# AUTOENCODER ARCHITECTURES FOR BEAM ANOMALY DETECTION AND STATISTICAL DISTRIBUTION CHARACTERIZATION

J. Szumega*[1,2], L. Bougueroua[3], B. Gkotse[1,4], P. Jouvelot[2], N. Minafra[1,5], F. Ravotti[1]
[1]European Organization for Nuclear Research, Geneva, Switzerland
[2]Mines Paris, PSL University, Paris, France
[3]Efrei Research Lab - Université Paris Panthéon Assas, Villejuif, France
[4]University of Wisconsin–Madison, Madison, United States
[5]University College of Dublin, Ireland

*Abstract*

In accelerator facilities, beam delivery control and assessment require capable monitoring systems, including both hardware and software components. In most accelerator beamlines, precise measurements and reliable beam delivery are critical factors in their operation. At the CERN IRRAD facility, the transverse beam profile contains essential information on the beam properties relevant to the irradiation of materials and components. Building upon the existing Beam Profile Monitor instrument at the CERN IRRAD Facility (IRRAD-BPM), we explore the possibilities of using Machine Learning techniques, with special focus on Autoencoder (AE) architectures. Dealing with a critical system that involves high-energy protons and extreme radiation conditions, we developed an AE-based anomaly-detection system. Its architecture, based on multiple parameters, is a result of hyperparameter optimisation aiming for the highest separation of anomalous samples. Additionally, to mitigate the existing limited IRRAD-BPM coverage that cannot capture the full extent of the beam tails, we perform a statistical inference using this AE architecture. Moreover, by using a Multi-Wire Proportional Chamber (MWPC) device also present on the IRRAD beamline, we improve the beam profile modeling within a data-fusion-like approach.

## INTRODUCTION

Experimental irradiation facilities must ensure the delivery of beams with the required characteristics and quality. As such, accurate assessment requires a reliable monitoring system capable of detecting unwanted deviations. During facility operation, deviation events can occur due to glitches during beam production, extraction or transport to the experimental facilities. Consequently, a highly precise anomaly detection needs to be an integral part of a monitoring system. This finally impacts all categories of experiments performed at IRRAD, from detector development and radiation-hardness assessment of electronic components, including both proton [1, 2] and heavy-ion [3] irradiation for space applications where the transverse beam profile properties need to be particularly precise.

The IRRAD facility at CERN uses four IRRAD-BPM sensors [4] to monitor the quality of the PS-extracted (Proton Synchrotron) 24 GeV/c proton beam. This is a grid of 39 metallic pads, 4x4 mm$^2$ each with 0.5 mm separation, resulting in a total effective coverage of ~6 cm$^2$ in the dense, central region or a total span of ~12 cm$^2$. In addition, nearby downstream the IRRAD-BPM on the T8 beamline of the experimental East Area, a Multi-Wire Proportional Chamber (MWPC) is also installed to monitor both proton and heavy-ion beams. It is equipped with 32 wires, with 6 mm separation, for each horizontal and vertical profile measurement. This device can cover a larger beam area of ~345 cm$^2$. The comparison between the acceptance of these two devices is presented in Fig. 1.

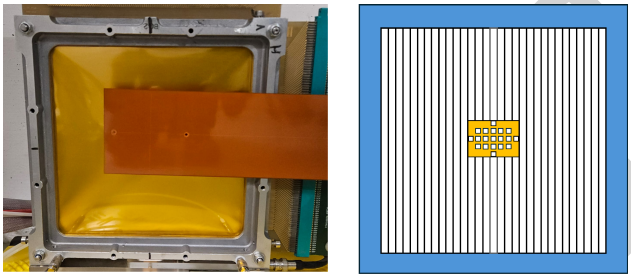


Figure 1: Comparison of the detector dimensions between the Multi-Wire Proportional Chamber (larger detector) and the Beam Profile Monitor. Figure a) shows both encased in protective coating so their internal structure remains hidden. Figure b) is a schema showing an in-scale comparison. While the coverage of the MWPC significantly exceeds that of the BPM, the latter provides better resolution in the beam center.

These two detectors differ in their construction and capabilities for measuring transverse profiles. One is an array of pixel-like pads, capable of 2D profile measurements, but on a narrow surface. The other offers more extensive coverage but reads two 1D profiles only - vertical and horizontal.

By studying autoencoder architectures and heterogeneous data from these detectors, their respective advantages are leveraged to support and improve the beam monitoring capabilities during irradiation experiments.

