Chapter III.12

Machine learning

Jochem Snuverink
Paul Scherrer Institute, Villigen, Switzerland

Machine learning (ML) is a powerful new mathematical toolbox that allows for new possibilities in a broad variety of research areas. It enables machines to fulfill tasks without being explicitly programmed. There is no reason to believe why machine learning cannot be applied to particle accelerators. After a short introduction to machine learning, in this chapter the existing applications to particle accelerators are reviewed with examples on the PSI accelerator complex.

When asking chatGPT [1] about machine learning, it answered (bold by the author):

Machine learning is not a hype, but rather a ***powerful tool*** *that has the potential to revolutionise many industries and fields, including healthcare, finance, transportation, and more.*

Machine learning has already made significant contributions to many applications, such as ***image recognition, speech recognition, natural language processing, and recommendation systems****. These applications have brought significant benefits to users, making their lives easier and more convenient.*

Moreover, the use of machine learning is ***not limited to specific industries or applications****. Its flexibility allows it to be applied to a* ***wide range of problems and challenges, making it an attractive tool for many researchers, engineers, and businesses****.*

However, it is also important to acknowledge that the use of machine learning ***requires significant expertise in mathematics, statistics, programming, and domain-specific knowledge****. Not all applications may benefit from machine learning, and some may require other techniques or methods.*

Therefore, while machine learning is not a hype, it is also not a panacea for all problems. Its success depends on many factors, including the quality and quantity of data, the expertise of the users, and the suitability of the problem for machine learning techniques.

Perhaps not surprisingly chatGPT is excited about the already made contributions, but at the same time notes it is only successful when the problem conditions are right. To understand the boundaries when machine learning can be successfully applied is active ongoing and expanding research area. Some examples from Swiss universities and research institutes:

– ETH-Z: `https://ml.inf.ethz.ch/`, `https://ethz.ch/en/news-and-events/eth-news/news/2020/10/pr-new-centre-for-ai-research.html`: 29 new professorships in 2020;

– EPFL: `https://www.epfl.ch/research/domains/ml/`;

– Swiss Data Science centre: `https://datascience.ch/`.

This chapter should be cited as: Machine learning, J. Snuverink, DOI: 10.23730/CYRSP-2024-003.2131, in: Proceedings of the Joint Universities Accelerator School (JUAS): Courses and exercises, E. Métral (ed.),
CERN Yellow Reports: School Proceedings, CERN-2024-003, DOI: 10.23730/CYRSP-2024-003, p. 2131.

It is important to understand the new toolbox and understand when it can be applied to accelerator physics. In this chapter a short introduction on machine learning is made and a few examples from the Paul Scherrer Institute (PSI) on applications to accelerator physics are presented.

III.12.1 Machine learning in one page

The seemingly sudden increase of machine learning and artificial intelligence is a combination of technology factors, which all happened in the last decade: the increase of computational capabilities in particular the development and wide availability of GPUs (graphics processing units) allows for more complicated models, faster "training" (optimising or fitting) of these models and larger data sets; availability of large datasets that can be easily shared e.g. via cloud computing; new network architecture and training paradigms, such as GANs (generative adversial networks); a better theoretical understanding of neural networks and optimisation methods. Applications such as self-driving cars and recommender systems (social networks) have driven advancement, both algorithmic and practical.

