

PAPER • OPEN ACCESS

# Surrogate modeling of the CLIC final-focus system using artificial neural networks

To cite this article: J. Ögren *et al* 2021 *JINST* **16** P05012

View the [article online](#) for updates and enhancements.

You may also like

- [Improved step-by-step chromaticity compensation method for chromatic sextupole optimization](#)  
Gang-Wen Liu, , Zheng-He Bai et al.
- [Analysis and correction of linear optics errors, and operational improvements in the Indus-2 storage ring](#)  
Riyasat Husain and A. D. Ghodke
- [Low emittance electron storage rings](#)  
E B Levichev



**PRIME**  
PACIFIC RIM MEETING  
ON ELECTROCHEMICAL  
AND SOLID STATE SCIENCE

HONOLULU, HI  
Oct 6–11, 2024

Abstract submission deadline:  
**April 12, 2024**

**Learn more and submit!**



**Joint Meeting of**

The Electrochemical Society  
•  
The Electrochemical Society of Japan  
•  
Korea Electrochemical Society



RECEIVED: February 2, 2021

REVISED: February 23, 2021

ACCEPTED: February 26, 2021

PUBLISHED: May 12, 2021

# Surrogate modeling of the CLIC final-focus system using artificial neural networks

J. Ögren,<sup>a,\*</sup> C. Gohil<sup>a,b</sup> and D. Schulte<sup>a</sup>

<sup>a</sup>European Organization for Nuclear Research (CERN),  
Geneva, Switzerland

<sup>b</sup>John Adams Institute, University of Oxford,  
Oxford, U.K.

E-mail: [jim.ogren@igp.uu.se](mailto:jim.ogren@igp.uu.se)

**ABSTRACT:** Artificial neural networks can be used for creating surrogate models that can replace computationally expensive simulations. In this paper, a surrogate model was created for a subset of the Compact Linear Collider (CLIC) final-focus system. By training on simulation data, we created a model that maps sextupole offsets to luminosity and beam sizes, thus replacing computationally intensive tracking and beam-beam simulations. This model was then used for optimizing the parameters of a random walk procedure for sextupole alignment.

**KEYWORDS:** Accelerator modelling and simulations (multi-particle dynamics; single-particle dynamics); Beam dynamics

ARXIV EPRINT: [2009.06454](https://arxiv.org/abs/2009.06454)

---

\*Corresponding author.



---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>The CLIC final-focus system</b>	<b>2</b>
2.1	Tuning	3
2.2	Sextupole misalignment	3
<b>3</b>	<b>Surrogate model</b>	<b>4</b>
3.1	Data generation	4
3.2	Model setup	5
3.3	Model training	7
3.4	How much data is needed?	7
3.5	Final model	9
<b>4</b>	<b>Model verification</b>	<b>9</b>
<b>5</b>	<b>Random walk optimization</b>	<b>10</b>
<b>6</b>	<b>Conclusions</b>	<b>12</b>

---

## 1 Introduction

The field of machine learning has grown immensely in the past decade and is now finding its way into a wide range of fields. Particle accelerators are controllable physical systems with rich complexities and offer interesting applications that can profit from different machine learning methods. Examples of applications in particle accelerators are anomaly detection in superconducting magnets [1], automatic control of collimators [2], electron bunch profile prediction [3, 4], electron beam size prediction in circular accelerators [5] and ion source prediction [6]. Different advanced techniques have also been used for optimization and automatic control in free-electron lasers [7, 8], longitudinal phase space control [9], dynamic aperture optimization [10] and optics correction [11, 12]. A more comprehensive list of machine learning applications in particle accelerators can be found in [13].

The field of particle accelerators relies to a large extent, especially in design and optimization, on high-fidelity simulations and there is an abundance of simulation tools. In many cases simulations are computationally expensive, which makes optimization time consuming and online modeling problematic. A surrogate model is an approximative model that can mimic the system and is particularly useful for applications that require extensive simulations. The surrogate model, which is fast to evaluate, can be used for optimization or as an online model. In this paper we make use of artificial neural networks and supervised deep learning to create a surrogate model for a part of the Compact Linear Collider (CLIC) with computationally expensive simulations: the final-focus system.

