

PAPER • OPEN ACCESS

## Fast kernel methods for data quality monitoring as a goodness-of-fit test

To cite this article: Gaia Grosso *et al* 2023 *Mach. Learn.: Sci. Technol.* 4 035029

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

You may also like

- [Fast neural network inference on FPGAs for triggering on long-lived particles at colliders](#)  
Andrea Coccaro, Francesco Armando Di Bello, Stefano Giagu *et al.*
- [Metric flows with neural networks](#)  
James Halverson and Fabian Rühle
- [Set-conditional set generation for particle physics](#)  
Nathalie Soybelman, Nilotpal Kakati, Lukas Heinrich *et al.*



The advertisement features a dark blue background with a large, stylized graphic of a hand holding a globe on the left. The globe is composed of several light blue and white curved segments. The ECS logo is positioned in the upper right, followed by the text 'The Electrochemical Society' and its tagline 'Advancing solid state & electrochemical science & technology'. The main event details are listed in white text: '247th ECS Meeting', 'Montréal, Canada', 'May 18-22, 2025', and 'Palais des Congrès de Montréal'. A call to action 'Showcase your science!' is placed below. A green circular callout in the bottom right corner contains the text 'Abstract submission deadline extended: December 20'. At the bottom left, the text 'ECS UNITED' is displayed in large white letters.

**ECS** The Electrochemical Society  
Advancing solid state & electrochemical science & technology

**247th ECS Meeting**  
Montréal, Canada  
May 18-22, 2025  
*Palais des Congrès de Montréal*

**Showcase your science!**

**Abstract submission deadline extended: December 20**

**ECS UNITED**



## PAPER

## OPEN ACCESS

## Fast kernel methods for data quality monitoring as a goodness-of-fit test

## RECEIVED

30 March 2023

## REVISED

11 July 2023

## ACCEPTED FOR PUBLICATION

28 July 2023

## PUBLISHED

25 August 2023

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Gaia Grosso<sup>1,2,3</sup> , Nicolò Lai<sup>1</sup> , Marco Letizia<sup>4,5,\*</sup> , Jacopo Pazzini<sup>1,2</sup> , Marco Rando<sup>4</sup> , Lorenzo Rosasco<sup>4,6,7</sup> , Andrea Wulzer<sup>8,9</sup>  and Marco Zanetti<sup>1,2</sup> 

<sup>1</sup> Dipartimento di Fisica e Astronomia, Università di Padova, Via Francesco Marzolo, 8, 35121 Padova, Italy

<sup>2</sup> INFN, Sezione di Padova, Via Francesco Marzolo, 8, 35121 Padova, Italy

<sup>3</sup> Experimental Physics Department, CERN, Espl. des Particules 1, 1211 Meyrin, Geneva, Switzerland

<sup>4</sup> MaLGa Center—DIBRIS, Università di Genova, Via Dodecaneso 35, 16146 Genova, Italy

<sup>5</sup> INFN, Sezione di Genova, Via Dodecaneso 33, 16146 Genova, Italy

<sup>6</sup> Istituto Italiano di Tecnologia, Via Morego 30, 16163 Genova, Italy

<sup>7</sup> CBMM—MIT, Cambridge, MA, United States of America

<sup>8</sup> Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology (BIST), Campus UAB, 08193 Bellaterra, Barcelona, Spain

<sup>9</sup> ICREA, Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys 23, 08010 Barcelona, Spain

\* Author to whom any correspondence should be addressed.

E-mail: [marco.letizia@edu.unige.it](mailto:marco.letizia@edu.unige.it)

**Keywords:** machine learning, high energy physics, data quality, monitoring, goodness-of-fit

## Abstract

We propose an accurate and efficient machine learning approach for monitoring particle detectors in real-time. The goal is to assess the compatibility of incoming experimental data with a reference dataset, characterising the data behaviour under normal circumstances, via a likelihood-ratio hypothesis test. The model is based on a modern implementation of kernel methods, nonparametric algorithms that can learn any continuous function given enough data. The resulting approach is efficient and agnostic to the type of anomaly that may be present in the data. Our study demonstrates the effectiveness of this strategy on multivariate data from drift tube chamber muon detectors.

## 1. Introduction

Modern high-energy physics experiments operating at colliders are extremely sophisticated devices consisting of millions of sensors sampled every few nanoseconds, producing an enormous throughput of complex data. Several types of technologies are employed for identifying and measuring the particles originated in the collisions; in all cases, the environmental conditions are severe, making the required performances challenging to achieve. Although the various subsystems are designed to offer redundancy, measurements can be undermined by malfunctions of parts of the experiment, either because of critical inefficiencies or because of possibly misinterpreted spurious signals. In addition to supervising the status (powering, electronic configuration, temperature, etc) of the various hardware components, data from all sources must thus be monitored continuously to assess their quality and to promptly detect any faults, possibly providing indications about their causes. Given the rate of tens of MHz at which data is gathered and the number of sensors to be checked, the monitoring process needs to be as automated as possible: approaches based on machine learning (ML) techniques are particularly appealing for this task and have started being employed by the experimental collaborations [1–4], complementing more traditional methods [5–9]. Data quality monitoring (DQM) consists, in essence, of comparing batches of data with corresponding reference samples gathered in nominal conditions; departures from the latter can then be analysed to identify their origin. The data processing must fit the computational constraints imposed by the frequency at which batches are delivered and by their size, with the latter depending on the granularity with which sensors are grouped and the statistical uncertainty aimed at.

In this work, we present the application of a methodology developed in the context of model-independent searches for new physics [10–12]—specifically of its kernel methods implementation [13] based on the Falkon [14] library—as an efficient and effective DQM tool. The method (dubbed New Physics Learning Machine (NPLM), see section 3) implements a goodness-of-fit statistical test in many dimensions. It leverages the ability of classifiers to infer the underlying data-generating distributions in order to estimate the likelihood ratio, eventually assessing goodness-of-fit by a hypothesis test based on the likelihood-ratio test statistics. At variance with anomaly detection methods, which are more common in the ML literature, NPLM is not only sensitive to the presence of outliers in the data. It searches for discrepancies in the statistical distribution of the data, relative to the expected (reference) distribution. It is thus also sensitive to data discrepancies that are due to an anomalous statistical concentration of multiple events, each of which is not an outlier but on the contrary very typical in the reference data population. The Falkon-based implementation of NPLM offers tremendous advantages in terms of training time compared to the one based on neural networks. It can thus be used for DQM.

Our proposed approach addresses two main limitations of conventional DQM methods. First, it introduces a sophisticated multivariate statistical data analysis strategy for the assessment of the data quality. Conventional methods typically employ one or a number of one-dimensional (1D) marginals, therefore their sensitivity to anomalies depends critically on the choice of input variables and it does not exploit anomalies in their correlations. A key feature of NPLM is instead the capability of examining the phase space as a whole, in many dimensions and exploiting correlations boosting the sensitivity to anomalous data. For the same reason, it will be possible to run NPLM on low-level quantities that require limited pre-processing. This can be advantageous for DQM, as it allows it to deal with almost raw data from the detectors' electronic front-ends, therefore limiting the bias introduced by further manipulations that could hide issues in the data.

Secondly, our proposal can seed a standardisation of the design of DQM methodologies. Current strategies are designed specifically for each detector system and monitoring task, while NPLM is a universal method with a wide range of applicability. An intelligent ad hoc selection of the relevant physical variables will still be beneficial for the sensitivity, but NPLM provides a universal framework to assess the agreement with the reference of their distribution. The availability of a standardised approach based on NPLM could be also exploited to validate the adequacy (or establish the inadequacy) of conventionally-designed approaches.