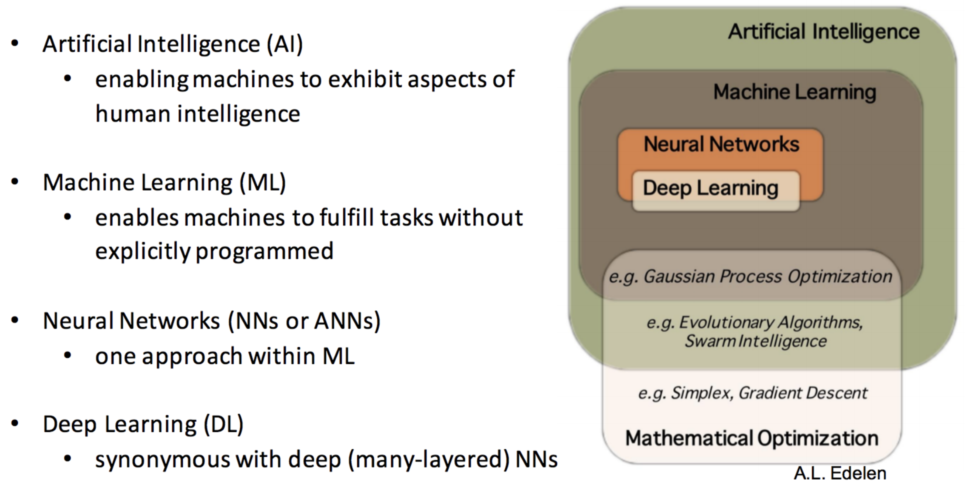

Fig. III.12.1: Definition of terms used in this chapter, from Auralee Edelen [2].

Since artificial intelligence and machine learning is often mixed, here the following definitions are used (see Fig. III.12.1).

In Fig. III.12.2, machine learning is clarified. In the traditional programming approach an algorithm or program is created from a set of commands and rules, generating output with a certain input. While in the machine-learning approach the program is created from input data and desired output data.

With the following conceptual example, linear regression, the ML approach can be explained (Fig. III.12.3). Let's say we have a set of points with Gaussian noise around $y = 4x + 3$. The input is x and desired output y. As a cost function the mean square error is used, which gives in this case an analytic fit based on linear equation: $y = 4.203x + 2.911$. The latter equation can be viewed as our ML model, which is a good approximation to the data. In some way ML is nothing else than "fitting a complicated function to data".

Now new input data can be evaluated and the expected output can be quickly generated. It should

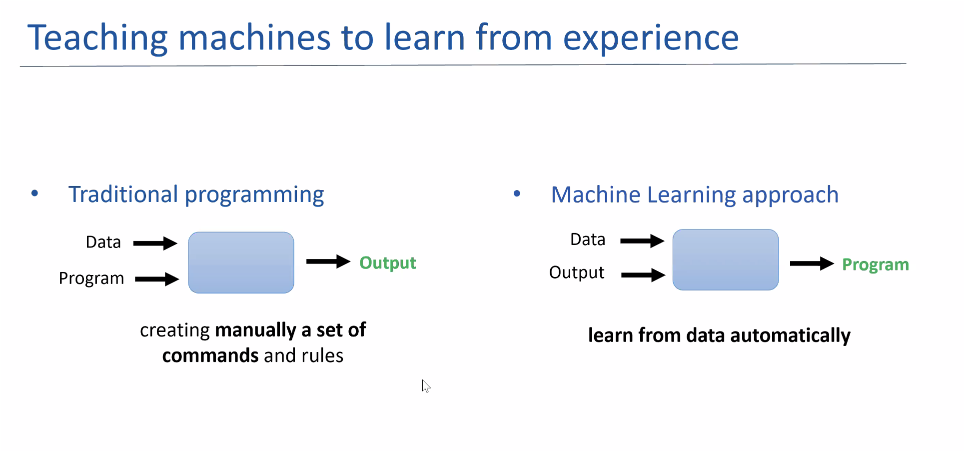

Fig. III.12.2: Conceptual difference between traditional programming or modelling approach and machine-learning approach [3].