**Table 1.** 380 GeV CLIC beam parameters.


Parameter	unit	value
Norm. emittance (end of linac) $\gamma\epsilon_x/\gamma\epsilon_y$	[nm]	900 / 20
Norm. emittance (IP) $\gamma\epsilon_x/\gamma\epsilon_y$	[nm]	950 / 30
Beta function (IP) $\beta_x^*/\beta_y^*$	[mm]	8.2 / 0.1
Target IP beam size $\sigma_x^*/\sigma_y^*$	[nm]	149 / 2.9
Bunch length $\sigma_z$	[ $\mu\text{m}$ ]	70
rms energy spread $\delta_p$	[%]	0.35
Bunch population $N_e$	[ $10^9$ ]	5.2
Number of bunches $n_b$		352
Repetition rate $f_{\text{rep}}$	[Hz]	50
Luminosity $L$	[ $10^{34}\text{cm}^{-2}\text{s}^{-1}$ ]	1.5
Peak luminosity $L_p$	[ $10^{34}\text{cm}^{-2}\text{s}^{-1}$ ]	0.9

## 2 The CLIC final-focus system

CLIC is a proposed linear electron-positron collider [14–16] at CERN. To produce high luminosities, linear colliders rely on ultra-small beam sizes at the interaction point. This requires emittance preservation along the whole machine and puts tight tolerances on alignment and other imperfections. Being able to reach nominal performance under imperfect conditions is crucial for the reliability of such a machine. Nominal parameters for the 380 GeV energy stage are summarized in table 1.

The final-focus system constitutes the final 780 m of each beamline, see figure 1. Two strong quadrupole magnets at the end of the system de-magnify the beam in both transverse directions. Since particles of different energies are bent differently in a magnetic field, a beam with an energy spread will not be focused to a single point. This effect is called chromaticity and without compensation the beam size at the interaction point will be greatly inflated with diminished luminosity as a result. The chromaticity correction involves nonlinear magnets and a delicate compensation of unwanted aberrations. Complex systems such as this require fine-tuning and are challenging to design.

The current version of the CLIC final-focus system has the final quadrupole doublet mounted outside the detector [17, 18] with  $L^* = 6$  m — distance between the final quadrupole and point of collision. The system consists of 20 quadrupole magnets, 6 sextupole magnets and 2 octupole magnets. Bending magnets generate dispersion and in combination with the sextupole magnets, an energy-dependent focusing is generated that compensates the chromaticity generated in the final doublet. The final-focus system follows a local chromaticity correction scheme [19] where two sextupoles are placed at the final doublet and then additional sextupoles are placed upstream with optics designed in such a way that unwanted geometric aberrations cancel. Figure 1 shows the beta functions ( $\beta_{x,y}$ ) and the dispersion profile ( $\eta_x$ ) of the final-focus system.



**Figure 1.** Optical functions for the 380 GeV CLIC final-focus system with  $L^* = 6$  m. Six sextupoles are placed in dispersive sections in a local chromaticity correction scheme.


## 2.1 Tuning

The final-focus system is sensitive to imperfections and the challenge is to achieve small beam sizes, and high luminosity, for a system with imperfections. Successful operation of a linear collider requires effective and fast tuning. Tuning studies have been done extensively for all parts of CLIC and for the final-focus system in particular. A recent report [20] considered single-beam tuning of the 380 GeV CLIC final-focus system including static imperfections such as magnet offsets, roll errors and strength errors. Single beam means that only one side, i.e. one of the two beamlines, of the final-focus system is tracked and the beam at the interaction point is mirrored for the beam-beam simulation.

In that study, the tuning procedure consisted of beam-based alignment that corrected the linear system, followed by sextupole alignment and tuning with sextupole and octupole knobs. Due to the nonlinearities of the system together with synchrotron radiation, small imperfections can have dramatic impact on beam size at the interaction point and consequently a substantial decrease in luminosity. The purpose of the tuning simulations is to find robust algorithms for achieving luminosity under imperfect conditions and thus ensuring the reliability of the collider. In all the simulations we use PLACET [21] for tracking the beam through the final-focus system lattice and GUINEA-PIG [22] for a full beam-beam simulation.