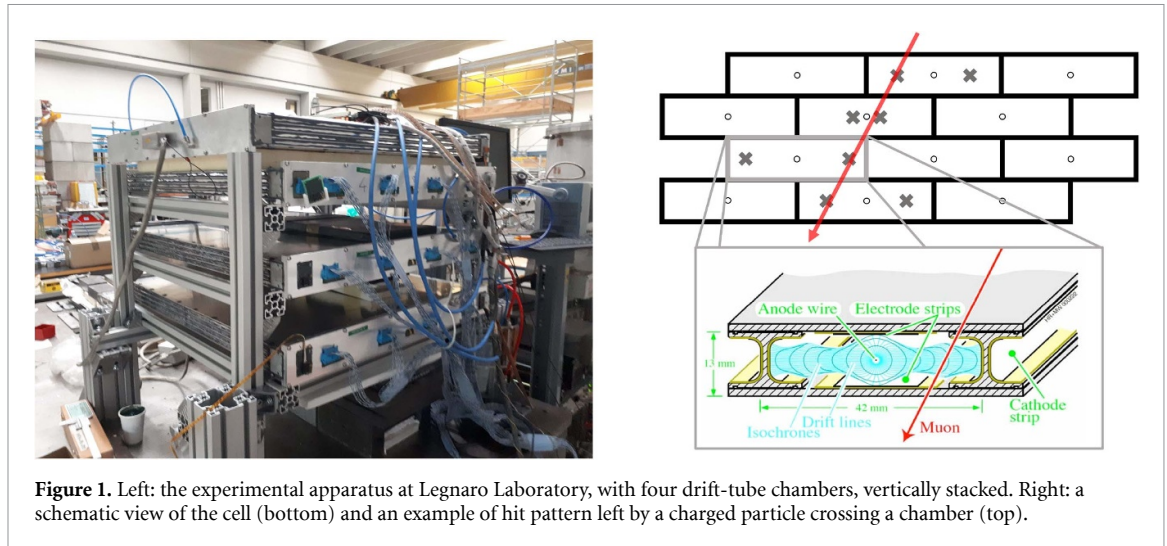
In order to test the effectiveness of NPLM for DQM, we exploit an experimental setup which we have full control of, consisting of a reduced-size version of the muon chambers installed in the Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC). The setup is operated as a cosmic muon telescope. As explained later, scaling tests are performed to assess the performances of the DQM algorithm in view of its possible deployment during standard LHC operations.

The paper is organised as follows. In the next section, we introduce the experimental setup and the algorithm input variables. These include a reference data set collected under standard conditions and smaller samples with anomalous controlled behaviours. The ML model and our core strategy are then described in section 3, whereas an overview of the results is given in section 4. Finally, the last section is devoted to conclusions and further developments.

## 2. Experimental setup

For this research, we exploited an experimental apparatus consisting of a set of drift tube (DT) chambers housed at the Legnaro INFN National Laboratory (figure 1, left). These chambers are a smaller in size copy of those deployed in the CMS experiment at the LHC [15]. The basic element of a DT chamber is a 70 cm long tube with a cross section of  $4 \times 2.1 \text{ cm}^2$  (figure 1, bottom right). Inside each tube, an electric field is produced by an anodic wire laid in the centre and two cathodic strips (cathodes) on the sides; the former is set at a voltage of 3.6 kV, the latter at  $-1.2 \text{ kV}$ . An additional pair of strips at 1.8 kV is placed above and below the wire to improve the homogeneity of the field. The tubes are filled with a mixture of argon and carbon dioxide gas (85%–15%) that gets ionised by charged particles passing through it. The produced electrons drift towards the wire at a constant velocity along the field lines, where they are collected. For each tube, the front-end electronics record the arrival time of the ions, amplify the signal, and filter out noise below a specific threshold (nominally 100 mV).

A DT chamber consists of 64 tubes arranged in four layers of 16 tubes each. The layers are staggered horizontally by half a cell. The setup at Legnaro records muons from cosmic rays, which occur at a rate of about 1 per minute per  $\text{cm}^2$  at sea level. Data acquisition occurs continuously at a rate of 40 MHz, without the need for any trigger logic. An external time reference is provided by plastic scintillators placed in between the DT chambers; the corresponding information is added to the data stream and used in the following analysis steps.



**Figure 1.** Left: the experimental apparatus at Legnaro Laboratory, with four drift-tube chambers, vertically stacked. Right: a schematic view of the cell (bottom) and an example of hit pattern left by a charged particle crossing a chamber (top).

Thanks to the homogeneity of the electric field, the particle's position within each tube (with a left-right ambiguity) is linearly dependent on the drift time. Namely, the distance of the muon track from the wire reads

$$x_{\pm} = \pm v_d (t_{\text{hit}} - t_0) = \pm v_d t, \quad (1)$$

with  $t_{\text{hit}}$  the time associated to each signal in a tube (called a hit). The two parameters are the drift velocity  $v_d$ , known by means of a calibration procedure (in our case,  $v_d = 53 \mu\text{m ns}^{-1}$ ), and the time pedestal  $t_0$ , which can be deduced from the timing information provided by the scintillators<sup>10</sup>. The drift time  $t$  is obtained by the difference between  $t_{\text{hit}}$  and the time pedestal.

The hits occurring in a time window of  $90 \mu\text{s}$  centred around the signal provided by the scintillators are grouped in quadruplets (with one hit pertaining to each of the four layers as in figure 1, right top). Then, a linear fit is performed on each of the quadruplets and the candidate muon track is obtained from the combination yielding the best  $\chi^2$ . In this way the trajectory of the muon in the plane transverse to the tubes is determined, with a precision on the position of about  $180 \mu\text{m}$  and on the slope of about 1 mrad. Tracks from various DT chambers can be combined to determine the 3D muon trajectory; in the following we will however consider only the 2D measurement.

If the detector conditions are anomalous, the efficiency and accuracy of the muon track reconstruction may be compromised. Ensuring the proper operation of the detector thus requires monitoring the quality of the recorded data. In what follows, we consider six basic quantities related to the passage of a muon through a DT chamber:

- Drift times  $t_i$ : the four drift times associated with the muon track. The drift time distribution is displayed in figure 2 in different ranges for the muon track angle  $\theta$  (or slope, see the next item), showing the correlations between these two variables. The  $t_i$  distributions are also reported in figures 3 and 4.
- Slope  $\theta$ : the angle formed by the muon track with the vertical axis. The chamber efficiency is expected to drop beyond  $|\theta| \sim 40^\circ$  as we see in figures 3 and 4.
- Number of hits  $n_{\text{Hits}}$ : the number of hits recorded in a time window of one second around the muon crossing time. Many spurious hits are present in addition to those due to the passage of a muon. The noise rate depends on the environmental conditions, with the one at the LHC orders of magnitude larger than that of our laboratory in Legnaro, but the recorded spurious hits rate can also be affected by issues related to the detector operation conditions.

The six variables  $x = \{t_1, \dots, t_4, \theta, n_{\text{Hits}}\}$  will be the input features of the NPLM algorithm for DQM, described in the next sections. Notice that the data are gathered from the subset of tubes in a single chamber that geometrically matches the scintillators, i.e. about three tubes per layer.

<sup>10</sup> In addition a mean-timer algorithm [16] is executed on the back-end board receiving the data. The timing information provided by that algorithm is currently not used in this analysis.

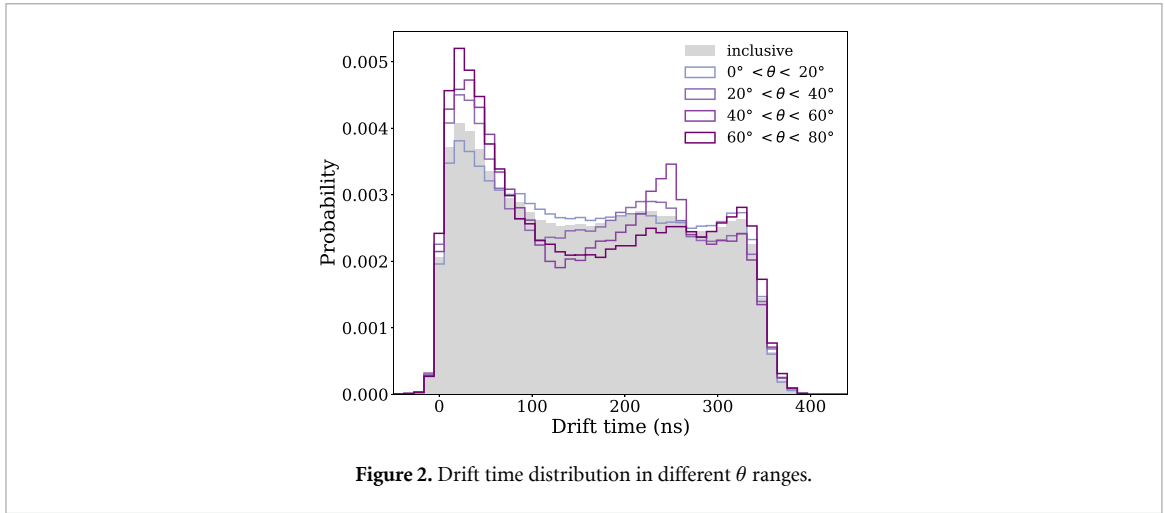


Figure 2. Drift time distribution in different  $\theta$  ranges.

