

Real-Time Anomaly Detection in Data Centers for Log-based Predictive Maintenance using an Evolving Fuzzy-Rule-Based Approach

Leticia Decker
University of Bologna
INFN Bologna
Bologna, Italy
leticia.deckerde@unibo.it

Daniel Leite
Federal University of Lavras
Department of Automatics
Lavras, Brazil
daniel.leite@ufla.br

Luca Giommi
University of Bologna
INFN Bologna
Bologna, Italy
luca.giommi3@unibo.it

Daniele Bonacorsi
University of Bologna
INFN Bologna
Bologna, Italy
daniele.bonacorsi@unibo.it

Abstract—Detection of anomalous behaviors in data centers is crucial to predictive maintenance and data safety. With data centers, we mean any computer network that allows users to transmit and exchange data and information. In particular, we focus on the Tier-1 data center of the Italian Institute for Nuclear Physics (INFN), which supports the high-energy physics experiments at the Large Hadron Collider (LHC) in Geneva. The center provides resources and services needed for data processing, storage, analysis, and distribution. Log records in the data center is a stochastic and non-stationary phenomenon in nature. We propose a real-time approach to monitor and classify log records based on sliding time windows, and a time-varying evolving fuzzy-rule-based classification model. The most frequent log pattern according to a control chart is taken as the normal system status. We extract attributes from time windows to gradually develop and update an evolving Gaussian Fuzzy Classifier (eGFC) on the fly. The real-time anomaly monitoring system has to provide encouraging results in terms of accuracy, compactness, and real-time operation.

Index Terms—predictive maintenance, anomaly detection, machine learning, evolving intelligent system, fuzzy logic

I. INTRODUCTION

A computing center (CC) is responsible for supporting a flexible, on-demand, dynamic, and computing-scalable cloud infrastructure, in which the resources are available directly or by means of services [1]. The complex CC infrastructure requires maintenance tools to keep itself operative, efficient, and reliable.

The maintenance of a CC is based on the complexity of the operation and idling time. It is usually classified as: (i) reactive; (ii) preventive; (iii) predictive; and (iv) advanced. The reactive maintenance refers to a set of procedures deployed after the fault occurrence, which aims at restoring the pristine behavior. Preventive maintenance is the collection of procedures performed to lessen the likelihood of a system failure. The predictive maintenance is designed to determine the status of running services, and predict events of interest. Advanced maintenance combines the other three paradigms in order to forecast and diagnose failures [2].

Usually, CC maintenance is based on offline statistical analysis of log records – in the preventive case, this is

based on fixed time intervals. Recently, online computational-intelligence-based systems, namely, evolving fuzzy modelling frameworks [3] [4] [5] [6] [7] [8] supported by fast incremental machine-learning algorithms, have been employed in general issues related to on-demand anomaly detection, forecasting, autonomous data classification, and predictive maintenance of a plethora of applications [9] [10] [11] [12] [13] [14].

Log records concern service-oriented unstructured data. Log data samples need to be ad-hoc processed by learning and modelling algorithms. The use of general-purpose solutions based on the content of log files has been a challenge over the years. In a log-based system, the data may be highly verbose such that it is hard to extract useful information from raw data. The amount of data is huge while a high percentage tends to be redundant. Any CC service run by a user generates log data using multiple files. After being processed, a reasonable amount of data for analysis is obtained.

Since all CC activities are recorded in log files, algorithms can track event occurrences through the data extracted from log files to monitor and predict the system status. In the predictive case, identification of anomalous behavior as an intermediate step using global attributes of log records is possible [15]. To reduce log-content processing, a common characteristic of the log files – the timestamp of each line writing – can be used. Furthermore, a reasonable assumption is that the system activity is proportional to the per-minute rate of lines written in a log file. Considering overall system faults, a direct impact on such rate of written log records is expected.

The background scenario of this study is the Tier-1 Bologna – the main Italian WLCG (Worldwide LHC Computing Grid) tier hold by the computing center of the Italian Institute of Nuclear Physics (INFN-CNAF). The WLCG involves 170 computing centers in over 42 countries, being the grid system that support the physics experiments performed at the biggest particle accelerator in the world, the CERN. It is organised in 4 layers – the tiers, from 0 (at CERN) to 3, in decreasing order of importance. The Tier-1 Bologna has approximately 40,000 CPU cores, 40 PB of disk storage, and 90 PB of tape storage.

It is connected to the Italian (GARR) and European (GÉANT) research networks, whose bandwidth of data transmission is over 200 Gbps. Currently, the Tier-1 has collected log data from 1,197 machines.

The INFN-CNAF provides a computing farming that accounts for all computing services of the Tier-1 Bologna. It acts as a service underlying the workload management system, allowing job scheduling to access directly the INFN-CNAF experiment data. On average, about 100 thousand batch jobs are executed per day at INFN-CNAF. The resources are continuously available, 24 hours a day, 7 days a week. The CC facility is based on a warehouse infrastructure for both storage and data transfer through a distributed system [16].

As Tier-1 Bologna CC concerns a dedicated infrastructure to support physics experiments [17], minimising the resources to maintain system operability is needed as log-data handling is a highly time and resource-consuming task. To achieve such computational cost minimisation, an approach is to identify which pieces of log data have processing priority aiming at maximising the likelihood to find useful information to the system maintenance. The present study addresses anomaly detection of the system behavior as an optimisation approach for predictive maintenance.

The rest of this paper is structured as follows. Section II presents related literature on anomaly detection and system maintenance at computing centers. Section III describes an evolving fuzzy-rule-based classification framework that is able to learn from summaries of log records and keep an updated representation of the spatial-temporal patterns related to the generation of log files. Section IV shows the methodology to perform the computational experiments. Classification results are given in Section V. Conclusions are outlined in Section VI.