## 2.2 Sextupole misalignment

Let us consider a single imperfection: transverse misalignment of sextupoles. Experience has shown that this imperfection has a large impact on luminosity. To quantify the impact we run simulations on a perfect machine and add transverse sextupole offsets assuming a normal distribution. For different rms values, we track 20 cases and run a beam-beam simulation to compute the luminosity. Similar to ref. [20] we consider the single-beam case. Figure 2 shows the average luminosity decreasing as



**Figure 2.** The average luminosity for the perfect machine with sextupole offsets of varying rms values. As the sextupole misalignments grow, luminosity quickly deteriorates. The shaded area show the maximum and minimum values.

the rms value of the transverse sextupole offset increases. Since transverse sextupole misalignments have a large impact, it is essential to have a quick and robust method for tuning the sextupoles. In the next section we build a surrogate model for the impact of sextupole misalignments on luminosity and beam size at the interaction point.

### 3 Surrogate model


We only consider transverse sextupole offsets and aim to develop a model that outputs: luminosity, peak luminosity, horizontal beam size and vertical beam size. The peak luminosity is the luminosity of collisions within 1% of the nominal energy. CLIC specifications require a minimum of 60% of collisions occur at the nominal energy, see table 1.

An analytical approach is not feasible due to the complexity of the system. The nonlinear fields of the sextupole magnets together with effects of synchrotron radiation make it difficult to compute the resulting beam distribution at the interaction point. Furthermore, the disruption parameter in the beam-beam interaction is high, which means that the beams have considerable effect on each other. In an electron-positron collider the transverse electromagnetic fields of one beam focuses the opposing beam (‘pinch effect’). In such circumstances analytical estimates of luminosity are not reliable.

Artificial neural networks are large connections of small units (artificial neurons) and by training on data the models can learn to perform different tasks such as classification, clustering or regression. In our case we make use of regression and train a model to map input to known outputs. In the context of machine learning, this is known as supervised learning. We use the machine learning system TensorFlow [23] interfaced via the Python library Keras [24] for training the models. The goal of a model is to make accurate predictions based on unseen data. Such models are said to generalize well. Large data sets are often needed to train artificial neural networks successfully.

#### 3.1 Data generation

Assuming normal distributions with rms values of 5, 10 and 20 μm, we added random offsets to the sextupoles. For each case we tracked  $10^5$  macroparticles through the final-focus system, ran a




**Figure 3.** Histograms of the data set showing the four outputs: luminosity (top left), peak luminosity (top right), horizontal beam size (bottom left) and vertical beam size (bottom right). The vertical line in the luminosity plot shows a cut that was used before training the models.

beam-beam simulation and saved sextupoles positions, beam sizes at the interaction point and the output from the beam-beam simulation: luminosity and peak luminosity. The data was generated in parallel using the CERN batch system and in a few days we had accumulated a data set of 450,000 cases. Each case took approximately 5 minutes to run on an ordinary desktop computer (8 cores, CPU computations only). Figure 3 shows histograms of the accumulated data. We use 100 bins in the histogram because we found that this was sufficient to resolve the shape of the distribution.

Since the numeric values of the inputs and outputs have very different ranges we must transform the data before training our models. We standardize the data set, i.e. apply transformations  $X_{\text{standard}} = (X - \mu)/\sigma$ , so that each column in our data set has zero mean and unit variance. The mean  $\mu$  and standard deviation  $\sigma$  are computed using the training data set. To the columns containing output data (luminosity, peak luminosity, horizontal and vertical beam sizes) we also apply a power transformation to make distributions closer to Gaussian, see figure 4. Standardization and power transformation are common practices to make the training of artificial neural networks numerically stable. Both these transformations are available in the Python library Scikit-learn [25]. When testing the model, the inverse transformations are applied to the output predictions.

### 3.2 Model setup

There are no clear guidelines of how to set up an artificial neural network and how to select a model architecture since it depends on the problem. In general, if the model has a complexity that is too



**Figure 4.** Histograms of the data set after power transformation and standardization: luminosity (top left), peak luminosity (top right), horizontal beam size (bottom left) and vertical beam size (bottom right). The data is also standardized to have zero mean and unit variance. The vertical line in the luminosity plot shows a cut that was used before training the models.

low it will not be able to mimic nonlinear behavior and if the model has too much complexity it can suffer from overfitting, where the model does not generalize well.

