

Calibrating for the Future: Enhancing Calorimeter Longevity with Deep Learning

S. Ali^{1*}, A. S. Ryzhikov^{1**}, D. A. Derkach^{1***}, F. D. Ratnikov^{1****}, and V. O. Bocharnikov^{1*****}

¹*HSE University, Moscow, Russia*

Received October 1, 2024; revised October 10, 2024; accepted October 20, 2024

Abstract—In the realm of high-energy physics, the longevity of calorimeters is paramount. Our research introduces a deep learning strategy to refine the calibration process of calorimeters used in particle physics experiments. We develop a Wasserstein GAN inspired methodology that adeptly calibrates the misalignment in calorimeter data due to aging or other factors. Leveraging the Wasserstein distance for loss calculation, this innovative approach requires a significantly lower number of events and resources to achieve high precision, minimizing absolute errors effectively. Our work extends the operational lifespan of calorimeters, thereby ensuring the accuracy and reliability of data in the long term, and is particularly beneficial for experiments where data integrity is crucial for scientific discovery.

Keywords: machine learning, deep learning, generative adversarial neural networks, calibration, high energy physics

DOI: 10.3103/S0027134924702047

1. INTRODUCTION

1.1. Calorimeter Architecture and Energy Measurement

A calorimeter [1–3] is a device used in particle physics and beyond to measure the energy of incident particles. By absorbing and measuring the total energy of an incoming particle, it provides crucial data about the particle's properties [4, 5]. Calorimeters are essential in high-energy physics experiments, nuclear reactors, and various applications in medical physics, such as radiation therapy [4, 5]. They are commonly found in particle detectors like those in the large hadron collider (LHC) and other major physics experiments worldwide [4].

Incorporating machine learning into high-energy physics (HEP) has increased for advancing our understanding of fundamental physics [6]. Machine learning (ML) techniques are transforming the way how data is analyzed and interpreted in high-energy physics experiments, especially in the wake of the

large hadron collider (LHC) and its upgraded version, the high luminosity LHC (HL-LHC). These experiments generate massive amounts of data, where traditional data analysis methods are insufficient for handling such large volumes, ML algorithms help in reducing data complexity and finding new features within large datasets, enhancing the overall physics performance of reconstruction and analysis algorithms [7]. ML algorithms, such as boosted decision trees (BDTs) [8] and neural networks (NNs), have become state-of-the-art for event and particle identification in HEP. These algorithms classify particles and events with greater efficiency and accuracy than traditional methods [9].

Accurate simulations are essential for comparing observed data with theoretical models. ML methods, particularly deep learning models like generative adversarial networks (GANs) [10] and variational autoencoders (VAEs), are used to generate realistic simulations of particle interactions and detector responses. These models can also accelerate computationally intensive tasks, such as event simulation and pattern recognition, which are critical for real-time data processing during experiments.

Many HEP experiments produce data at extreme rates that cannot be stored and processed offline. To manage these data volumes, real-time online data

*E-mail: thraaali@hse.ru

**E-mail: alyzhikov@hse.ru

***E-mail: dderkach@hse.ru

****E-mail: fratnikov@hse.ru

*****E-mail: vbocharnikov@hse.ru

compression, zero suppression, and filtering are necessary. This real-time data processing begins at the sensor readout and extends to data storage for offline analysis. Deploying ML algorithms in custom electronics systems and edge processing on local data center servers significantly enhances the efficiency of this process, ensuring timely and accurate data handling [11].

ML algorithms are also essential for distinguishing these overlapping signals and isolating the rare events of interest to physicists. Techniques like pile-up suppression use ML to improve signal clarity and reduce background noise, enabling more precise measurements [12].

1.2. Calorimeter Calibration

Calibration ensures that the output from each cell in the calorimeter is consistent when subjected to the same input. In an ideal scenario, identical particles hitting different parts of the calorimeter should yield the same response, ensuring accurate and reliable measurements [4]. However, due to manufacturing variations, construction misalignment and other factors each cell in a calorimeter might react slightly differently to the same input. To achieve uniformity across all cells, each one must be calibrated post-production. Calibration aligns the responses so that the measurement system can reliably report on particle interactions, which is essential for precise scientific outcomes [4, 5].

