

Received December 21, 2021, accepted January 22, 2022, date of publication February 16, 2022, date of current version March 29, 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3152206

Low-Rank Methods in Event Detection With Subsampled Point-to-Subspace Proximity Tests

JAKUB MARECEK¹, (Member, IEEE), STATHIS MAROULIS², VANA KALOGERAKI², AND DIMITRIOS GUNOPOULOS³

¹Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague, 121 35 Prague, Czech Republic

²Department of Informatics, Athens University of Economics and Business, 104 34 Athens, Greece

³Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, 157 84 Athens, Greece

Corresponding author: Jakub Marecek (jakub.marecek@fel.cvut.cz)

The work of J. Marecek was supported by the Operational Programme Research, Development and Education (OP RDE) Project “Research Center for Informatics” under Grant CZ.02.1.01/0.0/0.0/16_019/0000765. The work of D. Gunopoulos and V. Kalogeraki was supported in part by the EU H2020 TAILOR Project under Grant 952215.

ABSTRACT Monitoring of streamed data to detect abnormal behaviour (variously known as event detection, anomaly detection, change detection, or outlier detection) underlies many applications, especially within the Internet of Things. There, one often collects data from a variety of sources, with asynchronous sampling, and missing data. In this setting, one can detect abnormal behavior using low-rank techniques. In particular, we assume that normal observations come from a low-rank subspace, prior to being corrupted by a uniformly distributed noise. Correspondingly, we aim to recover a representation of the subspace, and perform event detection by running point-to-subspace distance query for incoming data. We use a variant of low-rank factorisation, which considers interval uncertainty sets around “known entries”, on a suitable flattening of the input data to obtain a low-rank model. On-line, we compute the distance of incoming data to the low-rank normal subspace and update the subspace to keep it consistent with the seasonal changes present. For the distance computation, we consider subsampling. We bound the one-sided error as a function of the number of coordinates employed. In our computational experiments, we test the proposed algorithm on induction-loop data from Dublin, Ireland.

INDEX TERMS Multidimensional signal processing, monitoring, matrix completion, point-to-subspace proximity, probably approximately correct learning.

I. INTRODUCTION

When detailed multivariate data are available in real time, it is highly desirable to monitor the appearance of “abnormal” behavior across the multivariate data, with guarantees on the performance of the monitoring procedure, but without the computational burden of processing the data set in its entirety. Across the Internet of Things, many examples abound [1]. To consider one example, many cities have been instrumented with large numbers of sensors capturing the numbers and average speeds of cars passing through the approaches of urban intersections (induction loops), volume of traffic (from CCTV data or aggregate data of mobile-phone operators), and speeds of public transport vehicles (e.g., on-board satellite positioning units in buses), but many still lack the infrastructure to detect traffic accidents prior to them being reported. This is to a large extent due to the limited utility of the

information from each of the sensors, e.g., maintaining statistics about traffic at a particular approach of an intersection. Only the combination of multivariate time series across multiple sensor types could allow the detection of events of interest in many applications.

More broadly, there is a monitoring component in most applications of the Internet of Things. In transportation applications, one may wish to detect traffic accidents [2]–[5], (imminent) aircraft engine failures [6], or deviations from a flight schedule [7]. In electric power distribution systems [5], [8]–[10], there may be reclosers and sectionalisers acting automatically upon a tree branch falling on an overhead power line, but the distribution system operator may not know about the event, until it is either detected from sparsely-deployed sensors or reported by customers. Similar techniques can be used [11] in water distribution networks. Likewise, Internet of Things (IoT) in manufacturing [12], [13] and environmental applications [14] crucially relies on monitoring, as does intrusion detection in IoT [15]–[17],

The associate editor coordinating the review of this manuscript and approving it for publication was Amin Zehtabian^{id}.

albeit the details of the model tend to be more involved and more application-specific. Correspondingly, there is a long history of work on monitoring and event detection (also known as anomaly detection or outlier detection), going back at least to [18], [19] in the univariate case. Outside of traditional methods, such as dimension reduction [20] and Gaussian processes [9], deep-learning methods [21]–[24] have been widely used recently. We refer to [25], [26] for excellent surveys.

Notice that processing heterogeneous sensor data in IoT applications poses several challenges: (1) One of the main challenges is, clearly, dealing with the velocity and, when accumulated, the volume of the data. A city can have thousands of sensors sampling at kHz rates. For example, in a network of 10, 000 sensors, sampling with 1-byte resolution at 1 kHz, one obtains close to 311 TB of data per year that needs to be analyzed to estimate what is normal. (2) The second challenge involves detecting an event in real-time. An automated event detection is useful in cases that the event is detected within seconds after it occurs, such as when a road is completely blocked before people start venting their frustration on social media or dialling rescue services. (3) Another common challenge is the missing values and failures of sensors. It is widespread for sensors to stop working or start reporting wrong values (*e.g.*, negative car flow). Distinguishing the mal-function of a single sensor from a genuine event shows the necessity of utilising multivariate data. (4) Finally, there is measurement noise. In field conditions, *e.g.*, an induction loop buried under inches of tarmac, or a traffic-volume estimate from a video feed captured in a rainstorm, does have a very limited accuracy. While there are methods for dealing with each of these challenges in isolation, one should like to address all four at the same time.

