Bayesian Neural Networks For Uncertainty Estimation In Particle Accelerator Applications

Aashwin Ananda Mishra
SLAC National Accelerator Laboratory
Menlo Park, CA-94025
aashwin@slac.stanford.edu

Auralee Edelen Linscott
SLAC National Accelerator Laboratory
Menlo Park, CA-94025

Adi Hanuka
SLAC National Accelerator Laboratory
Menlo Park, CA-94025

Christopher Mayes
SLAC National Accelerator Laboratory
Menlo Park, CA-94025

Abstract

Particle accelerators are used in a wide array of industrial, medical, and scientific applications. Accurate system models can be useful in experiment planning and in operation of these systems, for example, in model-based tuning schemes. To this end, machine learning is increasingly being applied for modeling particle accelerators, and neural networks are a popular choice of modeling paradigm as they can handle the large data sets and high-dimensional data that is common in this domain, such as particle beam images, spectra. However, for their deployment in high-regret and safety-critical systems, estimates of the predictive uncertainty are essential. Here we evaluate Bayesian Neural Networks to provide accurate predictions, with quantified uncertainties for particle accelerator systems. We select problems across different accelerator designs (a storage ring, a beam line for a free electron laser and an injector system). The cases also have diverse data volumes and formats, e.g. particle beam phase space images and scalar parameters. We show that Bayesian Neural Networks provide accurate predictions with reliable uncertainty estimates across diverse accelerator problems.

1 Introduction & Motivation

Particle accelerators use electromagnetic fields to accelerate beams of elementary particles, such as electrons or protons, and customize the beam shape in phase space to meet the requirements of different applications. Accelerators are used in a broad range of applications where safety and reliability are essential. For instance, numerous patients receive accelerator-based therapy each year for diseases like cancer [3]. Similarly, accelerators play an important role in security, including cargo inspection, nuclear non-proliferation treaty verification, etc [8]. High-power proton accelerators are used to produce radioisotopes [2], where tails of the beam distribution (or 'beam halo') can damage accelerator components if poorly controlled. Finally, accelerators are essential instruments to broad swaths of the scientific community; for example, accelerator based light sources are in high demand to provide scientific users with custom beams to image chemical, material, and biological samples.

Extended time spent in the tuning and control of accelerators is costly in terms of maximizing scientific output. In this vein, machine learning based surrogate models for accelerator problems are opportune. The existence of large data sets and the need to predict complicated outputs, such as beam images, make neural networks (NNs) an appealing approach for ML-based modeling. However, to be used reliably in particle accelerator applications for prediction and control, uncertainty estimates are needed along with point predictions. Deep learning models in particular have shortcomings where

Third Workshop on Machine Learning and the Physical Sciences (NeurIPS 2020), Vancouver, Canada.

their predictions may be overly confident, and do not inherently include prediction uncertainties. For instance, deterministic NNs are unable to recognize out-of-sample instances and make erroneous predictions for such cases with high confidence [4, 12]. Such uncertainty in predictions has had grave consequences while applying deep learning to high-regret applications. For instance, the first fatality in automated driving systems occurred due to the inability of a trained AI-agent to differentiate the hue of a trailer from the color of the sky [1]. Similarly, deep neural networks applied for facial recognition in law enforcement have exhibited critical errors [7].

Such epistemic uncertainty inherent to neural networks is exacerbated by the aleatoric and systemic uncertainties inherent to the modeling of particle accelerator applications. Interrelations between accelerator subsystems are complicated, involve large parameter spaces, and accelerator systems can be difficult to model *a priori*. Changes in system responses over time are also common due to drift, the existence of hidden variables, and transients (e.g. RF power fluctuations). Online tuning of accelerators to meet new custom beam requests also often involves entering previously unexplored setting combinations. Beyond this, the instrumentation to characterize the beam response is often limited due to cost constraints, resulting in limited measurement data. In addition, the instrumentation and controllable components also have different levels of sensitivity and inherent noise for different beam parameters, leading to heteroscedastic effects. This uncertainty is aggravated by the presence of compounding errors in the many individual beamline components.

