Towards automation of data quality system for CERN CMS experiment

View the article online for updates and enhancements.

Related content

Towards automation of data quality system for CERN CMS experiment

M Borisyak[1,2]**, F Ratnikov**[1,2] **and D Derkach**[1,2]**, A Ustyuzhanin**[1,2]

[1] National Research University - Higher School of Economics, 20 Myasnitskaya st., Moscow 101000, Russia
[2] Yandex School of Data Analysis, 11/2, Timura Frunze st., Moscow 119021, Russia

E-mail: `mborisyak@hse.ru`

Abstract. Daily operation of a large-scale experiment is a challenging task, particularly from perspectives of routine monitoring of quality for data being taken. We describe an approach that uses Machine Learning for the automated system to monitor data quality, which is based on partial use of data qualified manually by detector experts. The system automatically classifies marginal cases: both of good an bad data, and use human expert decision to classify remaining "grey area" cases.

This study uses collision data collected by the CMS experiment at LHC in 2010. We demonstrate that proposed workflow is able to automatically process at least 20% of samples without noticeable degradation of the result.

1. Introduction
Data Quality monitoring is a crucial task for every large scale High Energy Physics experiment. The challenge driven by the huge amount of data is the considerable amount of person power required for monitoring and classification. An automated system for data quality monitoring can thus save resources needed to keep high quality of collected physics data.

In these proceedings, we use data collected by the CMS experiment [1] at LHC in CERN. Currently, CMS data quality is certified by detector experts who base their judgment on the visual inspection of a set of pre-defined distributions. Our system is tested to monitor collision data [2, 3, 4] collected by the CMS experiment, however the only component of the system dependent from the particular experiment setup is the feature preprocesing step.

2. Automated Data Quality System
A detector measures physical properties of proton collisions products. When a subdetector exposes an abnormal behavior (e.g. part of subdetector becomes unresponsive), it is reflected in measured or reconstructed properties. Data Quality managers rely on some set of statistics and set of rules which describe normal values for these statistics: in case an anomaly occurs some of the statistics should show a considerable deviation from its normal values. Constructing these statistics requires an exhaustive knowledge of the detector properties and possible anomalies.

We, in contrast, follow an agnostic approach: instead of constructing a set of dedicated statistics, the system is based on physics properties of collected data, measured or reconstructed. This approach was chosen for a number of reasons. First of all, statistics, similar to those used

by experts, can be learned directly from data. This gives a possibility for automated detection of anomalies. Secondly, such system is easily adaptable to different experimental setups including modifications in the detector. Finally, an automated approach and the expert statistics are not mutually exclusive, and injection of expert statistics into the feature set is a good starting point for the improvement of the system.

The primary goal of the system is to assist Data Quality managers by filtering most obvious cases, both positive and negative. CMS data are aggregated into lumisections each corresponding to approximately 23 seconds of data taking. Lumisection is a data granularity for which the data quality flag is defined in the CMS software stack.

This task requires classification of all samples into three categories:

- definitely anomalous (black zone): decision can be made automatically, samples are marked as anomalous;

- definitely good (white zone): decision can be made automatically, samples are marked as good;

- ambiguous (gray zone): decision can not be made automatically, human intervention is required.

Initially, when no data is available, the system classifies all incoming samples as ambiguous. Samples from gray zone are passed for evaluation to the human experts. Evaluation results are then used for retraining the system. In this way, the systems learns to mimic the human expert.

Formally speaking, the systems objective is minimization of fraction of data samples passed for the human evaluation. i.e. the rejection rate:

$$\text{Rejection Rate} = \frac{\text{Rejected}}{\text{Total}} \to \min, \qquad (1)$$

under constraints:

$$\text{Loss Rate} = \frac{\text{False Negative}}{\text{True Positive} + \text{False Negative}} \le L_0, \qquad (2)$$

$$\text{Pollution Rate} = \frac{\text{False Positive}}{\text{False Positive} + \text{True Positive}} \le P_0, \qquad (3)$$

where:

- Rejected — total quantity of samples rejected by the system,

- Total — total quantity of processed samples,

- False Negative, False Positive, True Negative, True Positive — anomalous data classified as good, good data classified as anomalous, correctly classified good data and anomalous respectively.