II. RELATED LITERATURE

Because of the High-Luminosity LHC (HL-LHC) project, the major programmed upgrade at CERN, the used luminosity will increase by a factor of 10 from the original design. The luminosity is the rate of potential collisions per surface unit, which is proportional to the generated experimental data [18]. In this way, the amount of experimental and Monte Carlo analysis data will enlarge by, at least, the same factor, intensifying the maintenance complexity to keep the computing center quality of service (QoS).

For that reason, many efforts are being done at Tier-1 Bologna in order to create predictive maintenance tools using the log data. A first work based on the Elastic Stack Suite catalogues the log records and anomalies using an embedded unsupervised ML tool [19]. Another initiative uses supervised ML approaches to predict anomalies of system behavior in an ad-hoc solution [20]. Another work, also focused on a content-processing strategy, provides a clustering method used to characterize log records using Levenshtein distance [21]. In particular, it was created a prototype to identify a normal and an anomalous system behavior, in a binary classification, considering the log data generation rate and an One-class Support Vector Machine approach [16].

A. The StoRM logs use case

StoRM is the storage resource manager service (SRM) for generic disk-based storage systems adopted by the Tier-1 Bologna, providing high performance to parallel file systems.

SRM has a modular architecture made by two stateless components, Front-end (FE) and Back-end (BE), connected to database systems. The FE module manages user authentication, stores/retrieves database requests, interacting with the BE module [22].

In the other hand, the BE module is the core of StoRM service, executing all synchronous and asynchronous SRM functionalities and managing the Grid interactions. A simple StoRM architecture schema is presented in Fig. 1, showing the main module interactions. Typically, BE log files entries include the operator that has requested the action (DN), the involved files locations (SURL) and the result of the operation. A sample of its log messages is shown in Fig. 2.

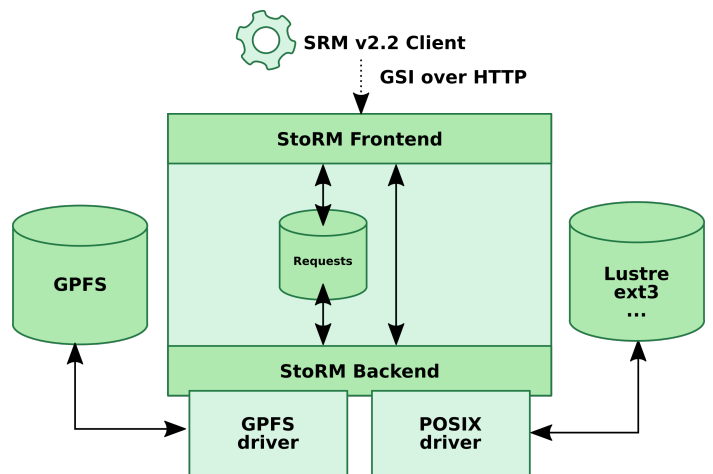


Fig. 1. A typical StoRM service architecture, with single back and front-end modules

In addition, the StoRM service at Tier-1 Bologna is used by high-energy physics experiments, in which each one has a different implementation of the structures and rules of the logging. In this work, the BE module log files from ATLAS implementation is chosen as input without any special reason.

III. EGFC: EVOLVING GAUSSIAN FUZZY CLASSIFIER

This section outlines eGFC, a semi-supervised evolving classifier derived from an online granular-computing framework [23] [24]. Although eGFC handles partially labeled data, we assume a fully-labeled log-file dataset in this paper. eGFC employs Gaussian membership functions to cover the data space with fuzzy granules, and associate new data samples to class labels. Granules are scattered in the data space wherever needed to represent local information. eGFC global response comes from the aggregation of local models. A recursive algorithm constructs a rule base, and updates local models to deal with changes. eGFC addresses issues such as unlimited amounts of data and scalability [3] [25].

```

00:00:00.615 - INFO [xmlrpc-75687] - srmLs: user </DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=atlpilo1/CN=614260/CN=Robot: ATLAS Pilot1> Request for [SURL: [srm://storm-fe.cr.cnaf.infn.it/atlas/atlasdatadisk/rucio/mc16_13TeV/63/ad/HITS.17978659._022059.pool.root.1.rucio.upload]] failed with: [status: SRM_FAILURE: All requests failed]
00:00:00.717 - INFO [xmlrpc-75643] - srmLs: user </DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=ddmadmin/CN=531497/CN=Robot: ATLAS Data Management> Request for [SURL: [srm://storm-fe.cr.cnaf.infn.it/atlas/atlasdatadisk/rucio/mc16_13TeV/de/a3/DAOD_HIGG2D1.15833676._000007.pool.root.1]] successfully done with: [status: SRM_SUCCESS: All requests successfully completed]
00:00:00.840 - INFO [xmlrpc-75683] - srmLs: user </DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=atlpilo1/CN=614260/CN=Robot: ATLAS Pilot1> Request for [SURL: [srm://storm-fe.cr.cnaf.infn.it/atlas/atlasdatadisk/rucio/mc16_13TeV/63/ad]] successfully done with: [status: SRM_SUCCESS: All requests successfully completed]
00:00:00.891 - INFO [Timer-4] - ADVANCED PICKER: dispatching 5 requests.

```

Fig. 2. Example of content of a *storm-backend.log* file

A. Preliminaries

Local models are created if the newest data are sufficiently different from the current knowledge. The learning algorithm expand, reduce, delete, and merge information granules. Rules are reviewed according to inter-granular relations. eGFC provides nonlinear, non-stationary, and fuzzy discrimination boundaries among classes [3] [23]. This paper particularly addresses a 4-class log-file classification problem.

Formally, let an input-output pair (\mathbf{x}, y) be related through $y = f(\mathbf{x})$. We seek an approximation to f to estimate the value of y given \mathbf{x} . In classification, y is a class label, a value in a set $\{C_1, \dots, C_m\} \in \mathbb{N}^m$, and f specifies class boundaries. In the more general, semi-supervised case, C_k may or may not be known when \mathbf{x} arrives. Classification of never-ending data streams involves pairs $(\mathbf{x}, C)^{[h]}$ of time-sequenced data, indexed by h . Non-stationarity requires evolving classifiers to identify time-varying relations $f^{[h]}$.

