

Machine learning in accelerator physics: Applications at the CERN Large Hadron Collider

**F.F. Van der Veken^{*a,b}, G. Azzopardi^a, F.H. Blanc^c, L.T.D. Coyle^{a,c}, E. Fol^{a,d},
M. Giovannozzi^a, T. Pieloni^{a,c}, S. Redaelli^a, B.M. Salvachua Ferrando^a,
M. Schenk^{a,c}, R. Tomas Garcia^a, G. Valentino^{a,b},**

^a*CERN, Beams Department
1211, Meyrin, Switzerland*

^b*University of Malta
MSD2080, Msida, Malta*

^c*Ecole Polytechnique Federale Lausanne,
1015, Lausanne, Switzerland*

^d*Johann-Wolfgang-Goethe University
60323, Frankfurt am Main, Germany*

E-mail: frederik.van.der.veken@cern.ch

With the advent of Machine Learning a few decades ago, Science and Engineering have had new powerful tools at their disposal. Particularly in the domain of particle physics, Machine Learning techniques have become an essential part in the analysis of data from particle collisions. Accelerator physics, however, only recently discovered the possibilities of using these tools to improve its analysis. In different laboratories worldwide, several activities are being carried out, typically in view of providing new insights to beam dynamics in circular accelerators. This is, for instance, the case for the CERN Large Hadron Collider, where, since a few years, exploratory studies are being carried out, covering a broad range of topics. These include the optimisation of the collimation system, the anomaly detection of beam position monitors, analysis of optimal correction tools for linear optics, lifetime and performance optimisation, and detection of hidden correlations in the huge data set of beam dynamics observables collected during the LHC Run 2. Furthermore, very recently, ML techniques are being scrutinised for the advanced analysis of numerical simulations data, in view of improving our models of dynamic aperture evolution.

*Artificial Intelligence for Science, Industry and Society, AISIS2019
October 21-25, 2019
Universidad Nacional Autónoma de México, Mexico City, México*

^{*}Speaker.

1. Introduction

For decades, Machine Learning (ML) techniques have been applied in several fields of Science and Engineering, with their successful application in experimental particle physics as a prime example. In the field of accelerator physics, however, they have been adopted at a much slower pace. Although the first attempts to use these techniques in beam diagnostics and beam control systems already date from a few decades ago [1, 2], it is only since a few years that some sizeable progress has been made. For example, several ML applications are under study or in full development at the CERN Large Hadron collider (LHC). From these we selected four different examples as a global overview of the field.

As a first use case, we show in Section 2 how supervised ML has reduced the time needed to align the ring collimators during beam commissioning. Then, in Section 3 we demonstrate how unsupervised ML can be used to clean optics measurements, and supervised ML to correct optics functions in the machine. Later, in Section 4, we investigate how unsupervised ML can be used in tracking simulations to detect outliers efficiently. Finally, it is shown in Section 5 how it is possible to create a model using supervised ML that directly relates measured beam lifetime to certain machine settings.

2. Alignment of the LHC Collimators

Accelerating a particle beam is never perfect, and unavoidable beam losses occur continuously. When these losses reach the coils of the superconducting magnets, they will eventually lead to a so-called quench [3] thus inducing a beam dump with consequent loss of time for physics. To protect against this scenario, the LHC relies on a collimation system to absorb these losses before they can inflict any damage or hamper operations. Such a system consists of a set of collimators, each made of two parallel blocks of specially-devised material that are placed symmetrically around the beam. When the machine goes into beam commissioning the exact beam position at the location of the collimators should be determined so to centre the collimators' jaws around the beam. All this is a necessary condition to push the LHC to nominal operation. Different alignment settings are needed for the different machine states, e.g. the beam energy. These settings are monitored throughout the year and different collimator setups are required when machine parameters are changed.