Quantities Rejection Rate, Loss Rate and Pollution Rate can also be measured as fractions of total luminosity rather than number of processed lumisections.

Constants L_0 and P_0 are to be set in advance and are driven by external requirements to the system. These constants define desired quality of label assignment: maximal fraction of 'lost' lumisections, i.e. good ones classified as anomalous (loss rate); and maximal fraction of anomalous lumisections classified as good ones (pollution rate). In this way, the system is forced to process automatically only the most obvious cases.

3. Data preprocessing

As mentioned above, CMS data samples are aggregated into lumisections each corresponding to approximately 23 seconds of data taking. Thus, a lumisection is a set of events, and each event corresponds to one beam crossing.

In the CMS pipeline, all taken events are split into 'streams' according to some criteria. In this work, we consider only the following streams:

- minimal bias stream: prescaled stream of all events [2];

- muon stream [3].

- photon stream [4];

Reconstructed high level objects in the event are may be divided into 'channels' depending on particles type or registering subsystem:

- muons;

- photons;

- particle flow jets;

- calorimeter jets.

Every object is characterized by its reconstructed physics properties. In this work, the following features are considered:

- p_T — traverse momentum;

- η, ϕ — pseudo-rapidity and angle between the transverse direction and the horizontal plane;

- f_x, f_y, f_z — coordinates of the reconstructed origin;

- m — reconstructed object mass for composed objects.

Since only the whole lumisection can be marked as good or anomalous, aggregation of physical features is performed.

All objects in a given channel in one event are sorted in the descending order by the momentum and $l = 5$ particles with indices: $0, \left\lfloor \frac{N}{l} \right\rfloor, \left\lfloor \frac{2N}{l} \right\rfloor, \ldots \left\lfloor \frac{(l-1)N}{l} \right\rfloor$ are selected to represent the channel for this event. In this way each event is characterized by fixed number of features. The last step of the preprocessing is to compute statistics for each feature for the entire lumisection. It was found experimentally, that among all considered combinations the best results are produced if we consider:

- mean and standard deviation;

- 1, 25, 50, 75, 99 one-sided percentiles.

Each lumisection is described by a fixed number of features, each feature having a hierarchical description:

- stream: minimum bias, photon or muon;

- channel: muons, photons, particle flow jets, calorimiter jets;

- quantile by momentum: $\frac{1}{5}, \frac{2}{5}, \ldots, 1$;

- physics property: p_T, η, ϕ, f_x, f_y, f_z, m;

- statistics for the lumisection: 1, 25, 50, 75, 99 percentiles, mean, standard deviation.

Additionally, features like the number of events in the lumisection and three components of the vector sum of momentum for all particles in the event, and others were also introduced.

4. Algorithm

In order to maximize number of automatically processed samples under constraints (2) and (3), a strong classifier is needed. We consider score function of the chosen classifier as a measure of certainty of the decision. Essentially, only two thresholds on classifier score, τ_L and τ_R, are of primary interest. These thresholds correspond to minimal and maximal score of a sample to be automatically classified as 'good' or 'anomalous', respectively. We estimate those thresholds on samples, already labeled by human experts. Independent scores, $\hat{y}$, are obtained by cross-validation procedure introduced in [5], and corresponding estimates for Loss Rate and Pollution Rate which are labeled as $\hat{L}_\tau(\hat{y}, y)$ and $\hat{P}_\tau(\hat{y}, y)$ respectively are evaluated for each threshold τ. Thresholds τ_L and τ_P are then selected as respectively maximal and minimal values of τ for which constraints (2) and (3) are satisfied.

In order to achieve maximal performance, the system memorizes each decision made by human expert, and the subsequent classifier is trained on all samples available. Note, that the system updates classifier every time when new sample has been processed by a human expert.

Pseudo-code for the system is provided in listing 1.