## RELATED WORK

Various neural network architectures are used to detect anomalies. At the Spallation Neutron Source Accelerator (SNS), a machine-learning-based infrastructure has been designed to predict errant beam and detect anomalies in

* jaroslaw.szumega@cern.ch

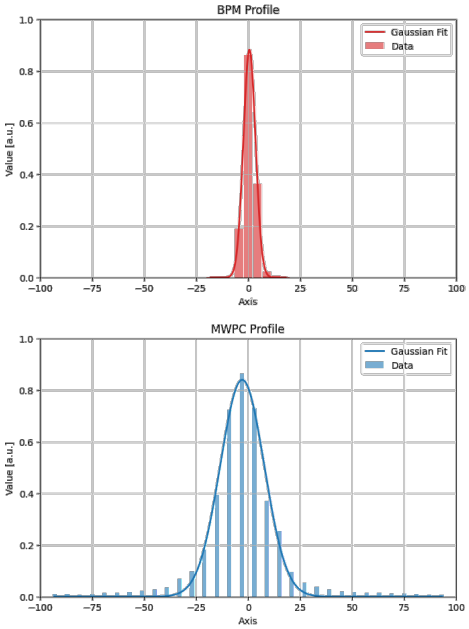


Figure 2: Comparison in captured beam profile data with corresponding Gaussian fits. The BPM measurement (red) is compared with the MWPC (blue). The heavier tails, exceeding the assumed Gaussian fit, become obvious when using the MWPC detector.

the target support cooling [5]. The training is done offline on historical data due to the rare occurrences of anomalies in online operation. In [6], the framework of a combined Variational Autoencoder and Bidirectional Long Short-Term Memory (BiLSTM) is used to detect anomalies in the multivariate time-series dataset from Jefferson Lab.

The common task of detecting out-of-distribution data points in particle physics is also supported with the use of autoencoders [7].

Machine-learning techniques are also used for beam monitoring itself. In [8], a fully connected neural network is used for non-destructive beam profile monitoring. Using the signal of eight pick-up electrodes, uniformly distributed around the beam, the network reconstructs the beam profile. Another study succeeds in six-dimensional phase-space reconstruction [9] and also employs a fully connected variant of an artificial neural network. With heterogeneous measurements, the high-dimensional beam features are reconstructed, thus allowing the resolution of complex beam phenomena.

## EXPERIMENTS

This paper presents a case study of the application of autoencoder architectures to 1) anomaly detection in 2D transverse beam profiles of IRRAD-BPM, and 2) manifold learning for beam distribution verification using data from both BPM and MWPC profiles (see the profile comparison in Fig. 2).

To fulfill the goal of Task 1, we build upon previous work [10] and use the beam profiles data recorded at IRRAD in the year 2024. As for Task 2, the general assumption for the proton beam profile at the LHC is that the profile follows a Gaussian distribution [11]. However, recent studies suggest that this is a highly simplified view, especially given that the actual tails of the bunch distributions are non-Gaussian [12]. Moreover, these effects are not exclusive to the LHC; they are also observed in the PS. The analysis of the profiles measured with the wire scanners suggests that their heavier tails correspond to a q-Gaussian distribution [13].

### *Methodology*

For both tasks, despite choosing different autoencoder architectures, we follow a similar experimental protocol. First, we define the data structure and the hyperparameters for the selected AE types, along with their corresponding values. We perform Hyperparameter Optimisation (HPO) to find a configuration that minimises or maximises the given objective function. The train-val-test splits are 0.7:0.15:0.15 for both datasets.

For anomaly detection, we use the 2D profile measurement obtained from the BPM. It allows us to treat each profile similarly to an image. Using the autoencoder architecture and a loss measure based on the Structural Similarity Index Measure (SSIM), we train the AE on only well-centred profiles. It allows the network to learn a hypothesis space biased towards seen data (inductive bias). Therefore, the reconstruction of centred profiles yields a higher SSIM value than the reconstruction of unseen off-centred anomalies. To assess its capabilities at each HPO iteration, we validate it using the separation metric based on Area Under the Curve (AUC). The goal is to maximise the separation, the value of AUC. Finally, we provide the Confusion Matrix and F1-score to assess the result of anomaly detection.