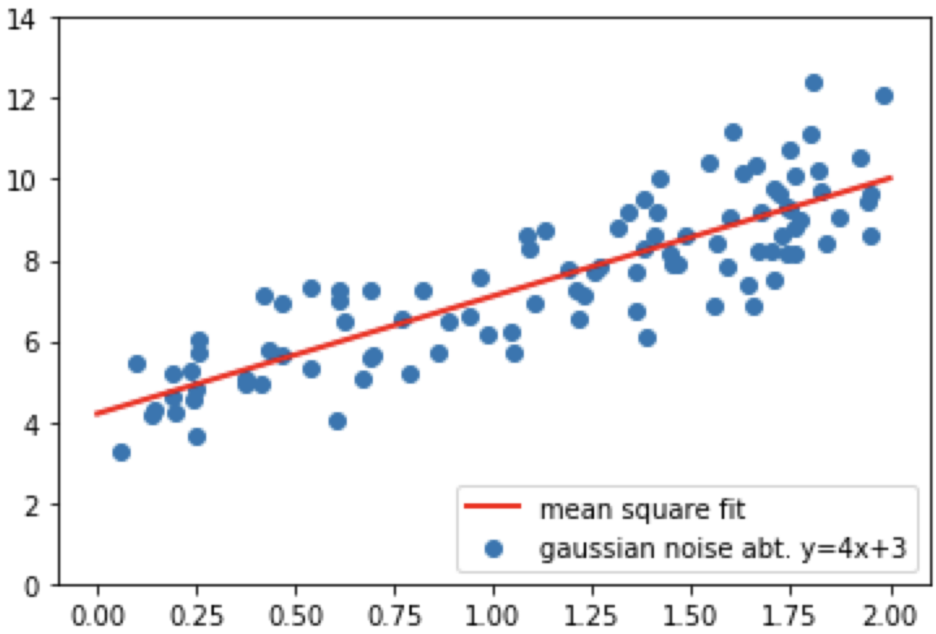

Fig. III.12.3: Linear regression fit to data points.

be noted that the ML model is typically only valid within the input domain. Outside this domain (e.g. $x > 2$) the model will give a value, but there is absolutely no guarantee that the return value is realistic. Especially, with multi-dimensional input it can be a challenge to have enough training data to cover the whole input domain.

Traditionally, the distinction is made between three types of ML, slightly tweaking Fig. III.12.2:

- Supervised learning: both input data and output are available, such as the example just discussed, or deciphering handwritten numbers (arguably the first ML application was made in the 70s by the US post, automatising this), or *improve and speedup the accelerator simulation?*
- Unsupervised learning: only data is available and no output. Still similarities or patterns can be

found (clustering). An example is: credit-fraud detection systems, generation of new images or *increase resolution of diagnostics?*
- Reinforcement learning: both input data and output are available, where the output is part of the cost function, such that after processing the input value a feedback (reward) is given. An example of this is learning (computer) games, or *find optimal accelerator settings?*

III.12.2 Machine learning for accelerator physics

There are several potential applications of machine learning to accelerator physics, including:

Beam optimization: Machine learning can be used to optimize the performance of particle accelerators by adjusting the parameters of the accelerator to achieve desired beam properties.

Accelerator control: Machine learning can help in designing controllers that ensure the stability of the particle beams, by learning the dynamics of the system.

Fault detection and diagnosis: Machine learning can detect and diagnose faults in accelerators, by learning from the data generated by the accelerator sensors.

Online monitoring: Machine learning can be used to develop real-time monitoring systems that detect deviations from normal accelerator operation and provide early warning of impending equipment failure.

Accelerator modeling: Machine learning can be used to improve accelerator modeling accuracy by developing models that can predict the behavior of the accelerator based on historical data.

Beam instrumentation: Machine learning can assist in improving beam instrumentation by analyzing the data acquired from the accelerator and extracting valuable insights.

Overall, machine learning has the potential to revolutionize accelerator physics by providing more efficient and accurate methods for controlling, optimizing, and monitoring the performance of particle accelerators.

Summarising, the following possible applications of ML in accelerator physics are identified:
- Tuning, optimisation and control: optimise or control accelerator parameters;
- Virtual diagnostics: surrogate model of destructive, imprecise, slow or broken diagnostics;
- Online modelling / Simulations: fast beam transport model;
- Anomaly detection and machine protection: preventive maintenance, interlock prediction;
- Advanced data analysis: enhanced phase-space measurements.

For each of these applications one can find many examples in the literature. Here for each category, except the latter, one or two PSI examples are briefly described.

III.12.2.1 Tuning, optimisation and control: optimisation at HIPA and SwissFEL

HIPA's beam power (1.3 MW) is limited by the beam losses, which cause component damage and activation. Optimisation is now mostly done empirically (manually). There is a large potential for automated optimisation and surrogate model construction. However, no accurate and fast physics model is available, and handling such a large power beam needs to be done safely.