Algorithm 1 A pseudocode for the automated data quality system.

    function TRAIN(X, y, L_0, P_0)
        compute scores $\hat{y}$ by k-fold cross-validation as in [5]
        $\tau_L = \max\{\tau \mid \hat{L}_\tau(\hat{y}, y) \leq L_0\}$
        $\tau_P = \min\{\tau \mid \hat{P}_\tau(\hat{y}, y) \leq P_0\}$
        return τ_L, τ_P, classifier trained on X, y
    end function

    function AUTOMATEDDATAQUALITY(L_0, P_0)
        $\tau_L, \tau_P \leftarrow 0, 1$
        classifier $\leftarrow \frac{1}{2}$
        $X_{\text{train}} = \varnothing$
        $y_{\text{train}} = \varnothing$
        for $i = 0, 1, \ldots, N$ **do**
            $x_i \leftarrow$ new sample
            $\hat{y}_i \leftarrow$ classifier(x_i)
            if $\hat{y}_i > \tau_L$ **then**
                classify x_i as good lumisection
            else if $\hat{y}_i < \tau_P$ **then**
                classify x_i as anomalous lumisection
            else
                $y_i \leftarrow$ label from human expert
                $X \leftarrow (X, x_i)$
                $y \leftarrow (y, y_i)$
                τ_L, τ_P, classifier $\leftarrow$ Train(X, y, L_0, P_0)
            end if
        end for
    end function

It is worth noting that the performance of the system changes over the course of learning, and it rapidly improves as more data is evaluated and labeled by the experts, and thus available for training. The performance is expected to have some intrinsic limit. An example of empirical learning curves are shown in figure 2.

5. Experiment

Performance of the system was evaluated on data collected by the CMS experiment at the LHC in the year 2010 and made available through the CERN OpenData portal ([2], [3], [4]). The data was preprocessed as described in section 3. The system is implemented with the Gradient Tree Boosting classifier [6] as an underlying classifier. 10-fold cross-validation scheme was used to estimate thresholds τ_L and τ_P *.

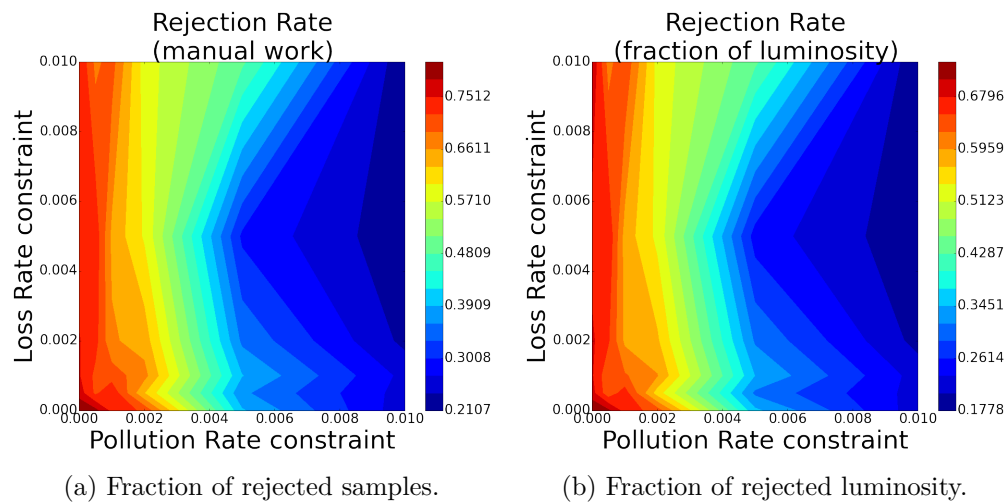(a) Fraction of rejected samples.   (b) Fraction of rejected luminosity.