The procedure to align the collimators is determined following a beam-based alignment (BBA), as established in [4]. Each collimator has a dedicated Beam Loss Monitoring (BLM) device positioned immediately downstream to detect beam losses generated when halo particles impact on the collimator jaws. During alignment, one of the jaws of a collimator is continuously moved towards the beam halo, whilst monitoring the measured beam loss signal. The jaw movement continues until a clear alignment spike is observed, indicating that the moving jaw touched the beam halo and is hence aligned. The procedure continues with the alignment of the other jaw, before moving to another collimator. In order to achieve a reliable alignment, one has to be able to correctly identify these alignment spikes. They consist of a steady-state signal before the spike, the loss spike itself, the temporal decay of losses, and a steady-state signal after the spike. The steady-state is a result of the continuous scraping of halo particles when the jaw positions are fixed. The further a jaw cuts into the beam halo the more the steady-state signal increases, as the density of the particles near




Figure 1: Difference between an alignment spike (left) and spurious spikes (right), from [5].

the jaw increases. Any other spikes that do not follow this pattern are classified as non-alignment spikes. They do not have a fixed structure and can even contain spurious high-amplitude spikes. Such non-alignment spikes arise due to other factors, i.e. beam instabilities or mechanical vibrations of the opposite jaw, thus indicating that the jaw has not yet touched the beam and must resume its alignment [5]. This difference is illustrated in Fig. 1.

As the alignment process is straightforward and repetitive, but tedious, it is a canonical example to be automated by ML algorithms. To this end, the process of spike recognition was cast as a classification problem, as explained in [5]. A dataset was assembled from previous alignment campaigns, in 2016 and 2018, to train several ML models to be able to recognise alignment spikes, and distinguish them from the non-alignment spikes discussed above. Fourteen manually-engineered features were extracted from this dataset and were analysed. The five most important features (one feature for spike height, three features for exponential decay, one feature for position of the jaw) were used to train six ML models for comparison (Logistic Regression, Neural Network, Support Vector Machine, Decision Tree, Random Forest, Gradient Boost). Each model was analysed in depth, optimised, and thoroughly tested on unseen data.

The ML model was incorporated into the operational BBA software [6], together with the necessary threshold-selection algorithm [7] and cross-talk analysis [8], in order to fully automate the alignment. This new alignment software was successfully used for all alignments throughout 2018, and decreased the alignment time of collimators at injection by 71.4% compared to the semi-automatic alignment in 2017, namely from 2.8 hours to 50 minutes [9, 10].

3. Optics Measurements and Corrections

For an efficient operation of any accelerator, it is vital to have good control over the optical functions, which are the observables stemming from the magnets' strength settings. This can be reconstructed during operation from the harmonic analysis of the readings of the Beam Position Monitors (BPMs) that record the position of the excited beam on a turn-by-turn basis. The optical functions thus retrieved typically differ from their design values, due to unavoidable magnetic errors; the difference thence is the basis for the computation of the magnets strength changes required to compensate the measured difference. Unfortunately, these BPMs are prone to producing faulty, i.e. nonphysical values, with around 10% of BPMs being flawed. These faulty BPMs produce non-physical outliers in the optical functions computed from their output. One hence needs to clean the BPM signal to get workable results.

The existing cleaning technique is based on threshold cuts and Singular Value Decomposition (SVD) [11]. We used unsupervised ML complementary to the existing technique to improve




Figure 2: Comparison between beta-beating computed before and after IF cleaning, which demonstrates that IF anomaly detection significantly reduces the number of nonphysical spikes. The optics is computed for Beam 2 in the horizontal plane with $\beta^* = 50$ cm.

the efficiency of BPMs data cleaning. The applied algorithm is performed on harmonic analysis obtained from SVD-cleaned BPM readings and uses some harmonic properties (tune, amplitude, and noise-to-amplitude ratio) as input, in order to detect faulty BPMs as anomalies before the optical functions are obtained from these data. Several clustering methods have been compared; the Isolation Forest (IF) algorithm [12] achieves the best results. Figure 2 shows a comparison of the β -beating, a typical optical function, computed from data cleaned using traditional techniques (blue, full marker) and additional cleaning with the IF algorithm (orange, hollow marker). There is a significant improvement when applying the IF algorithm after SVD. This method recently became a standard part of optics measurements at the LHC and has been successfully used during beam commissioning and machine development for different optics settings in 2018. The presented method, the statistics of its application in operation, and the performance evaluation of different algorithms on the simulated dataset are demonstrated in [13].