Methods of Calibration:

1. Hardware Calibration:

- **Internal Pulse Injection:** This involves injecting known electronic signals into the calorimeter to check its response [4].
- **Radioactive Source/Laser:** Using radioactive sources or lasers to provide a uniform and known energy input to calibrate the detector [4].

2. Test Beam Calibration:

- **Single Particles of a Known Type and Energy:** Single particles of a known type and energy are directed at the calorimeter to measure and adjust its response [4].

3. In Situ Calibration:

- **Well-Known Physics Samples:** Using particles from well-known physics processes, such as Z boson decaying to electron-positron pairs ($Z \rightarrow ee$) or W boson decaying to jet pairs ($W \rightarrow jj$), to calibrate the calorimeter in the actual experimental environment [4, 5].

However, over time, the aging of components and other effects can alter the characteristics of each cell within the calorimeter, necessitating periodic recalibration to maintain accuracy and reliability. This process often requires repeating multiple calibration methods to account for changes in detector response due to radiation damage [4].

Given the complexity and effort involved in optimizing a vast number of parameters during recalibration, we propose a software-based approach utilizing machine learning (ML). This method aims to automate the optimization process, reducing the need to repeat all calibration methods manually and minimizing potential issues, thereby enhancing efficiency and reliability in maintaining calibrated states of the calorimeter [4].

2. DATA DESCRIPTION-EXPERIMENTAL DATA

In our study, we conduct Monte Carlo simulations using the Geant4 toolkit [13] to simulate hadronic shower samples. Specifically, we generated 10 000 events with single 10 GeV pions. The simulations are carried out using a highly granular calorimeter geometry [14], which is typical for future collider experiments. The calorimeter employed a sampling design with simplified SiPM-on-tile technology, consisting of 20 mm copper plates as the absorber and polystyrene scintillator cells arranged in a 24×24 grid with a 5 mm thickness as the active layers. By analyzing the energy deposition in the active cells, we aim to evaluate the efficacy of our methods by accurately capturing the changes in behavior of hadronic showers. Additionally, we present synthetic aging in the simulations to mimic radiation damage of the detector. This involved assigning an aging coefficient A_i , to each readout cell i , which has a linear dependence on the integrated energy deposition in each cell over the entire simulated data set as in Fig. 1.

The aging coefficient A_i measures how much the response of a calorimeter cell changes over time due to factors such as radiation exposure. It reflects the difference between the cell's current state and its original, undamaged state. Specifically, for cell i , let $E_{\text{undamaged}_i}$ represent the correct energy read by the cell in its initial, undamaged state, and E_{damaged_i} denote the energy read by the same cell, from the same source of energy, after it has been damaged. The aging coefficient describes how much the current response deviates from the initial response and can be expressed as:

$$A_i = \frac{E_{\text{undamaged}_i}}{E_{\text{damaged}_i}}.$$




Fig. 1. Distribution of aging coefficients.




Fig. 2. Energy sum distribution in damaged vs. undamaged calorimeters.

If $A_i = 1$, it indicates no aging effects, meaning the cell's response has not changed. Conversely, if $A_i < 1$, it signifies a reduction in the cell's response due to damage, which necessitates calibration to ensure accurate energy measurements. By determining these coefficients for each cell, our model can make the necessary adjustments to correct the degraded responses, thereby ensuring that the calorimeter continues to provide precise energy measurements.

The dataset consists of two distinct sets: one representing energy readings from the undamaged calorimeter and the other from the damaged calorimeter. Unlike conventional machine learning applications, there is no need to split the data into separate training and test sets because our method is unsupervised. In unsupervised learning, the model is trained and tested on the same data. Our goal is to use both sets entirely to learn the aging coefficients that, when applied to one set, will match the other. Since the model does not have access to the true aging coefficients at any point and is applied in unsupervised setting, making it essential to use the full dataset to derive the best possible calibration.

