

CERN Summer Student Report

BSRH Auto-alignment

Lucie Kuncarova

Supervisors: Jan Pucek and Federico Roncarolo

Department: SY-BI-PM, CERN

September 5, 2025

European Organization for Nuclear Research (CERN)

Abstract

This report describes the work conducted during my internship at CERN on the automation of an alignment procedure for the Beam Synchrotron Radiation Halo monitor (BSRH). BSRH is an optical diagnostic instrument that monitors the transverse beam distribution in the LHC with a particular emphasis on the low-intensity halo. Its optical system has to be properly aligned for accurate measurements, with manual alignment being time-consuming and prone to drift.

To address this, a modular system of neural networks was developed, with each network responsible for aligning one optical path defined by a pair of targets. The models achieved high accuracy for the second and third optical paths, demonstrating the feasibility of data-driven alignment in such a complex optical system.

List of Acronyms

BSRH	Beam Synchrotron Radiation Halo monitor
CERN	European Organization for Nuclear Research
LHC	Large Hadron Collider
HL-LHC	High-Luminosity Large Hadron Collider
MD	Machine Development
NN	Neural Network
FFNN	Feedforward Neural Network

Contents

1	Introduction	3
2	The BSRH Instrument	5
2.1	Purpose in the LHC	5
2.2	Instrument Design	5
3	Alignment	8
3.1	Motivation	8
3.2	Alignment Algorithm Design	8
3.3	Data Collection	11
3.4	Modeling	13
3.4.1	First Optical Path	13
3.4.2	Second Optical Path	14
3.4.3	Third Optical Path	14
4	Conclusion	16

Chapter 1

Introduction

The CERN Large Hadron Collider (LHC) is the most powerful particle collider in the world, designed to collide particles at very high energies. To function efficiently and safely at these levels, many beam diagnostic devices are utilized to continuously measure properties of the beam and perform adjustments of the magnets' strength to keep the beam in closed orbit. These diagnostics are used for accelerator infrastructure protection and to maximize luminosity in the experiments.

One of the concerns of the machine protection panel is the beam halo, which is a collection of particles that reside in the tails of the transverse particle distribution. Though their relative intensity is low, the absolute halo charge could lead to increased particle loss, which could result in magnet quenches or excessive load on the collimation system. Monitoring the halo is therefore an important requirement for the planned High-Luminosity LHC (HL-LHC) upgrade.

The Beam Synchrotron Radiation Halo monitor (BSRH) is an optical diagnostic tool that is capable of fulfilling the functional specifications of the beam halo monitor. Imaging synchrotron radiation that originates from the bending of the trajectory of the circulating beams, it provides a non-invasive measurement method for both the beam core and halo. In order for the instrument to achieve the required precision (of charge in the halo region), the alignment of its optical system must be reliable. Currently used manual alignment is slow and prone to both short-term and long-term drifts, which motivated the development of an automated alignment procedure.

The main task of my internship was to design and implement an automatic alignment system for the BSRH instrument. It involved understanding the optical system, acquiring the necessary data, building data-driven models of the instrument, and developing control algorithms to keep the optical system correctly aligned during the operation in the LHC.

In addition to the alignment task, I also spent time completing several smaller related tasks. These included the creation of an auto-focus algorithm, which optimizes the sharpness of the beam picture through motor movements of the optical system. I also participated in a Machine Development (MD) session at the LHC, during which we were free to ask the operators to perform alterations of the beam size and tail content to perform a variety of dedicated measurements. Following the MD, I worked on the analysis of the collected data. For example, on charge-per-count (CPC) calibration scripts and bunch-by-bunch CPC studies.

Chapter 2

The BSRH Instrument

The Beam Synchrotron Radiation Halo monitor is a non-invasive optical diagnostic instrument used in the LHC. It is designed to measure the absolute charge of protons in the halo distribution the beam using synchrotron radiation that is emitted by the beam when its trajectory is deflected by the magnets in the LHC.

2.1 Purpose in the LHC

In high-intensity beams, a small but significant fraction of particles forms the so-called beam halo — the low-density tails of the transverse distribution. It is crucial to monitor these particles as they affect the losses around the accelerator that can lead to magnet quenches, damage of the collimators. The technique that is currently used to measure the halo content is called beam scraping, is invasive and cannot provide continuous monitoring. The BSRH aims to deliver a non-invasive and continuous measurement of the beam halo, thereby improving both machine protection and beam quality optimization in the HL-LHC era.

2.2 Instrument Design

The BSRH is based on a Cassegrain telescope, adapted to resolve faint halo structures in addition to the intense beam core. The instrument combines:

- **Optical Path:** Mirrors and lenses transport synchrotron light emitted from the beam pipe towards the detection system. The alignment of this optical chain is critical, as small deviations can distort halo measurements.
- **Detection System:** Cameras capture two-dimensional images of the beam. These must resolve the large intensity difference between the bright beam core and the much weaker halo (several orders of magnitude lower).