To overcome these challenges, we propose a novel framework that utilizes low-rank methods [27] to provide fast and accurate event detection on data from varied sources. Throughout, we consider uniformly-distributed measurement noise, but let us present the model in the noise-free case first in this paragraph. There, events correspond to points lying outside a certain subspace. To estimate the sub-space, we flatten the input data to a matrix and apply state-of-the-art low-rank matrix-factorization techniques. In particular, we factorize the original matrix into two smaller matrices, whose product approximates the original matrix. Subsequently, we develop a point-in-subspace membership test capable of detecting whether new samples are within the subspace spanned by the columns of one of the factors (smaller matrices). An affirmative answer is interpreted as an indication that the samples from the sensors present normal behavior. In the case of a negative answer, a point-to-subspace distance query can estimate the extent of abnormality of an event. Crucially, this point-in-subspace membership test can be sub-sampled, while still allowing for guarantees on its performance. The sub-sampling of, *e.g.*, one per cent of the data, allows for efficient applicability in IoT applications.

Our main contributions are the following:

- a general framework for representing what is an event and what is a non-event considering *heterogeneous data*, which are possibly *not sampled uniformly*, with *missing values* and *measurement errors*.
- a novel randomized event detection technique, implemented via a *point-to-subspace distance query*, with guarantees of *probably approximately correct* (PAC) learning [28] in *sublinear time*,
- an experimental evaluation on data from a traffic-control system in Dublin, Ireland, which shows that it is possible to process data collected from thousands of sensors over the course of one year within minutes, to answer point-to-subspace distance queries in milliseconds and thus detect even hard-to-detect events.

II. AN APPROACH

Our goal in this paper is to build a model of what is a non-event across many time series, possibly with *non-uniform sampling* across the time series, *missing values*, and *measurement errors* present in the values. We build a framework around this model and, in Section III, suggest algorithms for the individual components in this framework.

A. A MODEL

For example, one could consider applications in urban traffic management, where the number of vehicles passing over induction loops are measured, but often prove to be noisy, with the reliability of the induction loops and the related communication infrastructure limited. Subsequently, we aim at an online event detection mechanism, which would be able to decide whether multiple fragments of multiple incoming time-series present an event (abnormal behaviour) or not. In urban traffic management, for example, one aims at detecting a road accident, based on the evolution of the traffic volumes across a network of induction loops. Notice that an accident will manifest itself by some readings being low, due to roads being blocked, while other readings are high, due to re-routing, while no induction loop has to have its readings more than one standard deviation away from the long-run average, which renders univariate methods difficult to use. Such monitoring problems are central to many Internet-of-Things applications.

This pattern can be exploited by storing each day worth of data as a row in a matrix, possibly with many missing values. For multiple time series, we obtain multiple partial matrices, or a partial tensor. These can be flattened by concatenating the matrices row-wise to obtain one large matrix, as suggested in Figure 1. For D days discretised to T periods each, with up to S sensors available, the flattened matrix M is in dimension $n = TS$ and has $m = D$ rows.

Considering this flattened representation, it is natural to assume that each new day resembles a linear combination of r prototypical days, or rows in the flattened matrix in dimension $m \gg r$. Formally, we assume that there exists

$R \in \mathbb{R}^{r \times n}$, such that our observations $x \in \mathbb{R}^n$ are

$$x = cR + \mathcal{U}(-\Delta, \Delta), \quad (1)$$

possibly with many values missing, for some coefficients $c \in \mathbb{R}^r$ weighing the r vectors $\{e_1, e_2, \dots, e_r\}$ row-wise in R , with uniformly-distributed noise \mathcal{U} between $-\Delta$ and Δ .

We compute the matrix R , using low-rank approximation of the flattened matrix with an explicit consideration of the uniformly-distributed error in the measurements M_{ij} for $(i, j) \in M$. Considering the interval uncertainty set $[M_{ij} - \Delta, M_{ij} + \Delta]$ around each observation, this can be seen as matrix completion with element-wise lower bounds $X_{ij}^L := M_{ij} - \Delta$ for $(i, j) \in M$ and element-wise upper bounds $X_{ij}^U := M_{ij} + \Delta$ for $(i, j) \in M$.