Histograms in Fig. 2 of energy sum distributions for damaged and undamaged calorimeter reveal shifts in energy measurements, highlighting synthetic aging effects on calorimeter performance.

This work examines radiation damage that leads to lower signal amplitude in cells.

3. GENERATIVE ADVERSARIAL NETWORKS (GANs) IN CALORIMETER CALIBRATION

Generative adversarial networks (GANs) is a breakthrough in machine learning, particularly in the field of generative models. Introduced by Ian Goodfellow and his colleagues in 2014 [15], GANs consist of two neural networks: the generator (G) and the discriminator (D) [15]. These networks engage in a minimax game, where the generator aims to produce synthetic data that is indistinguishable from real data, and the discriminator strives to accurately distinguish between real and synthetic data.

3.1. Architecture and Mathematical Formulation

The GAN framework is built upon the adversarial relationship between G and D . The generator takes a random noise vector z from a latent space and transforms it into a synthetic data sample $G(z)$. The Discriminator, on the other hand, receives either a real data sample x or a synthetic sample $G(z)$ and outputs a probability $D(x)$ or $D(G(z))$ that the input is a real sample.

The objective of the GAN training process is described by the following value function $V(D, G)$ [15]:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))]. \quad (1)$$

Here, $p_{\text{data}}(x)$ represents the distribution of real data, and $p_z(z)$ is the distribution of the latent space from which noise vectors z are sampled. The discriminator tries to maximize the probability of correctly identifying real and fake samples, while the generator attempts to minimize the probability that the discriminator accurately classifies its outputs as fake.

3.2. Wasserstein GAN (WGAN)

While the original GAN framework has achieved remarkable success, it often suffers from training instability and issues such as mode collapse. To address these challenges, the Wasserstein GAN (WGAN) is introduced by Arjovsky et al. in 2017 [16]. WGAN modifies the GAN objective by employing the Earth mover's distance (also known as Wasserstein distance) as a measure of the divergence between the generated and real data distributions. This change




Fig. 3. Architecture of the WGAN-inspired network used for calibrating the calorimeter.




Fig. 4. Scatter plot of true aging coefficients versus predicted aging coefficients. The closer to line, the more consistent with true values the predictions are.

leads to a more stable training process and provides a meaningful loss metric that correlates well with the quality of generated samples.

In WGAN, the discriminator is replaced with a Critic that provides feedback on how to adjust the

synthetic data to better match the real data distribution, rather than simply classifying the data as real or fake. The WGAN objective function is expressed as:

$$\min_G \max_{C \in \mathcal{C}} V(C, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [C(x)] - \mathbb{E}_{z \sim p_z(z)} [C(G(z))], \quad (2)$$

where $\mathcal{C}$ denotes the set of 1-Lipschitz functions. The Critic C is constrained to be 1-Lipschitz, typically enforced by clipping the weights of the Critic within a small range $[-c, c]$, gradient penalize [17], or spectrally normalized [18]. This constraint ensures that the Critic provides a reliable estimate of the Wasserstein distance between the real and generated data distributions, which helps in guiding the generator towards producing high-quality samples.

The adoption of the Wasserstein distance in WGAN provides significant improvements in training dynamics, reducing the sensitivity to the architecture of the networks and hyperparameters, stabilizing the gradients, and often making the performance better compared to the original GAN framework. It also can be considered as a good distribution matching framework, where the WGAN Critic is responsible for the distributions (Earth mover's) distance assessment.