First we consider a simple case where our model consists of two hidden layers of twelve neurons each. We use a model with fully connected layers in a feedforward configuration (every neuron in one layer connects to every neuron in the subsequent layer). Figure 5 shows a schematic of the model. The 12-dimensional input vector can be written as


$$X = (x_1, \dots, x_{12})^T = (S1_x, S1_y, \dots, S6_x, S6_y)^T \quad (3.1)$$

where  $S1_x$  denotes the horizontal position of the first sextupole and so on. The 4-dimensional output can be written

$$Y = (y_1, \dots, y_4)^T = (L, L_p, \sigma_x, \sigma_y)^T \quad (3.2)$$

where  $L$  denotes the luminosity,  $L_p$  the peak luminosity and  $\sigma_x$  ( $\sigma_y$ ) the horizontal (vertical) beam size at IP. The input to a neuron is the sum of weighted inputs and a bias, and during the training of the model, weights and biases of all neurons are adjusted to minimize a loss function — in our case the mean squared error.





**Figure 5.** A schematic of an artificial neural network model that maps the horizontal and vertical positions of the six sextupole magnets in the CLIC final-focus system to the resulting luminosity, peak luminosity and horizontal and vertical beam sizes. This model consists of two fully connected hidden layers in a feedforward configuration.

### 3.3 Model training

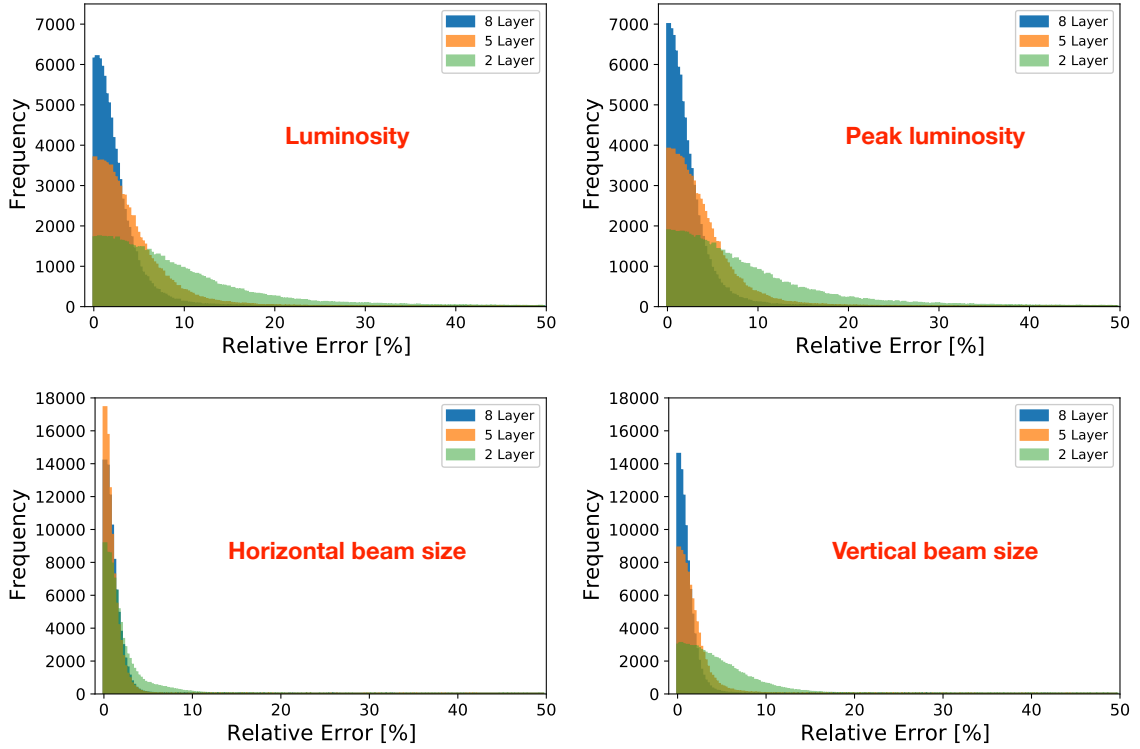
Following common practice, we split the complete data set and reserve 20% for testing. The remaining 80% is split into a training set (90%) and a validation set (10%). During training the model trains on the training set and evaluates on the validation set. The purpose of validation is to monitor overfitting. After training is done, i.e. when there is either no improvements or overfitting starts, the model performance is evaluated on the testing data — data never previously seen by the model.

