THE EUROPEAN
PHYSICAL JOURNAL C

Check for updates

Letter

Comparison of supervised and unsupervised anomaly detection in Belle II pixel detector data

Katharina Dort[1,a] , **Johannes Bilk**[1], **Stepahnie Käs**[1], **Jens Sören Lange**[1], **Marvin Peter**[1], **Timo Schellhaas**[1], **Benjamin Schwenker**[2], **Björn Spruck**[3]

[1] II. Physikalisches Institut, Justus-Liebig-University Gießen, 35392 Gießen, Germany
[2] II. Physikalisches Institut, Georg-August-Universität Göttingen, 37077 Göttingen, Germany
[3] Institut für Kernphysik, Johannes Gutenberg-Universität Mainz, 55128 Mainz, Germany

Abstract Machine learning has become a popular instrument for the search of undiscovered particles and mechanisms at particle collider experiments. It enables the investigation of large datasets and is therefore suitable to operate directly on minimally-processed data coming from the detector instead of reconstructed objects. Here, we study patterns of raw pixel hits recorded by the Belle II pixel detector, that is operational since 2019 and presently features 4 M pixels and trigger rates up to 5 kHz. In particular, we focus on unsupervised techniques that operate without the need for a theoretical model. These model-agnostic approaches allow for an unbiased exploration of data while filtering out anomalous detector signatures that could hint at new physics scenarios. We present the identification of hypothetical magnetic monopoles against Belle II beam background using self-organizing kohonen maps and autoencoders. These two unsupervised algorithms are compared to a Multilayer Perceptron and a superior signal efficiency of the Autoencoder is found at high background-rejection levels. Our results strengthen the case for using unsupervised machine learning techniques to complement traditional search strategies at particle colliders and pave the way to potential online applications of the algorithms in the near future.

1 Introduction

Machine learning has proven to be a valuable tool in reconstruction and analysis tasks for high-energy physics (HEP) [1]. In particular, the classification of signal and background using machine learning algorithms has sparked significant interest in recent times. The majority of these algorithms are trained in a supervised manner, and therefore rely on a prior definition of signal provided by a theoretical framework and simulations. However, detector signatures corresponding to elusive beyond the standard model physics (BSM) processes might be missed owing to a narrow signal definition or a mis-modelling of either signal or background.

Unsupervised and semi-supervised methods aim at the identification of signal features while minimizing predictions about signal or background. A data-driven approach is adopted allowing for a model-agnostic analysis that has the advantage of being independent from theoretical assumptions and therefore not confined to specific signal hypotheses and background modelling. In this document, we present two unsupervised machine learning methods with the objective to identify anomalous detector patterns that are potential indicators for unaccounted physics scenarios. The anomaly detection is performed such that it can be considered as a sophisticated filter mechanism that defines a potential signal region for further statistical analysis. In this paper, we focus on the filter itself.

The inner region of the Belle II detector is considered to demonstrate this filtering approach. The input data consists of pixel hits coming from the Belle II pixel detector presently featuring one layer of DEPFET (Depleted P-channel FET) [2] silicon sensors. The unsupervised data-driven machine learning algorithms are trained on beam-background data recorded by the pixel detector when a single beam was circulating in the collider. Simulated pixel hits by hypothetical long-lived magnetic monopoles serve as anomalous events to evaluate the performance of the presented algorithms. The unsupervised techniques are compared to a neural network, that is trained in a supervised manner.

The paper is structured as follows: in Sect. 2, the Belle II experiment and the pixel detector are introduced. Subsequently, the dataset and data preprocessing are described in

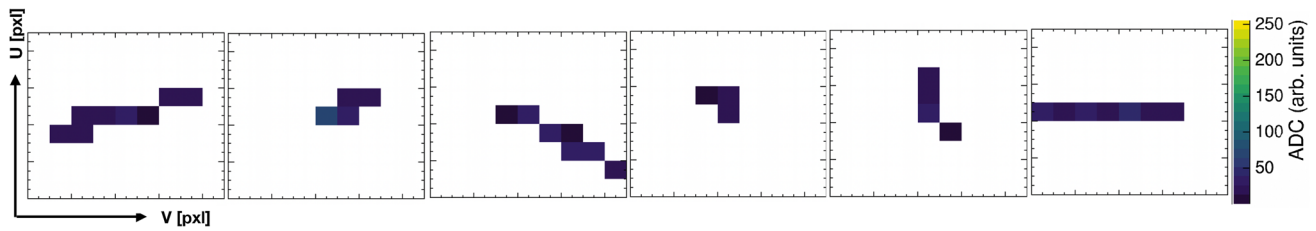a e-mail: katharina.dort@physik.uni-giessen.de (corresponding author)

Fig. 1 Example clusters of single-beam background particles from the inner region of the PXD (pixel dimensions: $50\,\mu m \times 55\,\mu m$). The seed pixel is located at the centre of the 9×9 pixel matrix. The colour scale represents the single-pixel charge in ADC values