Figure 1: Evaluation of the system using CMS experimental data published in the CERN open data portal. X and Y axis of plots correspond to Pollution Rate and Loss Rate respectively, color corresponds to rejection rate expressed as 1a fraction of rejected samples and 1b fraction of rejected luminosity.
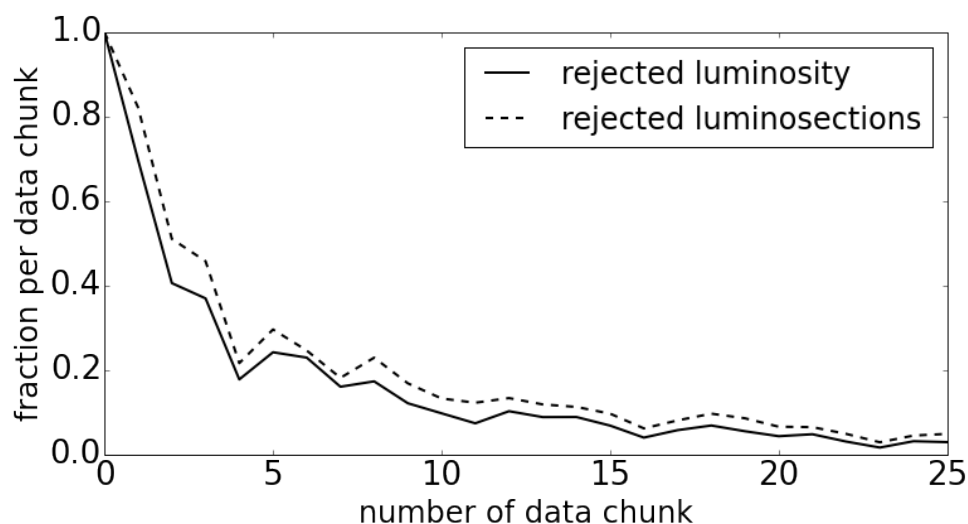

Figure 2: Example of learning curves: fractions of rejected luminosity and lumisections gradually decrease as classifier gets more labeled data.

To speed up the evaluation process the data was split randomly into 26 chunks and the system accepted a whole chunk at once rather than a single sample. This procedure might result in

* The code of the system is available at `https://github.com/yandexdataschool/cms-dqm/`.

slight underestimation of system performance, since in this case evaluation does not take into account the data from onr chunk (see Alg. 1). However, since a typical time scale of data taking in the CMS experiment is much larger than limitations of this experiment, the number of samples per chunk can be considerably decreased.

The evaluation was performed for constraint on Pollution Rate and Loss Rate set to 0, $1 \cdot 10^{-3}$, $2 \cdot 10^{-3}$, $5 \cdot 10^{-3}$, 10^{-2} as fractions of luminosity. Constraint were not violated with the exception of the point Loss Rate $= 0$, Pollution Rate $= 0$, where violations are present but negligible (less that 10^{-4}). The system was able to automatically process at least 20% of samples (which account for 30% of total luminosity). The rapid growth is observed in both quantities as restrictions become less strict.

6. Conclusion
In this work, we described an approach for automated data quality system. While developed with the CMS experiment in mind, we use an agnostic approach which allows the straightforward adaptation of the proposed algorithm to different experimental setups. We also define the clear strategy for improving performance using knowledge about detector specifics. Performance of the system was evaluated on the data collected by the CMS experiment at the LHC in 2010 and made available through the CERN OpenData portal. Experiments demonstrate that the system is able to automatically process at least 20% of samples and 30% of total luminosity keeping pollution and loss rates on negligible level, and with more relaxed restrictions on pollution and loss the performance of the system significantly increases.

References
[1] Chatrchyan S, Hmayakyan G, Khachatryan V, Sirunyan A, Adam W, Bauer T, Bergauer T, Bergauer H, Dragicevic M, Erö J *et al.* 2008
[2] CMS collaboration 2014 MinimumBias primary dataset in AOD format from RunB of 2010 (/MinimumBias/Run2010B-Apr21ReReco-v1/AOD) `http://doi.org/10.7483/OPENDATA.CMS.6BPY.XFRQ`
[3] CMS collaboration 2014 Mu primary dataset in AOD format from RunB of 2010 (/Mu/Run2010B-Apr21ReReco-v1/AOD) `http://doi.org/10.7483/OPENDATA.CMS.B8MR.C4A2`
[4] CMS collaboration 2014 Photon primary dataset in AOD format from RunB of 2010 (/Photon/Run2010B-Apr21ReReco-v1/AOD) `http://doi.org/10.7483/OPENDATA.CMS.QKAX.PSW6`
[5] Blum A, Kalai A and Langford J 1999 *Proceedings of the twelfth annual conference on Computational learning theory* (ACM) pp 203–208
[6] Friedman J H 2001 *Annals of statistics* 1189–1232