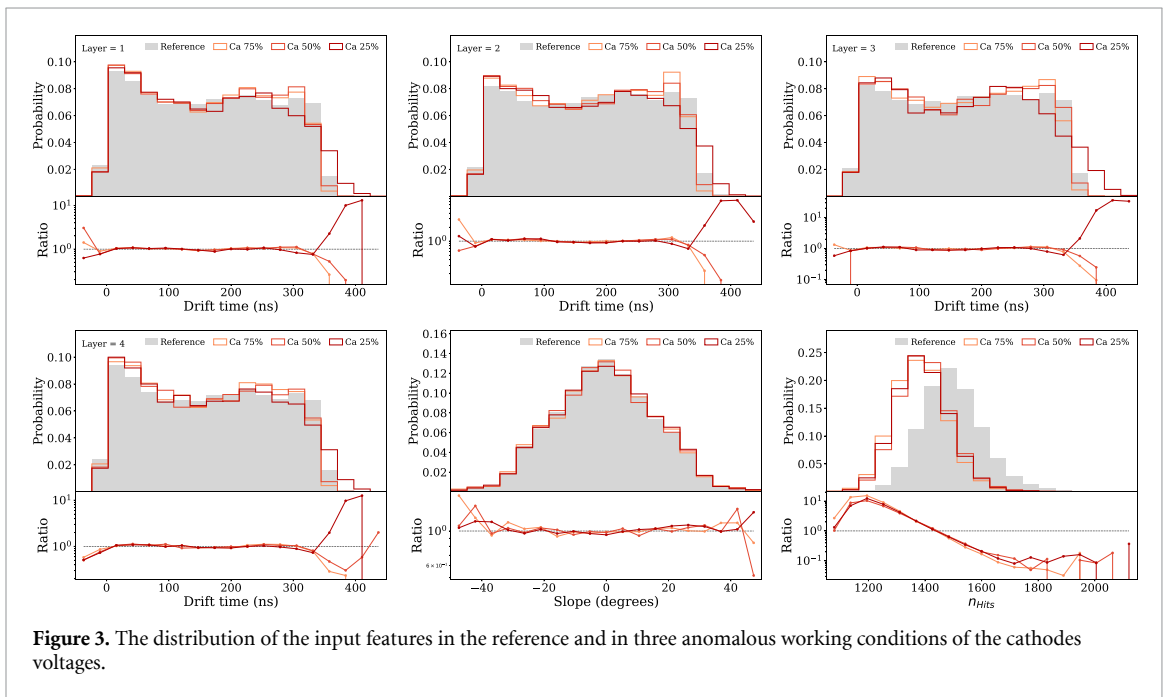
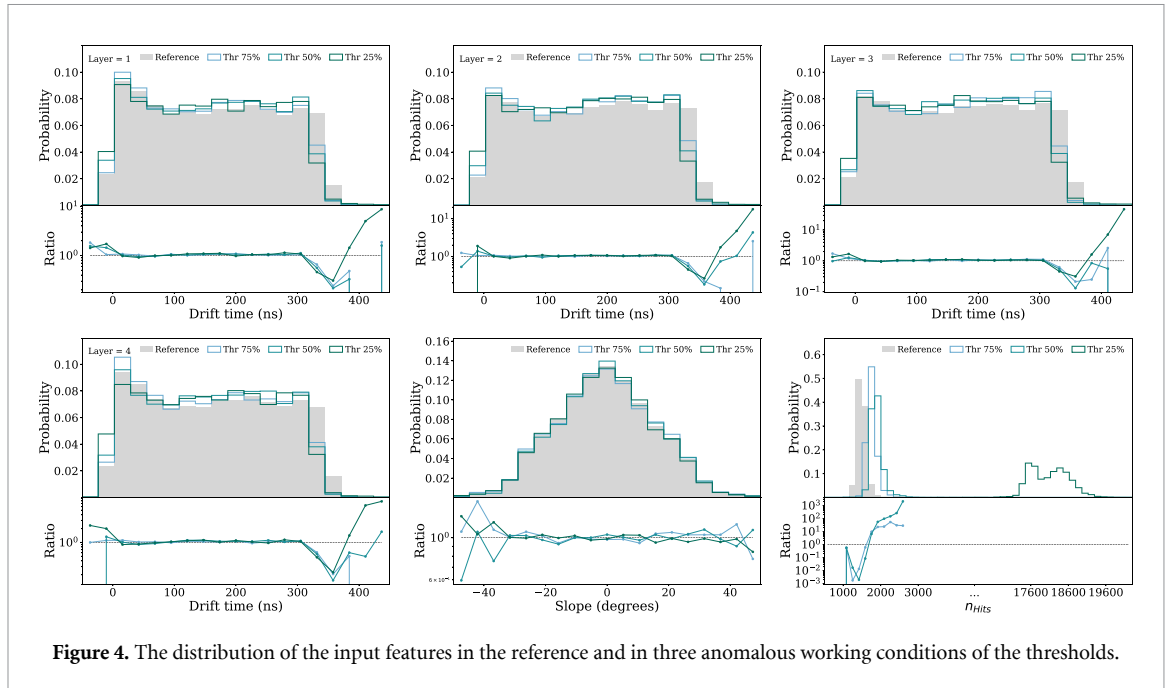


Figure 3. The distribution of the input features in the reference and in three anomalous working conditions of the cathodes voltages.

We collected the data by artificially inducing possible issues that can occur during detector operations. Specifically, we reduced the voltage of the cathodic strips to 75%, 50%, and 25% of their nominal value (−900 V, −600 V, and −300 V, respectively), and we lowered the front-end thresholds to 75%, 50%, and 25% of their nominal value (75 mV, 50 mV, and 25 mV, respectively). The former action distorts the electric field shape, whereas the latter mimics the sudden contribution of noise sources. We conducted a dedicated data acquisition campaign in these six anomalous configurations, collecting around  $10^4$  events for each configuration. We also collected around  $3 \times 10^5$  data points in the normal (or, reference) working conditions of the apparatus. The distribution of the six input features for the reference data and the data collected under the different anomalous conditions are shown in figure 3 (variation of the cathodes voltages) and figure 4 (variation of the thresholds). These data will be used to design and calibrate the DQM algorithm, as described in the following section.

### 3. Methodology

In the setup described in the previous section, we are interested in assessing the quality of individual batches of data collected by the apparatus, each of which denoted as  $\mathcal{D} = \{x_i\}_{i=1}^{N_{\mathcal{D}}}$ . Namely, we ask whether the statistical distribution of the data points in  $\mathcal{D}$  coincides or not with the one expected under *reference* working



**Figure 4.** The distribution of the input features in the reference and in three anomalous working conditions of the thresholds.

conditions,  $p(x|R)$ . We thus aim at performing what is known in statistics as a *goodness-of-fit test*. See [17] for references and a concise overview.

The reference distribution  $p(x|R)$  is not available in closed form. What is available is instead a second dataset  $\mathcal{R} = \{x_i\}_{i=1}^{N_{\mathcal{R}}}$  collected by the same apparatus when operated in the reference working conditions, such that the data in  $\mathcal{R}$  do follow the  $p(x|R)$  distribution. Our goodness-of-fit test is thus carried out by comparing the two datasets  $\mathcal{D}$  and  $\mathcal{R}$ , asking whether they are sampled from the same statistical distribution. The problem can then be formulated as a *two-sample test*, in which, however,  $\mathcal{D}$  and  $\mathcal{R}$  play asymmetric roles.

The data batch  $\mathcal{D}$  is what needs to be tested. Therefore its composition and its size,  $N_{\mathcal{D}}$ , are among the specification requirements of the DQM methodology we are developing.  $N_{\mathcal{D}} \sim 1000$  is in the ballpark of what is typically considered by DQM applications deployed at CMS.

The reference dataset  $\mathcal{R}$  is instead created within the methodology design, with mild or no limitation on its size,  $N_{\mathcal{R}}$ . A larger  $\mathcal{R}$  dataset offers a more faithful representation of the underlying reference statistical distribution and therefore a more accurate test. Furthermore, taking  $N_{\mathcal{R}}$  larger than  $N_{\mathcal{D}}$  reduces the effect of the  $\mathcal{R}$  dataset statistical fluctuation on the outcome of the test, leaving only those inherently due to the fluctuations of  $\mathcal{D}$ . This makes the outcome for a given data batch  $\mathcal{D}$  nearly independent on the specific instance of the set  $\mathcal{R}$  that is employed for the test, making the result more robust. In what follows, we will thus preferentially consider an unbalanced setup for the two datasets, with  $N_{\mathcal{R}} > N_{\mathcal{D}}$ . We will further exploit the availability of a relatively large volume of data collected under the reference working conditions for calibrating the test statistics variable and for selecting the hyperparameters, as discussed in the following.