Considering the factorization LR , where $L \in \mathbb{R}^{m \times r}$ and $R \in \mathbb{R}^{r \times n}$ obtained above (6), given an incoming $x \in \mathbb{R}^n$, the maximum likelihood estimate $\hat{c} \in \mathbb{R}^r$ of c in (1) is precisely the point minimizing $|x|_\infty$:

$$\min_{\hat{c} \in \mathbb{R}^r} \max_i |x_i - (\hat{c}R)_i| \quad (2)$$

whenever $\hat{c} \leq \Delta$. We refer to Section 7.1.1 of [29] for a discussion. In a linear program corresponding to (2), we consider a subset of coordinates of \mathbb{R}^n and prove a bound on the one-sided error when using the subset. This is the first use of a point-to-subspace query considering the supremum norm (ℓ^∞) in event detection.

B. THE MODULOR FRAMEWORK

This naturally leads to a framework comprising three main components, as illustrated in Figure 1:

- 1) Data flattener, which captures the unstructured raw input data coming from different sources and reformats them into a partial matrix, which is processed further.
- 2) Matrix-factorization component, which approximates the partial matrix obtained by the data flattener. The approximation consists of two matrices (known as factors), whose product is a low-rank approximation of the original one. Using the two factors, we are able to capture the most salient features of the original matrix in a compressed form. Sparsity of the partial matrix on the input makes the calculation of the matrix factorization possible with high accuracy within modest run time, while allowing for missing values in the input data. This component is further described in Section III-A.
- 3) Subsampled point-to-subspace proximity tester (or subspace-proximity tester for short) uses the output of the matrix-factorization component and estimates whether the current sensor readings present an abnormal behavior or not. This component is described in detail in Section III-B, but crucially, its run time is independent of the dimension. In experimental results with a history of sensor readings encoded in a partial matrix in $\mathbb{R}^{304 \times 299430}$ and current sensor readings encoded in a vector in \mathbb{R}^{299430} , for instance, it takes

only milliseconds to perform the test, as we illustrate in of Section V.

The framework can be utilised as follows: Data flattener collects all data from different sources and creates the corresponding matrices, e.g., a partial matrix with traffic volumes and speeds. Then the data structure thus produced is passed to the matrix-factorization component, which factorizes the data, and creates two matrices L and R . One of the factors (matrix R) is then passed to the subspace-proximity tester, which uses it to assess whether incoming sensor readings present abnormal behavior or not and report the results to the end user. Finally, subspace-proximity tester relays the data back to the matrix-factorization component to update the input matrix, replacing the oldest data present, and updating online [30]–[32], if needed.

We denote this framework MODULoR, where this backronym can stand for “MONitoring Distributed systems Using Low-Rank methods“, or more accurately as a “Method for Outlier Detection Using Low-Rank factorization and range-space subsampling”. We stress that the novelty lies in the subsampled point-to-subspace proximity tester, whose low run time makes the use of low-rank factorization practical. Without the subsampling, the point-to-subspace distance query in such an approach [30], [32]–[36] would be too demanding for online use.

III. THE ALGORITHMS

As outlined above, there are two key algorithms needed. The first one implements the matrix-factorization component. In our experiments, we chose the alternating parallel coordinate descent for *inequality-constrained matrix completion* to estimate the low-rank approximation of a partial matrix, either in an online or offline fashion. This makes it possible, rather uniquely, to be robust to uniformly-distributed measurement noise, while being able to detect sparse noise as abnormal (events, anomalies).

The second algorithm implements the subspace-proximity tester. In our experiments, we consider the test with the supremum norm, implemented as a *linear program*, which is subsampled. As an input, it uses the output of the matrix-factorization component and it is able to predict if an incoming time series presents normal or abnormal behavior. This second algorithm is run in an online fashion. We describe the two algorithms in more detail in the following two sections.

A. MATRIX-FACTORIZATION COMPONENT

To formalise the factorisation $M \approx LR$, let L_i and R_j be the i -th row and j -th column of L and R , respectively. With Frobenius-norm regularisation, the factorization problem we wish to solve reads:

$$\min_{L \in \mathbb{R}^{m \times r}, R \in \mathbb{R}^{r \times n}} f_{\mathcal{L}}(L, R) + f_{\mathcal{U}}(L, R) + \frac{\mu}{2} \|L\|_F^2 + \frac{\mu}{2} \|R\|_F^2 \quad (3)$$