In this context, obtaining quantified uncertainties for machine learning based models for particle accelerators is paramount if such models are to be applied in such high-regret and safety critical tasks. While different approaches are available to this end (such as Gaussian Processes [14], Bootstrap MAP ensembles [6]), in this investigation we evaluate Bayesian Neural Networks (BNNs) for particle accelerator applications. Herein, we select problems across different accelerator designs, with diverse data volumes and data formats. We evaluate the ability of BNNs to provide accurate predictions of the mean and reliable estimates for the predictive uncertainty. For inference, approximate Variational inference with the Bayes By Backprop algorithm [5] is used. The network architectures are selected using Bayesian Optimization [13]. To optimize the variational parameters, we utilize the Adam algorithm [11] wherein the rates are set using cross-validation. The activations across all neurons are Rectified Linear Units (ReLU), and all weights and biases are initialized with standard normal priors.

2 Results & Discussion

Emittance Prediction in a Storage Ring: SPEAR3 is a 3-GeV, high-brightness electron storage ring [9], operating with a beam current of 500 mA. The electron beam area in phase space (i.e. the *emittance*, ϵ) is an essential parameter that needs to minimized to produce a high-brightness photon beams for scientific users. In SPEAR3, 13 skew quadrupoles are adjusted to minimize the beam loss rate, which is a proxy measure for ϵ. Models of the relationship between the quadrupole settings and ϵ suffer from inherent epistemic uncertainty, as well as aleatoric uncertainty arising due to errors in beam current measurement (used to infer ϵ). For reliable predictions under uncertainty, we train BNNs using both experimental measurements and simulation data (see [10] for the simulation description). The experimental dataset has 650 measurements of the the beam loss rate (in mA/min) and corresponding parameters for the 13 skew quadrupoles. The simulation data has $3.5k$ samples of the beam loss rate with skew quadrupole settings. The network architecture was optimized as having 7 hidden layers with 8 neurons each. The input layers has 13 features and the output is a scalar. The results are outlined in Figure 1 for both these cases, using a randomized $80\% - 20\%$ train-test data split. In both the cases, the BNN mean predictions are comparable to deterministic NNs. For instance, the mean absolute error after training on the experimental data is 0.06 for the deterministic neural network and is 0.03 for the BNN for the test dataset. Additionally, the predictive uncertainty estimates are qualitatively consistent with the prediction error. For instance, in Figure 1 (c), for low values of beam loss rate the predictions are mostly accurate and the predicted standard error is low as well. However, for high values of beam loss, the predictions have appreciable discrepancy and the predicted standard error is accordingly higher.

Emittance Prediction in the LCLS Linac: In our second case, we examine modeling the transverse emittance (ϵ_x) of the Linac Coherent Light Source (LCLS) electron beamline. The LCLS is a free electron laser (FEL) based light source user facility providing customized photon beams for scientific experiments. The FEL process is extremely sensitive to variations in the electron beam phase space, which in turn is sensitive to a variety of accelerator settings and the impact of collective effects such as coherent synchrotron radiation. For this case, 4k data points are obtained using Bmad [16] simulations, which includes nonlinear collective beam effects. The data is uniformly sampled from a large operating range of the accelerating cavity phases and voltages (6 features total), which are commonly adjusted to optimize the beam's shape. However, uniform sampling in feature space does not translate to uniform sampling in target space, as is outlined in Figure 2 (b). Due to the paucity of samples at high values of emittance, data driven model predictions have significant discrepancy in these ranges, and reliable uncertainty estimates are essential. The validation and testing sets consist of 1k and 800 samples, respectively. The BNN architecture has 8 hidden layers. The input layer has 6 features and the output ϵ_x. It can be observed in Figure 2 (c), for low values of the emittance the