B. Gaussian Functions and Rule Structure

Learning in eGFC does not require initial rules. Rules are created and dynamically updated depending on the behavior of a system over time. When a data sample is available, a decision procedure may add a rule to the model structure or update the parameters of a rule.

In eGFC models, a rule R^i is

IF $(x_1 \text{ is } A_1^i)$ AND ... AND $(x_n \text{ is } A_n^i)$
THEN $(y \text{ is } C^i)$

in which x_j , $j = 1, \dots, n$, are attributes, and y is a class. The data stream is denoted $(\mathbf{x}, y)^{[h]}$, $h = 1, \dots$. Moreover, A_j^i , $\forall j$; $i = 1, \dots, c$, are Gaussian membership functions built from the available data; and C^i is the class label of the i -th rule. Rules R^i , $\forall i$, form the rule base. The number of rules, c , is variable, which is a notable characteristic of the approach since guesses on how many data partitions exist are needless [3].

A normal Gaussian membership function, $A_j^i = G(\mu_j^i, \sigma_j^i)$, has height 1 [26]. It is characterized by the modal value μ_j^i and dispersion σ_j^i . Characteristics that make Gaussians appropriate include: (i) easiness of learning and changing, i.e., modal values and dispersions can be captured and updated straightforwardly from a data stream; (ii) infinite support, i.e., since the data are priorly unknown, the support of Gaussians extends to the whole domain; and (iii) smooth surface of fuzzy granules, $\gamma^i = A_1^i \times \dots \times A_j^i \times \dots \times A_n^i$, in the n -dimensional

Cartesian space – obtained by the cylindrical extension of uni-dimensional Gaussians, and the use of the minimum T-norm aggregation [25] [26].

C. Adding Rules to the Evolving Fuzzy Classifier

Rules may not exist *a priori*. They are created and evolved as data are available. A new granule γ^{c+1} and the rule R^{c+1} that governs the granule are created if none of the existing rules $\{R^1, \dots, R^c\}$ are sufficiently activated for a sample $\mathbf{x}^{[h]}$. The learning algorithm assumes that $\mathbf{x}^{[h]}$ brings new information. Let $\rho^{[h]} \in [0, 1]$ be an adaptive threshold that determines if a new rule is needed. If

$$T\left(A_1^i(x_1^{[h]}), \dots, A_n^i(x_n^{[h]})\right) \leq \rho^{[h]}, \quad \forall i, \quad i = 1, \dots, c, \quad (1)$$

in which T is any triangular norm, then the eGFC structure is expanded. The minimum (Gödel) T-norm is used in this paper, but other choices are possible. If $\rho^{[h]}$ is equal to 0, then the model is structurally stable, and unable to capture concept shifts. In contrast, if $\rho^{[h]}$ is equal to 1, eGFC creates a rule for each new sample, which is not practical. Structural and parametric adaptability are balanced for intermediate values of $\rho^{[h]}$ (stability-plasticity trade-off) [27].

The value of $\rho^{[h]}$ is crucial to regulate how large granules can be. Different choices impact the accuracy and compactness of a model, resulting in different granular perspectives of the same problem. Section III-E gives a Gaussian-dispersion-based procedure to update $\rho^{[h]}$.

A new granule γ^{c+1} is initially represented by membership functions, A_j^{c+1} , $j = 1, \dots, n$, with

$$\mu_j^{c+1} = x_j^{[h]}, \quad (2)$$

and

$$\sigma_j^{c+1} = 1/2\pi. \quad (3)$$

We call (3) the Stigler approach to standard Gaussian functions, or *maximum approach* [25]. The intuition is to start big, and let the dispersions gradually shrink when new samples activate the same granule. This strategy is appealing for a compact model structure.

In general, the class C^{c+1} of the rule R^{c+1} is initially undefined, i.e., the $(c+1)$ -th rule remains unlabeled until a label is provided. If the corresponding output, $y^{[h]}$, associated to $\mathbf{x}^{[h]}$, becomes available, then

$$C^{c+1} = y^{[h]}. \quad (4)$$

Otherwise, the first labeled sample of the data stream that arrives after the h -th time step, and activates the rule R^{c+1} according to (1), is used to define its class, C^{c+1} .

In case a labeled sample activates a rule that is already labeled, but the sample's label is different from that of the rule, then a new (partially overlapped) granule and a rule are created to represent new information. Partially overlapped Gaussian granules tagged with different labels tend to have their dispersions reduced over time by the parameter adaptation procedure described in Section III-D. The modal values of the Gaussian granules may also drift, if convenient for a more suitable decision boundary.

With this initial rule parameterization, preference is given to the design of granules balanced along its dimensions, rather than granules with unbalanced geometry. eGFC realizes the principle of the balanced information granularity [28], but allows the Gaussians to find more appropriate places and dispersions through adaptation mechanisms.

D. Incremental Parameter Adaptation

Updating the eGFC model consists in: (i) reducing or expanding Gaussians $A_j^{i^*}$, $\forall j$, of the most active granule, γ^{i^*} , considering labeled and unlabeled samples; (ii) moving granules toward regions of relatively dense population; and (iii) tagging rules if labeled data are available. Adaptation aims to develop more specific local models in the sense of Yager [29], and provide pavement (covering) to the newest data.

A rule R^i is candidate to be updated if it is sufficiently activated by an unlabeled sample, $\mathbf{x}^{[h]}$, according to

$$\min \left(A_1^i(x_1^{[h]}), \dots, A_n^i(x_n^{[h]}) \right) > \rho^{[h]}. \quad (5)$$

Geometrically, $\mathbf{x}^{[h]}$ belongs to a region highly influenced by the granule γ^i . Only the most active rule, R^{i^*} , is chosen for adaptation in case two or more rules reach the $\rho^{[h]}$ level for the unlabeled $\mathbf{x}^{[h]}$. For a labeled sample, i.e., for pairs $(\mathbf{x}, y)^{[h]}$, the class of the most active rule R^{i^*} , if defined, must match $y^{[h]}$. Otherwise, the second most active rule among those that reached the $\rho^{[h]}$ level is chosen for adaptation, and so on. If none of the rules are apt, then a new one is created (Section III-C).

