Using Self-consistency to Determine Uncertainty in Particle Accelerator Diagnostic Measurements

Stephen D. Webb
RadiaSoft LLC
Boulder, CO 80301
swebb@radiasoft.net

Jonathan P. Edelen
RadiaSoft LLC
Boulder, CO 80301
jedelen@radiasoft.net

Abstract

Control of a particle accelerator relies on interpreting input diagnostics, and tuning the accelerator settings to achieve desired diagnostic values. Particle accelerator diagnostics can fail, slowly or abruptly, and currently a human operator is tasked with evaluating this. For truly autonomous operation of particle accelerators, we must be able to determine the reliability of a diagnostic using nothing more than the set of all diagnostic data. We present a method of quantifying the scale of error in diagnostic measurements based on gaussian process (GP) regression and the intuition that each diagnostic measurement from a beam position monitor (BPM) should be predictable given the set of all other BPM measurements.

1 Autonomous Accelerator Operation

Particle accelerator operation currently centers on a team of skilled human operators and some automated routines for optimizing the accelerator's operation against a set of performance metrics. Accelerators produce a large amount of data, that data is of variable reliability, and the optimization is typically multi-objective and spans a high-dimensional configuration space. This makes particle accelerators an ideal candidate for applying machine learning techniques [2]. Autonomous control is one of the key areas for progress for machine learning applications to particle accelerators [3].

There has been much progress in applying machine learning based optimization to accelerator controls. Scheinker *et al.* [7] demonstrated the ability for model-independent control to adapt an accelerator to deliver a certain longitudinal phase space in a free-electron laser linac based only on a handful of control knobs and a diagnostic read-out. Duris *et al.* [1] have demonstrated that Bayesian optimization with Gaussian process models can reach better optima faster than simplex models in computer simulations of free electron lasers. Other recent work by Scheinker *et al.* [8] demonstrated that a model-independent optimization algorithm can improve the average pulse energy in the LCLS SASE FEL by automatically tuningsettings with no foreknowledge of free electron laser physics.

These applications all show the capability of machine learning techniques to tune an accelerator to a desired operating point. However, they all involve human-in-the-loop operation – there are still operators in the control room. For industrial applications such as water treatment – where there might be two operators for ten accelerators, instead of ten operators for one – or for accelerators deployed for space applications – where there would be no operators present – the control systems must be truly autonomous. This means that the control system must be able to evaluate the reliability of each diagnostic measurement using only the set of all diagnostic measurements. Particle accelerator diagnostics can fail for numerous reasons, and either catastrophically or slowly, and any autonomous accelerator control system must be able to detect these and distinguish them.

In this paper we present a concept for determining the reliability of a set of beam position monitor (BPM) measurements using diagnostic self-consistency. In an ideal system with no errors, the BPM

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

measurements should be highly correlated. If the measurements have some noise on top of them, then a predictive model that captures uncertainty should have the scale of that uncertainty be proportional to the level of noise on the measurements. We can therefore use the predictive uncertainty of a model fit to the BPM data to determine the scale of that uncertainty with no outside input.

2 Diagnostic Self-Consistency

Consider a set of diagnostic measurements $\{X_i\}$. In many cases, there are physical reasons to believe that these measurements are correlated. For our case, where we analyze BPMs, beam dynamics suggests that the BPM measurements are very strongly correlated due to the existence of an underlying transfer map that relates the beam phase space coordinates at each BPM to each other. Given enough BPMs, we should be able to reconstruct that map in a model that allows us to predict the value measured at one BPM given all the values at the other BPMs. We refer to this as the existence of *diagnostic self-consistency*.

Formally, we expect that for each X_α measurement, given the set of inputs $\{X_{i\neq\alpha}\}$, there should be a function $X_\alpha = f_\alpha(\{X_{i\neq\alpha}\})$ in an ideal case with perfect diagnostic reliability for each i. In the case that f_α exists and is perfectly accurate for each α then we have diagnostic self-consistency.

In reality, the diagnostics will have some noise in their readings, due to electronic noise, resolution limits, etc. Thus, each X_α will have some best fit model $\overline{f}_\alpha$ which has some uncertainty in the prediction. By quantifying that uncertainty, we can extract the uncertainty in each diagnostic measurement and ascertain the variance in a diagnostic's measurement using only the set of all diagnostic measurements.