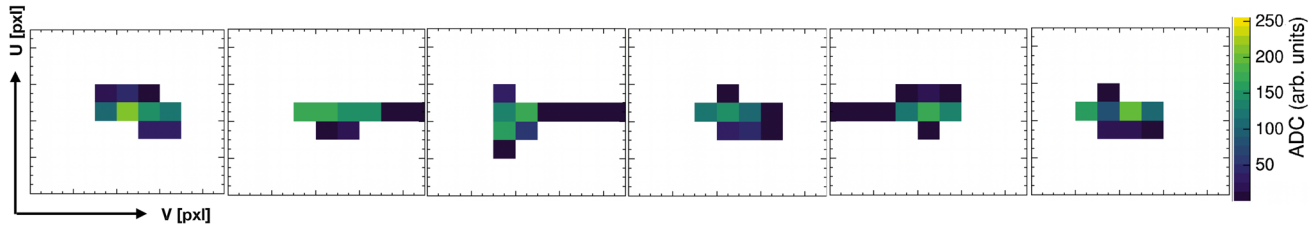**Fig. 2** Example clusters of simulated magnetic monopoles from the inner region of the PXD (pixel dimensions: $50\,\mu m \times 55\,\mu m$). The seed pixel is located at the centre of the 9×9 pixel matrix. The colour scale represents the single-pixel charge in ADC values

Sect. 3. The different machine learning algorithms are introduced in Sect. 4 and their performance is presented in Sect. 5.

2 The Belle II experiment

The Belle II experiment, located at the SuperKEKB accelerator, has started operation in spring 2019. SuperKEKB provides electron-positron collisions with a nominal centre-of-mass energy of 10.58 GeV and a design instant luminosity of $> 5 \times 10^{35}\,\mathrm{cm^{-2}\,s^{-1}}$ [3].

The Belle II detector is composed of several sub-detectors, that are arranged cylindrically around the interaction point [4]. This document only treats the inner-most detector: an all-silicon pixel detector that is part of the Belle II tracking system, as detailed below. The silicon vertex detector (SVD) based on double-sided silicon strip sensors and the central drift chamber (CDC), a cylindrical wire chamber, are the other two detectors making up the Belle II tracking system.

2.1 The Belle II pixel detector

The Belle II pixel detector (PXD) consists of pixelated DEPFET sensors, that are arranged in two layers at radii of 14 mm and 22 mm from the beam pipe [5,6]. Presently, only the first layer is installed. The PXD features nearly 4 million pixels with pixel sizes between $50\,\mu m \times 55\,\mu m$ and $50\,\mu m \times 85\,\mu m$ and a thickness down to $75\,\mu m$.

The data rate coming from the PXD is foreseen to reach about 20 GB/s, which necessitates an online reduction scheme [7]. The FPGA-based data-reduction system online selection node (ONSEN) is able to reduce the data rate by

a factor of 30 by using reconstructed tracks provided by the online event reconstruction, which are extrapolated to the PXD layers. A region-of-interest (ROI) is defined around the intercept of these tracks with the detector layers and only pixel hits within ROIs are considered [8]. However, this filtering mechanism relies on reconstructable particle tracks, as PXD hits outside a ROI are discarded by the ONSEN. In particular, particles with a low transverse momentum or high energy loss can escape tracking. As a consequence, no ROI is generated and the PXD data associated with these particles is lost.

To guarantee a high signal efficiency of particles with non-reconstructable tracks, a new veto system based on machine learning is proposed. A proof-of-principal for a veto system dedicated to the identification of slow pions has already been presented in the past [9]. In this document, the cluster rescue veto system is extended to exotic or *anomalous* particle signatures that do not generate a reconstructable particle track. To assess the efficiency of the cluster rescue mechanism, we simulate the creation of long-lived hypothetical magnetic monopoles in the particle collision. As a consequence of their high energy loss, magnetic monopoles are stopped in the inner layers of the Belle II detector [10]. The lack of hits in the outer sub-detectors inhibits the reconstruction of tracks, which also leads to the deletion of PXD data associated with a monopole by the ONSEN, once the ROI selection is switched on. The aim of the proposed veto system is to identify the relevant PXD data based on anomalous event signatures and tagging it to prevent deletion. We consider unsupervised machine learning algorithms to generate the veto, that could potentially run online during data-taking on FPGA-based systems.

Moreover, the tracking algorithms are currently relying on information from the CDC to form a track. However, with the rescued PXD data, a novel tracking approach using only the silicon-based PXD and SVD detectors could be envisioned.

3 Data generation

PXD background data was recorded in dedicated beam-background runs taken in 2020. For these runs, only a single particle beam circulated in the Belle II detector. Background generation mechanisms such as the interaction of the circulating beam with residue particles in the beam pipe are responsible for the background hits detected in these runs. The PXD hits generated by background are characterised by a small charge signal in each pixel, as shown by the example clusters displayed in Fig. 1. The hits are nevertheless detected by the PXD due to the high signal-to-noise ratio of the DEPFET pixel sensors. The v-coordinate is along the beam direction and the u-coordinate perpendicular to it. In view of future online applications of the investigated algorithms, the raw PXD hit information is used without applying an offline calibration.

Dedicated background samples are regularly produced in Belle II for background monitoring and as background-overlay files for simulations. It is thus possible to obtain updated background samples in the future, which allows for an adaptation of the classification algorithm to changing beam conditions.

The signal events are simulated using the official Belle II software framework `basf2` [11]. The creation of monopole-antimonopole pairs from electron-positron collisions is considered. The magnetic charge of the monopole is set to 68.5 e in accordance with the Dirac theory [12,13] and the mass to 3 GeV. A full detector simulation with `basf2` is performed, including the interaction of particles with the PXD layers. The simulated PXD hits associated with magnetic monopoles are shown in Fig. 2.