Another domain of application of ML techniques is in the correction of optical functions. When, in operation, the optical functions have been reconstructed from the BPM signals, a correction procedure should be applied, which implies acting on the strength of the most effective quadrupolar magnets. Currently, optics corrections in the LHC are performed in two steps, i.e. local corrections using the Segment-by-Segment technique [14] and global corrections using the Response Matrix technique [15]. We investigated the possibility to construct an optics model that aims to do both steps at once by reconstructing magnet errors from optics perturbations caused by these errors. We created a large dataset from accelerator simulations, producing input-output pairs to be used for supervised training. Each pair contains the errors in the magnet's circuits (quadrupoles powered in series to be adjusted) as input, and the horizontal and vertical phase advances of the simulated optics, perturbed with the given errors in the quadrupole circuits, as output. After training, several regression models are tested on simulated optics perturbed with single-quadrupole errors to examine the method performance on more realistic data. The comparison of β -beating corrections achieved with different regression models is given in Table 1. The results clearly demonstrate the

ability of ML-based regression models to compute the corrections needed to reduce the deviations from the design optics. The details on the presented study can be found in [16] and recent advances in [17].

β -beating %	peak	rms
Uncorrected	32 ± 10	11 ± 3
Response Matrix	11 ± 5	3 ± 2
Orthogonal Matching Pursuit	11 ± 2	3.5 ± 0.8
Convolutional Neural Network	11 ± 2	3.2 ± 0.5
Ridge regression	10 ± 2	2.9 ± 0.8
Linear regression	9 ± 2	2.6 ± 1.7

Table 1: Comparison of β -beating averaged over 100 simulations considering the standard deviation as uncertainty. Lower values imply better results. The optics measurements are simulated using the $\beta^* = 40$ cm optics for the 2016 LHC run for Beam 1.

4. Dynamic Aperture Studies

In the study of beam dynamics in general, and the development of new accelerators in particular, one needs to quantify the impact of the field quality of the magnetic elements of the accelerator's lattice on the accelerator performance. One of the important mathematical tools that aim to assess this is the concept of Dynamic Aperture (DA). It represents the size of the smallest connected volume in phase space where motion is bounded for a given amount of time. Though mainly a theoretical tool, it can easily be estimated from single-particle tracking simulations. In these, a set of initial conditions, representing particles, is tracked through the accelerator lattice for a given number of turns. While tracking is fully six-dimensional, it is in general computationally too expensive, in terms of CPU time, to provide an adequate number of initial conditions in six dimensions. It is therefore custom to set the coordinates and phases of the longitudinal plane at a given value, and use a distribution of transverse coordinates where their phases are either randomly assigned or kept at a fixed value. This gives a more manageable two-dimensional initial distribution, for which in our case a uniform polar grid is used. For every angle, we define the angular DA as the largest amplitude for which all particles survived. The total DA is then simply the integral over all angles. Finally, this total DA is calculated for 60 different random realisations of the machine (the so-called seeds), in order to account for the non-deterministic nature of the inhomogeneities in the magnetic field strengths and possible magnet misalignments.

One of the common issues in these calculations is the presence of outliers, where for a given seed and angle the DA differs a lot from the value obtained for the other seeds. It might arise due to a numerical error or file corruption, in which case a remedy should be found. It is however more plausible to be the result of a particular sensitivity of the seed to the underlying physics (e.g. due to special internal cancellations of magnetic errors or closeness to resonances). In that scenario one might argue that it is not necessarily representative of a real-life scenario, and it can safely be removed from the analysis. Another important consideration is that for one angle, a set of seeds can sometimes be clustered into two or more groups, and this is perfectly acceptable from a physics