The availability of a large set of data that are accurately labelled as having been collected under the reference detector conditions deserves further comments. These data are routinely available, in particular in high-energy physics experiments, and are in fact used for the design and calibration of regular DQM methods [5–9]. They are validated by a careful offline inspection, which typically requires human intervention. This validation process is way too demanding and slow to be employed as a DQM algorithm. The purpose of DQM is in fact to monitor the data quality online, i.e. while they are being collected. The offline validation is instead straightforwardly capable of producing labelled reference data samples that are way larger than individual data batches.

### 3.1. The NPLM method

We employ the NPLM method, which was proposed and developed by some of us [10–13] to address a similar problem in the different context of searches for new physical laws at collider experiments. The search for *new physics* is performed by comparing the measured data with a reference dataset whose statistical distribution is the one predicted by a *standard* set of physical laws that supposedly describe the experimental setup. The purpose of the comparison is not to assess the quality of the data like in DQM, but the quality of the distribution prediction and in turn to check whether the standard laws are adequate or, instead, new physical laws are needed to model the experimental setup. However, this conceptual difference does not have

practical consequences. The NPLM setup of  $\mathcal{D}$  versus  $\mathcal{R}$  data comparison is straightforwardly portable to DQM problems.

The NPLM method design is inspired by the classical approach to hypothesis testing based on the likelihood ratio [18]. A model  $f_{\mathbf{w}}(x)$  acting on the space of data  $x$ , with trainable parameters  $\mathbf{w}$ , is employed to define a set of alternatives to  $p(x|\mathcal{R})$  for the distribution of the data points in  $\mathcal{D}$ . Since the alternative hypothesis depends on  $\mathbf{w}$ , we denote it as  $H_{\mathbf{w}}$  and  $p(x|H_{\mathbf{w}})$  is the alternative distribution of  $x$ . In particular,  $f_{\mathbf{w}}(x)$  directly parametrises the logarithm of the ratio between  $p(x|H_{\mathbf{w}})$  and  $p(x|\mathcal{R})$ . The model  $f_{\mathbf{w}}(x)$  could be a neural network as in [10–12], or it could be built with kernel methods [13]. We will employ the latter option for reasons that will become clear soon. The model is trained by adjusting its parameters to best accommodate the observed data. Consequently, the trained parameters  $\hat{\mathbf{w}}$  define the best-fit hypothesis  $H_{\hat{\mathbf{w}}}$ . Following [18], the test statistic variable to be employed for the assessment of the quality of the data  $\mathcal{D}$  is <sup>11</sup>

$$t_{\hat{\mathbf{w}}}(\mathcal{D}) = 2 \sum_{x \in \mathcal{D}} \log \frac{p(x|H_{\hat{\mathbf{w}}})}{p(x|\mathcal{R})} = 2 \sum_{x \in \mathcal{D}} f_{\hat{\mathbf{w}}}(x). \quad (2)$$

In order to train the model we exploit a classical result of statistical learning: a continuous-output classifier trained to tell apart two datasets approximates—possibly up to a given monotonic transformation—the log ratio between the probability distribution of the two training sets. This property is proven explicitly in e.g. [10, 13] for the weighted logistic loss

$$\ell(y, f_{\mathbf{w}}(x)) = (1 - y)(1 + N_1/N_0) \log \left( 1 + e^{f_{\mathbf{w}}(x)} \right) + y(1 + N_0/N_1) \log \left( 1 + e^{-f_{\mathbf{w}}(x)} \right). \quad (3)$$

By assigning label  $y = 0$  to the data in  $\mathcal{R}$ , and  $y = 1$  to those in  $\mathcal{D}$ , the model  $f_{\hat{\mathbf{w}}}(x)$  trained with the loss in equation (3) approaches the logarithm of  $p(x|H_{\hat{\mathbf{w}}})/p(x|\mathcal{R})$  as it was needed in equation (2). The weight factors in equation (3), which depend on  $N_1/N_0 = N_{\mathcal{D}}/N_{\mathcal{R}}$ , are included because the two training datasets are unbalanced as previously explained.

A direct application of the classical theory of hypothesis testing [18] would actually suggest to employ a different loss function. In fact, the best-fit parameters  $\hat{\mathbf{w}}$  to be used in the definition of the test statistic (2) should be those that maximise the likelihood function. Minimising the logistic loss produces instead an estimate of the best-fit parameters that is different, *a priori*, from the maximum likelihood estimate. This can be remedied by employing a special loss function called ‘maximum likelihood loss, whose minimisation is equivalent to maximising the likelihood [10]. The maximum likelihood loss is not used in the kernel-based implementation of NPLM [13] and the logistic loss (3) is preferred for practical reasons. No strong performance degradation has been observed using the logistic loss in place of the maximum likelihood loss in the tests of the NPLM method performed so far.

Using the elements above, the design of the NPLM method for DQM works as follows. We first pick up a model for  $f_{\mathbf{w}}(x)$  and select its hyperparameters. The hyperparameters selection strategy is described in the next section for the kernel-based implementation of NPLM. Next, we need to calibrate the test statistics variable (2) in order to be able to associate its value  $t(\mathcal{D})$  to a probability  $p[t(\mathcal{D})]$ , the *p-value*. This probability will be the output of the DQM algorithm. Based on its value, the analyser will eventually judge the quality of each data batch  $\mathcal{D}$ . For instance, the analyser might define a probability threshold, below which the data batch is discarded or set apart for further analyses. Above the threshold the batch could be retained as a good batch.

It should be noted that the selected hyperparameters and the *p-value* do depend on the detailed setup of the DQM problem under consideration. For instance, different hyperparameters will be used in section 4 for the setup with five input features and data batch size  $N_{\mathcal{D}} = 1000$  than in the case of six features and  $N_{\mathcal{D}} = 500$ . The *p-value* calibration function  $p[t]$  will be also different. However, once these elements are made available for a given setup, they can be used to evaluate the quality of all the  $\mathcal{D}$  batches in that setup. The only operation that the DQM algorithm has to perform at run-time is one single training of  $\mathcal{D}$  against  $\mathcal{R}$ , out of which  $t(\mathcal{D})$  is obtained and in turn  $p[t(\mathcal{D})]$ .

Calibration is performed as follows. The test statistics (2) is preferentially large and positive if the best-fit alternative distribution  $p(x|H_{\hat{\mathbf{w}}})$  accommodates the data better than the reference distribution  $p(x|\mathcal{R})$  does, signalling that the data batch is likely not thrown from  $p(x|\mathcal{R})$ . Large  $t(\mathcal{D})$  should thus correspond to a small probability. The precise correspondence is established by comparison with the typical values that  $t$  attains

<sup>11</sup> Unlike in NPLM applications to new physics searches, the total number of data points in  $\mathcal{D}$  is not a random variable, but is fixed to the data batch size. The regular likelihood for i.i.d. data is thus employed rather than the extended likelihood. Correspondingly, the test statistic contains one term less than in [10–13].

when the data batch is instead a good batch. We thus compute the distribution,  $p(t|R)$ , that the  $t$  variable possesses when the data follow the reference statistical distribution and the  $p$ -value is defined as

$$p[t] = \int_t^{\infty} dt' p(t'|R). \quad (4)$$

The physical meaning of  $p[t(\mathcal{D})]$  is the probability that a good data batch gives a value of  $t$  that is more unlikely (i.e. larger) than the value  $t(\mathcal{D})$  produced by the batch  $\mathcal{D}$ . If a threshold is set on  $p$ , this threshold measures the frequency at which good data batches are not recognised as such by the algorithm.

The  $p(t|R)$  distribution is straightforwardly estimated empirically, thanks to the availability of reference-distributed labelled data points. We create several artificial data batches—called *Toy* datasets—of the same size  $N_{\mathcal{D}}$  as the true batches. We run the training and compute  $t$  on each of them. Each Toy dataset should be statistically independent, and independent from the reference dataset  $\mathcal{R}$  that is employed for training. A very large sample of reference-distributed data is thus used in order to produce both the Toy batches and the reference dataset. By histogramming the values of  $t$  computed on the Toys we could easily obtain an estimate of  $p(t|R)$  and hence of  $p[t]$ . A different procedure is adopted here, exploiting the empirical observation [13] that  $p(t|R)$  is well approximated by a chi-squared ( $\chi^2$ ) distribution. The number of degrees of freedom of the  $\chi^2$  depends on the setup but can be determined by fitting to the empirical distribution of the  $t$  values computed on the Toys. The survival function (one minus the cumulative) of the corresponding  $\chi^2$  distribution will be used as an estimate of  $p[t]$ . It should be noted that by proceeding in this way we will be formally able to compute very small  $p$ -values that correspond to highly-discrepant data batches with very large  $t(\mathcal{D})$ . However, the agreement of  $p(t|R)$  with the  $\chi^2$  cannot be verified in the high  $t$  region, which the Toys do not populate, and there is no theoretical reason to expect that this agreement will persist in that region. Our quantification of the  $p$ -value is thus only accurate in the region that the Toys statistically populate. For instance, if 300 Toys are thrown, only  $p$ -values larger than around  $1/300$  are accurately computed. If  $t(\mathcal{D})$  falls in a region where our determination of  $p$  is much smaller than that, ours should be regarded as a reasonable estimate that is particularly useful to compare the level of discrepancy of different batches, but it cannot be directly validated. However, in those cases we will be able to ensure that  $p[t(\mathcal{D})] \lesssim 1/300$  by directly comparing with the  $t$  values on the Toys.