Because we are fitting the relation between ~ 10s of diagnostics, and need to have a variance in those predictions, we will use gaussian process (GP) regression [6] to compute the correlation. GP regression has already been used effectively for optimization [1] as well as in building physics-based models for improved optimization [5].

3 The ATR Beamline: A Test Case

To generate test data, we use MAD-X simulations of the ATR beamline at Brookhaven National Laboratory, fig. (1). The ATR beamline is used to extract beam from the Alternating Gradient Synchrotron and transport it to the Relativistic Heavy Ion Collider. This beamline has a number of vertical and horizontal bends, and a number of correctors and BPMs along the line.

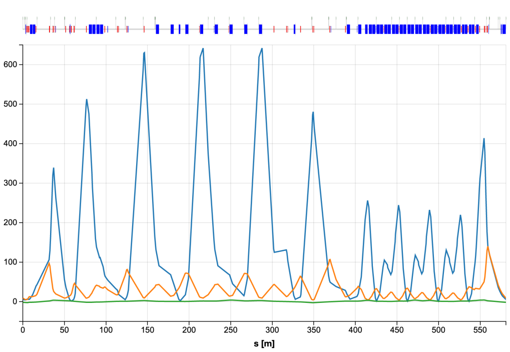

Figure 1: Lattice layout and Twiss parameters for the ATR beamline.

For what follows, we generated a dataset of 500 different beams with randomly offset initial positions and angles to the ideal axis. To compute the $\overline{f}_\alpha$ for each BPM, we train the GP models using a training set of 400 randomly selected beams, with the additional 100 held out as a test set, which we will use to infer the error in the models.

3.1 Baseline – Ideal Diagnostics

To test if we are able to achieve diagnostic self-consistency with the ATR beamline, we begin with simulation data with no noise on the beamline. For our self-consistency model we used a polynomial kernel with degree 1 for gaussian process regression using GPflow [4] (version 2.0.5). Because the underlying physics can be described exactly by a matrix transform, it is sensible that a linear kernel would perform best.

In this case, we find that BPMs 0, 14, and 22 do not give reliable fits, with BPM 0 having the worst fitting. We show some examples in fig. (2) – BPMs other than 0, 14, and 22 have the same ideal predictions as those shown in figs. (2b, 2c).

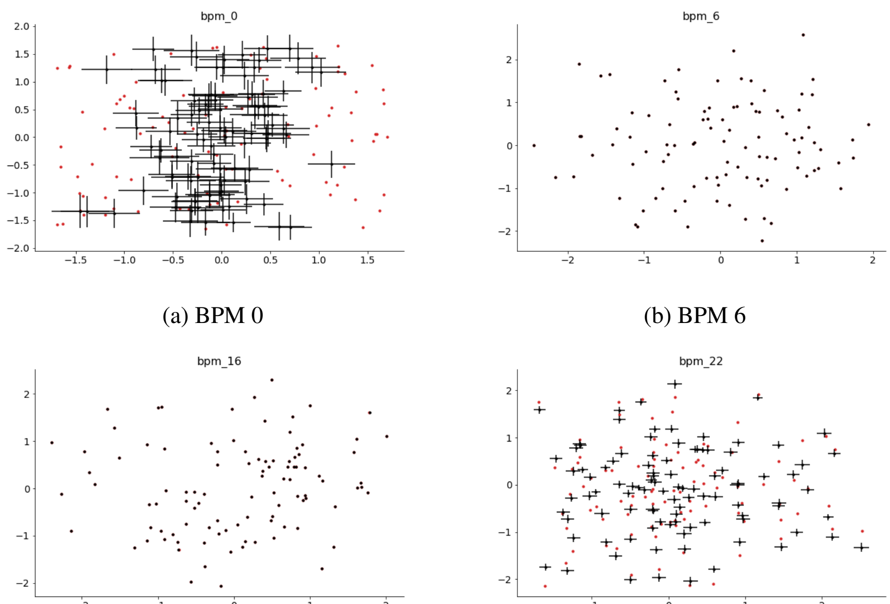

(a) BPM 0

(b) BPM 6

(c) BPM 16

(d) BPM 22

Figure 2: Example comparisons of test set data (red) versus predictions with uncertainty (black).