To include $\mathbf{x}^{[h]}$ in R^{i^*} , eGFC's learning algorithm updates the modal values and dispersions of the corresponding membership functions $A_j^{i^*}$, $j = 1, \dots, n$, from

$$\mu_j^{i^*}(\text{new}) = \frac{(\varpi^{i^*} - 1)\mu_j^{i^*}(\text{old}) + x_j^{[h]}}{\varpi^{i^*}}, \quad (6)$$

and

$$\sigma_j^{i^*}(\text{new}) = \left(\frac{(\varpi^{i^*} - 1)}{\varpi^{i^*}} \left(\sigma_j^{i^*}(\text{old}) \right)^2 + \frac{1}{\varpi^{i^*}} \left(x_j^{[h]} - \mu_j^{i^*}(\text{old}) \right)^2 \right)^{1/2}, \quad (7)$$

in which ϖ^{i^*} is the number of times the i^* -th rule was chosen to be updated. Notice that (6)-(7) are recursive and, therefore, do not require data storage. As σ^{i^*} defines a convex region

of influence around μ^{i^*} , very large and very small values may induce, respectively, a unique or too many information granules per class. An approach is to keep $\sigma_j^{i^*}$ between a lower, $1/4\pi$, and the Stiegler, $1/2\pi$, limits.

E. Dispersion-Based Time-Varying ρ -Level

Let the activation threshold, $\rho^{[h]} \in [0, 1]$, be time-varying, similar to [23] [6]. The threshold assumes values in the unit interval according to the overall average dispersion

$$\sigma_{avg}^{[h]} = \frac{1}{cn} \sum_{i=1}^c \sum_{j=1}^n \sigma_j^{i[h]}, \quad (8)$$

in which c and n are the number of rules and attributes, so that

$$\rho(\text{new}) = \frac{\sigma_{avg}^{[h]}}{\sigma_{avg}^{[h-1]}} \rho(\text{old}). \quad (9)$$

As mentioned, rules' activation levels for an input $\mathbf{x}^{[h]}$ are compared to $\rho^{[h]}$ to decide between parametric or structural changes of an eGFC model. In general, eGFC starts learning from an empty rule base, and without knowledge about the properties of the data. Practice suggests $\rho^{[0]} = 0.1$ as starting value. The threshold tends to converge to a proper value after some time steps if the classifier structure and parameters achieve a level of maturity and stability. Non-stationarities and new classes guide $\rho^{[h]}$ to values that better reflect the needs of the current environment. A time-varying $\rho^{[h]}$ avoids assumptions about how often the data stream changes.

F. Merging Similar Granules

Similarity between two granules with the same class label may be high enough to form a unique granule that inherits the essential information conveyed by the merged granules. Analysis of inter-granular relations requires a distance measure between Gaussian objects. Let

$$d(\gamma^{i_1}, \gamma^{i_2}) = \frac{1}{n} \left(\sum_{j=1}^n |\mu_j^{i_1} - \mu_j^{i_2}| + \sigma_j^{i_1} + \sigma_j^{i_2} - 2\sqrt{\sigma_j^{i_1} \sigma_j^{i_2}} \right) \quad (10)$$

be the distance between the granules γ^{i_1} and γ^{i_2} . This measure considers Gaussians and the specificity of information, that is, in turn, inversely related to the Gaussians' dispersion [30]. For example, if the dispersions $\sigma_j^{i_1}$ and $\sigma_j^{i_2}$ differ one from another, rather than being equal, the distance between the underlying Gaussians is larger.

eGFC may merge the pair of granules that presents the smallest value of $d(\cdot)$ for all pairs of granules. Both granules must be either unlabeled or tagged with the same class label. The merging decision is based on a threshold value, Δ , or expert judgement regarding the suitability of combining such granules to have a more compact model. For data within the unit hyper-cube, we suggest $\Delta = 0.1$ as default, which means that the candidate granules should be quite similar and, in fact, carry the same information.

A new granule, say γ^i , which results from γ^{i_1} and γ^{i_2} , is built by Gaussians with modal values

$$\mu_j^i = \frac{\frac{\sigma_j^{i_1}}{\sigma_j^{i_2}} \mu_j^{i_1} + \frac{\sigma_j^{i_2}}{\sigma_j^{i_1}} \mu_j^{i_2}}{\frac{\sigma_j^{i_1}}{\sigma_j^{i_2}} + \frac{\sigma_j^{i_2}}{\sigma_j^{i_1}}}, \quad j = 1, \dots, n, \quad (11)$$

and dispersion

$$\sigma_j^i = \sigma_j^{i_1} + \sigma_j^{i_2}, \quad j = 1, \dots, n. \quad (12)$$

These relations take into consideration the uncertainty ratio of the original granules to determine an appropriate location and size of the resulting granule. Merging granules reduces the number of rules and redundancy [23] [30].

G. Deleting Rules

A rule is removed from the eGFC model if it is inconsistent with the current environment. In other words, if a rule is not activated for a number of iterations, say h_r , then it is deleted from the rule base. However, if a class is rare, then it may be the case to set h_r to infinity and keep the inactive rules. Removing rules periodically helps to keep the knowledge base updated in some applications.

H. Semi-Supervised Learning from Data Streams

The semi-supervised learning procedure to construct and update eGFC models along their lifespan is given in the next column.

IV. METHODOLOGY

We describe a dynamic control chart approach we propose for attribute extraction and log data tagging. We give details about the data-set and evaluation measures.