3.3. WGAN-Inspired Approach to Calorimeter Calibration

WGAN works by minimizing the Wasserstein distance (Earth mover's distance) between the generated data distribution and the real data distribution. This is achieved through having the loss for the critic as the Wasserstein distance. Inspired by this approach, we employed WGAN architecture to calibrate a calorimeter response, aiming to align the distribution of data from a undamaged calorimeter with that from damaged calorimeter. In our application, the generator's weights represent the predicted aging coefficients, and the synthetic data processed by the generator corresponds to the corrected data from the undamaged calorimeter. The critic is trained to estimate the Wasserstein distance, while the generator iteratively updates its weights to reduce this distance, learning the appropriate aging coefficients. This process results in a corrected data distribution that closely matches the damaged calorimeter data. Unlike traditional WGANs focused on generating new data, our WGAN-inspired approach emphasizes aligning existing data distributions. To ensure the model has the best potential to learn the aging effects, we conducted the training specifically on the most affected cells located in the center of the calorimeter in our dataset, while not focusing on the cells at the borders of the calorimeter. The architecture of our WGAN network, including the generator and critic components, is illustrated in Fig. 3.




Fig. 5. Mean absolute error (MAE) of aging coefficients. The MAE decreases consistently, stabilizing around 0.0074 after 100 epochs, with R^2 value of 0.88, demonstrating the model's potential to tune aging coefficients.

4. RESULTS AND DISCUSSION

The results of the predicted aging coefficients by our WGAN-inspired approach are illustrated in Fig. 4. The scatter plot of the true values versus the predicted values shows that the points are closely aligned along the diagonal line. This indicates that the model has a strong capability to predict the aging factors accurately. Moreover, Fig. 5 demonstrates the mean absolute error (MAE) of the aging factors throughout the training process. The MAE decreases consistently, starting from 0.018 and stabilizing around 0.0074 after 100 epochs. This consistent reduction in error showcases the model's effectiveness in tuning the aging coefficients.

Additionally, the alignment of the scatter plot points along the diagonal line suggests that the model captures the underlying distribution of the data. The continuous decrease in MAE further reinforces the reliability of the model's predictions over time. These results demonstrate the potential of the inspired-WGAN approach to effectively calibrate the calorimeter by adjusting the aging coefficients to closely match the distribution of data from the undamaged calorimeter. The improvement in MAE over epochs highlights the model's learning capability and the ability to reduce prediction errors, thus providing potential in its application for real-world scenarios.

Figure 6 illustrates the performance of our model using the predicted aging coefficients to calibrate the same dataset used for training. Initially, there is a visible shift between the two distributions, indicating significant differences due to the damage. After applying our predicted aging coefficients and calibrating the damaged calorimeter, the distribution of the calibrated data aligns much more closely with that of the undamaged calorimeter, which also demonstrates the potential of our calibration method. Calibration is performed at the cell level, and the damaged




Fig. 6. Energy sum distribution for the undamaged and damaged calorimeters. As Fig. 2 shows a visible shift between the two distributions before calibration while here we see the distributions after calibrating the damaged calorimeter with our predicted aging coefficients, indicating a better alignment. Calibration is done at the cell level.

calorimeter shows fewer changes in its distribution postcalibration, indicating a successful adjustment of the aging coefficients.

5. CONCLUSIONS

Generative machine learning models have demonstrated their ability to address degradation effects in particle detectors, providing an innovative approach to calibration challenges. Our WGAN-inspired model has shown significant potential in identifying calibration coefficients. In our study, the model achieved a mean absolute error (MAE) of 0.0074, using a dataset of 5000 events for each damaged and undamaged states.

The efficiency of our WGAN model translates into substantial benefits, potentially reducing the cost and resource intensity of calibration experiments. By enabling calibration with a reduced amount of experimental data, this approach offers a more practical and economical solution for maintaining the accuracy of particle detectors over time.

Moreover, the presented method is universal and adaptable, capable of being applied to detectors of various shapes and based on different physics principles. This versatility underscores the broad applicability and impact of our approach, paving the way for its use in a wide range of detector calibration scenarios. Our work highlights the promise of generative models in enhancing the longevity and performance of particle detectors, contributing to the advancement of experimental physics.

FUNDING

This work is supported by the HSE Basic Research Fund. The computation for this research was performed using the computational resources of HPC facilities at the HSE University.