- **Control:** Motorized components (mirrors, lenses, filters, cameras) enable alignment of the optical path. Automatic alignment strategies, such as the one developed in this project, are important for operational stability.




Figure 2.1: BSRH schema




Figure 2.2: BSRH during assembly in the LHC

Chapter 3

Alignment

3.1 Motivation

The accuracy of the BSRH depends greatly on the alignment of its optical system. Synchrotron radiation from the circulating beam must pass through a succession of mirrors and lenses before it reaches the camera. Small displacements in the position or angle of these optical components will shift or distort the beam image, reducing the precision of the halo measurement.

Until now, BSRH alignment has been done manually. Users typically adjust motors while visually observing the camera image to verify whether the beam spot is centered. The procedure is time-consuming, depends on operators' skill, and is susceptible to drifts caused by mechanical instabilities or environmental conditions. During extended LHC operation, such drifts may accumulate and degrade data quality unless alignment is occasionally re-established.

The motivation behind this project was therefore to design and validate the architecture of an automatic alignment process that has the potential to align the optical system repeatedly without human intervention. Not only does an auto-alignment system improve efficiency in terms of reducing setup time, it also ensures reproducibility and stability of measurements across different fills and machine settings.

3.2 Alignment Algorithm Design

The alignment of BSRH is defined by the relation between the motors' positions and the beam position on the cameras. Aligning the optical system means finding motor positions such that the optical path of the incoming beam is well-defined (centered on all targets) and visible on the main camera at the end of the system.

This mapping is nonlinear and difficult to model analytically. The optical chain consists of multiple mirrors and lenses, each introducing coupled effects: moving one motor can shift the beam not only in one direction, but in a way that depends on the current configuration of all other elements. In addition, the system exhibits mechanical backlash and noise from the beam and cameras. Because of this, deriving explicit equations for alignment is not practical.

Neural networks provide a natural solution for this problem because:

- **Data-driven modeling:** Neural networks can learn the complex, non-linear relationships between motor movements and beam spot positions directly from measured data, without requiring an analytical model of the optics.
- **Generalization:** The trained network can interpolate within the range of beam and motor positions covered by the training data, allowing it to predict alignment corrections even for unseen configurations that are still within this range. However, extrapolation beyond the training domain remains a limitation, highlighting the importance of broad and representative data collection.
- **Multi-input, multi-output capability:** Neural networks handle multiple motors and coupled degrees of freedom simultaneously, making them suitable for describing an optical system such as BSRH.

Instead of attempting to align the whole optical system in one go, the problem can be broken down into individual straight optical lines, each described by two points or targets in our case. Every path has its own neural network assigned. This modularization offers several advantages:

- **Decomposition of complexity:** Training one big model for the entire optical system would require huge amounts of good-quality data and could be unstable. Breaking the problem down into small sub-networks makes both training and testing easier and more stable.
- **Better error propagation regulation:** Errors do not get out of control by correcting one optical path at a time. Each network ensures its part is corrected before moving onto the next one.
- **Reusability:** When an optical element changes (e.g. one mirror is replaced), then only the affected network needs to be retrained.
- **Better interpretability:** Operators can observe alignment step by step if needed.

In summary, a system of neural networks, each trained to align one optical line defined by two points was used. As seen in the BSRH schema, there are three such paths.

1. The first straight optical line consists of one mirror coupled to three motors: x-axis translation, x-tilt, and y-tilt. Two cameras capturing their

respective targets. And a mask that does not directly influence the beam position, so it is not used during the alignment procedure.




Figure 3.1: First optical path

In the figure above, the oval shape is the motorized mirror and the two dots are the targets/first and second camera of the optical line.

2. The second straight optical line consists of one mirror coupled to four motors: x-axis, y-axis, x-tilt, and y-tilt. Two cameras capturing their respective targets. Lastly, a special lens, a filter wheel, and one other filter are used that do not directly influence the position of the beam, so they are not used during the alignment procedure.




Figure 3.2: Second optical path

3. The third optical line consists of one mirror coupled to two motors: x-tilt and y-tilt. Two cameras capture their respective targets, one of which is the main camera at the end of the line. Lastly, two more lenses and

the second camera in case it was not after the alignment of the previous optical path. For the second and third optical path, it is also true that if the beam is visible on the second camera, then it is always visible on the first one as well.

The quality of the alignment model heavily depends on how well the training data covers the space of possible motor and beam positions. If the dataset is too limited, the neural networks may perform well only in regions the beam has been measured in before but fail when exposed to new positions or slightly different beam conditions.

The following are the beam position coverage distributions that were acquired for each corresponding optical path and its neural network.