Another feature of the NPLM approach is the possibility of exploiting the function  $f_{\tilde{w}}$  learned during the training task to characterise anomalous batches of data. The function  $f_{\tilde{w}}$  represents the log-ratio between  $p(x|H_{\tilde{w}})$  and  $p(x|R)$  and, hence, can be used to deform and adapt the reference distribution to the data by reweighting, according to the following expression

$$p(x|H_{\tilde{w}}) = e^{f_{\tilde{w}}(x)} p(x|R). \quad (5)$$

The function  $\exp(f_{\tilde{w}}(x))$  will be close to one if the data are well-described by the reference distribution, while it will depart from it otherwise. One should therefore be able to gain additional information about the anomalous batch by inspecting this quantity as a function of the input variables, or any combination of them, even when not explicitly provided as an input feature for the training. Having access to this kind of information is a valuable element in the context of the search for new physics [10, 11, 13], since the physics-motivated variables that one might want to inspect to explain a potential anomalous score could be some type of nontrivial combination of the input features with a clear physical meaning, such as the invariant mass of a many-body final state. For DQM applications, this analysis is less relevant since a direct visual inspection of the ratio between the binned data and reference marginal distributions is already quite informative and the user might not be interested in exploring specific high-level features in the first place. On the other hand, one can still exploit the possibility of reconstructing the data distribution using  $f_{\tilde{w}}$  as a debugging tool, namely to check whether the learning model correctly recognises if the data deviates from the reference and how.

Moreover, somewhat aside from the main goal of the present article, the output of the NPLM-DQM application could be exploited to study those data batches that display significant deviations from the reference and, depending on the characteristic of the departures, to classifying them into different anomalous categories. Further investigations on a possible extension of the application in this respect are left for future work.

### 3.2. Falkon-based NPLM

Applying NPLM to the DQM problem is simpler than using it for new physics searches. For new physics searches one needs to worry about imperfections in the reference data that stem from the mismodelling of the reference distribution based on the underlying standard physical laws. Including these effects in NPLM is



possible but requires dedicated work and domain-specific expertise [12]. Mismodelling is not a concern in DQM problems because no modelling is required at all: the reference-distributed data are merely collected from the same experimental apparatus and not simulated. NPLM algorithms for DQM can thus be designed more easily and systematically without the need for extremely specialised domain knowledge.

DQM applications are, however, much more computationally demanding than new physics searches. For new physics searches there is typically only one dataset  $\mathcal{D}$  to be analysed. For DQM, a large flow of data batches needs to be analysed online. We will see in section 5 that, for instance, order 10 s are needed to the CMS muon system to collect one data batch. Our DQM algorithm must respond on a competitive timescale in order to be applicable to that problem. The relevant operation time is the one needed for a single training, as previously explained. The original implementation of NPLM based on neural networks is vastly incompatible with this requirement. On the other hand, the one based on kernel methods is much faster to train on problems of comparable scale [13]. It could thus match the specification requirements for applications to LHC detectors.

The performance of the kernel-based version of NPLM stems from those of the Falkon [14] library, the core algorithm powering our implementation. A sketch of the basic theoretical and algorithmic ideas implemented in Falkon, developed in [19–21], are reported below.

With kernel methods, one learns functions of the following form

$$f_{\mathbf{w}}(x) = \sum_{i=1}^N w_i k_{\sigma}(x, x_i), \quad (6)$$

with  $N = N_0 + N_1$  the total size of the training dataset. Here  $k_{\sigma}(x, x_i)$  is the kernel function and  $\sigma$  some hyperparameter. We consider the Gaussian kernel

$$k_{\sigma}(x, x') = e^{-\|x-x'\|^2/2\sigma^2}, \quad (7)$$

so that  $f_{\mathbf{w}}$  is a linear combination of Gaussians of fixed width  $\sigma$ , centred at the training data points. The optimisation of the model parameters  $\mathbf{w}$  is achieved by minimising the empirical risk  $\hat{L}(f_{\mathbf{w}})$ , plus a regularisation term

$$\hat{L}_{\lambda}(f_{\mathbf{w}}) = \hat{L}(f_{\mathbf{w}}) + \lambda R(f_{\mathbf{w}}). \quad (8)$$

The empirical risk in our case is the one associated with the logistic loss (3)

$$\hat{L}(f_{\mathbf{w}}) = \sum_{i=1}^N \ell(y_i, f_{\mathbf{w}}(x_i)). \quad (9)$$

The regularisation term is given by

$$R(f_{\mathbf{w}}) = \sum_{ij} w_i w_j k_{\sigma}(x_i, x_j). \quad (10)$$

Its relative importance in the optimisation target (8) is controlled by the hyperparameter  $\lambda$ .

Kernel methods are non-parametric approaches, in the sense that the number of parameters  $\mathbf{w}$  in equation (6) increases automatically with the total number of data points. Gaussian kernel methods are universal, meaning that they can recover any continuous function in the large sample limit [22, 23]. However, optimising the function in equation (6), with the target in equation (8), requires handling an  $N \times N$  matrix—the *kernel matrix*—with entries  $k_{\sigma}(x_i, x_j)$ . The computational complexity of the optimisation thus scales cubically in time and quadratically space with respect to the number of training points  $N$  [14, 19]. These costs prevent the application to large-scale settings, and some approximation is needed.

Within the Falkon library, the problem of minimising equation (8) is formulated in terms of an approximate Newton method (see algorithm 2 of [14]). The algorithm is based on the Nyström approximation, which is used twice. First, to reduce the size of the problem, by considering solutions of the form

$$f_{\mathbf{w}}(x) = \sum_{i=1}^M w_i k_{\sigma}(x, \tilde{x}_i), \quad (11)$$

where  $\{\tilde{x}_1, \dots, \tilde{x}_M\} \subset \{x_1, \dots, x_N\}$  are called Nyström centres and are sampled uniformly at random from the input data. The number of centres  $M \leq N$  is a hyperparameter to be chosen. Then, Nyström approximation is again used to derive an approximate Hessian matrix

$$\tilde{\mathbf{H}} = \frac{1}{M} T \tilde{D} T^\top + \lambda \mathbf{I}. \quad (12)$$

Here,  $T$  is such that  $T^\top T = \tilde{K}$  (Cholesky decomposition), with  $\tilde{K} \in \mathbb{R}^{M \times M}$  the kernel matrix subsampled with respect to both rows and columns.  $\tilde{D} \in \mathbb{R}^{M \times M}$  is a diagonal matrix s.t. the  $i$ th element is the second derivative of the loss  $\ell''(y_i, f_w(x_i))$ , with respect to its first variable. Equation (12) is then used as a preconditioner to perform conjugate gradient descent. With this strategy, the overall computational cost to achieve optimal statistical bounds is  $\mathcal{O}(N)$  in memory and, of particular importance for our scope,  $\mathcal{O}(N\sqrt{N}\log N)$  in time. It is known in the literature [24, 25] that the effect of the projection in the subspace determined by the centres is a form of regularisation. On the other hand, the stochasticity of the projection can potentially lead to a subspace that does not guarantee stability. From this point of view, the inclusion of a further explicit penalty term can be used to ensure stability as needed. Indeed, the regularisation level is determined by both the penalty parameter and the number of centres. These ideas are formalised and made quantitative in [24]. The reader can find more details about the Falkon algorithm, including comparisons with alternative approaches, in [14].

### 3.2.1 Hyperparameters selection

The selection of the three Falkon hyperparameters  $M$ ,  $\sigma$  and  $\lambda$  follows the prescriptions of [13], with one minor modification described below. The hyperparameters selection employs data collected under the reference working condition, and proceeds as follows.