To assess the model performance we compare the predictions to the known outputs in the test data set. As a figure of merit we use the mean absolute percentage error (MAPE). From figure 3 it is clear there are a few outliers with low luminosity. This results in large errors which can inflate the MAPE. Therefore, we make a cut in our testing data and only consider cases with luminosities above  $5 \times 10^{31} \text{ cm}^{-2}\text{s}^{-1}$ .

Figure 6 shows histograms of the performance of the three different models. We tested a different number of hidden layers (2-9) and the plots show three cases: two, five and eight layers. It is clear that a model with two layers does not have sufficient complexity and gives rather large relative errors. There is a substantial improvement when we increase the number of layers from two to five and eight. Table 2 shows the mean of the absolute percentage error. Eight layers gave the lowest MAPE and there was no improvement from increasing the number of layers beyond eight. We also tried other network architectures, for example: we increased the number of nodes per layer; tried reducing the number of nodes in subsequent layers and tried architectures with a large number of nodes in the first layer followed by smaller layers. In the end it did not seem possible to lower the MAPE further.

### 3.4 How much data is needed?

The full data set used in the previous section consisted of 450,000 samples containing misalignments drawn from a random distribution. 360,000 samples were randomly selected for the training dataset




**Figure 6.** Histograms of the absolute percentage error for the four different outputs: luminosity (top left), peak luminosity (top right), horizontal beam size (bottom left) and vertical beam size (bottom right). Three different models with two, five and eight layers are compared.

**Table 2.** Model evaluation on testing data. Mean absolute percentage error for the four model outputs for three different model architectures.

	2 layers	5 layers	8 layers
Luminosity [%]	8.87	4.01	2.47
Peak luminosity [%]	8.42	3.77	2.29
Horizontal beam size [%]	2.06	0.913	1.02
Vertical beam size [%]	4.74	1.68	1.04

and the remaining 90,000 samples were used for testing. To investigate the influence of number of data points, we train the model with eight layers on training sets of different sizes. Subsets of [1000, 5000, 50000, 100000, . . . , 300000] are drawn randomly from the full training set. Each case is trained ten times and tested on the same full testing set used in previous section. Figure 7 shows the mean and standard deviation of the MAPE for the models trained on different number of data points. It is clear that more data is better but with diminishing returns. Already at 50,000 data points the model performance approaches the final result. For a model valid for luminosities above  $5 \times 10^{31} \text{ cm}^{-2}\text{s}^{-1}$ , 50,000 data points seems sufficient. On the other hand, in order to increase the range of the model more data and more data over a wider range is likely required.



**Figure 7.** Mean absolute percentage error (MAPE) for models trained on different number of training samples. Each case is trained ten times and the markers show the mean values and the shaded areas the standard deviation. It is clear that more data is better but with diminishing returns.

### 3.5 Final model


The final model consisted of eight fully connected layers in a feedforward configuration, cf. figure 5. Each layer had 12 neurons that used the rectified linear unit as an activation function. The model was trained on the full training set of 360,000 data points using the Adam optimizer [23] — an adaptive learning method implemented in TensorFlow — and mean squared error as loss function. We used a batch size (number of training samples used in every iteration) of 50 and trained for 400 epochs (the number of times the complete training data set is passed through). The final model took a total of 8-10 hours to train on an ordinary desktop computer (8 cores, CPU computation). The results from testing are presented in table 2.

Surely there is potential room for improvement. For instance, the hyperparameter optimization was not done exhaustively and other model architectures might be better. However, finding the optimum model is not within the scope of this paper. Instead, we want to show that the model was good enough to be useful. To assess the final model’s usability we compared it to full simulations.

## 4 Model verification

To verify the model performance we compare it to the full simulation. A nice feature of the tracking code PLACET is it can easily interface Python. This made it possible to seamlessly load and interface the trained machine learning model in our simulations. At each iteration we ran the full beam-beam simulation and at the same time checked the prediction of the machine learning model for comparison.