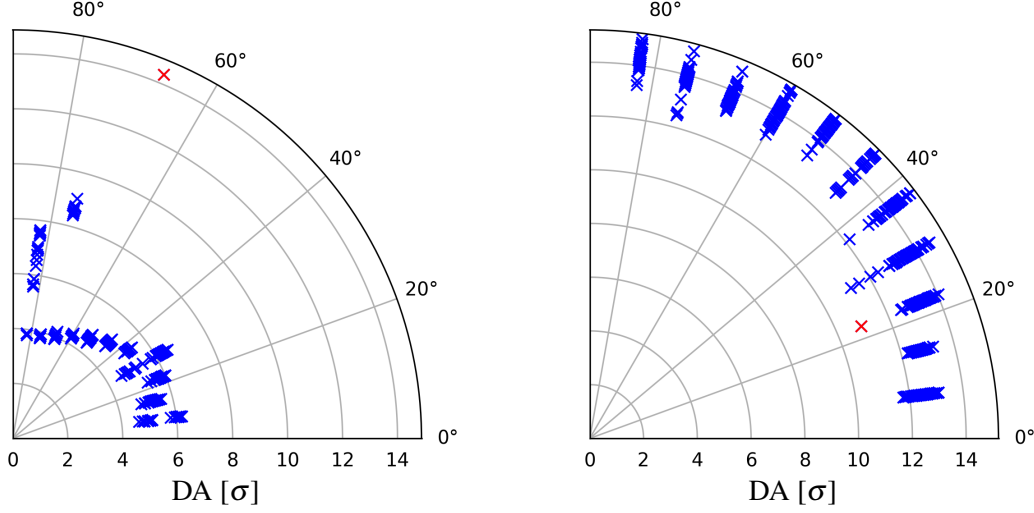


Figure 3: Left: Example of a DA computation where an outlier is correctly flagged. Right: Example of a potential false positive, where further investigation is needed.

point of view. It is of course of paramount importance that this scenario is distinguished from a set of true outliers.

We chose to tackle these issues by cutting down the outlier detection algorithm into several steps. First, for each angle the DA values for the different seeds are rescaled between the minimum and maximum values, as to quantify the outlier strength in function of the density of the seeds. These points are then clustered with help of the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm [18], where a cluster is defined as including at least three points. For those points that are not assigned to any cluster after applying the algorithm, the Local Outlier Factor (LOF) is calculated to quantify their outlier strength. Finally, a minimum threshold on the distance between an outlier and the nearest cluster is imposed, to avoid being overly strict on what is recognised as an outlier. For practical reasons, a DA value can only be an outlier when it is at a minimum or a maximum for that angle, as single points in between (but far away from) two clusters would not be considered outliers. The algorithm is tuned with bias towards false positives rather than false negatives. False positives will be removed anyway, because so far our algorithm is meant to be supplemented by a manual review.

The application of the algorithm is shown in Fig. 3; on the left side we see an example of a point that is clearly an outlier, and correctly identified as such by the algorithm. On the right side, we see an example of a point that seems to be an outlier for this specific angle. However, when comparing it to its neighbouring angles this might not necessarily be true anymore. If one or both minima of the neighbouring angles are points from the same seed and have a value similar to the outlier, it might be a false positive instead and could be kept. On the other hand, the governing dynamics that defines the DA can be very different from angle to angle. Even when the neighbouring DA values are similar, the value might still be an outlier, e.g. due to resonance structures. It is clear that particular care needs to be taken in these cases before drawing any conclusions, and additional investigations might be advisable. Another important point is to investigate the dependence of the

number of outliers on the angular distribution and on the seed number. For most studies, we indeed observe that outliers are far more present at high angles and that certain seeds contain far more outliers than others. This potentially gives us more insight into the sensitivity of the underlying physics; investigations on this matter are still on-going.

A second investigation of the usage of ML techniques in DA studies concerns the evolution of DA over time. It is known from theory that over time the DA shrinks, following well-defined scaling laws [19, 20, 21, 22]. These laws can potentially be used to extrapolate a CPU-intensive simulation to much longer, and hence more realistic, timescales. So far this has led to very promising results, however, the fitting procedure is rather sensitive, making it less reliable for prediction. We conducted some very preliminary research into using a Recurrent Neural Network (RNN) [23] to extrapolate the DA. This is of course an overly simplistic, brute-force approach as it contains no information on the theoretical knowledge about the actual scaling law. The results indicate that the RNN is not really able to make an accurate prediction. One very interesting alternative could be to use ML techniques to improve the above-described fitting procedure, e.g. by finding a set of optimal fitting weights for the deterministic models. This is currently being investigated.