We can therefore conclude that we can achieve diagnostic consistency with some subset of the BPMs. In testing, we have found that the first BPM in a sequence, regardless of how many BPMs we have available, always has this poor level of fitting. This is likely due to insufficient information upstream of the first BPM. The error in BPMs 14 and 22, is likely related to the coupling in the horizontal and vertical motion introduced in the W line. In all cases, we can identify these BPMs where diagnostic consistency does not occur in simulation data.

3.2 Gaussian noise on diagnostic data

We then tested our approach to diagnostic consistency by adding Gaussian noise to each of the 23 BPM measurements. Here each BPM is assigned a random σ that represents the variance in the noise level. Thus, each BPM has a different scale of noise. We then train each BPM's f_α with this noisy data, and compare the predictions to the test set values as before.

Once noise is introduced, the model $\overline{f}_\alpha$ is no long a perfect predictor of test data, compare figs. (3b, 3c) to figs. (2b, 2c). However, for the $\overline{f}_\alpha$ other than BPMs 0, 14, and 22, the test data points are still within the error predictions from the GP model. We will use the errors in the predictions of the test data as a proxy for estimating the uncertainty in the individual BPMs.

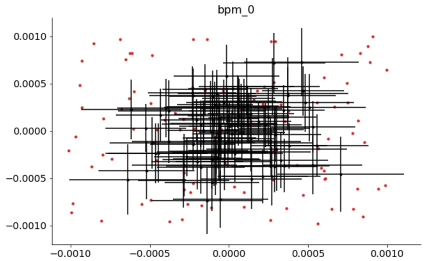

(a) BPM 0

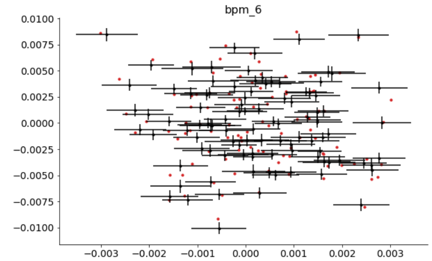

(b) BPM 6

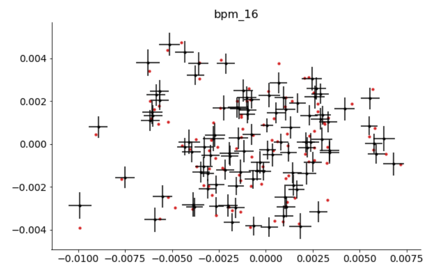

(c) BPM 16

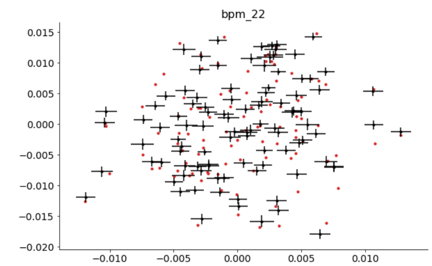

(d) BPM 22

Figure 3: Example comparisons of test set data (red) versus predictions with uncertainty (black) with noise in the data.

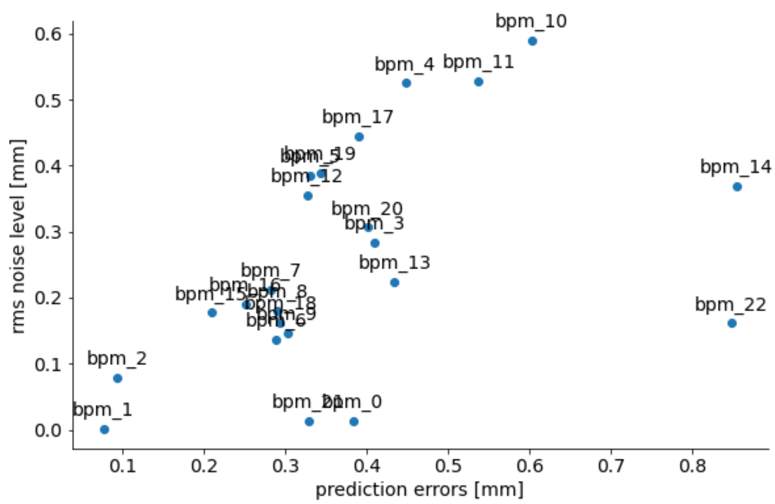

Figure 4: Variance in the GP model test set predictions versus the noise σ applied to the dataset.