At SwissFEL one would like to increase the FEL pulse intensity. Here manual tuning is time consuming and often inefficient due to the many available knobs (about 40).

Bayesian optimisation (BO) tries to combine the strengths of human optimisation (learning and experience) with the strengths of numerical optimisation (juggling many things at once and fast decisions). Bayesian Optimisation is a flexible, data-driven approach for global optimisation with noisy feedback. Instead of fitting a single function, a statistical regression model of the data is fitted, typically a Gaussian process. On each iteration the model looks for evaluation points that reduce the uncertainty of the unknown target function. This can be naturally extended with a safety function (e.g. losses recorded by various detectors), modelled as a statistical regression model as well, which is evaluated and its model is also updated, taking into account the uncertainties of the safety function this defines the safe-set marked in green, see Fig. III.12.4.

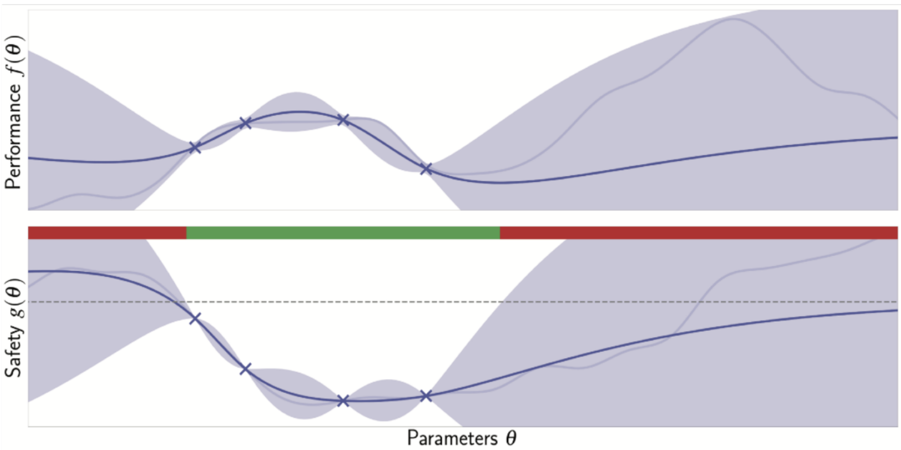

Fig. III.12.4: For each evaluation point (indicated with a cross) the performance (e.g. FEL pulse energy) is measured and the model is updated. Additionally, the safety function (e.g. losses recorded by various detectors) is evaluated and its model is also updated, taking into account the uncertainties of the safety function this defines the safe-set marked in green.

In Fig. III.12.5 it is shown that at the HIPA accelerator the safe variants are competitive with the non-safe methods that create interlocks (violate constraints), which proves that the safety constraints are working [4].

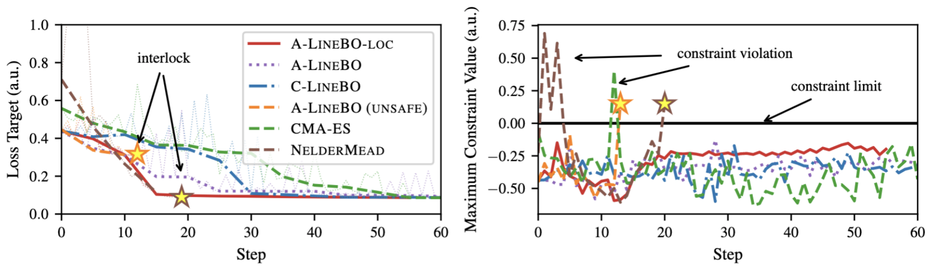

Fig. III.12.5: Safe Bayesian loss optimisation with different algorithms on the HIPA accelerator.

III.12.2.2 Virtual diagnostics: SINQ target diagnostics

The SINQ target has two diagnostic systems that are protecting the target from damage. The VIMOS system monitors the SINQ target beam spot with a metal grid. If the beam is focussed too much or changes too fast interlocks are triggered.

This VIMOS grid is degrading over time and cannot be replaced. Also the temperature sensors can break or degrade over time. Therefore, we would like to develop a model that predict the images and temperatures from other data.