The PXD information used as input for the machine learning algorithms is extracted as follows: For each simulated event and for the background events, the charge values of a 9×9 pixel matrix are considered around the PXD hit with the highest charge value (*seed pixel*). The matrix size is sufficiently large to capture the entire cluster for the majority of events and small enough to guarantee a fast convergence of the investigated algorithms. In addition, the global position of the seed pixel within the PXD is extracted. In total, 84 features are considered, which are normalized to the range [0, 1] to avoid dominance of a single parameter.

4 Machine learning techniques

We propose a sophisticated filter based on unsupervised machine learning algorithms to identify anomalous signatures in the PXD. The filter operates on a matrix-by-matrix basis and labels each 9×9 matrix as anomalous or normal based on an *anomaly score*. While the scope of the anomaly score depends on the selected algorithm, we adopt the definition that low values represent normal and high values anomalous events.

4.1 Performance metrics

For all algorithms, the receiver operating characteristics (ROC) is obtained by scanning the signal efficiency ϵ_S and recording the background rejection ϵ_B. The area-under-curve (AUC) is commonly used as a figure of merit for the performance of classifiers [14]. For anomaly detection, a high background rejection is particularly desirable. Therefore, the signal efficiency at three different operation points featuring a high background-rejection level are studied as well, i.e. the signal efficiencies $\epsilon_S(\epsilon_B = 10^{-2})$, $\epsilon_S(\epsilon_B = 10^{-3})$ and $\epsilon_S(\epsilon_B = 10^{-4})$ are extracted.

The uncertainty is extracted by repeating the training and evaluation five times with random shuffling of input vectors i.e. a vector in the training set in the first iteration can be assigned to the evaluation set in the second one. In each iteration, the performance metrics are determined. Their mean represents the nominal value and the quadratic sum of deviations the uncertainty.

4.2 Multilayer perceptron (MLP)

First, a supervised multilayer perceptron (MLP) is considered, to which the unsupervised learning approaches are compared.

The supervised training is performed using 350k background and signal events each. After each training epoch, a dedicated testing set is presented to the MLP containing additional 150k events for both classes. The training is stopped automatically once the reduction of the predicted error (*loss*) from the testing set is only marginal. An evaluation set comprising 500k events for each class is considered to assess the performance of the algorithm.

4.3 Self-organizing maps (SOM)

A self-organizing map (SOM) is an unsupervised machine learning technique enabling the transformation of a high-dimensional dataset to a low-dimensional discrete grid, while keeping the topological structure [15,16]. After training, vectors that are close in the high-dimensional input space are represented by adjacent grid points in the low-dimensional

Table 1 Hyperparameters for the multilayer perceptron

Parameter	Value
Activation function	Rectified linear unit (ReLu)
Loss function	Cross entropy
Optimizer	Stochastic gradient descent (SGD)
Learning rate	1e−4
Momentum	0.9
Trainable parameters	15 k

Table 2 Hyperparameters for the self-organizing maps

Parameter	Value
Nodes	100
Neighbourhood function	Gaussian
Gaussian width	16
Learning rate	0.03
Trainable parameters	8 k

Table 3 Hyperparameters for the autoencoder

Parameter	Value
Activation function	Rectified linear unit (ReLu)
Loss function	Mean squared error
Optimizer	Adaptive moment estimation
Learning rate	1e−4
Trainable parameters	13 k
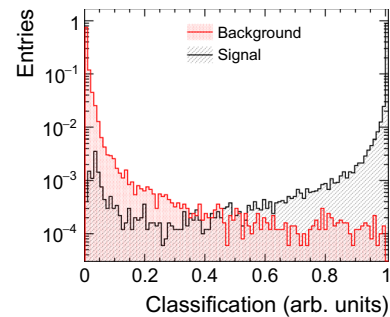
Fig. 3 MLP classification distribution for signal and background
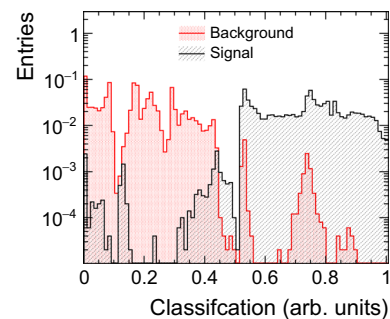
Fig. 4 SOM classification distribution for signal and background

space. The same training, testing and evaluation sets as for the MLP are used.

4.4 Autoencoders (AE)

An autoencoder (AE) is a feed-forward multilayer neural network that aims to reproduce the input vector without using an identity mapping. It consists of two parts: an encoder and decoder. While the encoder compresses the input to a lower-dimensional vector, the decoder reconstructs the original input from the reduced representation. The latent space in the centre of the AE is an information bottleneck that enforces the selection of relevant patterns from the input data.

During training, only background events are presented to the AE and their reconstruction error is minimized, making the AE specialized in the reproduction of background events. In the evaluation phase, the AE is able to recognize background events by a low reconstruction error. Signal events appear anomalous to the AE and are characterized by a high error, that can therefore serve as an anomaly/classification score [17].

5 Results

In the following, the hyperparameters of the three algorithms are presented and the identification performance of hypothetical magnetic monopoles against beam background is studied. The evaluation for other example signals is presented in Appendix E.

5.1 Hyperparameters

The hyperparameters of the three algorithms are listed in Tables 1, 2 and 3, and the architectures of the MLP and the AE is shown in Appendix A. All algorithms are optimised by performing a grid search of possible hyperparameters and selecting the ones yielding the highest signal efficiency at low background levels of 10^{-4}.