3.3 Relation of measurement noise to gaussian process model errors

We can now compare the σ of the noise we applied to each BPM with the RMS errors in prediction for each BPM's GP model applied to a test dataset

$$\varepsilon_\alpha = \sqrt{\sum_i \left(\overline{f}_\alpha(X_i) - Y_i\right)^2} \tag{1}$$

ε_α can be computed from existing data with no reference to the level of noise. As we see in fig. (4), there is a linear relationship between σ and ε, with the problem BPMs 0, 14, and 22 identified in the baseline case studied above as the outliers. This relationship, combined with the ability to identify problem diagnostics using ideal simulation data, suggests that this technique can be used to infer the uncertainty of a given BPM's measurements using only the BPM data itself.

4 Conclusion and Outlook

We have presented a technique for inferring the uncertainty in a diagnostic's measurement using gaussian process regression and an assumption of diagnostic self-consistency. We tested this technique using simulations of the ATR beamline BPMs, which showed that diagnostic self-consistency is not uniform across all BPMs in a beamline. When we then added noise to the BPM data and found that, except for the problem BPMs identified in the ideal dataset, the RMS error on a holdout dataset can be accurately predicted for almost all of the BPMs. Although we do not yet understand the particular failures for the handful of BPMs that do not perform well in computing a self-consistency fit, we can detect them clearly in ideal simulations, and therefore can choose to exclude them in practical applications. Two key developments remain for this work: (1) an application to rings rather than linacs, where there is a periodicity in the diagnostics and no "first BPM" and (2) using the computed diagnostic uncertainty in optimization routines for the accelerator, in such a way that a diagnostic that is unreliable is given less weight in tuning than a more reliable diagnostic, but in a smooth way.

Broader Impact

Autonomous operation of particle accelerators would open up a variety of applications, ranging from medical accelerators for cancer therapy to high power electron linacs for water treatment facilities to accelerator-based experiments in space. Safely operating such an accelerator in a fully autonomous mode, with minimal input from human operators, requires that the autonomous control system be able to evaluate the accuracy of the diagnostic inputs the system is tuning against. This work represents the first steps towards quantifying diagnostic reliability using only the set of available diagnostic measurements, without any external information supplied.

Acknowledgements

The authors would like to thank Auralee Edelen, Joe Duris, and Adi Hanuka (SLAC) for helpful discussions.

References

[1] J. Duris et al. "Bayesian optimization of a free-electron laser". In: *Phys. Rev. Lett.* 124.124801 (2020). DOI: 10.1103/PhysRevLett.124.124801.

[2] A. L. Edelen et al. "Neural Networks for Modeling and Control of Particle Accelerators". In: *IEEE Transactions on Nuclear Science* 63.2 (Apr. 2016), pp. 878–897. ISSN: 0018-9499. DOI: 10.1109/TNS.2016.2543203.

[3] A. Edelen, C. Mayes, and D. Bowring. *Opportunities in Machine Learning for Particle Accelerators*. arXiv:1811.03172 [physics.acc-ph]. 2018.

[4] A. G. de G. Matthews et al. ""GPflow: A Gaussian process library using TensorFlow"". In: *Journal of Machine Learning Research* 18.40 (apr 2017), 1–6.

[5] A. Hanuka et al. *Physics-informed Gaussian Process for Online Optimization of Particle Accelerators*. arXiv:2009.03566 [physics.comp-ph].

[6] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. THE MIT Press, 2006.

[7] A. Scheinker et al. "Demonstration of Model-Independent Control of the Longitudinal Phase Space of Electron Beams in the Linac-Coherent Light Source with Femtosecond Resolution". In: *Phys. Rev. Lett.* 121.044801 (2018).

[8] A. Scheinker et al. "Model-independent tuning for maximizing free electron laser pulse length". In: *Phys. Rev. Acc. Beams* 22.082802 (2019).