*The number of centres  $M$*  controls the expressive power of the model and therefore it should be as large as possible not to compromise the sensitivity to anomalous distributions with intricate shapes. It must also be at least as large as  $\sqrt{N}$  in order to achieve statistically optimal bounds of the training convergence. At the same time, training is faster if  $M$  is smaller. The experiments performed in [13] show that any value of  $M$  above around the data batch size  $N_{\mathcal{D}}$  does not compromise sensitivity.

*The Gaussian width  $\sigma$*  is selected as the 90th percentile of the pairwise distance between reference-distributed data points. Notice that the model (11) acts on an input vector  $x$  whose input features are standardised to have zero mean and unit variance on reference-distributed data. The same standardisation is applied before computing the distances.

*The regularisation parameter  $\lambda$*  is kept as small as possible while keeping training stable, i.e. avoiding large training times or non-numerical outputs. A number of reference-distributed Toy data batches is employed for this study, each trained against the reference sample  $\mathcal{R}$ . Some of the experiments performed in this paper employ quite smaller data batches (e.g.  $N_{\mathcal{D}} = 250$ ) than those considered in [13]. In these new conditions we observe that the compatibility of the test statistic distribution with a  $\chi^2$  (see the end of section 3.1) is violated for very small  $\lambda$ . In these cases, we raise  $\lambda$  until when the agreement with the  $\chi^2$  is restored.

The hyperparameters selected with the above criteria, in the different setups for DQM considered in this paper, are reported in table 1. It should be emphasised that the hyperparameters selection problem for NPLM is rather different than for regular applications of Falkon or other types of classifiers. The hyperparameters for regular classifiers can be optimised based on the performances in the specific classification task under examination. NPLM aims instead at goodness-of-fit, namely at attaining good sensitivity to a wide class of anomalous data distributions that are not known or specified *a priori*. Hence, a prior reasonable choice of the hyperparameters must be performed and cannot be re-optimised *a posteriori*. In particular, no re-optimisation can be or has been performed to enhance the sensitivity to the specific types of anomalies that are considered in this paper to demonstrate the sensitivity of the method.

### 3.3. Alternative approaches

Goodness-of-fit and two-sample test problems are of interest in several domains of science. Many approaches exist, and developing new strategies is an active area of research. One heuristic reason to choose NPLM for DQM, among the many different options, is that it has been developed in the challenging context of new physics searches. Prior experimental and theoretical knowledge suggests that new physics is elusive. The target for new physics searches is thus to spot out minor departures of the actual data from the reference distribution. These departures could emerge either as small corrections to the distribution shape or as relatively large corrections like sharp peaks, which however only account for a very small fraction of the experimental data. Detecting such small effects requires precisely comparing the reference distribution with

**Table 1.** NPLM algorithm parameters configuration for the five-dimensional and six-dimensional experiments considered in this work. The numbers of degrees of freedom of the  $\chi^2$  that best approximates  $p(t|R)$  is reported in the last column.

	$N_{\mathcal{R}}$	$N_{\mathcal{D}}$	$M$	$\sigma$	$\lambda$	dof
5D	2000	250	2000	4.5	$10^{-6}$	40
		500			$10^{-7}$	83
		1000			$10^{-8}$	171
6D	2000	250	2000	4.8	$10^{-6}$	58
		500				78
		1000				109

large datasets, which NPLM is designed to perform. Using NPLM for DQM could thus enable a more accurate monitoring of the data offering sensitivity to more subtle failures of the apparatus. The number of input features in the data that are typically relevant for new physics searches ranges from few to tens, which is an adequate number also for the monitoring of individual detectors and detector systems fully exploiting the correlations among the variables. For comparison, methods to assess the quality of generated images target instead order thousand-dimensional input data. They could be less performant for DQM as they are designed to address a radically different problem.

These heuristic considerations suggest that NPLM is a reasonable starting point for the development of novel DQM algorithms based on advanced multivariate goodness-of-fit or two-sample test methods, which we advocate in this paper. On the other hand, no comprehensive comparative study of the NPLM performances is currently available. Such comparison is beyond the scope of this paper. However, the DQM problems and datasets we study will be useful benchmarks for future work in this direction.

Work has initiated [26, 27] to compare NPLM with a certain class of methods, called ‘classifier-based’ methods. The classifier-based approaches [28] are all those that entail training a classifier to tell apart  $\mathcal{D}$  from  $\mathcal{R}$  and using the trained classifier to construct a test statistic for the hypothesis test. A simple implementation [29] employs the classification accuracy as test statistics. Following the standard pipeline for classifiers, the model is trained on part of the  $\mathcal{D}$  and  $\mathcal{R}$  datasets (the training set), while the accuracy is evaluated on the remaining data (the test set). The idea is that while the accuracy will be poor (around random guess) if  $\mathcal{D}$  and  $\mathcal{R}$  follow the same distribution, it will be higher if their distributions differ.

NPLM is technically a classifier-based method. Its major peculiarities are the choice of the likelihood ratio test statistic in equation (2) and the fact that the entire datasets are employed both for training and for the evaluation of the test statistics. None of these choices is motivated from the viewpoint of the theory of classification, while they are both natural or in fact required from the perspective of the theory of hypothesis testing that underlies the NPLM approach. Performance studies in [27] show that these choices are beneficial for the sensitivity. These results partly contradict [26], which however employs different classification models, different criteria for hyperparameters selection and uses permutation tests for the estimate of the sensitivity rather than computing it empirically as in NPLM. These differences are evidently responsible for the different findings and more work is needed for a conclusive assessment.

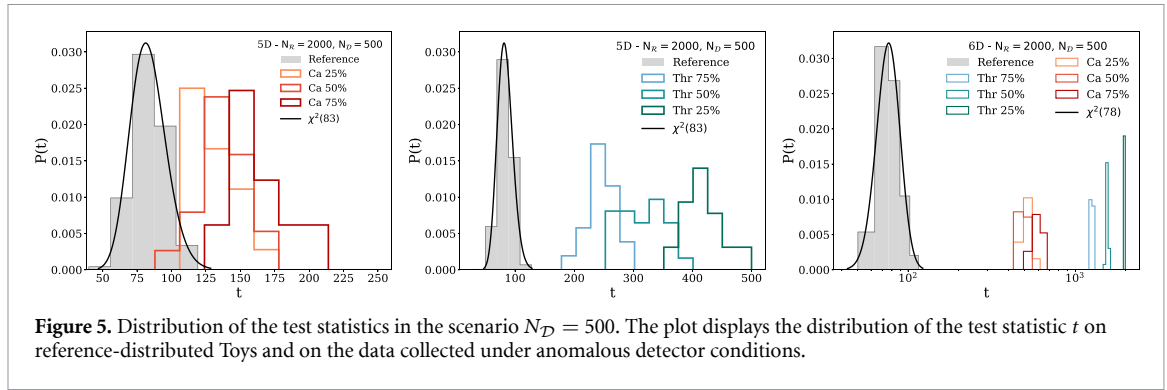
## 4. Results

In this section, we present the application of the NPLM strategy for DQM to the DT chambers data described in section 2<sup>12</sup>. We will consider monitoring data batches of variable size  $N_{\mathcal{D}} = 250, 500$  and  $1000$ , by employing a reference dataset of fixed size  $N_{\mathcal{R}} = 2000$ .

The input data consists of six features: the four drift times, the muon angle and the number of hits. As shown in the bottom-right plots of figures 3 and 4, the number of hits,  $n_{\text{Hits}}$ , is highly discriminant for the anomalies we considered in our study, and in particular for the ones affecting the thresholds (the lower the threshold, the higher the noise). At the LHC, however, that quantity also depends on the luminosity delivered to the experiment, which could vary greatly even during a single run. Not being necessarily a proxy to a detector issue, it is worth considering also the case where only the other five variables are provided to the algorithm; as an additional benefit, this will allow assessing the ability of the NPLM DQM approach to exploit correlations between variables and detect anomalies even when their effect is unexpected and not straightforwardly evident.

The left and middle panels of figure 5 show the test statistics distribution in the five-dimensional (5D) problem, for data batches size  $N_{\mathcal{D}} = 500$ . The grey histograms display the distribution of  $t$  in the reference

<sup>12</sup> Code available at <https://github.com/FalkonHEP>.



**Figure 5.** Distribution of the test statistics in the scenario  $N_D = 500$ . The plot displays the distribution of the test statistic  $t$  on reference-distributed Toys and on the data collected under anomalous detector conditions.

**Table 2.** Median  $p$ -values for different anomalies and data batches size. Five input features are considered, excluding  $n_{\text{hits}}$ .