Figure 8 shows a random walk tuning of the transverse sextupole position. Random sextupole offsets were added to a perfect machine and using a random walk the luminosity was optimized within 350 iterations. At each iteration the beam distribution was tracked and the luminosity was computed with a full beam-beam simulation. At each step we also checked the prediction of the machine learning model (plotted in solid and dotted lines). We plot luminosity normalized to nominal luminosity,




**Figure 8.** Example of a random walk optimization on the full simulation.

$L_0$ , which is the luminosity of the unperturbed lattice. Note that sometimes the luminosity in the random walk tuning is higher than the nominal value since sextupole offsets can introduce waist shifts that enhance luminosity via the pinch effect and the nominal luminosity does not include this effect. Over the range of almost two orders of magnitude there was excellent agreement between the machine learning model and the full simulation. The MAPE in this case was 2.6% for luminosity, 3.0% for peak luminosity, 1.8% for horizontal beam size and 2.6% for vertical beam size. The MAPE was higher in this simulation compared to the testing dataset, which showed a MAPE of approximately 2% for luminosity and 1% for beam size. This is due to the small sample size in this simulation which leads to a higher variance in the MAPE calculated. The random walk method always used the luminosity from the full simulation and the machine learning model was only evaluated parasitically.

The converse was also tested, where the output of the machine learning model guided the random walk optimization and the full simulation was done parasitically for comparison. Both cases showed similar results.

## 5 Random walk optimization

Running a single case on the full simulation, i.e. tracking  $10^5$  macroparticles and running the full beam-beam simulation, on a normal desktop computer takes a few minutes. Making a single prediction with the machine learning model takes a few milliseconds. Such a speed-up in performance



**Figure 9.** Grid search of gain and subset used in a random walk procedure. The points show the average values and shaded area the standard deviation. For each setting a 100 random cases were tested 100 times.

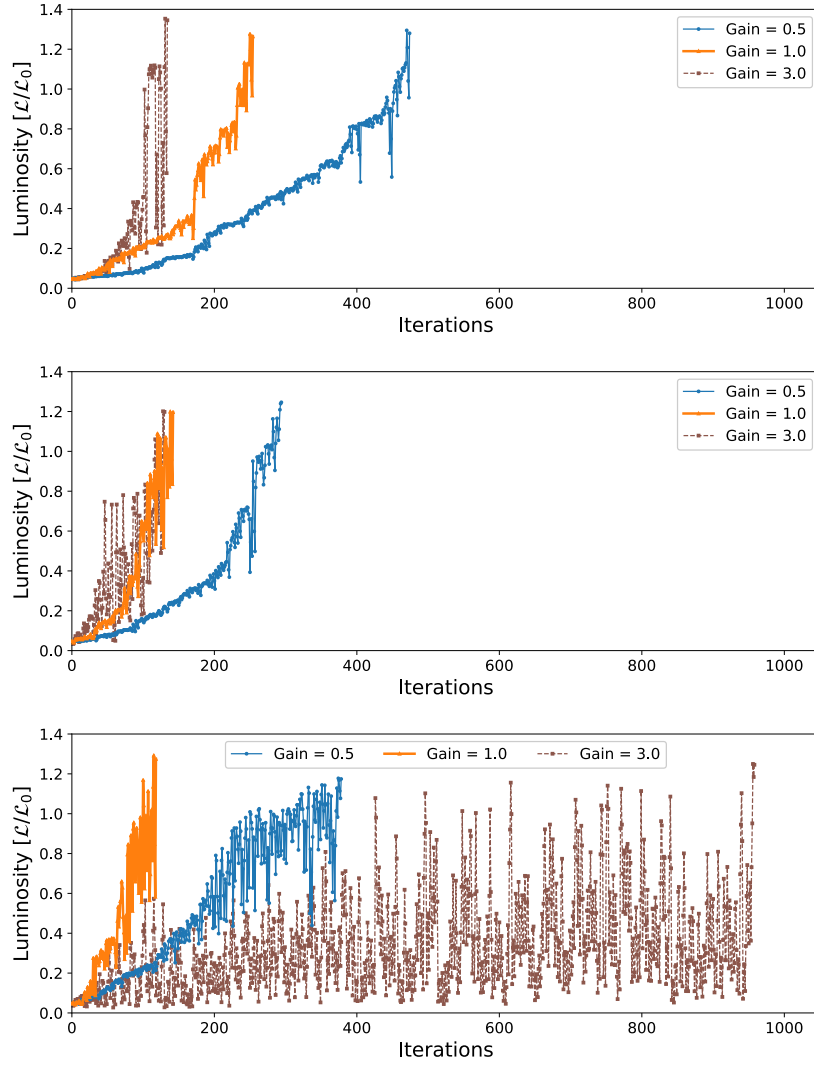
allows for systematic studies of optimization otherwise not feasible. As an example, we use our machine learning model to optimize a simple random walk procedure.