For the SOM, the dimension of the low-dimensional grid space is set to one to allow for a comparable performance evaluation as for the other two machine learning techniques. The low-dimensional representation will therefore span only a single line, with the aim to have grid points responding to background cluster on the one end and signal on the other end.
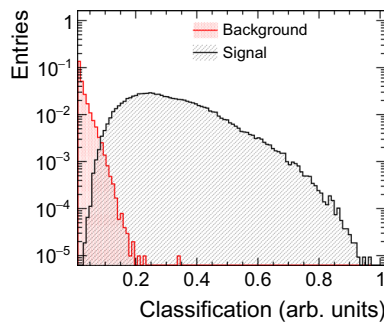
Fig. 5 AE classification distribution for signal and background
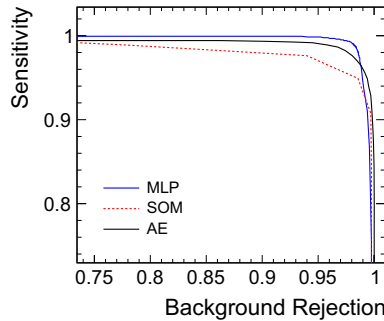

Fig. 6 ROC curves for the three machine learning algorithms: a supervised MLP – Multilayer Perceptron, an unsupervised SOM – Self-Organizing Map and an unsupervised AE – Autoencoder

5.2 Classification performance

The one-dimensional classification distribution is presented in Figs. 3, 4 and 5 for the three algorithms.

For the MLP, signal/background is suppressed by approximately three orders of magnitude for low/high classification values. In case of the SOM, the trained one-dimensional grid serves as classification axis and exhibits similar suppression factors as the MLP. For the AE, a signficant overlap region is visible at low classification scores, but a high signal purity is achieved at high values. The reconstruction of example matrices for both signal and background is shown Appendix B.

The ROC curves for the three algorithms are presented in Fig. 6 and the performance results are listed in Table 4.

The AUC of the MLP reaches $99.69^{+0.01}_{-0.01}\%$, which indicates an excellent classification performance. The AUC of the AE is about 1% lower and the one for the SOM about 3% lower. At high background-rejection levels of

$\epsilon_S(\epsilon_B = 10^{-4})$, the signal efficiency of the MLP deteriorates to $39.5^{+3.6}_{-2.3}\%$. The efficiency for the AE still reaches $60.1^{+3.3}_{-2.7}\%$ showing that the AE outperforms the other two algorithms if a high signal purity is demanded. The origin of the different classification performance of the MLP and the AE is investigated in Appendix D and the robustness of the AE against changing signal and background sets is shown in Appendix C.

6 Summary and outlook

The unsupervised identification of anomalous pixel-detector data using Self-Organizing Maps and Autoencoders was presented. To exemplify the approach, hypothetical magnetic monopoles at the Belle II pixel detector were simulated and identified against beam-background data. The two unsupervised algorithms have shown a similar performance as a supervised multilayer perceptron. The Autoencoder outperforms the multilayer perceptron, if high background-rejection levels are required.

This study is an essential cornerstone for future online applications of anomaly detection at the Belle II pixel detector in order to further improve its sensitivity to undiscovered physics. Moreover, the identification of anomalous data could be considered for experiment protection and data-quality monitoring as well, since anomalies in the data could hint at unsatisfactory beam stability or malfunctioning detector components.

Table 4 Performance metrics of the investigated machine learning techniques for the evaluation set

Algorithm	AUC [%]	$\epsilon_S(10^{-1})$ [%]	$\epsilon_S(10^{-2})$ [%]	$\epsilon_S(10^{-3})$ [%]	$\epsilon_S(10^{-4})$ [%]
MLP	$99.69^{+0.01}_{-0.01}$	$99.99^{+0.01}_{-0.01}$	$96.9^{+0.4}_{-0.2}$	$83.4^{+0.9}_{-1.2}$	$39.5^{+3.6}_{-2.3}$
SOM	$96.30^{+0.12}_{-0.24}$	$89.3^{+0.3}_{-0.2}$	$80.3^{+1.1}_{-1.4}$	$49.9^{+4.2}_{-4.1}$	$28.8^{+5.1}_{-6.2}$
AE	$98.86^{+0.05}_{-0.02}$	$98.2^{+0.1}_{-0.3}$	$95.9^{+1.2}_{-0.4}$	$87.4^{+1.4}_{-1.2}$	$60.1^{+3.3}_{-2.7}$

Table 6 Network architecture of the AE

Layer	Nodes	Activation
Linear	84	ReLU
Linear	64	ReLU
Linear	16	ReLU
Linear	8	ReLU
Linear	16	ReLU
Linear	64	ReLU
Linear	84	ReLU

Appendix A: Neural network architecture

In Tables 5 and 6, the architectures for the MLP and the AE are shown that were used to obtain the results presented in this paper.

Table 5 Network architecture of the MLP

Layer	Nodes	Activation
Linear	84	ReLU
Linear	128	ReLU
Linear	32	ReLU
Linear	16	ReLU
Linear	1	ReLU

Appendix B: AE reconstruction performance

Example clusters from beam background before and after reconstruction are depicted in Fig. 7. The upper row shows the original input matrices of background events. In the lower row, the output of the AE is displayed with matrices in the same column belonging to the same event. While the overall shape is faithfully reproduced, the AE has a tendency to add more pixel hits to the matrix especially directly adjacent to the original pixel hits.