A. Control Chart: Tagging Log Data

A control chart is a time-series graph to monitor the evolution of a process, phenomenon or variable based on the Central Limit theorem [31]. The main idea is that the mean, $\mu(u)$, of an independent random variable, u , with unknown distribution, follows a normal distribution.

Let

$$\mathbf{u}_j = [u_1 \dots u_i \dots u_n], \quad (13)$$

be a sequence of values that represent the number of log entries in a log file over a time window $w_j = [\underline{w}_j \ \overline{w}_j]$; $\mathbf{u}_j \in \mathbb{N}^n$. The time interval from u_1 to u_n coincides with the window boundaries, \underline{w}_j and \overline{w}_j . Additionally, let μ_j be the mean of \mathbf{u}_j , thus

$$\mu_j = \frac{1}{n} \sum_{i=1}^n u_i, \quad u_i \in [\underline{w}_j \ \overline{w}_j]. \quad (14)$$

A time series of means, with cardinality m , is

$$\mu = [\mu_1 \dots \mu_j \dots \mu_m]. \quad (15)$$

As μ follows a normal distribution, a sample μ_j can be tagged by means of a control chart, see Fig. 3. The mean of the time

eGFC: Online Semi-Supervised Learning

- 1: Initial number of rules, $c = 0$;
 - 2: Initial meta-parameters, $\rho^{[0]} = \Delta = 0.1$, $h_r = 200$;
 - 3: Read input data sample $\mathbf{x}^{[h]}$, $h = 1$;
 - 4: Create granule γ^{c+1} (Eqs. (2)-(3)), unknown class C^{c+1} ;
 - 5: **FOR** $h = 2, \dots$ **DO**
 - 6: Read $\mathbf{x}^{[h]}$;
 - 7: Calculate rules' activation degree (Eq. (1));
 - 8: Determine the most active rule R^{i^*} ;
 - 9: Provide estimated class C^{i^*} ;
 - 10: // Model adaptation
 - 11: **IF** $T(A_1^i(x_1^{[h]}), \dots, A_n^i(x_n^{[h]})) \leq \rho^{[h]} \ \forall i, \ i = 1, \dots, c$
 - 12: **IF** actual label $y^{[h]}$ is available
 - 13: Create labeled granule γ^{c+1} (Eqs. (2)-(4));
 - 14: **ELSE**
 - 15: Create unlabeled granule γ^{c+1} (Eqs. (2)-(3));
 - 16: **END**
 - 17: **ELSE**
 - 18: **IF** actual label $y^{[h]}$ is available
 - 19: Update the most active granule γ^{i^*} whose class C^{i^*} is equal to $y^{[h]}$ (Eqs. (6)-(7));
 - 20: Tag unlabeled active granules;
 - 21: **ELSE**
 - 22: Update the most active γ^{i^*} (Eqs. (6)-(7));
 - 23: **END**
 - 24: **END**
 - 25: Update the ρ -level (Eqs. (8)-(9));
 - 26: Delete inactive rules based on h_r ;
 - 27: Merge granules based on Δ (Eqs. (10)-(12));
 - 28: **END**
-

series of means μ is

$$\bar{\mu} = \frac{1}{m} \sum_{j=1}^m \mu_j. \quad (16)$$

The k -th upper and lower horizontal lines in relation to $\bar{\mu}$ refer to the k -th standard deviation,

$$\sigma_k(\mu) = k \sqrt{\frac{1}{m} \sum_{j=1}^m (\bar{\mu} - \mu_j)^2}, \quad (17)$$

such that if a sample

$$\mu_j \in [\bar{\mu} - \sigma_k(\mu_j \forall j), \bar{\mu} + \sigma_k(\mu_j \forall j)], \quad (18)$$

for $k = 1$, it is tagged as 'Class 1' (normal system condition). Otherwise, if (18) holds for $k = 2, 3$, and 4, respectively, μ_j is tagged as 'Class 2', 'Class 3', and 'Class 4', which mean low, medium, and high-severity anomaly. The greater the value of k , the greater the severity of the anomalous behavior.

Control charts are widely used in quality monitoring to identify anomalies according to the control lines calculated from a stream of means. The probability that a sample μ_j is within the different class boundaries are 68.3%, 27.1%, 4.3%, and 0.3%, respectively. Therefore, the online data classification problem is unbalanced.

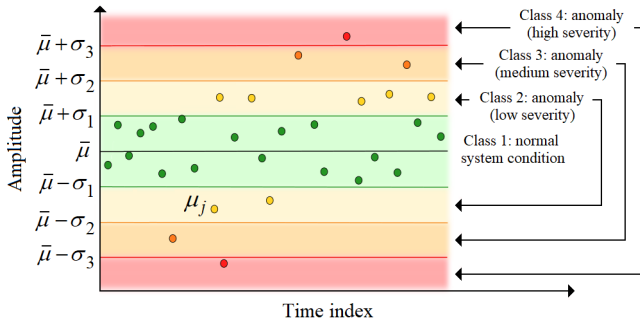


Fig. 3. Control chart used to tag mean log data within a time window

B. About the Data-set

A stream of time-indexed log records is generated by the StoRM service. Each log entry is composed by the timestamp in which it was written, and the message itself. Analysis of the message type and its content is out of the scope of this paper.

We extract relevant attributes from the original log data stream by analysing constant sliding time windows. Transformed data are provided as 5-attribute vectors

$$\mathbf{x} = [x_1 \ x_2 \ x_3 \ x_4 \ x_5], \quad (19)$$

whose elements evaluated in time window w_j are $\bar{\mu}$, $\sigma(\mu_j \ \forall j)$, $\min(\mu_j \ \forall j)$, $\max(\mu_j \ \forall j)$, and $\max(\Delta\mu_j \ \forall j)$. The latter means the maximum difference of the amplitude of two consecutive μ_j , belonging to the time window w_j .

A vector $\mathbf{x}^{[h]}$ is associated to a class label $C = \{1, 2, 3, 4\}$ that, in turn, indicates the system behavior. The true label C is available after an estimation \hat{C} provided by the eGFC model. The pair $(\mathbf{x}, C)^{[h]}$ is used by the eGFC online learning algorithm for an updating step.