By using the accelerator diagnostics data, such as beam position, loss and current monitors as input and the VIMOS or temperature sensors as desired output a ML model can be trained, see Fig. III.12.6.

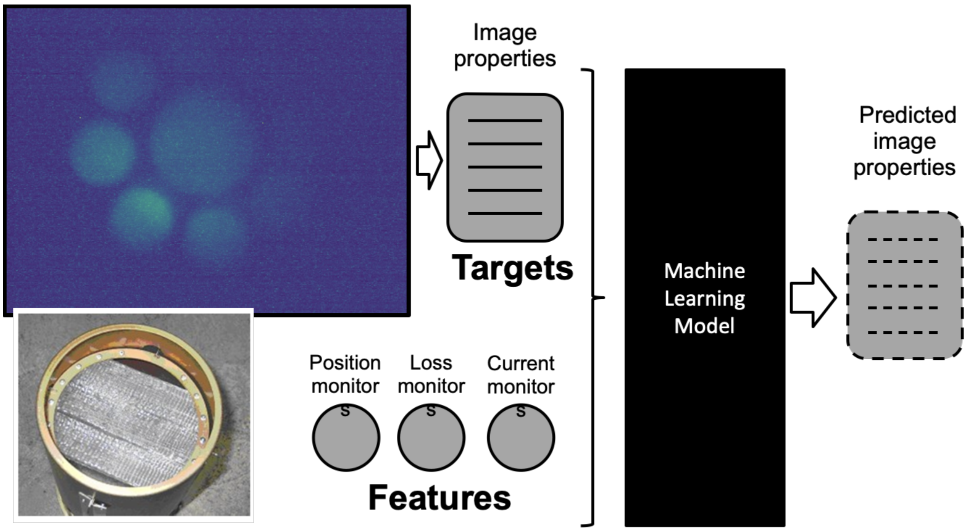

Fig. III.12.6: Conceptual virtual diagnostic model on the VIMOS system.

A so-called random forest model (an ensemble of boosted decision trees) was trained (fitted) to the 2018 temperature data. In Fig. III.12.7 it can be seen that the model is generally within 1 degree of the real value.

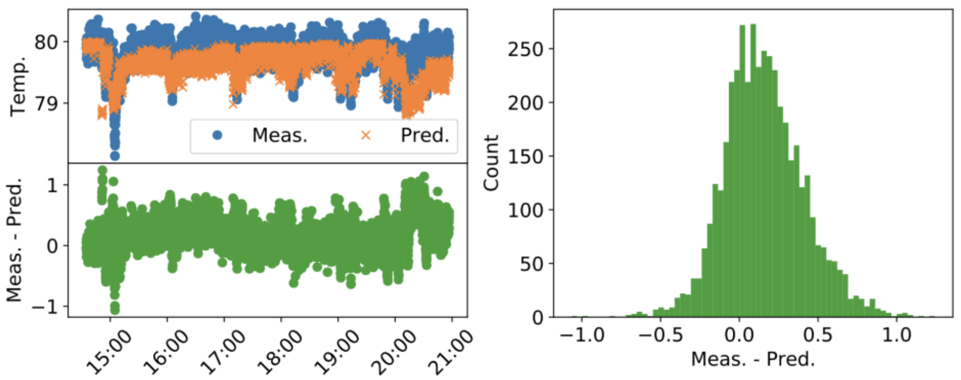

Fig. III.12.7: Performance of the Random Forest model on SINQ temperature data.

A word of caution, while the performance of such models is typically very impressive, one should note that the model will only work if all required input data is still there in the future or that the operating conditions are similar (within the input domain). ML models that incorporate physics knowledge are generally better guarded against such performance issues.

III.12.2.3 Virtual diagnostics: ultra-fast pump-probe experiments

FELs are successfully used for the investigation of ultrafast phenomena. One classical FEL experiment is the pump-probe experiment. The experimental setup is excited (started) by a short “pump” pulse, and then probed by a second “probe” pulse after a certain delay time. By repeating the measurement at different delay times the full dynamics of the process can be obtained. In this experiment it is important to know delay time precisely. See also https://www.psi.ch/de/lmn/ultra-fast-pump-probe-experiments.