Example matrices of signal before and after reconstruction are depicted in Fig. 8. The unfamiliar shape and in particular the high ADC values pose a problem for the AE. While some shape characteristics are still preserved, the AE distorts the matrices in a way that they appear to be more background-like, i.e. instead of the concentrated pixel hit distribution, the distributions become long-stretched and the high ADC values are transformed to low ones.
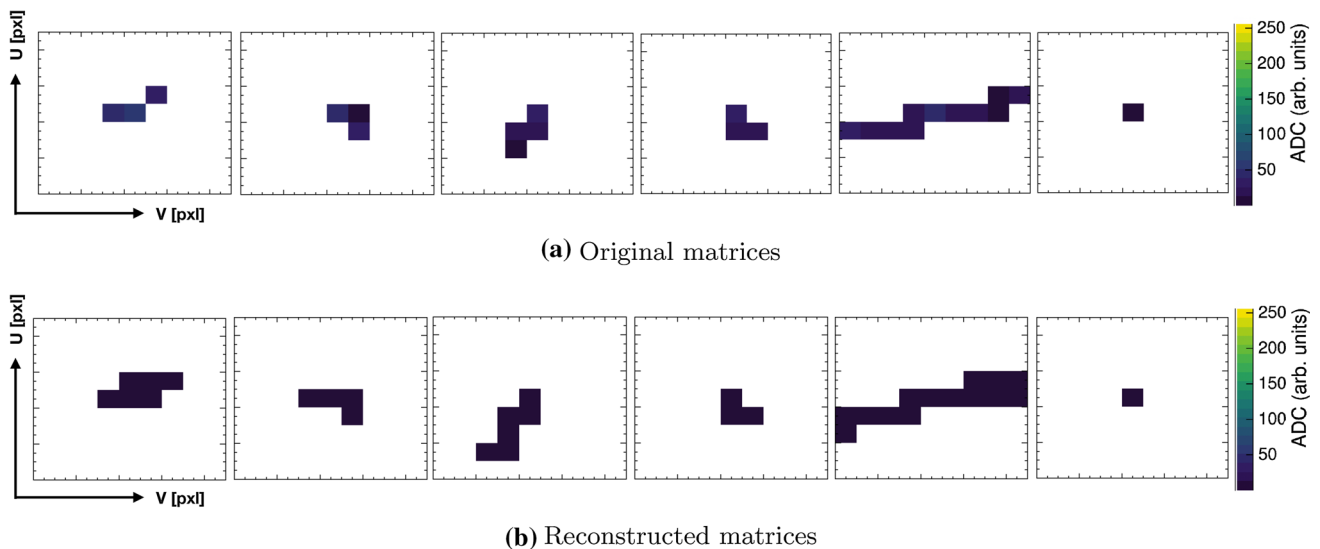

(a) Original matrices

(b) Reconstructed matrices

Fig. 7 **a** Input and **b** reconstructed pixel matrices associated with background particles

(a) Original matrices
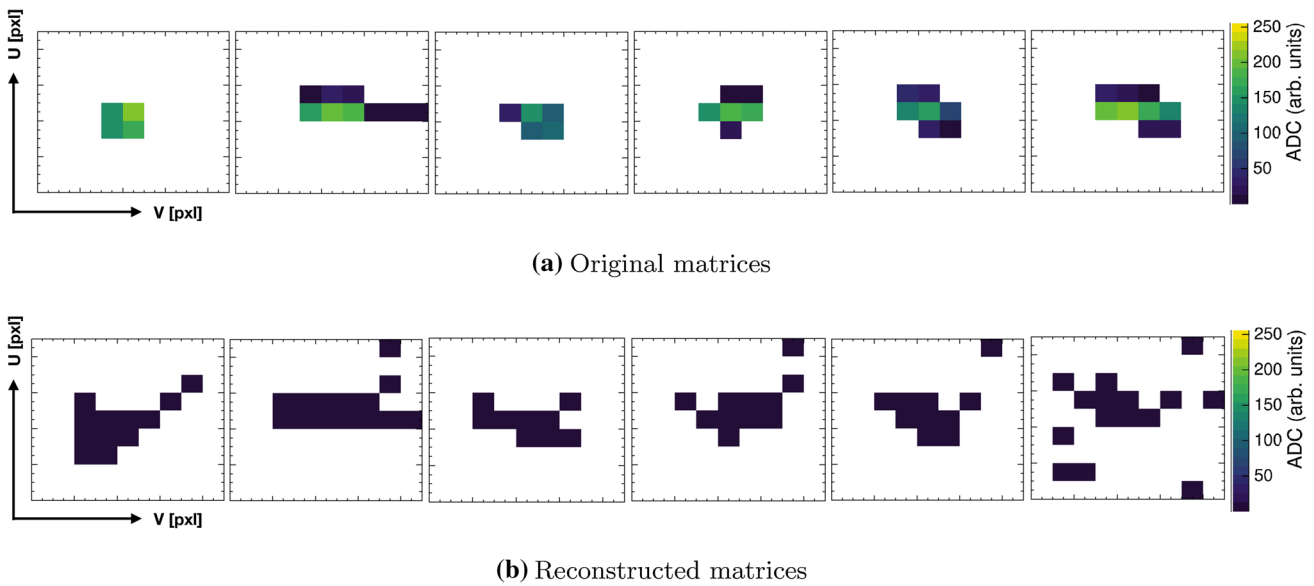

(b) Reconstructed matrices

Fig. 8 **a** Input and **b** reconstructed pixel matrices associated with signal particles

Appendix C: Robustness

In the following, the robustness of the AE with respect to different input datasets is discussed.