Figure 3.4: First optical line - beam position before motor movements




Figure 3.5: First optical line - beam position after motor movements

Acquiring data on the first optical line was challenging as not many data points could be acquired during the operation of the LHC. This plot shows around 730 data points.

Acquiring data on the second optical path was easier, and around 15,000 were collected. The bounding boxes in the figure show the full visible camera target for each camera, respectively. Although the second camera has very good coverage, the first is only covered around its central part. This is not a problem as when the beam hits this uncovered area of the first camera, the beam is not visible at all on the second camera. This is not a situation the neural network of the second optical part has to solve. It is also important to know that even though the second optical path has four movable motors, the tilts and the axis movements are coupled, so only the tilts were enough to cover the camera spaces. Therefore, the x-axis and y-axis motors were not used at all during the data acquisition, nor during the model training.

For the last optical path, data from the first camera was not collected at all




Figure 3.6: Second optical line - beam position before motor movements




Figure 3.7: Second optical line - beam position after motor movements

as it required moving the first camera target in and out during acquisition, making the acquisition lengthy and thus avoided as the mirror could move only in the two tilts. For the neural network of the second optical line, this data was enough. Around 20,000 samples were collected.

3.4 Modeling

All models were trained on their respective datasets using the Adam optimizer and Huber loss as it is quite robust to outliers but still sensitive to small errors. The models were trained as feedforward neural networks with slightly different architectures and parameters, depending on the specifics of the corresponding optical path. Since the ranges of the different motors varied, a standard scaler was used for both the inputs and the outputs of the neural networks.

3.4.1 First Optical Path

The motor ranges for the first optical path were measured as:

- **motor x-axis:** 2,640 steps
- **motor x-tilt:** 10,000 steps
- **motor y-tilt:** 6,000 steps

The best model reached a validation error of 213.07 steps on average per motor. The respective errors were:

- **motor x-axis:** 155.9 steps, 5.9%
- **motor x-tilt:** 162.6 steps, 1.6%




Figure 3.8: Third optical line - beam position before motor movements




Figure 3.9: Third optical line - beam position after motor movements

- **motor y-tilt:** 320.7 steps, 5.3%

The FFNN architecture was 8-64-64-32-32-3. As stated above, this model would benefit from more data points and can be improved in the future.

3.4.2 Second Optical Path

The motor ranges for the second optical path were measured as:

- **motor x-tilt:** 1,050 steps
- **motor y-tilt:** 1,350 steps

The best model reached a validation error of 6.93 steps on average per motor. The respective errors were:

- **motor x-tilt:** 8.09 steps, 0.8%
- **motor y-tilt:** 5.78 steps, 0.4%

The FFNN architecture was 4-128-128-64-64-32-2. The models using only the tilts performed much better than if all motors were involved.

3.4.3 Third Optical Path

The motor ranges for the third optical path were measured as:

- **motor x-tilt:** 720 steps
- **motor y-tilt:** 600 steps

The best model reached a validation error of 1.99 steps on average per motor.
The respective errors were:

- **motor x-tilt:** 1.45 steps, 0.2%
- **motor y-tilt:** 2.53 steps, 0.4%

The FFNN architecture was 4-128-64-64-32-2.

Chapter 4

Conclusion

During this internship, I worked on the development of an automated alignment system for the Beam Synchrotron Radiation Halo monitor in the LHC. The project demonstrated that a modular system of neural networks, each responsible for aligning one optical path, can successfully model the nonlinear relationship between mirror motors and beam spot positions. The trained networks achieved good accuracy, particularly for the second and third optical paths, where errors were reduced below 1% of the motor ranges. This suggests a potential application of machine learning to replace manual alignment with a reproducible and efficient automated procedure.

In addition to the alignment project, I contributed to several smaller tasks, including the implementation of an auto-focus function, the integration of auto-focus and auto-alignment, and data analysis from a Machine Development session on the LHC. These activities provided important practical experience with beam diagnostics, data handling, and control system integration.

There are still challenges to address. The first optical path was limited by the small amount of data available, which affected model accuracy. Expanding the dataset in the future, improving data coverage, and refining the neural network architectures would help to further increase robustness.

Overall, the work carried out provides a foundation for deploying automated alignment in the BSRH, reducing operational overhead, and ensuring stable and reliable halo measurements in the High-Luminosity LHC era.

List of Figures

2.1	BSRH schema	6
2.2	BSRH during assembly in the LHC	7
3.1	First optical path	10
3.2	Second optical path	10
3.3	Third optical path	11
3.4	First optical line - beam position before motor movements	12
3.5	First optical line - beam position after motor movements	12
3.6	Second optical line - beam position before motor movements . . .	13
3.7	Second optical line - beam position after motor movements . . .	13
3.8	Third optical line - beam position before motor movements . . .	14
3.9	Third optical line - beam position after motor movements	14