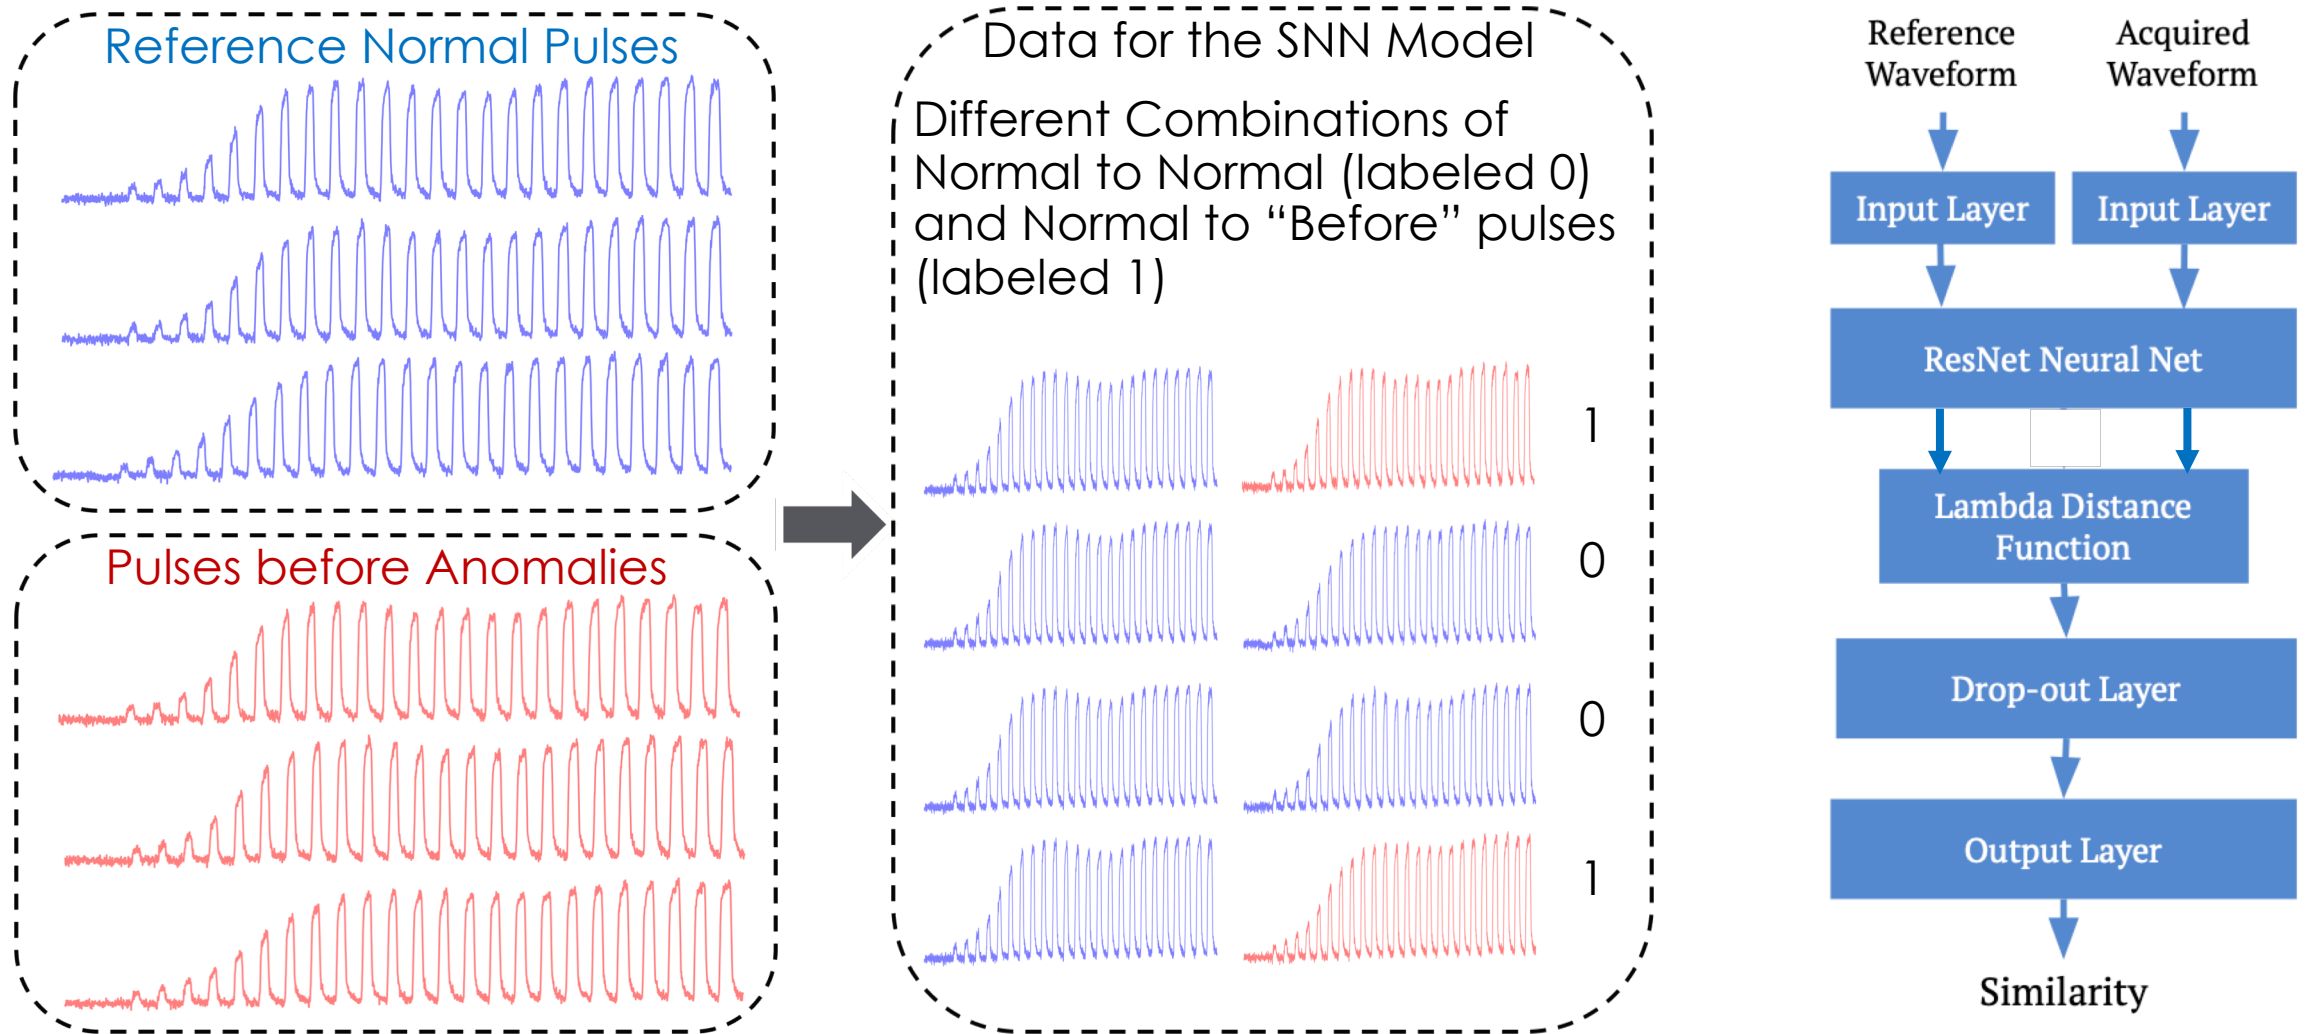
Siamese Neural Network (SNN) Model

Traditional classification models vs Siamese model

- Traditional DL classification models fails to identify unseen anomalies (OOD)
- Similarity based models can correctly classify unseen anomalies. Ex Siamese model
- Siamese model does not explicitly model the classification but focuses on the similarities
- It learns twin embedding models to transform inputs into a latent space
- Distance measures are applied at latent space to compute the similarity