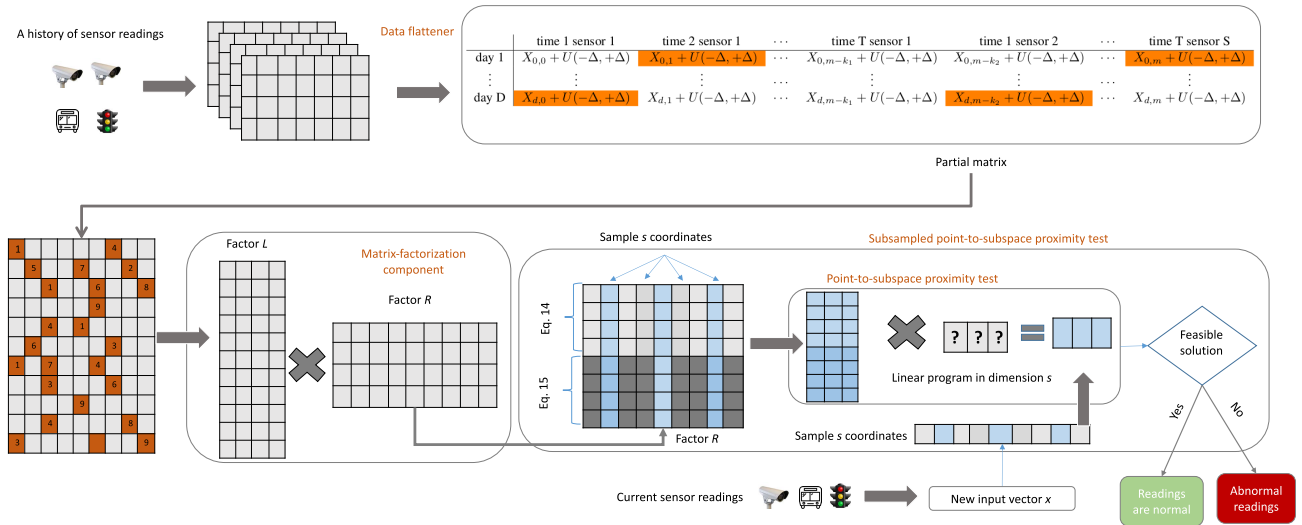


FIGURE 1. A schematic illustration of the MODULO-R framework: a history of sensor readings is processed into a partial matrix, which is factorized. One of the factors is subsampled and the corresponding subsampling is applied also to incoming sensor readings. This makes it possible to run a point-to-subspace test in an online fashion, with a constant run time and a small one-sided error.

Algorithm 1 Matrix Factorization via Alternating Parallel Coordinate Descent, cf. [27]

Input: $\mathcal{E}, \mathcal{L}, \mathcal{U}, X^{\mathcal{E}}, X^{\mathcal{L}}, X^{\mathcal{U}}, \text{rank } r$
Output: $m \times n$ matrix

- 1: choose $L \in \mathbb{R}^{m \times r}$ and $R \in \mathbb{R}^{r \times n}$
- 2: **for** $k = 0, 1, 2, \dots$ **do**
- 3: choose a random subset $\hat{S}_{row} \subset \{1, \dots, m\}$
- 4: **for** $i \in \hat{S}_{row}$ **in parallel do**
- 5: choose $\hat{r} \in \{1, \dots, r\}$ uniformly at random
- 6: compute $\delta_{i\hat{r}}$ using formula (11)
- 7: update $L_{i\hat{r}} \leftarrow L_{i\hat{r}} + \delta_{i\hat{r}}$
- 8: **end for**
- 9: choose a random subset $\hat{S}_{column} \subset \{1, \dots, n\}$
- 10: **for** $j \in \hat{S}_{column}$ **in parallel do**
- 11: choose $\hat{r} \in \{1, \dots, r\}$ uniformly at random
- 12: compute $\delta_{\hat{r}j}$ using (12)
- 13: update $R_{\hat{r}j} \leftarrow R_{\hat{r}j} + \delta_{\hat{r}j}$
- 14: **end for**
- 15: **end for**
- 16: **return** (L, R)

where

$$f_{\mathcal{L}}(L, R) := \frac{1}{2} \sum_{(ij) \in \mathcal{L}} (X_{ij}^{\mathcal{L}} - L_i R_j)^2_+, \quad (4)$$

$$f_{\mathcal{U}}(L, R) := \frac{1}{2} \sum_{(ij) \in \mathcal{U}} (L_i R_j - X_{ij}^{\mathcal{U}})^2_+, \quad (5)$$

where $\xi_+ = \max\{0, \xi\}$, calligraphic \mathcal{L} is used for bounds from below and \mathcal{U} for bounds from above. Notice that this is a non-convex problem, whose special case of $\Delta = 0$ is NP-hard [37], [38].

The matrix completion under interval uncertainty can be seen as a special case of the inequality-constrained matrix completion of [27]:

$$\min\{f(L, R) : L \in \mathbb{R}^{m \times r}, R \in \mathbb{R}^{r \times n}\}, \quad (6)$$

where

$$f(L, R) := f_{\mathcal{E}}(L, R) + f_{\mathcal{L}}(L, R) + f_{\mathcal{U}}(L, R) + \frac{\mu}{2} \|L\|_F^2 + \frac{\mu}{2} \|R\|_F^2 \quad (7)$$

$$f_{\mathcal{E}}(L, R) := \frac{1}{2} \sum_{(ij) \in \mathcal{E}} (L_i R_j - X_{ij}^{\mathcal{E}})^2, \quad (8)$$

$$f_{\mathcal{L}}(L, R) := \frac{1}{2} \sum_{(ij) \in \mathcal{L}} (X_{ij}^{\mathcal{L}} - L_i R_j)^2_+, \quad (9)$$