In the case of manifold learning, we leverage the presence of heterogeneous detectors. As the MWPC can measure only two 1D profiles, we match this capability by extracting the two main axes from BPM measurements. To compensate for the difference in their resolution, we pad the shorter vectors. The autoencoder is trained to reconstruct the input from the latent representation. By minimising the reconstruction Mean Squared Error (MSE), we aim to learn the manifold of selected measurements, i.e., the autoencoder's latent space that faithfully represents the beam profiles. We use this model to verify the beam's distribution: first, we calculate the fitting parameters of the x-y profiles for the Gaussian, q-Gaussian and Student's t distributions. Then, using these parameters, we generate samples from the corresponding distributions and attempt to reconstruct them with the AE. Since the network was trained using real measurements, the lowest reconstruction error on the test data will indicate the closest match to the true distribution.

### *Results and Discussion*

**Anomalous Beam Detection** The 2D BPM dataset contains almost 1,400,000 samples, with a valid-to-anomalous beam profile ratio of 0.85:0.15. The task of anomaly detection is well justified, given the strong class imbalance and the unspecified nature of future, unobserved anomalies.

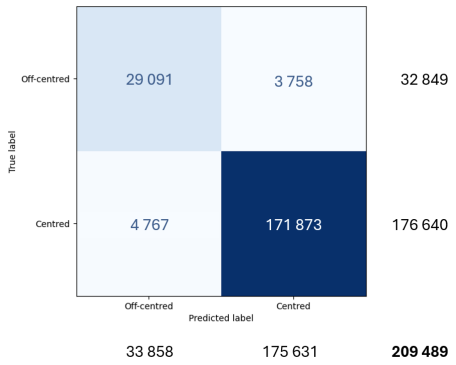


Figure 3: Classification results of the Fully Connected Autoencoder. The detection threshold is set to $T_{SSIM} = 0.89$. Profiles below the threshold are classified as anomalies. The F1-score is 0.872.

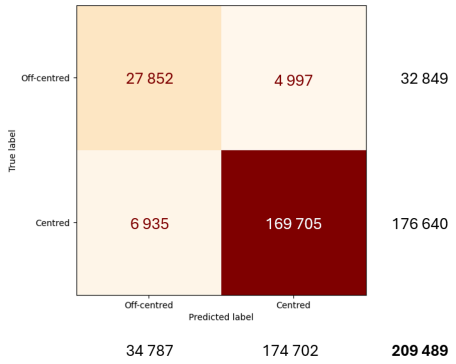


Figure 4: Beam profile classification results of the Variational Autoencoder. The threshold was determined and set to $T_{SSIM} = 0.93$. The F1-score for VAE is 0.824.

We trained a Fully Connected (Dense) Autoencoder and a Variational Autoencoder and obtained good separation between normal (well-centred) and anomaly (off-centred) classes. The two AE architectures can be briefly described as `[ENC layers] -> [Latent] -> [DEC layers]`:

- Dense AE:
  [128, 4, 64] -> 4 -> [64, 4, 128];
- VAE:
  [256, 128, 64, 32, 16] -> 2 -> [16, 32, 64, 128, 256].

The Dense AE is smaller and achieved better separation results for anomaly detection. Yet, VAE succeeded in finding a more compact latent space. The confusion matrices for two architectures are shown in Figs. 3 and 4.

**Manifold Learning** To train the AE for beam distribution verification, we chose an LSTM-based autoencoder. With proper padding and alignment of the BPM and MWPC data, we treat each set of x-y profiles as a 2-vector sequence, 32 elements each. The HPO parameter search was set to include the bidirectional LSTM as well. The result was a BiLSTM Autoencoder with a single layer of 256 hidden units and a latent space represented as a 16-element vector.

We used test data to obtain corresponding samples of three distributions. The distribution that matches the manifold learned by the Autoencoder should yield the lowest error. The results are presented in the Table 1.

Table 1: Comparison of AE-Based Reconstruction Errors for Different Distributions

| Distribution | MSE | RMSE |
|---|---|---|
| Gaussian | $1.38 \times 10^{-4}$ | $1.17 \times 10^{-2}$ |
| q-Gaussian | $2.83 \times 10^{-5}$ | $5.32 \times 10^{-3}$ |
| Student's t | $2.83 \times 10^{-5}$ | $5.32 \times 10^{-3}$ |

As we verified the fit parameters, e.g., for q-Gaussian, the $q$ parameter was such that $q > 1$ (around 1.575). This strongly implies heavier tails, which in the case of BPM-only measurements were not possible to observe. On the other hand, we observed a Student's t's $\nu$ parameter = 2.48. Considering the relation between the Student's t and q-Gaussian distributions

$$q = \frac{\nu + 3}{\nu + 1}, \tag{1}$$

then:

$$\nu \approx 2.48 \;\Rightarrow\; q = \frac{2.48 + 3}{2.48 + 1} \approx 1.575. \tag{2}$$

Thus, the inferred Student's t parameters correspond to the observed q-Gaussian fitting. This result confirms the q-Gaussian nature of the data distribution, which is also reflected in the calculated errors.