5. Beam Lifetime Optimisation

The so-called beam lifetime τ is the time after which the beam intensity $N(t)$ of an accelerator would have decreased down to a fraction $1/e$ of its initial value. When considering only the intensity loss due to the protons that are lost in collisions, i.e. the burn-off, this beam lifetime is inversely proportional to the probability per second of a single proton colliding with an opposite-moving proton. The intensity evolution is then governed by the simple equation

$$N(t) = N_0 e^{-\frac{t}{\tau}}. \quad (5.1)$$

With the nominal design specifications of the LHC, the proton collision probability is around $3 \cdot 10^{-6} \text{ s}^{-1}$ giving rise to a beam lifetime of $\tau_{\text{bo}} \sim 90 \text{ h}$ [3]. However, several other processes affect the beam lifetime as well, like proton-residual gas collisions, Coulomb scattering, and magnetic field imperfections, just to name a few. The total beam lifetime is then simply equal to the inverse sum of all parts, $1/\tau_{\text{tot}} = 1/\tau_1 + 1/\tau_2 + \dots$

The beam lifetime can be acquired in real-time during operation by measuring the normalised instantaneous losses at the main collimators, and extrapolating those to find the value of t at $N(t) = 1/e N_0$. This measurement of instantaneous beam lifetime is a valuable indicator of the overall quality of the beam for a given fill of the machine. During the many years of operation of the LHC, it has been observed that its beam lifetime can vary drastically depending on the operational settings at hand; the main aim behind the research presented in this section is to link the former to the latter. The system hence developed would be capable of determining the optimal set of operational parameters, that maximises the lifetimes of both beams given a specific machine configuration. In this way, it could be used to assist operators in the control room during the decision-making process.

To make maximal use of the large amount of data acquired through the many instrumentation systems, the surrogate LHC model is constructed using a supervised-learning framework. The




Figure 4: Beam lifetime as a function of the LHC working point as measured for Beam 1. Red dot: Lifetime-optimised working point, as determined by the model. The model prediction is close, but not exactly equal, to the measured maximum.

various inputs of the model are the operational parameters of the machine, such as tunes, chromaticities, and magnet currents, while its outputs are the associated lifetimes of each of the two beams. The data used cover an entire operational year, but we restrict ourselves to a small part of the complete machine cycle, in order to be able to process the huge data volume and to simplify the input/output relationship. This part corresponds to the first phase of each LHC fill where the energy is kept at a low level, just after completing the filling of the machine with protons from CERN’s injectors chain. Using this part of the complete cycle is of course but a first step, and in order to create a complete model at a later stage, one has to consider data from the full machine cycle. Several supervised learning models were trained and compared, and the best performance was achieved with a Gradient Boosted Decision Tree model [24]. Once the surrogate model is trained, it can be paired with a variety of optimisers in order to extract the optimal machine configuration from the trained lifetime response. In our case, an off-the-shelf simplex optimiser [25] was used to this end.

An important first observation was that the distribution of the input data was far from ideal. This is to be expected, as the operators, rightfully, do not explore the input parameter space during physics runs, but keep their settings at values that are expected to give optimal results. However, this creates several internal correlations between the different input parameters, which might severely hamper the predictive power of the surrogate model. To address this, the input parameter space was explored during a dedicated machine experiment, in which multiple random tune scans were performed over varying machine configurations. The data collected during this experiment

are used to benchmark and supplement the current model.

A second observation is that we uncovered a number of instabilities that increase the beam emittances. This is to be avoided, as in general, emittance blow-up reduces the machine performance significantly, in terms of luminosity. Unfortunately, such instabilities are so far not taken into account by the model. Even though the surrogate model is obviously falling short by not including the former, it demonstrates inherent value when restricting ourselves to the naive lifetime optimisation problem. In a vertical vs horizontal tune diagram, the model finds the optimal region in terms of lifetime, see Fig. 4, although it falls short of the maximum.