$$f_{\mathcal{U}}(L, R) := \frac{1}{2} \sum_{(ij) \in \mathcal{U}} (L_i R_j - X_{ij}^{\mathcal{U}})^2_+, \quad (10)$$

where for $(i, j) \in \mathcal{U}$ we have an element-wise upper bound $X_{ij}^{\mathcal{U}}$, for $(i, j) \in \mathcal{L}$ we have an element-wise lower bound $X_{ij}^{\mathcal{L}}$, for $(i, j) \in \mathcal{E}$ we know the exact value $X_{ij}^{\mathcal{E}}$, and $\xi_+ = \max\{0, \xi\}$.

A popular heuristic for matrix completion considers a product of two matrices, $X = LR$, where $L \in \mathbb{R}^{m \times r}$ and $R \in \mathbb{R}^{r \times n}$, obtaining $X = LR$ of rank at most r , cf. [39]. In particular, we use a variant of the alternating parallel coordinate descent method for matrix completion introduced by [27] under the name of ‘‘MACO’’, summarized in Algorithm 1. It is based on the observation that while f is not convex jointly in (L, R) , it is convex in L for fixed R and in R for fixed L . In Steps 3–8 of the algorithm, we fix R , choose random \hat{r} and a random set \hat{S}_{row} of rows of L , and update, in parallel, for $i \in \hat{S}_{row}$: $L_{i\hat{r}} \leftarrow L_{i\hat{r}} + \delta_{i\hat{r}}$. Following [27], we use

$$\delta_{i\hat{r}} := -\langle \nabla_L f(L, R), E_{i\hat{r}} \rangle / W_{i\hat{r}}, \quad (11)$$

where the computation of $\langle \nabla_L f(L, R), E_{i\hat{r}} \rangle$ can be simplified as suggested in Figure 2. In Steps 9–14, we fix L , choose random \hat{r} and a random set \hat{S}_{column} of columns of R , and update, in parallel for $j \in \hat{S}_{column}$: $R_{\hat{r}j} \leftarrow R_{\hat{r}j} + \delta_{\hat{r}j}$.

$$\delta_{\hat{r}j} := -\langle \nabla_R f(L, R), E_{\hat{r}j} \rangle / V_{\hat{r}j}, \quad (12)$$

where the computation of $\langle \nabla_R f(L, R), E_{\hat{r}j} \rangle$ can, again, be simplified as suggested in Figure 2.

$$\begin{aligned} \langle \nabla_L f(L, R), E_{i\hat{r}} \rangle &= \mu L_{i\hat{r}} + \sum_{j: (ij) \in \mathcal{E}} (L_{i:R:j} - X_{ij}^{\mathcal{E}}) R_{\hat{r}j} + \sum_{j: (ij) \in \mathcal{U}L_{i:R:j} < X_{ij}^{\mathcal{U}}} (L_{i:R:j} - X_{ij}^{\mathcal{U}}) R_{\hat{r}j} + \sum_{j: (ij) \in \mathcal{L}L_{i:R:j} > X_{ij}^{\mathcal{L}}} (L_{i:R:j} - X_{ij}^{\mathcal{L}}) R_{\hat{r}j} \\ \langle \nabla_R f(L, R), E_{\hat{r}j} \rangle &= \mu R_{\hat{r}j} + \sum_{i: (ij) \in \mathcal{E}} (L_{i:R:j} - X_{ij}^{\mathcal{E}}) L_{i\hat{r}} + \sum_{i: (ij) \in \mathcal{L}L_{i:R:j} < X_{ij}^{\mathcal{L}}} (L_{i:R:j} - X_{ij}^{\mathcal{L}}) L_{i\hat{r}} + \sum_{i: (ij) \in \mathcal{U}L_{i:R:j} > X_{ij}^{\mathcal{U}}} (L_{i:R:j} - X_{ij}^{\mathcal{U}}) L_{i\hat{r}}. \end{aligned}$$

FIGURE 2. Multiplication by the gradients in (11) and (12) of Algorithm 1 can be simplified considerably.

We should also like to comment on the choice of Δ and ϵ . A sensible approach seems to be based on cross-validation: out of the historical data (or out of L), one can pick one row, and compute the Δ needed. The maximum of Δ for any row seems to be a good choice. We refer to [27] for a discussion of the choice of the parameter $\mu > 0$.