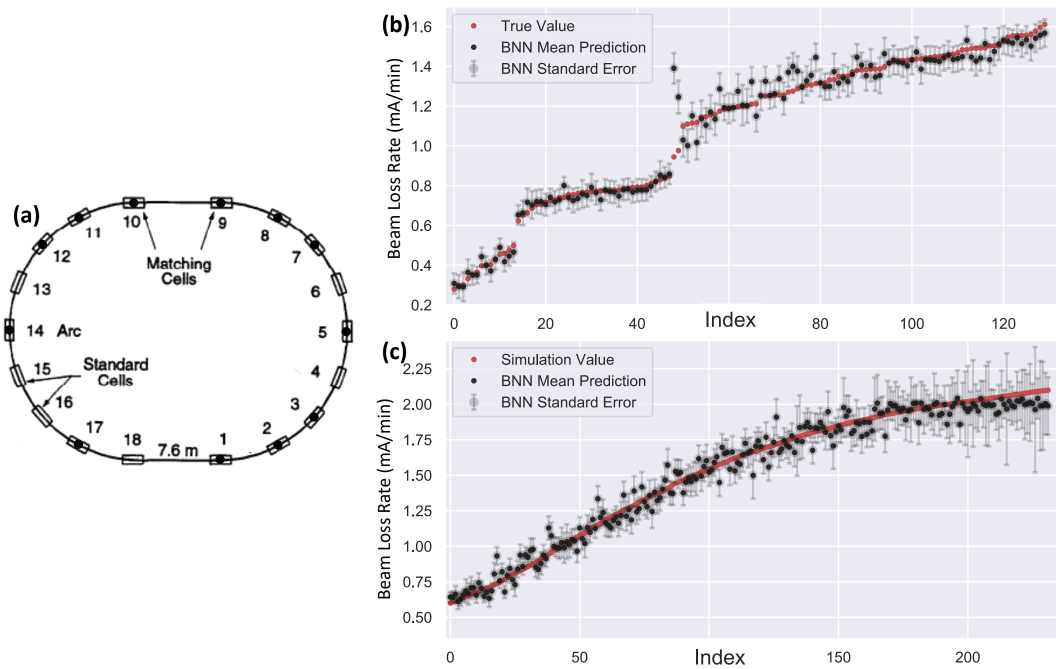

Figure 1: (a) Schematic of SPEAR3 Ring, BNN mean predictions and predicted standard error for (b) SPEAR3 measurements, (c) simulation data. Here we sort the simulation values by magnitude and show corresponding predictions. The beam loss rate is a proxy measure for the beam emittance.

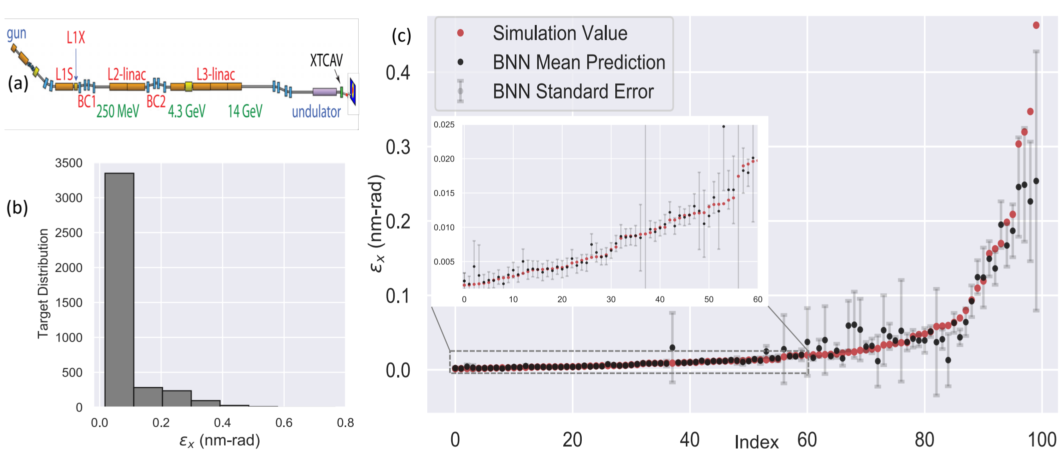

Figure 2: A) Schematic of the LCLS Beamline, b) Distribution of target emittance (ϵ_x) in the learning data sampled uniformly in feature space, c) BNN mean predictions and standard errors for a sample from the test set, with an inset view of the lower ranges.

BNN mean predictions are accurate and the uncertainty bounds are largely negligible. In the few cases where the BNN prediction is erroneous, the predicted standard error reflects this. At high values of the emittance, where data is sparse, the BNN mean predictions have discrepancy comparable to a deterministic neural network. However, the predicted standard error is accordingly higher and reflects the discrepancy in predictions. The BNN provides prediction accuracy comparable to a deterministic neural network (MAE=0.02 for both), but the uncertainty bounds make this a reliable model for such applications.