The presented research is only a first approach to using data-driven ML techniques to develop a surrogate model that relates the LHC beam lifetime to its operational settings. The current surrogate model has some important limitations, but is nonetheless able to generate optimal values that are close to the experimental measurements. To overcome the prevailing limitations, this naive approach will need to take the emergence of instabilities into account, and address these in parallel to the lifetime optimisation using a multi-objective optimisation framework.

6. Conclusions

With these selected examples we have demonstrated that ML is steadily finding its place in the field of accelerator physics. We have presented some applications of ML techniques at the CERN Large Hadron Collider, using a variety of algorithms based on both supervised and unsupervised learning.

Unsupervised learning has been shown to be adequate in two of our study cases, both concerning the detection of outliers. In one of these, it is used to identify outliers in optics measurements. This is important to avoid being dominated by faulty BPMs when extracting the values of optical functions from measurement data. In the other study, unsupervised learning algorithms are used to detect outliers in DA results from beam tracking simulations. This aids in spotting numerical errors or remove unlikely realisations of the acceleration lattice. For both studies, the algorithms work as expected and will be fine-tuned further during their usage.

Supervised learning has been used in three of our study cases: to automatise the alignment of collimators during beam commissioning, to correct optical functions also during the beam commissioning stage, and to predict the beam lifetime based on the operational settings of the machine. For the collimator alignment, the time gain as compared to manual alignment has been impressive enough to make the ML implementation the default one. A continued cross-talk analysis will allow in the future to perform more alignments in parallel. Concerning the optics corrections, the basic Neural Network implementation produced already very interesting and encouraging results. A larger dataset is being generated and more error sources and non-linearities are being added in order to create a more general model. Finally, the beam lifetime ML model showed satisfactory results, once the parameter space was uncorrelated by a dedicated machine experiment. The naive model currently created will be used as a basis to build upon to further improve the framework.

References

- [1] E. Bozoki and A. Friedman, *Neural networks and orbit control in accelerators*, in *Proc. 4th European Particle Accelerator Conf. (EPAC'94)*, London, England, (1994) 1589.

- [2] E. Meier, Y. E. Tan and G. LeBlanc, *Orbit Correction Studies using Neural Networks*, in *Proc. 3rd Int. Particle Accelerator Conf. (IPAC'12)*, New Orleans (LA), USA, (2012) 2837.
- [3] O. S. Brüning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole et al., *LHC Design Report*, CERN Yellow Reports: Monographs. CERN, Geneva, 2004, [10.5170/CERN-2004-003-V-1](#).
- [4] R. W. Aßmann, E. B. Holzer, J.-B. Jeanneret, V. Kain, S. Redaelli, G. Robert-Demolaize et al., *Expected performance and beam-based optimization of the LHC collimation system*, in *Proc. 9th European Particle Accelerator Conf. (EPAC'04)*, Lucerne, Switzerland, (2004) 1825.
- [5] G. Azzopardi, G. Valentino, B. Salvachua and A. Muscat, *Automatic spike detection in beam loss signals for LHC collimator alignment*, *Nucl. Instrum. Methods Phys. Res.* **A934** (2019) 10.
- [6] G. Azzopardi, G. Valentino, B. Salvachua, S. Redaelli and A. Muscat, *Software Architecture for Automatic LHC Collimator Alignment using Machine Learning*, in *Proc. 17th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'19)*, New York (NY), USA, (2019).
- [7] G. Azzopardi, G. Valentino, B. Salvachua, S. Redaelli and A. Muscat, *Automatic Beam Loss Threshold Selection for LHC Collimator Alignment*, in *Proc. 17th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'19)*, New York (NY), USA, (2019).
- [8] G. Azzopardi, B. Salvachua and G. Valentino, *Data-driven cross-talk modeling of beam losses in LHC collimators*, *Phys. Rev. Accel. Beams* **22** (2019) 083002.
- [9] G. Azzopardi, A. Muscat, S. Redaelli, B. Salvachua and G. Valentino, *Operational results of LHC collimator alignment using machine learning*, in *Proc. 10th Int. Particle Accelerator Conf. (IPAC'19)*, Melbourne, Australia, (2019) 1208.
- [10] G. Azzopardi, B. Salvachua, G. Valentino, S. Redaelli and A. Muscat, *Operational results on the fully automatic LHC collimator alignment*, *Phys. Rev. Accel. Beams* **22** (2019) 093001.
- [11] R. Calaga and R. Tomás, *Statistical analysis of RHIC beam position monitors performance*, *Phys. Rev. Accel. Beams* **7** (2004) 042801.
- [12] F. T. Liu, K. M. Ting and Z.-H. Zhou, *Isolation forest*, in *Proc. 8th IEEE Int. Conf. on Data Mining (ICDM'08)*, IEEE Computer Society, (2008) 413.
- [13] E. Fol, J. M. Coello de Portugal and R. Tomás, *Unsupervised Machine Learning for Detection of Faulty Beam Position Monitors*, in *Proc. 10th Int. Particle Accelerator Conf. (IPAC'19)*, Melbourne, Australia, (2019) 2668.
- [14] M. Aiba, S. Fartoukh, A. Franchi, M. Giovannozzi, V. Kain, M. Lamont et al., *First β -beating measurement and optics analysis for the CERN Large Hadron Collider*, *Phys. Rev. ST Accel. Beams* **12** (2009) 081002.
- [15] T. Persson, F. Carlier, J. Coello de Portugal, A. Garcia-Tabares Valdivieso, A. Langner, E. H. Maclean et al., *LHC optics commissioning: A journey towards 1% optics control*, *Phys. Rev. Accel. Beams* **20** (2017) 061002.
- [16] E. Fol, J. M. Coello de Portugal, G. Franchetti and R. Tomás, *Optics corrections using machine learning in the LHC*, in *Proc. 10th Int. Particle Accelerator Conf. (IPAC'19)*, Melbourne, Australia, (2019) 3990.
- [17] E. Fol, J. M. Coello de Portugal, G. Franchetti and R. Tomás, *Application of Machine Learning to Beam Diagnostics*, in *Proc. 39th Int. Free Electron Laser Conf. (FEL'19)*, Hamburg, Germany, (2019).