B. SUBSAMPLED POINT-TO-SUBSPACE PROXIMITY TESTER

As suggested previously, instead of computing the distance of an incoming time-series to each one of those already available per-day time-series, classified as event or non-event, we consider a point-to-subspace query in the infinity norm:

$$\min_{\hat{c} \in \mathbb{R}^r} \max_i |x_i - (\hat{c}R)_i|, \quad (13)$$

and test whether the distance (13) is less than or equal to Δ . As we described in Section II for uniform noise, the supremum norm (ℓ^∞) gives the maximum likelihood estimate. The infinity norm is sometimes seen as difficult to work with, due of the lack of differentiability. However, note that it (13) can be recast as a test of the feasibility of a linear programming problem:

$$\min_{\hat{c} \in \mathbb{R}^r} 1 \quad \text{s.t. } x_i - (\hat{c}R)_i \leq \Delta, \quad (14)$$

$$(\hat{c}R)_i - x_i \leq \Delta. \quad (15)$$

Alternatively, this is an intersection of hyperplanes, also known as a hyper-plane arrangement. As we will show in the following section, this geometric intuition is useful in the analysis of the algorithms.

In Algorithm 2 we present a test, which considers only a subset S , $|S| \ll n$ of coordinates, picked uniformly at random. As we show in the following section, this test has only a modest one-sided error.

IV. AN ANALYSIS

Before we present the main result, let us remark on the convergence properties of Algorithm 1, which has been proposed and analyzed by [27] and [32]. A simple convergence result of [27] states that the method is monotonic and, with probability 1, $\lim_{k \rightarrow \infty} \inf \|\nabla_L f(L^{(k)}, R^{(k)})\| = 0$, and $\lim_{k \rightarrow \infty} \inf \|\nabla_R f(L^{(k)}, R^{(k)})\| = 0$. This applies in our case as well. See [32] for details of the rate of convergence.

Our main analytical result concerns the statistical performance of the point-to-subspace query. Informally, the randomized point-to-subspace distance query in Algorithm 2 has one-sided error: If the distance between the vector x and $\text{span}(R)$ is no more than Δ in ℓ^∞ , we never report otherwise. If, however, the distance actually is more than Δ in ℓ^∞ ,

Algorithm 2 Subsampled Point-to-Subspace Proximity Tester With the Supremum Norm

Input: $R \in \mathbb{R}^{r \times n}$, $x \in \mathbb{R}^n$, $s, \Delta \in \mathbb{R}$

Output: true/false

- 1: choose $S \subset \{1, \dots, n\}$, $|S| = s$, uniformly at random
- 2: initialise a linear program P in variable $v \in \mathbb{R}^s$
- 3: **for** $i \in S$ **do**
- 4: add constraint $x_i - (\text{proj}_S(L)v)_i \leq \Delta$
- 5: add constraint $x_i - (\text{proj}_S(L)v)_i \geq -\Delta$
- 6: **end for**
- 7: **if** $\exists v \in \mathbb{R}^s$ such that the constraints are satisfied **then**
- 8: **return true**
- 9: **else**
- 10: **return false**
- 11: **end if**

considering only a subset S of coordinates may ignore a coordinate where the distance is larger, and hence mis-report that the vector is within distance Δ in ℓ^∞ , with a certain probability, depending on the number of constraints that are actually violated. For example, to achieve the one-sided error of ϵ with probability of 1/3 or less, this test needs to solve a linear program in dimension $O(\frac{r \log r}{\epsilon} \log \frac{r \log r}{\epsilon})$. Notice that this bound is independent of the ‘‘ambient’’ dimension n .

Formally:

Theorem 1: (i) When the distance (13) is $D \leq \Delta$, Algorithm 2 never reports the point is outside the sub-space. (ii) When the distance (13) is $D > \Delta$, because there are en coordinates i such that for all \hat{c} , there is $|x_i - (\hat{c}R)_i| \geq \Delta$, then for any $\delta \in (0, 1)$, when Algorithm 2 considers s coordinates

$$O\left(\frac{1}{\epsilon} \log \frac{1}{\delta} + \frac{r \log r}{\epsilon} \log \frac{r \log r}{\epsilon}\right)$$

sampled independently uniformly at random, the point is inside the subspace with probability $1 - \delta$.

To prove this theorem, we need several definitions from discrete geometry. We follow the notation of [40] and [41] and use calligraphic fonts to distinguish from symbols in the remainder of the text (e.g. S from \mathcal{S}).

Definition 2 (Range Space of [40]): A range space \mathcal{S} is a pair $(\mathcal{X}, \mathcal{R})$, where \mathcal{X} is a set and \mathcal{R} is a family of subsets of \mathcal{X} , $\mathcal{R} \subset 2^{\mathcal{X}}$. Members of \mathcal{X} are called elements or points of \mathcal{S} and members of \mathcal{R} are called ranges of \mathcal{S} . \mathcal{S} is finite if \mathcal{X} is finite.

Notice that the range space is a (possibly infinite) hypergraph.

Definition 3 (Shattering of [40]): Let $\mathcal{S} = (\mathcal{X}, \mathcal{R})$ be a range space and let $\mathcal{A} \subset \mathcal{X}$ be a finite set. Then $\Pi_{\mathcal{R}}(\mathcal{A})$

denotes the set of all subsets of \mathcal{A} that can be obtained by intersecting \mathcal{A} with a range of \mathcal{S} . If $\Pi_{\mathcal{R}}(\mathcal{A}) = 2^{\mathcal{A}}$, we say that \mathcal{A} is shattered by \mathcal{R} .