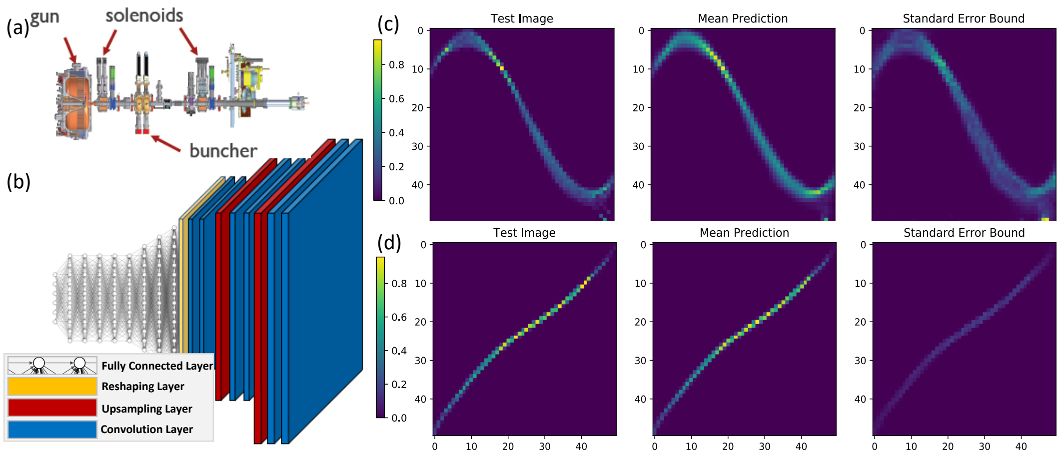

Figure 3: (a) LCLS-II Injector Schematic, (b) BNN architecture, and BNN mean predictions and predicted standard error for the LCLS-II injector longitudinal phase space projections, which show the time vs. energy histogram of the bunch (here shown as 50×50 binned images). Here we are showing two randomly-selected representative examples (c) and (d)

Prediction of Phase Space Image Projections in the LCLS-II Injector: In addition to scalar quantities, prediction of beam phase space projections are often used to provide additional information about the beam. The beam itself is a collection of six-dimensional information (3 positions and 3 momenta for each particle), and 2D projections of the phase space can also be directly measured. Thus, for NN models of accelerator systems, it is highly desirable to predict both these 2D projections and the associated uncertainty. In addition to being useful for general accelerator modeling, ML-based image prediction is useful for providing non-invasive estimates of the beam phase space in cases where it cannot be continuously measured. Here, we focus on prediction of the longitudinal phase space images of the LCLS-II injector. To generate the dataset, simulations of the injector using the IMPACT code [15] were carried out. The scalar inputs were generated by randomly sampling 5 settings of interest (the injector cavity phase, two solenoid strengths, and a buncher cavity amplitude and phase). The network architecture consisted of encoder and decoder sections, outlined in Figure 3 (b). The encoder section consists of 9 densely connected layers consisting of 20 ($\times$ 5), 100, 200, 600 and 10k neurons respectively. The vector output of the encoder section is reshaped into a higher-order tensor before being fed into the decoder section. The decoder section consists of sets of convolutional layers, followed by upsampling layers. Here, the upsampling factor for the rows and columns was 2. The convolutional layers had 16 filters, except for the last layer having 1 filter. The kernel sizes over the six convolutional layers were $4, 4, 4, 3, 2, 1$ respectively. The training dataset consisted of 15k pairs of scalar inputs and image outputs. The validation and testing dataset consisted of 2k pairs of scalar inputs and image outputs each. Representative predictions on the test set are shown in Figure 3. For each instance, we report the test image from the simulation, the mean prediction from the BNN and the standard error predicted by the BNN. The standard error highlights regions where there is significant discrepancy between the mean prediction and the simulation output, for instance in Figure 3 (c) . In cases where the mean prediction is in close agreement with the simulation, for instance in Figure 3 (d), the standard error is correspondingly lower.