At SwissFEL, the delay time is measured with a special diagnostic, however for specific pulses with low signal strength no time is measured. To recover the delay time for these pulses, a ML model has been built that predicts the delay time from other signals, see Fig.III.12.8.

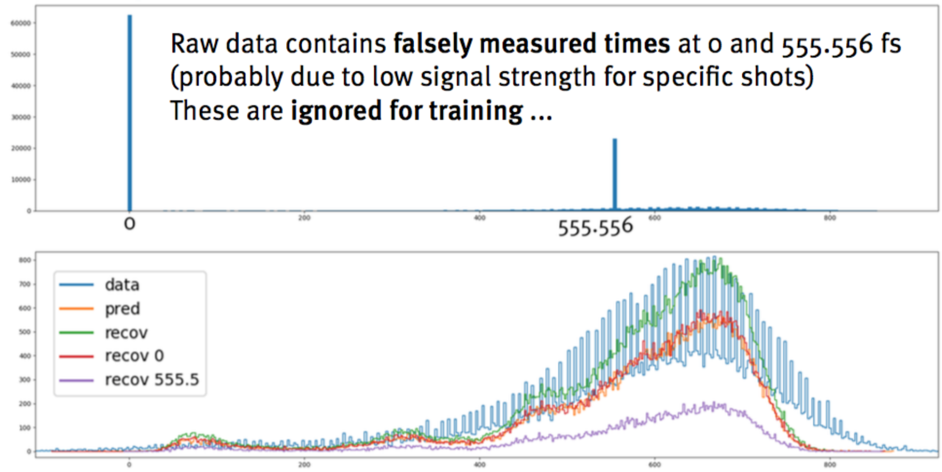

Fig. III.12.8: With the help of an ML model wrongly measured delay times can be recovered (red and purple graphs).

III.12.2.4 Online model / simulations: dynamic-aperture optimisation for SLS 2

The dynamic aperture is the stability region of phase space in a circular accelerator. It is an important parameter to estimate efficiency and stable operation. The traditional method for computing dynamic aperture involves the use of a tracking code and tracking thousands of turns (in case of electrons, for protons many more up to billions of turns might be needed), which can be a time and CPU-consuming process. For the SLS 2 study the strengths of the magnets are the design parameters. With a multi-objective optimisation algorithm and the traditional method, many satisfiable designs were found. To speed up, the search process a surrogate model based on a neural network was made. Using this model with the same optimisation method, many more promising solutions were found. However, it turned out that the found solutions by the surrogate model were not good when verified by the traditional accelerator code. Retraining (refitting) the surrogate model during the optimisation process overcame this model degradation and finally delivered fast and good results [5].

III.12.2.5 Anomaly detection: HIPA loss monitor

According to Hawkins [6] an anomaly is: *An observation which deviates so much from other observations as to arouse suspicions that it was generated by a different mechanism.*

So, anomaly detection is trying to find observations or sequences that deviate from the "normal behaviour". Experts would recognize these anomalous patterns easily, but cannot be monitoring the huge amount of data some systems produce. Examples are credit card fraud detection, intrusion detection in cybersecurity, or fault diagnosis in industry.

At HIPA a loss monitor broke down without anyone noticing. With an anomaly detection model it was shown that at least half an hour before the monitor broke down, an early warning could have been raised, see Fig. III.12.9 [7].