Definition 4 (Dimension of [40]): The Vapnik-Chervonenkis dimension of \mathcal{S} is the smallest integer d such that no $\mathcal{A} \subset \mathcal{X}$ of cardinality $d + 1$ is shattered by \mathcal{R} . If no such d exists, we say the dimension of \mathcal{S} is infinite.

Definition 5 (ϵ -net of [42]): An ϵ -net of a finite subset of points $P \subseteq \mathcal{X}$ is a subset $\mathcal{N} \subseteq P$ such that any range $\nabla \in \mathcal{R}$ with $|\nabla \cap P| \geq \epsilon|P|$ has a non-empty intersection with \mathcal{N} .

With these definitions, recall the celebrated result:

Theorem 6 ([41]–[44]): Let $(\mathcal{X}, \mathcal{R})$ be a range space of Vapnik-Chervonenkis dimension d . Let $\epsilon, \delta \in (0, 1)$. If \mathcal{S} is a set of

$$O\left(\frac{1}{\epsilon} \log \frac{1}{\delta} + \frac{d}{\epsilon} \log \frac{d}{\epsilon}\right)$$

points sampled independently from a finite subset of \mathcal{X} , then \mathcal{S} is an ϵ -net for the finite subset with probability at least $1 - \delta$.

Now, let us present the proof of Theorem 1.

Proof: To see (i), consider the linear program constructed in Algorithm 2 and notice that its constraints are a subset of those in (14). If (14) is feasible, then any subset of constraints will be feasible. To see (ii), we show that a set related to the polyhedron of feasible x has a small Vapnik-Chervonenkis (VC) dimension, which makes it possible to apply classical results from discrete geometry presented above. In particular, we proceed in four steps:

- 1) denote by \mathcal{S}_1 the range space for all possible constraints added in Line 4 and by \mathcal{S}_2 the range space for all possible constraints added in Line 5.
- 2) The VC dimension of each of $\mathcal{S}_1, \mathcal{S}_2$ is at most $r + 1$.
- 3) The VC dimension of $\mathcal{S}_1 \cup \mathcal{S}_2$ is $O(r \log r)$.
- 4) Subsequently, we apply Theorem 6 of [41]–[44].

A. STEP 1

The range spaces \mathcal{S}_1 and \mathcal{S}_2 will share the same set of points, namely $[n] := 1, 2, \dots, n$, and feature very similar ranges: \mathcal{S}_1 will feature the hyperplanes $x_i - (\hat{c}\mathcal{R})_i \leq \Delta$ corresponding to the first set of constraints in the LP (14), while \mathcal{S}_2 will feature the hyperplanes $(\hat{c}\mathcal{R})_i - x_i \leq \Delta$. We keep them separate, so as to allow for the hyperplanes to be in a generic position.

Alternatively, one could construct a single range space, with the same set of points and ranges given by the subspaces given by the intersections of $x_i - (\hat{c}\mathcal{R})_i \leq \Delta$ and $(\hat{c}\mathcal{R})_i - x_i \leq \Delta$ for $i \in [n]$. This would, however, complicate the analysis, somewhat.

B. STEP 2

The VC dimension of each of $\mathcal{S}_1, \mathcal{S}_2$ is at most $r + 1$. For range spaces, where the ranges are hyper-planes, this is a standard result. We refer to Section 15.5.1 of [45] for a very elegant proof using Radon’s theorem. Notice that r would suffice, if there were no vertical hyperplanes.

C. STEP 3

The VC dimension of $\mathcal{S}_1 \cup \mathcal{S}_2$ is $O(r \log r)$. This follows by the counting of the possible ranges and Sauer-Shelah lemma, a standard result. We refer to Lemma 15.6 in [45].

D. STEP 4

The intuition is that if there is a large-enough subset, a large-enough random sample will intersect with it. The surprising part of Theorem 6 on the existence of ϵ -nets is that the bound of the large-enough does not depend on the number of points of the ground set, but only on the VC dimension established above. In particular, we sample coordinates $S, |S| = s$ in Line 1. This corresponds to sampling from \mathcal{X} in $\mathcal{S}_1 \cup \mathcal{S}_2$. Because we assume there are ϵn coordinates i such that such that for all \hat{c} , there is $|x_i - (\hat{c}\mathcal{R})_i| \geq \Delta$, an ϵ -net will intersect these by Theorem 6. \square

Next, let us consider the run-time of Algorithm 2, which is dominated by the feasibility test of a linear program P in Line 7. Using standard interior-point methods [46], if there is a feasible solution to the linear program P , an ϵ -accurate approximation to the can be obtained in $O(\sqrt{s} \ln(1/\epsilon))$ iterations, wherein each iteration amounts to solving a linear system. This yields an upper bound on the run-time of

$$O\left(\frac{r^{3.5} \log^{3.5} r}{\epsilon^{3.5}} \log^{3.5} \frac{r}{\epsilon}\right),$$

which could be improved considerably by exploiting the sparsity in the linear program’s constraint matrix. The same iterations make it possible to detect infeasibility using the arguments of [47], although the homogeneous self-dual approach of [48] with a worse iteration complexity may be preferable in practice. Either way, a solver generator [49], [50] allows for excellent performance.