C. Performance Measure

Classification accuracy, $Acc \in [0, 1]$, is computed recursively from

$$Acc(new) = \frac{h-1}{h} Acc(old) + \frac{1}{h} \tau, \quad (20)$$

in which $\tau := 1$ if $\hat{C}^{[h]} = C^{[h]}$ (right estimation). Otherwise, $\tau := 0$ (wrong class estimation).

The average number of granules or rules over time, c_{avg} , is a measure of model concision. Recursively,

$$c_{avg}(new) = \frac{h-1}{h} c_{avg}(old) + \frac{1}{h} c^{[h]}. \quad (21)$$

V. RESULTS

We evaluate the evolving Gaussian fuzzy classification system. No prior knowledge about the data is assumed. Classification models are developed from scratch based on information extracted from an online log data stream.

A. eGFC Results

We look for an evolving classifier based on the newest input data. The default meta-parameters are used (see *eGFC Learning Algorithm*). Table I summarizes the results averaged over 5 runs for shuffled data-sets extracted from log records. Four data-sets were produced using the same data, but different lengths of time windows, namely 60, 30, 15, and 5-minute time windows. Larger time windows impose a higher-order low-pass filter effect, and tends to isolate the trend component of the time series from cyclical and random (stochastic) components. Each data-set consists of 1,436 samples and 5 attributes. Four classes are possible, namely ‘normal operation’, ‘low severity’, ‘medium severity’, and ‘high severity’.

TABLE I
eGFC PERFORMANCE IN MULTI-CLASS CLASSIFICATION OF SYSTEM ANOMALIES (99% OF CONFIDENCE)

Lenght (min)	Acc(%)	# Rules	Time (s)
60	92.48 ± 1.21	13.42 ± 4.32	0.36 ± 0.10
30	88.01 ± 4.96	17.22 ± 2.59	0.45 ± 0.04
15	82.57 ± 5.64	18.13 ± 4.79	0.49 ± 0.10
5	81.97 ± 5.02	16.09 ± 2.51	0.41 ± 0.06

Table I shows that analysis of larger 60-minute windows facilitates the eGFC learning algorithm to detect and classify spatial-temporal patterns, which represent the anomaly classes. Notice that using a more compact model structure (13.42 fuzzy rules on average along the learning steps), the eGFC model produced an average accuracy of 92.48%. The CPU time in a quad-core i7-8550U with 1.80GHz and 8GB of RAM are similar in all scenarios.

Figure 4 gives a typical example of evolution of the ρ -level, accuracy, and number of eGFC rules. Four dimensions of the final Gaussian granules, at $h = 1436$, are also shown. Notice that data from Class 2 and Class 3 (low and medium-severity anomalies) spread in a nonlinear way over the data space. These classes require more than one granule and rule to be represented, whereas the remaining classes are generally given in a common region. Class-4 data (high-severity anomaly) belong to a more compact region than the data of other classes and, therefore, are represented by a single granule. A higher number of granules to represent a class, in general, provides larger non-linearity of decision boundaries, which improves classification accuracy.

Figure 5 emphasizes the multi-dimensional ellipsoidal geometry of eGFC granules. This contour lines representation confirm the spreading characteristic specially related to Class-2 data, showing large overlapping regions of Class-1 and Class-2. Figure 6 shows the confusion matrix for a 94.08%-accuracy scenario. Notice that confusion happens in the neighbourhood of a target class, which means that if a higher number of streaming samples are further available, the eGFC model may improve its accuracy by fine tuning its decision boundaries. Class 1 (normal operation) and Class 2 (low severity) are those responsible for a larger reduction of the overall accuracy.

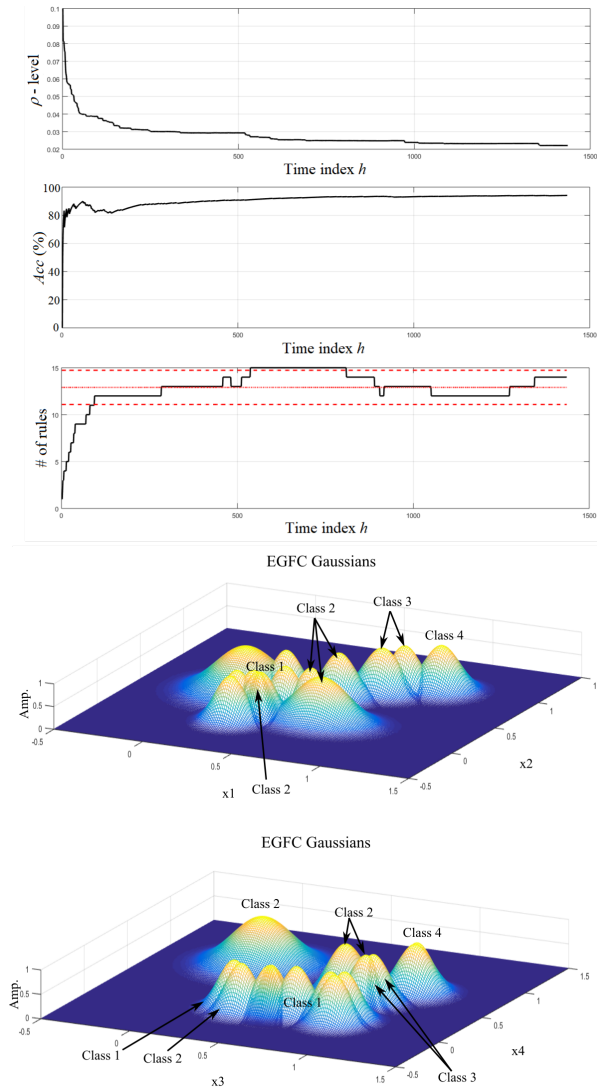


Fig. 4. The time evolution of the evolving factors: granulation ρ and number of rules, and the model accuracy until the convergence at the 3 first graphics. At the last 2 graphics the eGFC Gaussian classes.