Anomaly	$N_D = 250$	$N_D = 500$	$N_D = 1000$
Cathode 75%	0.0034	$1.1 \times 10^{-6}$	$<10^{-7}$
Cathode 50%	0.029	$3.4 \times 10^{-4}$	$<10^{-7}$
Cathode 25%	0.14	0.0019	$<10^{-7}$
Threshold 75%	$2.8 \times 10^{-7}$	$<10^{-7}$	$<10^{-7}$
Threshold 50%	$<10^{-7}$	$<10^{-7}$	$<10^{-7}$
Threshold 25%	$<10^{-7}$	$<10^{-7}$	$<10^{-7}$

**Table 3.** Median  $p$ -values in the setup  $N_D = 500$ .

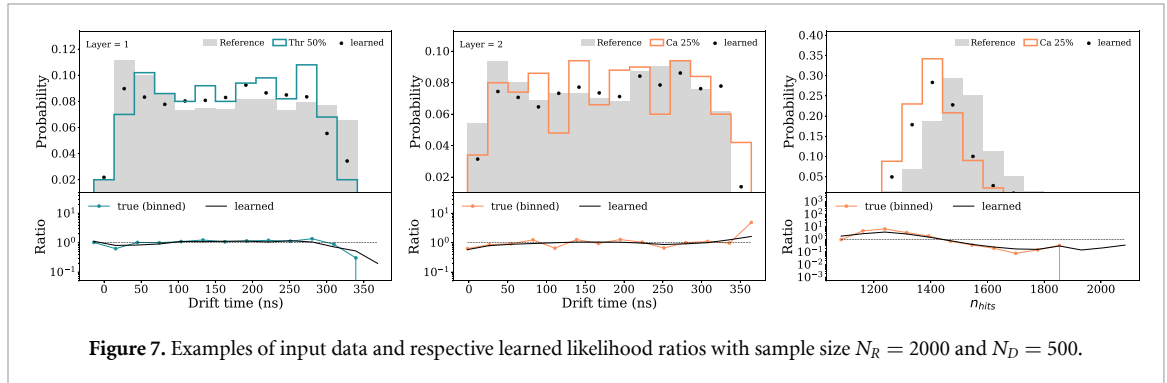
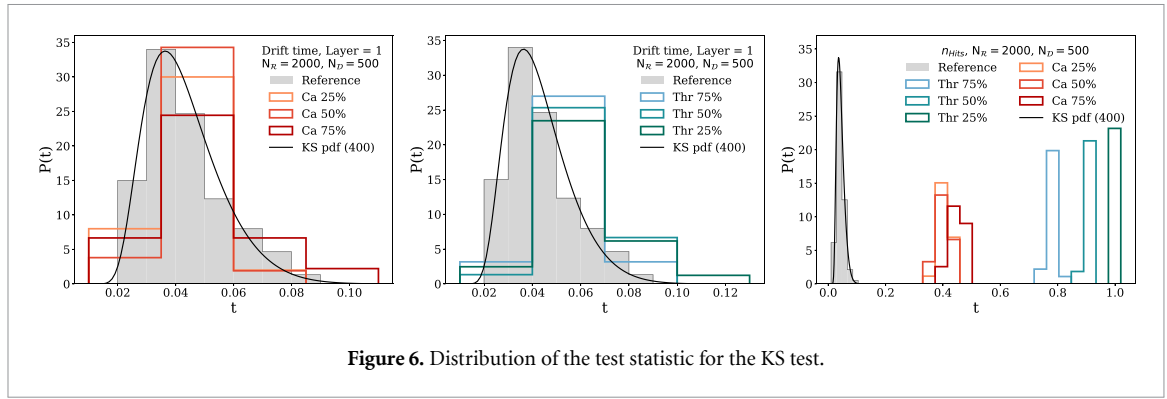
Anomaly	NPLM (5D)	KS ( $t_1$ )	KS ( $t_2$ )	KS ( $t_3$ )	KS ( $t_4$ )	KS ( $\phi$ )
Cathode 75%	$1.1 \times 10^{-6}$	0.50	0.41	0.43	0.40	0.42
Cathode 50%	$3.4 \times 10^{-4}$	0.47	0.27	0.47	0.37	0.41
Cathode 25%	0.0019	0.45	0.44	0.13	0.45	0.50
Threshold 75%	$<10^{-7}$	0.23	0.14	0.16	0.14	0.48
Threshold 50%	$<10^{-7}$	0.09	0.10	0.06	0.17	0.42
Threshold 25%	$<10^{-7}$	0.11	0.07	0.04	0.11	0.66

working conditions,  $P(t|R)$ . This is obtained empirically by processing reference-distributed Toy data batches, and fitted to a  $\chi^2$  distribution as explained in section 3.1. The different distributions of the test statistic associated with the anomalous batches shown in the coloured histograms are very well separated from the reference distribution, meaning that anomalous data are very likely to be identified as such by the algorithm. This is quantified by the median  $p$ -value of the anomalous batches, reported in the central column of table 2. The table also reports the median  $p$ -value for larger ( $N_D = 1000$ ) and smaller ( $N_D = 250$ ) batches. The sensitivity to the anomaly increases with  $N_D$ , as expected.

For a comparative assessment of the performance, we computed a Kolmogorov–Smirnov (KS) test on each individual feature for the same data used to train the NPLM model. The KS median  $p$ -values are reported in table 3 and compared with the ones obtained with the 5D NPLM test. We see that individual variables have a very limited power to discriminate the anomalous batches. The NPLM method instead is sensitive to correlated discrepancies in the different distributions and discriminates the anomalies effectively. For illustrative purposes, we show in the left and middle panels of figure 6 the distribution of the 1D KS statistic computed on the drift time of the first layer ( $t_1$ ) for reference and anomalous batches. By comparison with figure 5, it is easy to recognise the advantage of the NPLM strategy.

We now turn to the study of the complete six-dimensional problem, including the variable  $n_{\text{hits}}$ . The reference and anomalous test statistic distributions are shown on the right panel of figure 5. By comparing with the other panels of the figure we can appreciate the tremendous discriminating power of the  $n_{\text{hits}}$  variable: including  $n_{\text{hits}}$  all the anomalies can be detected with very high significance. Therefore, using this variable alone for the NPLM DQM test, or running a regular KS test (as shown in the right panel of figure 6), is sufficient to identify the anomalies, as previously mentioned.

We conclude this section by showing some examples of the data marginal distribution reconstructed by the model. The three plots reported in figure 7 are produced by reweighting each event of the reference sample used for the training by an exponential factor  $e^{\beta w(x)}$ , as explained in equation (5); both the reweighted reference and the data samples are binned, and their ratio with respect to the original reference sample is



shown in the bottom panels. By comparing the data-versus-reference ratio (labelled as ‘true’) with the reconstructed one (‘learned’) we can appreciate the correctness of the model in understanding the nature of the anomaly and, hence, trust the results of the ML task.

All the numerical experiments presented in this paper have been performed on a single machine equipped with a NVIDIA Titan Xp GPU with 12 GB of VRAM. We tested the performances of the algorithm in terms of execution time; the training time for a single 5D classification task is approximately 0.5 s, with no significant dependency on the nature of the data and the size of the sample.

## 5. Conclusions and outlook

We presented the test of a powerful ML-based algorithm, NPLM, as a tool to monitor the quality of the data originated by a typical detector used for measuring particles at high energy colliders. NPLM compares collected measurements with a reference dataset describing the standard detector readout, performing a multidimensional likelihood-ratio hypothesis test.

The study demonstrated the capability of the algorithm to detect anomalous detector conditions, with a much greater discriminating power than simpler traditional methods, like KS test. Traditional strategies for DQM are specifically designed for individual detectors and monitoring needs. We will thus be in the position of performing a full-fledged sensitivity comparison only after implementing our strategy in a concrete realistic setup with an established standard alternative. Nevertheless, the strong observed advantage in performances suggests that major sensitivity improvements will be attained thanks to the deployment of a sophisticated goodness-of-fit method such as NPLM.

Although conducted on simplified experimental conditions, the test presents figures appropriate for a typical monitoring system of a detector operating at the LHC; in particular, the number of channels and the size of the datasets are of the same order of magnitude as the corresponding CMS DQM application. The amount of data we consider for each batch can be gathered much more quickly at the LHC than in a cosmic stand like the one used here, anyhow the rate at which possible issues should be detected is not larger than one in a minute<sup>13</sup>; the time requested by NPLM to run—less than a second—makes the algorithms suitable to be executed online.

<sup>13</sup> Failures potentially leading to catastrophic consequences that requires a much prompt reaction are typically controlled by hardware interlock systems.