CONFLICT OF INTEREST

The authors of this work declare that they have no conflicts of interest.

REFERENCES

1. A. A. Alves, Jr., L. M. Andrade Filho, A. F. Barbosa, I. Bediaga, G. Cernicchiaro, G. Guerrier, H. P. Lima, Jr., A. A. Machado, J. Magnin, F. Marujo, J. M. de Miranda, A. Reis, A. Santos, A. Toledo, K. Akiba, S. Amato, B. de Paula, L. de Paula, T. da Silva, M. Gandelman, J. H. Lopes, B. Maréchal, D. Moraes, E. Polcarpo, F. Rodrigues, J. Bal-lansat, Y. Bastian, D. Boget, I. De Bonis, V. Coco, P. Y. David, D. Decamp, P. Delebecque, C. Dran-court, N. Dumont-Dayot, C. Girard, B. Lieunard, M. N. Minard, B. Pietrzyk, T. Rambure, G. Rospabe, S. T'Jampens, Z. Ajaltouni, G. Bohnert, R. Bonnefoy, D. Borras, C. Carloganu, H. Chanal, E. Conte, R. Cornat, M. Crouau, E. Delage, O. Deschamps, P. Henrard, P. Jacquet, C. Lacan, J. Laubser, J. Lecoq, R. Lefèvre, M. Magne, M. Martemyanov, M.-L. Mercier, S. Monteil, V. Niess, P. Perret, et al. (The LHCb Collaboration), *J. Instrum.* **3**, S08005 (2008).
<https://doi.org/10.1088/1748-0221/3/08/S08005>
2. K. Aamodt, A. A. Quintana, R. Achenbach, S. Acounis, D. Adamová, C. Adler, M. Aggarwal, F. Agnese, G. A. Rinella, Z. Ahammed, A. Ahmad, N. Ahmad, S. Ahmad, A. Akindinov, P. Akishin, D. Aleksan-drov, B. Alessandro, R. Alfaro, G. Alfaro, A. Alici, J. Alme, T. Alt, S. Altinpinar, W. Amend, C. Andrei, Y. Andres, A. Andronic, G. Anelli, M. Anfreville, V. Angelov, A. Anzo, C. Anson, T. Antić, V. Anto-nenko, D. Antonczyk, F. Antinori, S. Antinori, P. An-tonioli, L. Aphecetche, H. Appelshäuser, V. Aprodu, M. Arba, S. Arcelli, A. Argentieri, N. Armesto, R. Ar-naldi, A. Arefiev, I. Arsene, A. Asryan, A. Augustinus, T. C. Awes, J. Åysto, M. D. Azmi, S. Bablock, A. Badalà, S. K. Badyal, et al. (The ALICE Collabo-ration), *J. Instrum.* **3**, S08002 (2008).
<https://doi.org/10.1088/1748-0221/3/08/S08002>
3. G. Aad, E. Abat, J. Abdallah, A. A. Abdelalim, A. Abdesselam, O. Abdinov, B. A. Abi, M. Abolins, H. Abramowicz, E. Acerbi, B. S. Acharya, R. Achen-bach, M. Ackers, D. L. Adams, F. Adamyan, T. N. Addy, M. Aderholz, C. Adorisio, P. Adragna, M. Aharrouche, S. P. Ahlen, F. Ahles, A. Ahmad, H. Ahmed, G. Aielli, P. F. Åkesson, T. P. A. Åkesson, A. V. Akimov, S. M. Alam, J. Albert, S. Albrand, M. Aleksa, I. N. Aleksandrov, M. Aleppo, F. Alessan-dria, C. Alexa, G. Alexander, T. Alexopoulos, G. Al-imonti, M. Aliyev, P. P. Allport, S. E. Allwood-Spiers, A. Aloisio, J. Alonso, R. Alves, et al. (The ATLAS Collaboration), *J. Instrum.* **3**, S08003 (2008).
<https://doi.org/10.1088/1748-0221/3/08/S08003>
4. Ch. W. Fabjan and F. Gianotti, *Rev. Mod. Phys.* **75**, 1243 (2003).
<https://doi.org/10.1103/revmodphys.75.1243>
5. S. S. Korpachev and M. V. Chadeeva, *Bull. Lebedev Phys. Inst.* **45**, 322 (2018).
<https://doi.org/10.3103/S1068335618100093>
6. P. Calafiura, D. Rousseau, and K. Terao, *Artifi-cial Intelligence for High Energy Physics*, Ed. by P. Calafiura, D. Rousseau, and K. Terao (World Sci-entific, 2022).
<https://doi.org/10.1142/12200>