C.1: Size of pixel matrix

The nominal input features consist of a 9×9 pixel matrix as well as the global hit coordinates of the seed pixel within the PXD. To investigate the impact of the size of the pixel matrix, larger matrices are extracted from the same data and simulation sets. The training, testing and evaluation steps of the AE are repeated and three ROC curves belonging to matrix sizes of 9×9, 15×15 and 25×25 are computed, as illustrated in Fig. 9. An improvement for larger matrix sizes at low background-rejection levels is observable. As the improvement is comparably small, the increase in computation time is not justified for this study and therefore the 9×9 matrix is kept.
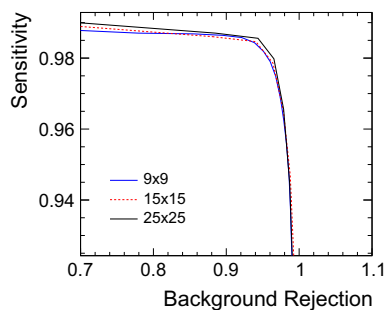
C.2: Signal in training data

We consider the hypothetical scenario of having a background training set that is contaminated with signal events. In this case, the AE is already exposed to signal during training and will consequently adapt to it, potentially resulting in a lower reconstruction error for signal and therefore a weaker classification performance. To evaluate the impact of a contaminated data set, 0.01%/0.1%/1% of the training events are replaced by signal and the analysis is repeated.

The resulting ROC curves are presented in Fig. 10. The performance of the AE deteriorates for 1% signal in the data but is not affected for the lower contamination values. It can be concluded that a signal component in the training set has to be sufficiently rare to guarantee the optimal performance of the AE.
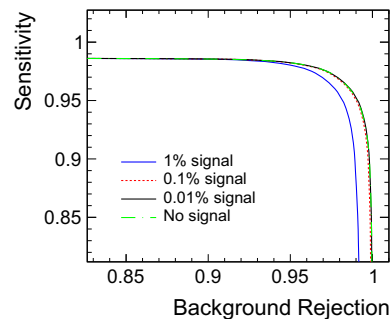

Fig. 9 ROC curves using different matrix sizes for the input data

Fig. 10 ROC curves for different signal levels in the training sample
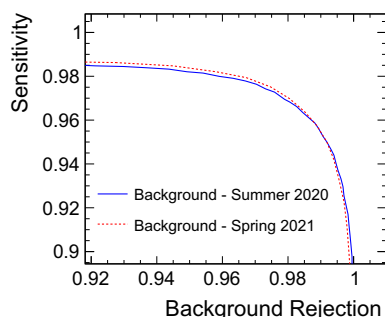
Fig. 11 ROC curve for different beam-background samples recorded in different run periods

Fig. 12 ROC curves for simulated and real beam background. In both cases, the same signal set is used

C.3: Stability over time

As the luminosity of SuperKEKB is still increasing steadily, changing beam conditions in different run periods are expected. To investigate the stability of the AE over time, beam background samples from different run periods were tested and the resulting ROC curves are presented in Fig. 11. The results for both run periods are similar, which confirms that the different background conditions are negligible for the identification of the magnetic monopole signal. If a stronger change in background conditions occurs in the future, the data-driven training of the AE would ensure a relatively fast adaptation of the classifier.

C.4: Simulated beam background

Data from the PXD is subject to various systematic effects that were not explicitly accounted for in this study such as particularly noisy pixels or pixel-by-pixel variations resulting in a slightly different AUC value for the same energy deposition. In online applications, correcting for these effects is not easily feasible. It is therefore beneficial to examine their impact on the classification performance. To this end, dedicated beam-background simulations are performed. The extraction of pixel matrices is conducted in the same manner as for data and an AE is trained and evaluated on the simulation.

The beam-background simulation is performed with `basf2`. The framework uses the strategic accelerator design (SAD) software [18] and Geant4 [18] for the generation of beam-background particles and subsequently a Geant4-based approach to propagate the particles through the detector. A detailed comparison of simulated beam background and data for the inner tracking detectors can be found elsewhere [19].

The resulting ROC curve using simulated background is compared to the one from data in Fig. 12. Good agreement between the two curves is achieved, which implies a sufficiently good background modelling in simulation. In
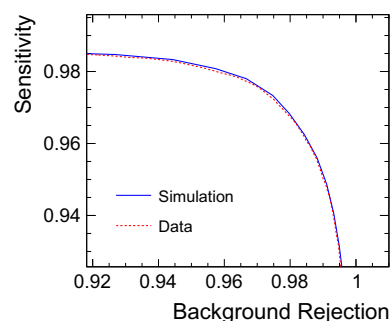
addition, remaining differences between data and simulation related to noisy pixels or pixel-by-pixel variations do not affect the results. Although noisy pixels have the potential to disturb the training, the results from the previous section imply that a disturbance of about 1% of the considered matrices is necessary to affect the AE. While there are uncertainties in the ADC values due to pixel-by-pixel variations, these uncertainties are small compared to the large difference in energy deposition between background and signal.