For the sextupoles, there are twelve degrees of freedom: six sextupoles and two transverse directions. We devise a random walk procedure in the following way. Out of the twelve degrees of freedom a subset of the degrees of freedom is selected randomly. In this subset a direction is selected at random, i.e. we specify a random linear combination of the degrees of freedom in the subset. Four points along the selected direction,  $g[-1, -0.5, 0.5, 1]$  with gain  $g$ , are checked. If the highest luminosity of the four points exceeds the luminosity from the previous iteration, that points is selected. Otherwise the procedure moves to the next iteration and selects a new subset and direction. The whole procedure is iterated until a certain luminosity is reached. There are two parameters to be optimized: gain and size of the subset.

We performed a grid search of gains  $[0.5, 1.0, 1.5, 2.0, 2.5, 3.0]$  and subsets 1-12 on the eight layer artificial neural network model. For each setting we tested 100 different initial conditions with random sextupole offsets. For each initial condition we ran the random walk procedure 100 times and recorded the required number of iterations to reach the target luminosity. Figure 9 shows the average number of iterations needed for the different parameter settings. For a smaller subset a higher gain decreases the number of iterations whereas for a larger subset a smaller gain is preferable. The optimum seems to be a small subset of 2 or 3 and a higher gain.

To verify the optimum parameters from the grid search, we tested the random walk procedure using a full simulation. Figure 10 shows three different cases: subsets 2, 6 and 12. For each case, three different gains were tested. The results are in good agreement with the grid search using the machine learning model. Using a higher gain is better when the subset is small whereas a high gain and a large subset lead to large fluctuations and requires more iterations to reach target.

This random walk procedure was rather simple and could surely be optimized further. In a real machine there are other constraints to consider as well and in most cases a slow smoothly varying optimization might be preferred over a quicker procedure with large fluctuations. The point of this exercise was to show the utility of the surrogate model and that results from optimization on the surrogate model translated to the full simulation.



**Figure 10.** Random walk procedure tested on the full simulation. Three different gains were tested using subset = 2 (top), subset = 6 (middled) and subset = 12 (bottom). A higher gain is better when the subset is small and a smaller gain when subset is large.

## 6 Conclusions

In this paper we presented a surrogate model that maps sextupole offsets to resulting luminosities and beam sizes in the CLIC final-focus system. We trained an artificial neural network on a large set of simulation data. The model was verified by comparing its predictions to the full simulation with a mean absolute percentage error of a few percent. We used the model to optimize two parameters of a random walk procedure and compared the results to the full simulation, with good agreement.

We considered only a small subset of the system (transverse sextupole offsets) and the single-beam case. A more complete surrogate model could be realized by adding more imperfections to the system. Naturally, as the degrees of freedom increase, training the model becomes more

challenging and may require a larger data set but the principle remains the same. Nonetheless, as we showed with our example of optimizing the random walk procedure: even a model with limited scope can already be very useful. In our case it helped improving the sextupole alignment procedure and reduced the overall tuning time.