3 Summary & Future Outlook

In this investigation, we show that BNNs can provide accurate predictions and uncertainty estimates for several different kinds of accelerator systems, using data from both physics simulations and measurements. We include predictions of scalar data and image data that describe relevant aspects of the particle beam with respect to changing accelerator settings. Evaluating methods for incorporating uncertainty estimates into neural network-based models of particle accelerator systems is an essential step in ensuring that they can be reliably put to use in accelerator operations. Examples of use cases includes experiment planning, model-based control, and online prediction of beam parameters that cannot be continuously measured (i.e. virtual diagnostics) for use in control and data analysis.

Impact Statement

Particle accelerators are essential for a variety of high-regret and safety-critical tasks, such as cancer treatment, nuclear non-proliferation treaty verification, food sterilization, etc. In addition, optimization of high-power accelerators used for isotope production and high-energy physics experiments is often difficult due to the potential damage that can be caused by the high-power beam halo. In such applications, accelerators need to consistently ensure safety and reliability. To support these ends, machine learning models for accelerator applications need to be accurate, to provide high fidelity predictions, and flexible, to process high data volumes and different formats. However, such machine learning models also need to be “uncertainty aware" to support safety and reliability in decisions made using the machine learning model. In this investigation, we show that Bayesian Neural Networks meet these criteria across diverse and diametric accelerator problems. Testing these models with quantified uncertainties in online operation at large scientific user facilities is paving the way toward transitioning to novel safety-critical and high-regret applications of particle accelerators, like proton therapy.

Acknowledgement

This work was supported under the U.S. Department of Energy/ Stanford University Contract for Management and Operation of SLAC National Accelerator Laboratory under Contract No. `DE-AC02-76SF00515`. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. `DE-AC02-05CH11231`.

References

[1] National Highway Traffic Safety Administration. Tesla crash preliminary evaluation report ODI resume. investigation: PE 16-007, 2017.

[2] Jose R. Alonso, Roger Barlow, Janet M. Conrad, and Loyd Hoyt Waites. Medical isotope production with the isodar cyclotron. *Nature Reviews Physics*, 1(9):533–535, 2019. doi: 10.1038/s42254-019-0095-6. URL `https://doi.org/10.1038/s42254-019-0095-6`.

[3] Ugo Amaldi. The importance of particle accelerators. *Europhysics News*, 31(6):5–9, 2000.

[4] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.

[5] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. In *International Conference on Machine Learning*, pages 1613–1622, 2015.

[6] John G Carney, Pádraig Cunningham, and Umesh Bhagwan. Confidence and prediction intervals for neural network ensembles. In *IJCNN'99. International Joint Conference on Neural Networks. Proceedings (Cat. No. 99CH36339)*, volume 2, pages 1215–1218. IEEE, 1999.

[7] Laurence Dodds. Chinese businesswoman accused of jaywalking after ai camera spots her face on an advert. *The Telegraph*, 2018.

[8] Michael Fazio, George Laramore, Suresh Pillai, Ahmed Badruzzaman, Harry Martz, Jeff Buchsbaum, David Jaffray, Mary-Keara Boss, Andrea Schmidt, and Jeff Calame. Basic research needs workshop on compact accelerators for security and medicine: Tools for the 21st century, may 6-8, 2019. 2019.

[9] R Hettel. The completion of SPEAR 3. Technical report, Stanford Linear Accelerator Center (SLAC), Menlo Park, CA, 2005.

[10] Xiaobiao. Huang. Beam-based correction and optimization for accelerators, 2020. URL `https://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=2330060`.

[11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. URL `http://arxiv.org/abs/1412.6980`. cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.

[12] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015.

[13] Tom O'Malley, Elie Bursztein, James Long, François Chollet, Haifeng Jin, Luca Invernizzi, et al. Keras Tuner. `https://github.com/keras-team/keras-tuner`, 2019.

[14] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.

[15] Robert Ryne and Salman Habib. Impact: An integrated-map & particle accelerator tracking code for linac simulation. *APS*, pages 9P–105, 1997.

[16] David Sagan. The bmad reference manual. *Cornell University*, 2008. URL `https://www.classe.cornell.edu/bmad/`.