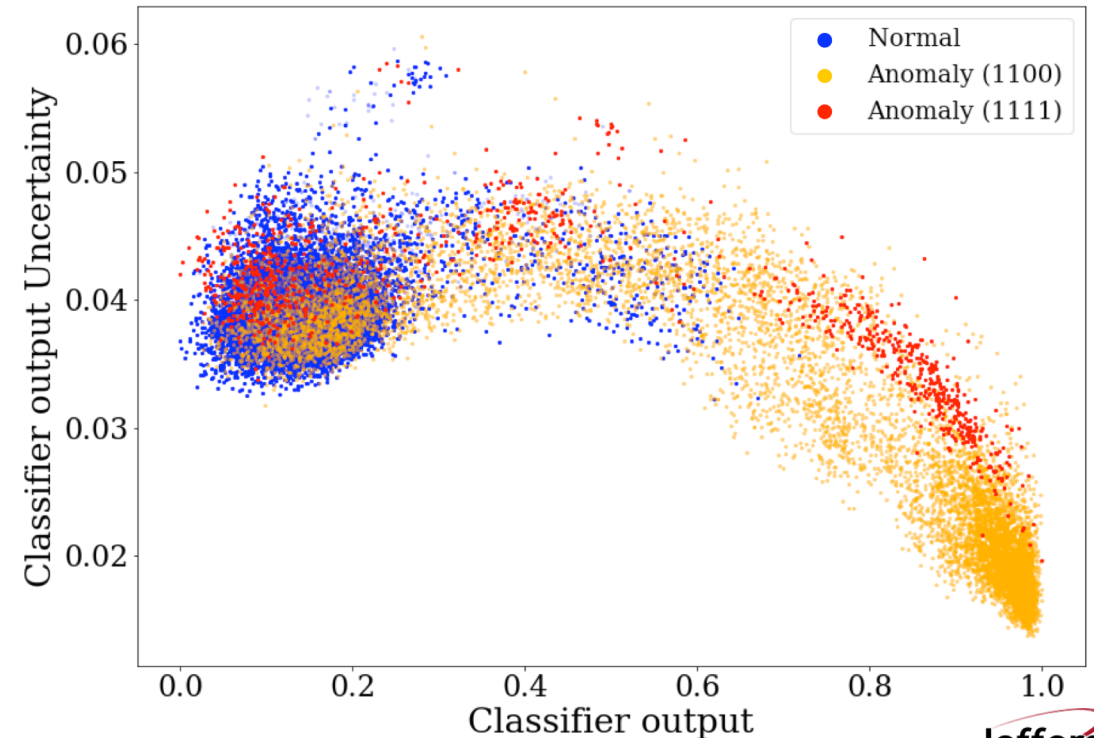
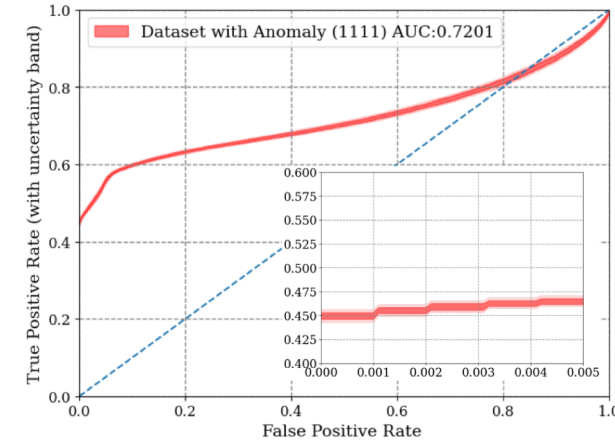

Data Preparation for SNN Model Training



Uncertainty aware Siamese model


- We enhanced our Siamese model by adding GP layer providing an uncertainty estimate
- Results from similarity model showed a ~4x improvement in performance over previously published results, it is also much better than a vanilla Auto-encoder
- The ROC curves shows true fault detection rate above 60% while keeping the false alarms below 0.5% (not optimized)
- We introduced an out-of-domain anomaly, labelled 1111 (red), the UQ-based model performed similar in classifying the anomalies and indicated high uncertainty (as expected)

After a fault is predicted, is it possible to associate with a particular equipment failure?