## References

- [1] M. Wielgosz, A. Skoczeń and M. Mertik, *Using LSTM recurrent neural networks for monitoring the LHC superconducting magnets*, *Nucl. Instrum. Meth. A* **867** (2017) 40 [[arXiv:1611.06241](#)].
- [2] 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.
- [3] C. Emma, A. Edelen, M.J. Hogan, B. O’Shea, G. White and V. Yakimenko, *Machine learning-based longitudinal phase space prediction of particle accelerators*, *Phys. Rev. Accel. Beams* **21** (2018) 112802.
- [4] A. Scheinker and S. Gessner, *Adaptive method for electron bunch profile prediction*, *Phys. Rev. ST Accel. Beams* **18** (2015) 102801.
- [5] S.C. Leemann et al., *Demonstration of machine learning-based model-independent stabilization of source properties in synchrotron light sources*, *Phys. Rev. Lett.* **123** (2019) 194801.
- [6] Y.B. Kong, M.G. Hur, E.J. Lee, J.H. Park, Y.D. Park and S.D. Yang, *Predictive ion source control using artificial neural network for RFT-30 cyclotron*, *Nucl. Instrum. Meth. A* **806** (2016) 55.
- [7] N. Bruchon, G. Fenu, G. Gaio, M. Lonza, F.A. Pellegrino and L. Saule, *Free-electron laser spectrum evaluation and automatic optimization*, *Nucl. Instrum. Meth. A* **871** (2017) 20.
- [8] A. Scheinker et al., *Model-independent tuning for maximizing free electron laser pulse energy*, *Phys. Rev. Accel. Beams* **22** (2019) 082802.
- [9] A. Scheinker, A. Edelen, D. Bohler, C. Emma and A. Lutman, *Demonstration of model-independent control of the longitudinal phase space of electron beams in the linac-coherent light source with femtosecond resolution*, *Phys. Rev. Lett.* **121** (2018) 044801.
- [10] Y. Li, W. Cheng, L.H. Yu and R. Rainer, *Genetic algorithm enhanced by machine learning in dynamic aperture optimization*, *Phys. Rev. Accel. Beams* **21** (2018) 054601.
- [11] Y. Li, R. Rainer and W. Cheng, *Bayesian approach for linear optics correction*, *Phys. Rev. Accel. Beams* **22** (2019) 012804 [[arXiv:1904.08214](#)].
- [12] E. Fol, J.M.C. De Portugal, G. Franchetti and R. Tomás, *Optics corrections using machine learning in the LHC*, in *Proceedings of the 10<sup>th</sup> Int. Particle Accelerator Conf.*, [THPRB077](#), Australia (2019).
- [13] A. Edelen et al., *Opportunities in machine learning for particle accelerators*, [arXiv:1811.03172](#).
- [14] M. Aicheler et al. eds., *A multi-TeV linear collider based on CLIC technology: CLIC conceptual design report*, [CERN-2012-007](#), CERN, Geneva, Switzerland (2012).
- [15] CLIC and CLIC<sub>DP</sub> collaboration, *Updated baseline for a staged Compact Linear Collider*, [CERN-2016-004](#), CERN, Geneva, Switzerland (2016) [[arXiv:1608.07537](#)].
- [16] CLIC ACCELERATOR collaboration, *The Compact Linear Collider (CLIC) — project implementation plan*, [CERN-2018-010-M](#), CERN, Geneva, Switzerland (2018) [[arXiv:1903.08655](#)].
- [17] N. Alipour Tehrani et al., *CLICdet: the post-CDR CLIC detector model*, [CLICdp-Note-2017-001](#), CERN, Geneva, Switzerland (2017).

- [18] F. Plassard, A. Latina, E. Marin, R. Tomás and P. Bambade, *Quadrupole-free detector optics design for the Compact Linear Collider final focus system at 3 TeV*, *Phys. Rev. Accel. Beams* **21** (2018) 011002.
- [19] P. Raimondi and A. Seryi, *A novel final focus design for future linear colliders*, *Phys. Rev. Lett.* **86** (2001) 3779.
- [20] J.A. Ogren, A. Latina, R. Tomas Garcia and D. Schulte, *Tuning of the CLIC 380 GeV final-focus system with static imperfections*, *CERN-ACC-2018-0055*, CERN, Geneva, Switzerland (2018).
- [21] *The tracking code placet*, <https://clicsw.web.cern.ch/clicsw/>.
- [22] D. Schulte, *Study of electromagnetic and hadronic background in the interaction region of the TESLA collider*, Ph.D. thesis, *DESY-TESLA-97-08*, Universität Hamburg, Hamburg, Germany (1996) [TESLA-97-08].
- [23] M. Abadi et al., *TensorFlow: large-scale machine learning on heterogeneous systems*, software available from <https://www.tensorflow.org/>, (2015).
- [24] F. Chollet et al., *Keras*, <https://keras.io>, (2015).
- [25] F. Pedregosa et al., *Scikit-learn: machine learning in Python*, *J. Machine Learning Res.* **12** (2011) 2825 [[arXiv:1201.0490](https://arxiv.org/abs/1201.0490)].