To sum up, using the evolving fuzzy classification methodology and the sliding window control-chart-based approach, a CC maintenance system can accurately identify time windows that require further analysis in terms of text content. The evolving methodology supports data and information mining to assist predictive maintenance. Overall system status can be modelled as Gaussian granules of the log activity rate, and status changing can be noticed visually from the control charts. In addition, the stream of system status can be used to diagnose the context of the current log status, and to predict the next status. Since eGFC preserves its accuracy in non-stationary environment, the approach has shown to be a reliable solution to the predictive maintenance problem.

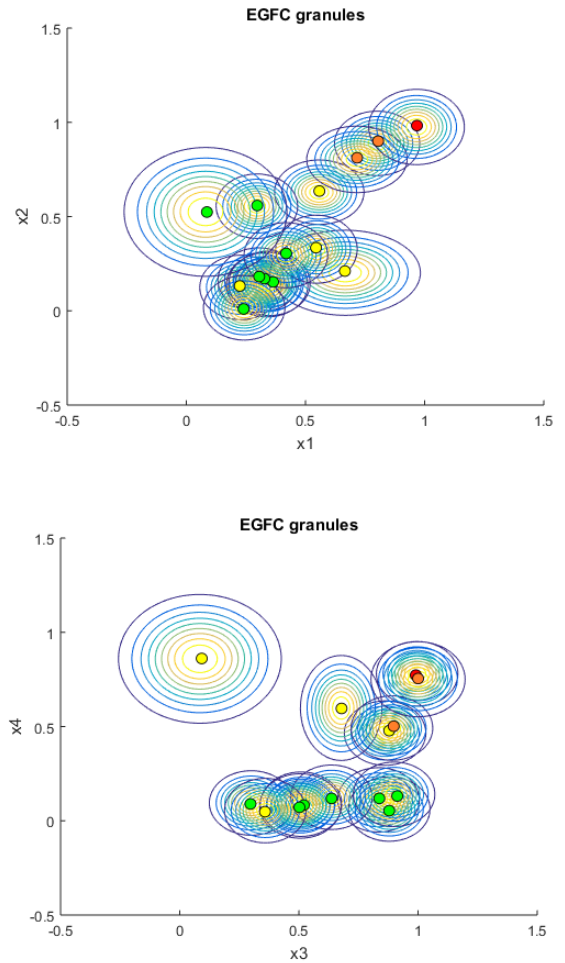


Fig. 5. The multi-dimensional ellipsoidal geometry of eGFC granules using the first four attributes of the log stream. The colours of the centers refer to the control chart of Fig. 3, i.e., green: normal system condition; yellow, orange and red: low, medium and high anomaly severity

VI. CONCLUSION

We described a real-time evolving general-purpose solution, namely, eGFC, to the log-based anomaly detection problem considering time-varying data from the Tier-1 Bologna computer center. eGFC models achieved an average accuracy of $92.48\% \pm 1.21$ with a confidence interval of 99% using a 60-minute sliding time window. Since the anomaly detection issue is context-sensible, the eGFC approach provides a strategy to update and evolve information granules and the parameters and structure of a fuzzy rule-based classifier in real-time. Multi-dimensional Gaussian granules are placed and sized autonomously in the data space aiming at constant improving the classification performance.

Fuzzy information granulation gives flexible and smooth boundaries to the classification model such that a wide variety of computer-center behaviors related to the same class label – even occurring in a conflicting region with overlapped classes – could be captured. This way, the eGFC approach, as a data-stream-oriented method, has shown to be highly

		Target Class				
		1	2	3	4	
Estimated Class	1	1086 75.6%	10 0.7%	0 0%	0 0%	99.1% 0.9%
	2	59 4.1%	186 13.0%	5 0.3%	0 0%	74.4% 25.6%
	3	1 0.1%	0 0%	66 4.6%	2 0.1%	95.7% 4.3%
	4	0 0%	1 0.1%	6 0.4%	14 1.0%	66.7% 33.3%
		94.8% 5.2%	94.4% 5.6%	85.7% 14.3%	87.5% 12.5%	94.2% 5.8%

Fig. 6. Example of confusion matrix provided by a 94.2%-accuracy eGFC model

applicable to a large range of classification issues concerning large log records from computing centers such as the Tier-1 Bologna, which supports the high-energy physics experiments at the Large Hadron Collider. Additionally, the autonomous sliding-window-based tagging strategy using control charts was successfully applied to the anomaly detection problem in question. Hand-labelling large volumes of online data (a key research issue in the machine learning community) is usually infeasible. Therefore, the chart-based approach seems quite promising to lead accuracy improvement in evolving classification frameworks.

The present study provides basis for extracting information from log content and identifying the best components to be text-processed, which minimise computational resource consumption. In the future we shall identify the type of message associated to anomalous time windows, and investigate autonomous feature extraction procedures.

REFERENCES

- [1] B. Furht and A. Escalante, *Handbook of Cloud Computing*. Springer Publishing Company, Incorporated, 1st ed., 2010.
- [2] F. Trojan and R. Marçal, "Proposal of maintenance-types classification to clarify maintenance concepts in production and operations management," *Journal of Business Economics*, vol. 8, pp. 562–574, 07 2017.
- [3] I. Škrjanc, J. Iglesias, A. Sanchis, D. Leite, E. Lughofer, and F. Gomide, "Evolving fuzzy and neuro-fuzzy approaches in clustering, regression, identification, and classification: A survey," *Inf. Sci.*, vol. 490, pp. 344–368, 2019.
- [4] L. A. Cordovil Júnior, P. H. Coutinho, I. Bessa, M. D'Angelo, and R. Palhares, "Uncertain data modeling based on evolving ellipsoidal fuzzy information granules," *IEEE Transactions on Fuzzy Systems*, 2019. DOI: 10.1109/TFUZZ.2019.2937052.
- [5] G. Casalino, G. Castellano, and C. Mencar, "Data Stream Classification by Dynamic Incremental Semi-Supervised Fuzzy Clustering," *Int J Artif Intell T*, vol. 28, no. 08, p. 26p, 2019.
- [6] C. Garcia, D. Leite, and I. Škrjanc, "Incremental missing-data imputation for evolving fuzzy granular prediction," *IEEE T FUZZY SYST*, p. 15p, 2019. DOI: 10.1109/TFUZZ.2019.2935688.