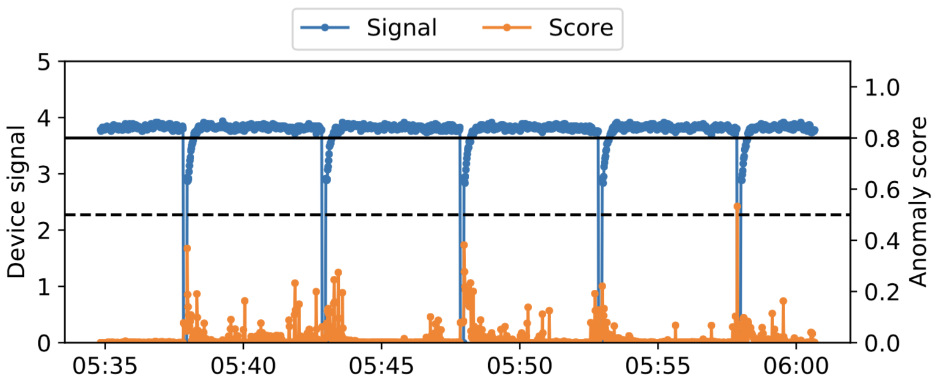

Fig. III.12.9: Loss monitor signal HIPA and anomaly score, hours before it broke down (11:20).

III.12.2.6 Machine protection: interlock prediction at HIPA

At the HIPA accelerator, every interlock turns off the beam for 25 s. With around 30-40 interlocks a day, this accounts for some minutes of lost beamtime every day. At other proposed accelerators, especially for ADS systems, only a few interlocks per year can be sustained. Some of the interlocks could be prevented, e.g. by lowering the current, if the interlocks can be predicted in advance.

At HIPA taking into account 450 accelerator input channels with several months of data are analysed and two different approaches with neural networks were tried to classify at each time step if an interlock would happen in the immediate future [8, 9].

Classification is always balance between good (correct hit) and bad predictions (false alarms). In most classification models a threshold can be set to balance between a good and a bad prediction. Since in the accelerator case the stable state is much more common only a very low false positive rate can be tolerated. An example of a successful interlock prediction on training data can be viewed in Fig. III.12.10.

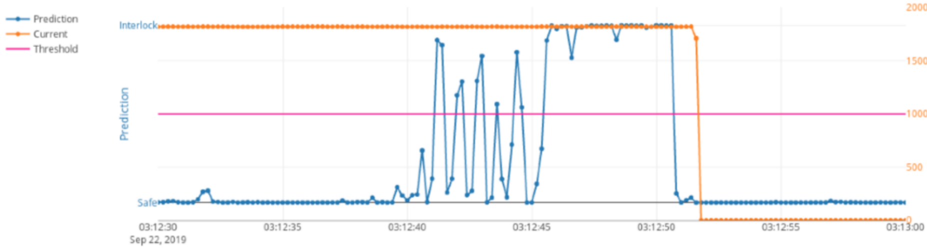

Fig. III.12.10: Example of successful interlock prediction on training data.

III.12.2.7 Summary and words of caution

Machine learning is another useful toolbox to improve accelerator operation. A very brief introduction with a few examples on the PSI accelerator complex have been described. Its usage should be focused on problems where no "physics" or analytical solution is available. Machine learning works especially well when lots of good data (from operation or simulations) is available and when the environment is static, which means no external influences (outside of the data) and not changing over time. Possible challenges that calls for special attention are: the limited ability of ML models to extrapolate outside of the data range; when the input to the ML model has deficiencies, the output cannot be trusted (garbage in, garbage out); ML models might not generalise well (suffer from "overfitting"). And in general if surprises come up, one is well advised to be sceptical.

References

[1] OpenAI, ChatGPT Default (GPT-3.5), 2023

[2] A. Edelen, presentation at ICFA Machine Learning Workshop 2018

[3] E. Fol, CERN BE seminar, 2022

[4] J. Kirschner *et al.*, Tuning particle accelerators with safety constraints using Bayesian optimization, Phys. Rev. Accel. Beams 25 6 2022,

[5] M. Kranjčević *et al.*, Multiobjective optimization of the dynamic aperture using surrogate models based on artificial neural networks, Phys. Rev. Accel. Beams 24 1 2021

[6] D. M. Hawkins, Identification of outliers, 1980, Springer.

[7] J. Coello de Portugal and J. Snuverink, Experience with anomaly detection using ensemble models on streaming data at HIPA, Nucl. Instr. Meth. A, 2021,

[8] S. Li, M. Zacharias *et al.* A Novel Approach for Classification and Forecasting of Time Series in Particle Accelerators, Information 2021, 12, 121.

[9] S. Li *et al.* Forecasting Particle Accelerator Interruptions Using Logistic LASSO Regression, 2023