Gradient Class Activation Mapping (GradCAM)

- GradCAM provides mapping between the the model output to the features in the input that the model thinks are the most relevant
- Extracts the most active features in the last convolutional layer and maps them back to the input



Original Image



Grad-CAM 'Cat'

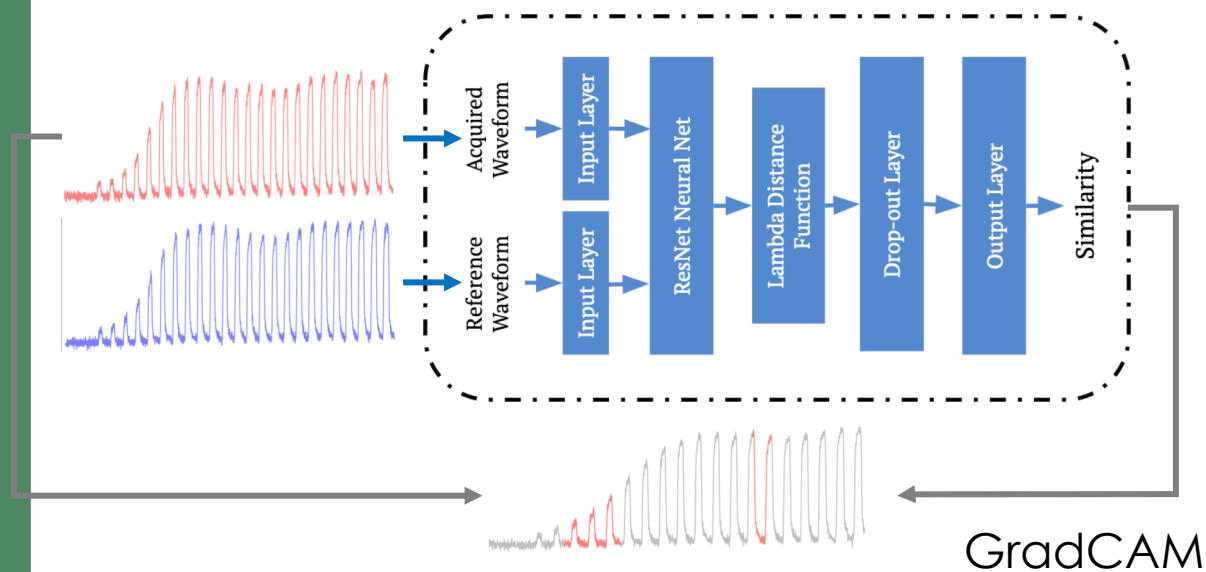


Grad-CAM 'Dog'


<https://arxiv.org/abs/1610.02391>

Equipment Fault Classification using GradCAM

- Applied GradCAM on SNN model trained to predict Errant Beam Pulses
- It identified sections of the waveform most relevant for a particular decision from the model




Stacked Relevance Vectors ("Before" Pulses)



Equipment Fault Classification using GradCAM

- The salient feature vectors are reduced to 2-dimensional space using UMAP*
- Studying how the cluster location from anomalies relate to specific equipment failures



* <https://arxiv.org/abs/1802.03426>

Path Forward

- Currently deployed deterministic SNN model will be replaced by Uncertainty Aware SNN
- Continue the study on equipment fault classification via GradCAM and SNN model
- Because SNN focuses on similarity, it is sensitive to changes in the beam configuration
- Ongoing work to add beam configuration to the SNN model as conditional inputs

Outline


- Uncertainty Quantification in Deep Learning
- Errant Beam Prediction at SNS Accelerator
 - Introduction
 - Data Processing
 - Siamese Model
 - Uncertainty Aware Siamese Classifier
 - Equipment Fault Classification
 - Future Roadmap
- Uncertainty Aware Booster Surrogate for FNAL
 - Uncertainty Aware Deep Regression with single inference