7. K. Albertsson, P. Altøe, D. Anderson, J. Ander-son, M. Andrews, J. P. A. Espinosa, A. Aurisano, L. Basara, A. Bevan, W. Bhimji, D. Bonacorsi, B. Burkle, P. Calafiura, M. Campanelli, L. Capps, F. Carminati, S. Carrazza, Y.-F. Chen, T. Childers, Ya. Coadou, E. Coniavitis, K. Cranmer, C. David, D. Davis, A. De Simone, J. Duarte, et al., *arXiv Preprint* (2018).
<https://doi.org/10.48550/arXiv.1807.02876>
8. Y. Coadou, in *Artificial Intelligence for High En-ergy Physics*, Ed. by P. Calafiura, D. Rousseau, and K. Terao (World Scientific, 2022), p. 9.
https://doi.org/10.1142/9789811234033_0002
9. HEP ML living review (2022). <https://iml-wg.github.io/HEPML-LivingReview/>.
10. A. Butter and T. Plehn, in *Artificial Intelligence for High Energy Physics*, Ed. by P. Calafiura, D. Rousseau, and K. Terao (World Scientific, 2022), p. 191.
https://doi.org/10.1142/9789811234033_0007
11. P. Harris and N. Tran, in *Artificial Intelligence for High Energy Physics*, Ed. by P. Calafiura, D. Rousseau, and K. Terao (World Scientific, 2022), p. 265.
https://doi.org/10.1142/9789811234033_0009
12. J. Duarte and J.-R. Vlimant, in *Artificial Intelli-gence for High Energy Physics*, Ed. by P. Calafiura, D. Rousseau, and K. Terao (World Scientific, 2022), p. 387.
https://doi.org/10.1142/9789811234033_0012
13. S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytrcek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Di-etrich, R. Enami, A. Feliciello, C. Ferguson, H. Fe-sefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J. J. Gómez Cadenas, I. González, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, et al., *Nucl. Instrum. Meth-ods Phys. Res., Sect. A* **506**, 250 (2003).
[https://doi.org/10.1016/s0168-9002\(03\)01368-8](https://doi.org/10.1016/s0168-9002(03)01368-8)

14. M. Arratia, K. Barish, L. Blanchard, H. Z. Huang, Zh. Ji, B. Karki, O. Long, R. Milton, A. Paul, S. J. Paul, S. Preins, B. Schmookler, O. Tsai, and Zh. Xu, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1047**, 167866 (2023).
<https://doi.org/10.1016/j.nima.2022.167866>
15. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, Sh. Ozair, A. Courville, and Yo. Bengio, in *Advances in Neural Information Processing Systems*, Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Curran Associates, 2014), Vol. 27, p. 2672.
https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf.
16. M. Arjovsky, S. Chintala, and L. Bottou, *arXiv Preprint* (2017).
<https://doi.org/10.48550/arXiv.1701.07875>
17. I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, in *Advances in Neural Information Processing Systems*, Ed. by I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, 2017), Vol. 30, p. 5767.
<https://doi.org/10.48550/arXiv.1704.00028>
18. T. Miyato, T. Kataoka, M. Koyama, and Yu. Yoshida, *arXiv Preprint* (2018).
<https://doi.org/10.48550/arXiv.1802.05957>

Publisher's Note. Allerton Press, Inc. remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

AI tools may have been used in the translation or editing of this article.