- [18] M. Ester, H.-P. Kriegel, J. Sander and X. Xu, *A density-based algorithm for discovering clusters in large spatial databases with noise*, in *Proc. 2nd Int. Conf. on Knowledge Discovery and Data Mining (KDD'96)*, Portland (OR), USA, (1996) 226,
<https://www.aaai.org/Papers/KDD/1996/KDD96-037.pdf>.
- [19] M. Giovannozzi and F. F. Van der Veken, *Description of the luminosity evolution for the CERN LHC including dynamic aperture effects. Part I: the model*, *Nucl. Instrum. Methods Phys. Res. A* **905** (2018) 171 [1806.03058].
- [20] M. Giovannozzi and F. F. Van der Veken, *Description of the luminosity evolution for the CERN LHC including dynamic aperture effects. Part II: application to Run 1 data*, *Nucl. Instrum. Methods Phys. Res. A* **908** (2018) 1 [1806.03059].
- [21] F. F. Van der Veken and M. Giovannozzi, *Scaling Laws for the Time Dependence of Luminosity in Hadron Circular Accelerators based on Simple Models of Dynamic Aperture Evolution*, in *Proc. 61st ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (HB'18)*, Daejeon, Korea, (2018) 260.
- [22] A. Bazzani, M. Giovannozzi, E. H. Maclean, C. E. Montanari, F. F. Van der Veken and W. Van Goethem, *Advances on the modeling of the time evolution of dynamic aperture of hadron circular accelerators*, *Phys. Rev. Accel. Beams* **22** (2019) 104003 [1909.09516].
- [23] G. Valentino. Personal conversation.
- [24] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma et al., *Lightgbm: A highly efficient gradient boosting decision tree*, in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan et al., eds., pp. 3146–3154, Curran Associates, Inc., (2017),
<http://papers.nips.cc/paper/6907-lightgbm-a-highly-efficient-gradient-boosting-decision-tree.pdf>.
- [25] J. A. Nelder and R. Mead, *A Simplex Method for Function Minimization*, *The Computer Journal* **7** (1965) 308
[<http://oup.prod.sis.lan/comjnl/article-pdf/7/4/308/1013182/7-4-308.pdf>].