C.5: Standard model background

The performance of the AE is assessed against different standard model (SM) background sources, which are simulated in `basf2`. The AE is trained on beam background and subsequently applied to discriminate magnetic monopoles from other SM background particles. The resulting classification distributions and ROC curves are shown in Figs. 13 and 14, respectively. For all particle species, the anti-particle was included in the simulation set as well. In contrast to beam background, the electrons, muons and protons are created at the interaction point and thus enter the PXD under a steeper angle with respect to the PXD surface. As a result, the clusters are smaller and have less resemblance with the long-stretched clusters of magnetic monopoles, which facilitates identifica-
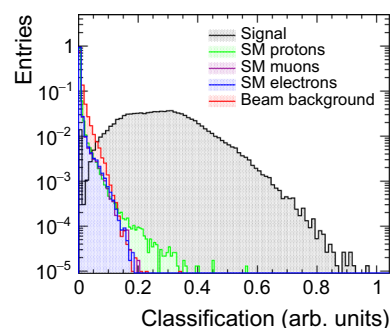

Fig. 13 AE classification distribution for signal and different SM (standard model) background
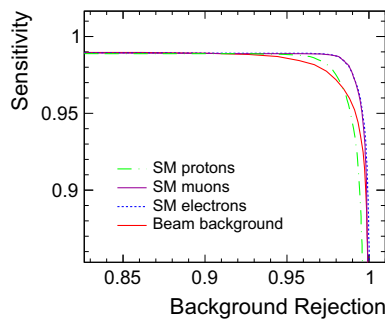
Fig. 14 ROC curve for different SM (standard model) background

tion and results in an improved identification performance for electrons and muons. The degraded performance for the protons is related to their higher energy deposition compared to the lighter particles, which gives rise to high-charge clusters resembling magnetic-monopole clusters.

Appendix D: Difference in classification performance

In Sect. 5, it was shown that the AE exhibits a lower AUC compared to the MLP but a higher signal efficiency at low background levels. The lower AUC is found to be related to beam-background matrices with uncommon cluster shapes that lead to a high anomaly score. Example matrices are depicted in Fig. 15. In comparison to the majority of beam-background matrices (cf. Fig. 1), broader clusters with a higher number of pixel hits are characteristic for these matrices. At low/intermediate background-rejection levels, they

are wrongly classified as signal by the AE, leading to the lower AUC compared to the MLP.

At high background-rejection levels, we considered those signal matrices with a high AE classification score, but a comparably low MLP score, which are responsible for the lower signal efficiency of the latter. Example matrices are depicted in Fig. 16. As these clusters have little resemblance with background clusters that the AE was trained on, the AE assigns a high classification score marking them as signal. For the MLP, the classification is hampered by the uncommon signal features when compared to the majority of signal matrices (cf. Fig. 2). As these features were rare in the training set as well, the MLP has not picked them up correctly leading to a lower classification score compared to the AE and thus a lower signal efficiency.

Appendix E: Different HIP signals

The results for the three algorithms are cross checked with two other HIP candidates: Standard Model antideuterons and hypothetical dyons, which carry both magnetic and electric charge.

Antideuterons First, antideuterons with a momentum up to 1 GeV are identified against the beam-background data. Antideuterons are particularly interesting in cosmological dark-matter searches [20]. Studying their properties at collider experiments delivers information about the formation of antideuterons, which is a crucial input for these searches. At Belle II, low-momentum antideuterons have a limited range in the detector and their identification is hampered by the
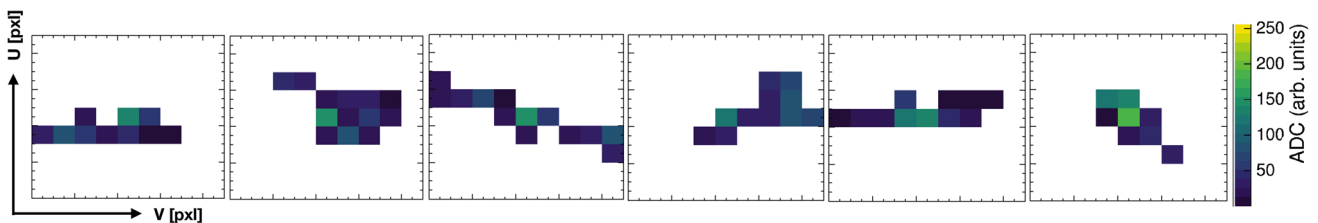

Fig. 15 Example beam-background matrices with a low MLP and a comparably high AE classification score
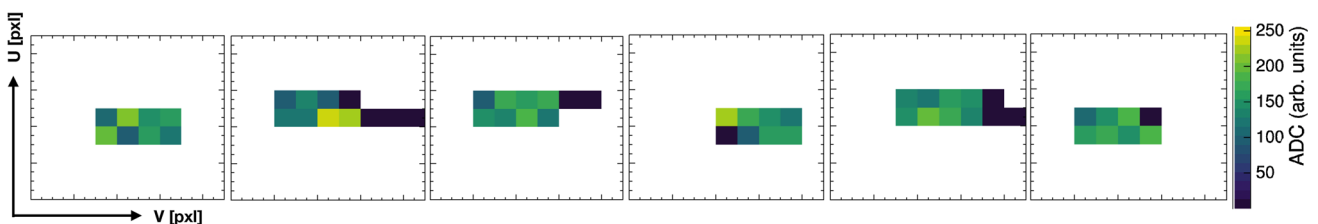

Fig. 16 Example magnetic-monopole matrices with a high AE and a comparably low MLP classficiation score
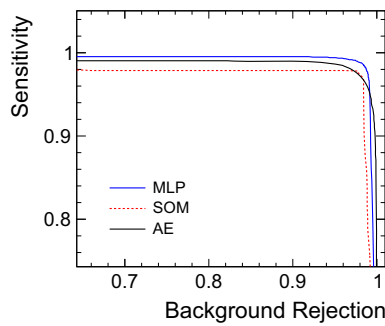
Fig. 17 ROC curve for the classification of antideuterons (signal) and beam background for three different algorithms. *MLP* multilayer perceptreon, *SOM* self-organising map, *AE* autoencoder