Our results open the door to the systematic deployment of NPLM for DQM, in particular to CMS detector subsystems in real LHC data-taking conditions, potentially advancing the standard practice in two main aspects. First, by boosting the sensitivity to detector malfunctioning exploiting a more refined multivariate statistical methodology and employing low-level variables reducing the bias introduced by further manipulations of the data. This will ultimately improve the quality of the data used for physics analyses. Second, by introducing a universal problem-independent approach for a more efficient and nearly automated design of DQM strategies.









## Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://zenodo.org/record/7128223>.

## Acknowledgments

L R, M L and M R acknowledge the financial support of the European Research Council (Grant SLING 819789). L R acknowledges the financial support of the AFOSR Projects FA9550-18-1-7009, FA9550-17-1-0390 and BAA-AFRL-AFOSR-2016-0007 (European Office of Aerospace Research and Development), the EU H2020-MSCA-RISE Project NoMADS—DLV-777826, and the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216. G G is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 772369). A W is supported by the Grant PID2020-115845GB-I00/AEI/10.13039/501100011033.

## ORCID iDs

Gaia Grosso  <https://orcid.org/0000-0002-8303-3291>  
Nicolò Lai  <https://orcid.org/0000-0001-9973-6509>  
Marco Letizia  <https://orcid.org/0000-0001-9641-4352>  
Jacopo Pazzini  <https://orcid.org/0000-0002-1118-6205>  
Marco Rando  <https://orcid.org/0009-0008-3839-1429>  
Lorenzo Rosasco  <https://orcid.org/0000-0003-3098-383X>  
Andrea Wulzer  <https://orcid.org/0000-0002-4523-1940>  
Marco Zanetti  <https://orcid.org/0000-0003-4281-4582>

## References

- [1] Pol A A, Cerminara G, Germain C and Pierini M 2022 Data quality monitoring anomaly detection *Artificial Intelligence for High Energy Physics* (World Scientific) ([https://doi.org/10.1142/9789811234033\\_0005](https://doi.org/10.1142/9789811234033_0005))
- [2] Pol A A, Cerminara G, Germain C, Pierini M and Seth A 2019 Detector monitoring with artificial neural networks at the CMS experiment at the CERN Large Hadron Collider *Comput. Softw. Big Sci.* **3** 3
- [3] Azzolini V, Andrews M, Cerminara G, Dev N, Jessop C, Marinelli N, Mudholkar T, Pierini M, Pol A and Vlimant J-R (CMS Collaboration) 2019 Improving data quality monitoring via a partnership of technologies and resources between the CMS experiment at CERN and industry *EPJ Web Conf.* **214** 01007
- [4] Adinolfi M, Archilli F, Baldini W, Baranov A, Derkach D, Panin A, Pearce A and Ustyuzhanin A 2017 LHCB data quality monitoring *J. Phys.: Conf. Ser.* **898** 092027
- [5] Rovere M (CMS Collaboration) 2015 Rovere and CMS Collaboration the data quality monitoring software for the CMS experiment at the LHC *J. Phys.: Conf. Ser.* **664** 072039
- [6] Azzolini V et al 2019 The data quality monitoring software for the CMS experiment at the LHC: past, present and future *EPJ Web Conf.* **214** 02003
- [7] Marantis A (ATLAS Collaboration) 2019 The ATLAS fast tracker—architecture, status and high-level data quality monitoring framework *Universe* **5** 32
- [8] Kaur S (ATLAS Muon Spectrometer System Collaboration) 2022 Online data monitoring of the ATLAS Muon system and commissioning of the new small wheel (NSW) data quality system *PoS ICHEP2022* 1013
- [9] Aad G et al (ATLAS Collaboration) 2020 ATLAS data quality operations and performance for 2015–2018 data-taking *J. Instrum.* **15** 04003
- [10] D'Agnolo R T and Wulzer A 2019 Learning new physics from a machine *Phys. Rev. D* **99** 015014
- [11] D'Agnolo R T, Grosso G, Pierini M, Wulzer A and Zanetti M 2021 Learning multivariate new physics *Eur. Phys. J. C* **81** 89
- [12] D'Agnolo R T, Grosso G, Pierini M, Wulzer A and Zanetti M 2022 Learning new physics from an imperfect machine *Eur. Phys. J. C* **82** 275
- [13] Letizia M, Losapio G, Rando M, Grosso G, Wulzer A, Pierini M, Zanetti M and Rosasco L 2022 Learning new physics efficiently with nonparametric methods *Eur. Phys. J. C* **82** 879
- [14] Meanti G, Carratino L, Rosasco L and Rudi A 2020 Kernel methods through the roof: handling billions of points efficiently *Advances in Neural Information Processing Systems* vol 33 (Curran Associates, Inc) pp 14410–22 (available at: <https://proceedings.neurips.cc/paper/2020/hash/a59afb1b7d82ec353921a55c579ee26d-Abstract.html>)

- [15] Chatrchyan S et al (CMS Collaboration) 2008 The CMS experiment at the CERN LHC *J. Instrum.* **3** S08004
- [16] Migliorini M, Pazzini J, Triossi A, Zanetti M and Zucchetta A 2022 Muon trigger with fast neural networks on FPGA, a demonstrator *J. Phys.: Conf. Ser.* **2374** 012099
- [17] Cousins R D 2016 On goodness-of-fit tests (available at: [www.physics.ucla.edu/~cousins/stats/ongoodness6march2016.pdf](http://www.physics.ucla.edu/~cousins/stats/ongoodness6march2016.pdf))
- [18] Neyman J and Pearson E S 1933 On the problem of the most efficient tests of statistical hypotheses *Phil. Trans. R. Soc. A* **231** 289
- [19] Rudi A, Carratino L and Rosasco L 2017 Falkon: an optimal large scale kernel method *Advances in Neural Information Processing Systems* vol 30 (Curran Associates, Inc) (available at: [https://papers.nips.cc/paper\\_files/paper/2017/hash/05546b0e38ab9175cd905eebcc6ebb76-Abstract.html](https://papers.nips.cc/paper_files/paper/2017/hash/05546b0e38ab9175cd905eebcc6ebb76-Abstract.html))
- [20] Marteau-Ferey U, Bach F and Rudi A 2019 Globally convergent newton methods for ill-conditioned generalized self-concordant losses *Advances in Neural Information Processing Systems* vol 32 (Curran Associates, Inc) (available at: [https://papers.nips.cc/paper\\_files/paper/2019/hash/60495b4e033e9f60b32a6607b587aadd-Abstract.html](https://papers.nips.cc/paper_files/paper/2019/hash/60495b4e033e9f60b32a6607b587aadd-Abstract.html))
- [21] Marteau-Ferey U, Ostrovskii D, Bach F and Rudi A 2019 Beyond least-squares: fast rates for regularized empirical risk minimization through self-concordance *Proc. 32nd Conf. on Learning Theory (25–28 June 2019) (Proc. of Machine Learning Research)* vol 99 (PMLR) pp 2294–340 (available at: <https://proceedings.mlr.press/v99/marteau-ferey19a.html>)
- [22] Micchelli C A, Xu Y and Zhang H 2006 Universal kernels *J. Mach. Learn. Res.* **7** 2651–67 (available at: <http://jmlr.org/papers/v7/micchelli06a.html>)
- [23] Christmann A and Steinwart I 2008 *Support Vector Machines* (Springer)
- [24] Rudi A, Camoriano R and Rosasco L 2015 Less is more: Nyström computational regularization *Advances in Neural Information Processing Systems* vol 28 (available at: [https://proceedings.neurips.cc/paper\\_files/paper/2015/file/03e0704b5690a2dee1861dc3ad3316c9-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2015/file/03e0704b5690a2dee1861dc3ad3316c9-Paper.pdf))
- [25] Manzhos S and Ihara M 2023 Rectangularization of Gaussian process regression for optimization of hyperparameters *Machine Learning with Applications* **13** 100487
- [26] Chakravarti P, Kuusela M, Lei J and Wasserman L 2021 Model-independent detection of new physics signals using interpretable semi-supervised classifier tests (arXiv:2102.07679)
- [27] Grosso G, Letizia M, Pierini M and Wulzer A 2023 Goodness of fit by Neyman-Pearson testing (arXiv:2305.14137)
- [28] Friedman J H 2003 On multivariate goodness of fit and two sample testing *eConf C030908* THD002 (available at: <https://www.slac.stanford.edu/econf/C030908/papers/THPD002.pdf>)
- [29] Lopez-Paz D and Oquab M 2017 Revisiting classifier two-sample tests *5th Int. Conf. on Learning Representations (ICLR 2017) (Toulon, France, 24–26 April 2017) (Conf. Track Proc.)* (OpenReview.net) (available at: <https://openreview.net/forum?id=SJKXfE5xx>)