Uncertainty Aware Booster Surrogate


Aim:

Reduce beam losses in the FNAL Booster by developing a Machine Learning (ML) model that provides an optimal set of actions for GMPS regulator

FNAL Accelerator Complex:



Courtesy: Christian Herwig



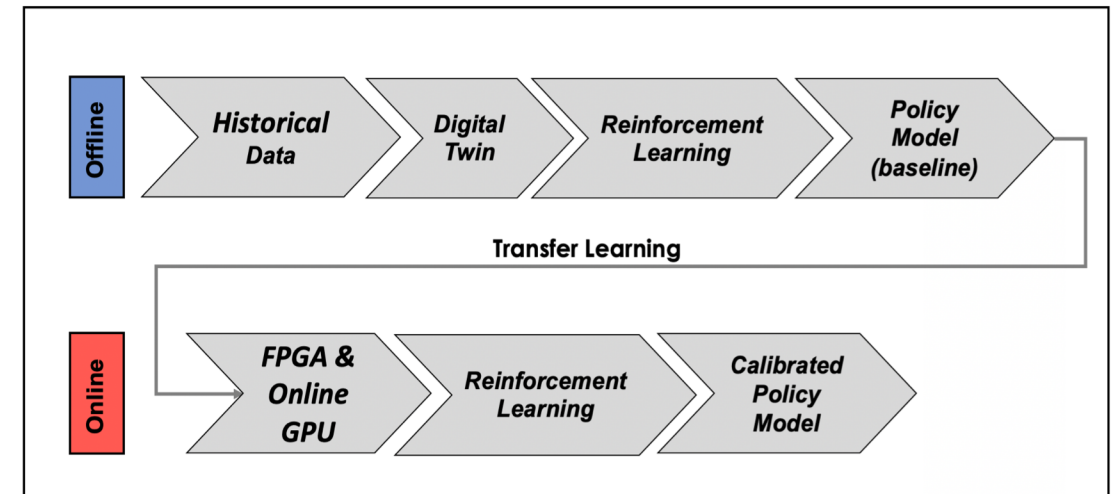
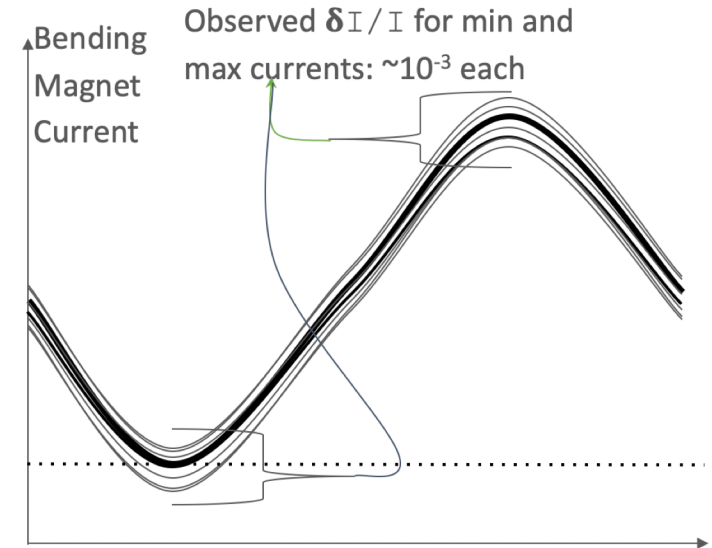
The Booster receives the 400 MeV (kinetic energy) beam from the Linac

It is then accelerated to 8 GeV with the help of booster cavities and Combined-function bending and focusing electromagnets known as gradient magnets.