## CONCLUSION AND FUTURE WORK

In this paper, we studied the applications of autoencoder architectures in beam anomaly detection and beam distribution characterisation. We have shown that the Fully Connected Autoencoder and Variational Autoencoder successfully separate the normal beam profile measurements from the anomalous samples. Using a BiLSTM network, we analysed samples from heterogeneous sources with different data dimensions. The LSTM-based network succeeded in learning the beam-profile manifold. The distribution study strongly suggests that the PS-extracted beam, subsequently propagated to the IRRAD facility on the T8 beamline at CERN, is better described by a q-Gaussian distribution.

With these promising results, our future work will focus on research towards even more detailed beam analysis and possible application of attention-oriented mechanisms. The novel beam instrumentation and support of ML-based algorithms also extend our research interest towards creating a digital twin of the IRRAD facility at CERN.

## REFERENCES

[1] P. Pelissou *et al.*, "IRRAD/CHARM, a CERN Irradiation Facility for Accelerator and Experiments Radiation Hardness Qualification", in *2024 IEEE Radiation Effects Data Workshop (REDW) (in conjunction with 2024 NSREC)*, Jul. 2024. doi:10.1109/REDW61286.2024.10759157

[2] F. Faccio, "ASIC survival in the radiation environment of the LHC experiments: 30 years of struggle and still tantalizing", *Nucl. Instrum. Meth. A*, vol. 1045, p. 167569, 2023. doi:10.1016/j.nima.2022.167569

[3] D. Söderström *et al.*, "Readiness of the HEARTS@CERN facility for space electronics high-energy heavy-ion testing", *JACoW*, vol. IPAC2025, TUPB006, 2025. doi:10.18429/JACoW-IPAC2025-TUPB006

[4] I. Mateu *et al.*, "A Beam Profile Monitor for High Energy Proton Beams Using Microfabrication Techniques", in *9th International Beam Instrumentation Conference*, TUPP37, 2020. doi:10.18429/JACoW-IBIC2020-TUPP37

[5] A. Kasparian *et al.*, "Machine learning at the Spallation Neutron Source accelerator and target", *JACoW*, vol. NA-PAC2025, MOP057, 2026. doi:10.18429/JACoW-NAPAC2025-MOP057

[6] A. Verma, "Variational autoencoder and bilstm-based anomaly detection for beam stability in spallation neutron sources", *Authorea Preprints*, 2025. doi:10.36227/techrxiv.175615627.78320502/v1

[7] V. Belis, P. Odagiu, and T. K. Aarrestad, "Machine learning for anomaly detection in particle physics", *Rev. Phys.*, vol. 12, p. 100091, 2024. doi:10.1016/j.revip.2024.100091

[8] Z. Omarov and S. Haciomeroglu, "Machine learning assisted non-destructive beam profile monitoring", *Nucl. Instrum. Meth. A*, vol. 1026, p. 166132, 2022. doi:10.1016/j.nima.2021.166132

[9] R. Roussel *et al.*, "Efficient six-dimensional phase space reconstructions from experimental measurements using generative machine learning", *Phys. Rev. Accel. Beams*, vol. 27, no. 9, p. 094601, 2024. doi:10.1103/PhysRevAccelBeams.27.094601

[10] J. Szumega, L. Bougueroua, B. Gkotse, P. Jouvelot, and F. Ravotti, "Machine Learning for the Anomaly Detection and Characterization of the 24 GeV/c Proton Beam at CERN IRRAD Facility", *JACoW*, vol. IPAC2025, THPM110, 2025. doi:10.18429/JACoW-IPAC2025-THPM110

[11] G. Trad *et al.*, "Status of the beam profile measurements at the LHC", in *6th Evian Workshop on LHC beam operation*, pp. 163–170, 2016.

[12] A. Babaev *et al.*, "Impact of beam–beam effects on absolute luminosity calibrations at the CERN Large Hadron Collider", *Eur. Phys. J. C*, vol. 84, no. 1, p. 17, 2024. doi:10.1140/epjc/s10052-023-12192-5

[13] M. Bozatzis *et al.*, "Tail population studies in the CERN Proton Synchrotron", *JACoW*, vol. IPAC2024, THPC15, 2024. doi:10.18429/JACoW-IPAC2024-THPC15