- [7] R. Hyde, P. Angelov, and A. Mackenzie, "Fully online clustering of evolving data streams into arbitrarily shaped clusters," *Inf. Sci.*, vol. 382, p. 41p, 2016.
- [8] H. J. Sadaei, P. C. de Lima e Silva, F. G. Guimarães, and M. H. Lee, "Short-term load forecasting by using a combined method of convolutional neural networks and fuzzy time series," *Energy*, vol. 175, pp. 365–377, 2019.
- [9] R. Venkatesan, M. Er, M. Dave, M. Pratama, and S. Wu, "A novel online multi-label classifier for high-speed streaming data applications," *Evolving Systems*, pp. 303–315, 07 2016.
- [10] J. L. Lobo, I. Lana, J. D. Ser, M. N. Bilbao, and N. Kasabov, "Evolving spiking neural networks for online learning over drifting data streams," *Neural Networks*, vol. 108, pp. 1–19, 2018.
- [11] P. Angelov and X. Gu, "Deep rule-based classifier with human-level performance and characteristics," *Inf. Sci.*, vol. 463–464, pp. 196–213, 10 2018.
- [12] M. Pratama, E. Lughofer, C. Lim, W. Rahayu, T. Dillon, and A. Budiyo, "pclass+: A novel evolving semi-supervised classifier," *INT J FUZZY SYST*, vol. 19, pp. 863–880, 07 2016.
- [13] L. Decker, D. Leite, F. Viola, and D. Bonacorsi, "Comparison of evolving granular classifiers applied to anomaly detection for predictive maintenance in computing centers," in *IEEE Conference on Evolving and Adaptive Intelligent Systems (EAIS), Bari*, p. 8p.
- [14] L. D. Sousa et al., "Event detection framework for wireless sensor networks considering data anomaly," in *IEEE Symposium on Computers and Communications (ISCC)*, pp. 500–507, 2012.
- [15] L. D. Sousa et al., "Big data analysis for predictive maintenance at the INFN-CNAF data center using machine learning approaches," in *Conf of Open Innovations Association (FRUCT), Helsinki*, pp. 448–451, 2019.
- [16] F. Minarini, "Anomaly detection prototype for log-based predictive maintenance at INFN-CNAF," Master's thesis, U. of Bologna, 2019.
- [17] A. Di Girolamo et al., "Operational intelligence for distributed computing systems for exascale science," in *Int Conf on Computing in High Energy and Nuclear Physics (CHEP), AU*, p. 8p, 2020.
- [18] W. Herr and B. Muratori, "Concept of luminosity," in *CAS - CERN Accelerator School: Course on Accelerator Physics*, pp. 361–378, 2006.
- [19] T. Diotallevi et al., "Collection and harmonization of system logs and prototypal analytics services with the elastic (ELK) suite at the INFN-CNAF computing centre," in *Int Symposium on Grids & Clouds (ISGC), Taiwan: Proceedings of Science*, p. 15p, 2019.
- [20] L. Giommi et al., "Towards predictive maintenance with machine learning at the INFN-CNAF computing centre," in *International Symposium on Grids & Clouds (ISGC), Taipei, Taiwan: Proceedings of Science*, p. 17p, 2019.
- [21] S. R. Tisbeni, "Big data analytics towards predictive maintenance at the INFN-CNAF computing centre," Master's thesis, U. of Bologna, 2019.
- [22] *The StoRM project*. <https://italiangrid.github.io/storm/index.html>.
- [23] D. Leite, R. Ballini, P. Costa Jr, and F. Gomide, "Evolving fuzzy granular modeling from nonstationary fuzzy data streams," *Evolving Systems*, vol. 3, pp. 65–79, 06 2012.
- [24] D. Leite, L. Decker, M. Santana, and P. Souza, "EGFC: Evolving Gaussian fuzzy classifier from never-ending semi-supervised data streams - with application to power quality disturbance detection and classification," in *IEEE World Congress on Computational Intelligence (WCCI - FUZZ-IEEE), Glasgow*, p. 8p, 2020.
- [25] D. Leite, G. Andonovski, I. Škrjanc, and F. Gomide, "Optimal rule-based granular systems from data streams," *IEEE Transactions on Fuzzy Systems*, 2019. doi: 10.1109/TFUZZ.2019.2911493.
- [26] W. Pedrycz and F. Gomide, *An Introduction to Fuzzy Sets: Analysis and Design*. NetLibrary, 01 2000.
- [27] D. Leite and I. Škrjanc, "Ensemble of evolving optimal granular experts, owa aggregation, and time series prediction," *Inf. Sci.*, vol. 504, pp. 95–112, 07 2019.
- [28] W. Pedrycz, *Granular Computing : An Introduction*, vol. 45, pp. 309–328. Springer, 01 2000.
- [29] R. R. Yager, "Measures of specificity," in *Computational Intelligence: Soft Computing and Fuzzy-Neuro Integration with Applications* (O. Kaynak, L. A. Zadeh, B. Türkşen, and I. J. Rudas, eds.), (Berlin, Heidelberg), pp. 94–113, Springer Berlin Heidelberg, 1998.
- [30] E. A. Soares, H. A. Camargo, S. J. Camargo, and D. F. Leite, "Incremental gaussian granular fuzzy modeling applied to hurricane track forecasting," *2018 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, p. 8p, 2018.
- [31] P. Qiu, *Introduction to Statistical Process Control*. Wiley: India, 2014.