These magnets are powered by the gradient magnet power supply (GMPS)

Reinforcement Learning for Booster Control

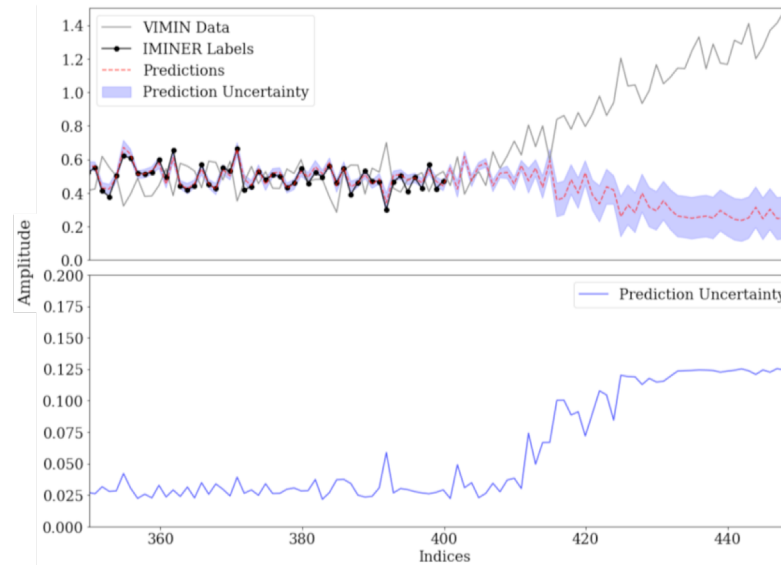

- Other high-current, high-power electrical loads near GMPS varies in time
- Causing unwanted fluctuations of the actual GMPS electrical current and thus fluctuations of the magnetic field in the Booster gradient magnets
- This spread in B-field degrades the beam quality
- A GMPS regulator is required to calculate and apply small compensating offsets in the GMPS driving signal
- Use of RL to improve the existing PID based regulator
- Policy model is focused on controlling the regulator to reduce the error
- **This invokes a need of Surrogate model to build the RL environment**



Uncertainty Aware DL Regression Model

Why Care so much about Uncertainty in Digital Model?

- Uncertainty Quantification can help determine how well a region of a phase space is modeled by the surrogate
- RL agent might spend most of the time in unstable region (might not be able to avoid local minima) or learn suboptimal policy, UQ can help in guiding the RL agent training
- We developed Deep Gaussian Process Approximation (DGPA) method to quantify the regression uncertainties for a DL model
- Unlike other methods, DGPA does not require multiple inferences and does not require offline calibrations making it easy to deploy in online settings



Thank You!

References:

Uncertainty aware anomaly detection to predict errant beam pulses in the SNS accelerator

W Blokland, K Rajput, M Schram, T Jeske, P Ramuhalli, C Peters, Y Yucesan, A Zhukov... - <https://arxiv.org/abs/2110.12006>

Uncertainty Aware ML-based surrogate models for particle accelerators: A Study at the Fermilab Booster Accelerator Complex

M Schram, K Rajput, K Somayaji, P Li, JS John... - <https://arxiv.org/abs/2209.07458>

Umap: Uniform manifold approximation and projection for dimension reduction

L McInnes, J Healy, J Melville - arXiv preprint arXiv:1802.03426, 2018 - <https://arxiv.org/abs/1802.03426>

Grad-cam: Visual explanations from deep networks via gradient-based localization

RR Selvaraju, M Cogswell, A Das... - Proceedings of the ..., 2017 - openaccess.thecvf.com

Simple and principled uncertainty estimation with deterministic deep learning via distance awareness

J Liu, Z Lin, S Padhy, D Tran... - Advances in ..., 2020 - proceedings.neurips.cc

Real-time Artificial Intelligence for Accelerator Control: A Study at the Fermilab Booster

Jason St. John, Christian Herwig, Diana Kafkes, Jovan Mitrevski, William A. Pellico, Gabriel N. Perdue, Andres Quintero-Parra, Brian A. Schupbach, Kiyomi Seiya, Nhan Tran, Malachi Schram, Javier M. Duarte, Yunzhi Huang, Rachael Keller

Developing Robust Digital Twins and Reinforcement Learning for Accelerator Control Systems at the Fermilab Booster

D. Kafkes, M. Schram - <https://arxiv.org/abs/2105.12847>