Table 7 Performance metrics of the investigated machine learning techniques for simulated antideuterons

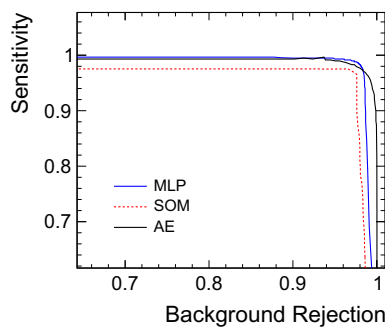| Algorithm	AUC [%]	$\epsilon_S(10^{-4})$ [%]
MLP	$97.18^{+0.06}_{-0.04}$	$10.4^{+5.6}_{-4.1}$
SOM	$91.22^{+0.22}_{-0.48}$	$8.1^{+4.5}_{-4.3}$
AE	$96.10^{+0.12}_{-0.24}$	$32.1^{+2.3}_{-1.4}$

Fig. 18 ROC curve for the classification of hypothetical dyons (signal) and beam background for three different algorithms. *MLP* multilayer perceptreon, *SOM* self-organising map, *AE* autoencoder

current PXD readout scheme, similar to other HIPs such as the hypothetical magnetic monopoles.

While the architecture of the three machine learning algorithms remains the same for the antideuteron studies, the MLP and SOM require re-training due to the novel signal definition. The results for antideuterons are depicted in Fig. 17 and the AUC as well as the signal efficiency at $\epsilon_S(10^{-4})$ are listed in Table 7. Similar to the results for the magnetic monopoles, the MLP exhibits the highest AUC, but the AE has the best signal efficiency at high background-rejection levels.

Table 8 Performance metrics of the investigated machine learning techniques for simulated hypothetical dyons

Algorithm	AUC [%]	$\epsilon_S(10^{-4})$ [%]
MLP	$99.06^{+0.03}_{-0.03}$	$33.9^{+2.6}_{-2.3}$
SOM	$96.75^{+0.09}_{-0.08}$	$25.3^{+5.3}_{-4.9}$
AE	$98.73^{+0.07}_{-0.05}$	$62.1^{+1.3}_{-1.5}$

Dyons Dyons are hypothetical particles carrying both magnetic and electric charges that were first hypothesised by Schwinger [21] and appear in several theories such as grand unified theories [22] or string theory [23]. Here, we simulate dyons with an electric charge of 10 e and a magnetic charge of 68.5 e corresponding to the unit Dirac charge. Results are depicted in Fig. 18 and the performance metrics are given in Table 8. The overall trend is found to be compatible with the other two HIP candidates as well.

References

1. K. Albertsson et al., J. Phys.: Conf. **1085**, 022008 (2018)
2. H.G. Moser, D. Collaboration et al., Nucl. Instrum. Methods A **831**, 85 (2016)
3. K. Akai et al., Nucl. Instrum. Methods A **907**, 188 (2018)
4. T. Abe et al., (2010). arXiv:1011.0352
5. C. Marinas et al., Nucl. Instrum. Methods A **731**, 31 (2013)
6. O. Alonso et al., IEEE Trans. Nucl. Sci. **60**, 1457 (2013)
7. D. Levit, I. Konorov, D. Greenwald, S. Paul, IEEE Trans. Nucl. Sci. **62**, 1033 (2015)
8. T. Geßler, W. Kühn, J.S. Lange, Z. Liu, D. Münchow, B. Spruck, J. Zhao, IEEE Trans. Nucl. Sci. **62**, 1149 (2015)
9. S. Bähr, O. Sander, M. Heck, C. Pulvermacher, M. Feindt, J. Becker, J. Phys. Conf. Ser. **664**, 092001 (2015)
10. S. Cecchini, L. Patrizii, Z. Sahnoun, G. Sirri, V. Togo, (2016). arXiv:1606.01220
11. T. Kuhr, C. Pulvermacher, M. Ritter, T. Hauth, N. Braun, Comput. Softw. Big Sci. **3**, 1 (2019)
12. P.A.M. Dirac, Proc. Math. Phys **133**, 60 (1931)
13. K.A. Milton, Rep. Prog. Phys. **69**, 1637 (2006)
14. J.A. Hanley, B.J. McNeil, Radiology **143**, 29 (1982)
15. T. Kohonen, Biol. Cybern. **43**, 59 (1982)
16. K. Kiviluoto, Proc. Int. Jt. Conf. Neural Netw. **1**, 294 (1996)
17. J. An, S. Cho, Spec. Lect. IE **2**, 1 (2015)
18. S. Agostinelli et al., Nucl. Instrum. Methods A **506**, 250 (2003)
19. H. Tanigawa et al., Nucl. Instrum. Methods A **982**, 164580 (2020)
20. T. Aramaki et al., Phys. Rep. **618**, 1 (2016)
21. J. Schwinger, Science **165**, 757 (1969)
22. W. l'Yi, Y. Park, I.G. Koh, Y. Kim, Phys. Rev. Lett. **49**, 1229 (1982)
23. R. Dijkgraaf, E. Verlinde, H. Verlinde, Nucl. Phys. B **484**, 543 (1997)