Alternatively, however, one may consider:

Theorem 7: There is an algorithm that can pre-process a sample of s coordinates such that the point-in-subspace membership query can be answered in time $O(\log s)$ in the worst case. The expected run-time of the pre-processing is $O(s^{r+\epsilon})$, $\epsilon \geq 0$, where the expectation is with respect to the random behaviour of the algorithm, and remains valid for any input.

Proof (Sketch): Notice that one can replace the test of feasibility of a linear program P with a point-location problem in a hyperplane arrangement. We refer to [51], [52] for a very good introduction to hyperplane arrangements, but to provide an elementary intuition: An alternative geometric view of Algorithm 2 is that we have a subspace $P \subseteq \mathbb{R}^s$, initialise $P = \mathbb{R}^s$ in Line 2, and then intersect it with hyperplanes on Lines 4–5. Equally well, one may consider a hyper-plane arrangement P , initialise it to an empty set in Line 2, and then add hyperplanes on Lines 4–5. Our goal is not to optimise a linear function over P , but rather to decide whether there exists a point within P , the intersection of the hyperplanes, which corresponds to one cell of the arrangement. The actual

result follows from the work of [43], [44] on hyperplane arrangements. \square

While the use of solver-generator [49], [50] may be preferable in many IoT applications, there may be large-scale use cases, where the asymptote of the run-time of the algorithm of [44] does matter and the sampling of the coordinates may be reused.

V. EXPERIMENTAL EVALUATION

To evaluate our approach, we have implemented our matrix-factorization component in Apache Spark [53], in order to ensure its scalability, and the subspace proximity tester in Python, using Numpy [54] for numerical linear algebra and `multiprocess` for parallel processing. Our open-source code is available at <https://github.com/stathisq/modulor>.

A. THE DATA

To validate the ability of our proposed method to detect events, we evaluated it on both synthetic and real-world datasets. Considering the limitations of the benchmarks in the literature [55], we used data from traffic monitoring collected by the Sydney Coordinated Adaptive Traffic System (SCATS) system of Dublin City Council (DCC) from intersections in Dublin, Ireland, between January 1 and November 30, 2017. Therein, each time series is obtained by one induction loop at an approach to an intersection with sensors at stop-lines and irregular intervals from stop-lines. Overall, our data contains readings from 3432 such sensors, distributed across the city. To use a realistic data set, reflecting the asynchronous operations of the system, we record the samples as they arrive asynchronously and do not impute any missing values. In particular, each intersection operates asynchronously, with all predefined phases changing, in turn, within a cycle time varying between 50 and 120 seconds both across the intersections and over time. Whenever an intersection's cycle time finishes, we record the flow over the cycle time. Within any given period, e.g., 2 minutes, we receive vehicle count data from only a fraction of the sensors. For each day, we consider data between 7 a.m. and 10 p.m., which are of particular interest to traffic operators.

Altogether, the data from 3,432 sensors recorded with sampling period of 2 minutes, or shorter, are flattened to a partial matrix $X \in \mathbb{R}^{304 \times 299,430}$, where there were 38,767,895 zeros out of the 91,026,720 elements, representing 42% sparsity. This is due to a large part to the asynchronicity of the sensor readings, and to a lesser part due to actual sensor failures. To evaluate our approach, we have created several matrices from X : matrix Y with a small amount of noise, which represents normal behaviour, and matrix G with additional noise, which represents events. We have repeated this process in a repeated six-fold cross validation (out-of-sample testing) and we report the mean and standard deviation of the performance measures across the six runs.

In each run, using rows of matrix X , we have created matrices Y and Y' in the following way. First, we have constructed the matrix $Y \in \mathbb{R}^{1,200 \times 299,430}$ representing normal behaviour

in several steps. In the first step, each row of Y has been initialised with one row sampled uniformly at random (with repetition) from the 304 rows of matrix X . In the second step, we have multiplied each row with a random scalar sampled (independently) from the uniform distribution over $(0, 2)$. In the third step, we have applied a perturbation by an independently identically uniformly distributed noise on $[-0.8\Delta, 0.8\Delta]$. Thus constructed matrix Y represents 1,200 time series of normal behaviour. Next, we have introduced events Y , obtaining $Y' \in \mathbb{R}^{1,200 \times 299,430}$, or rather five variants thereof. In particular, we have sampled 200 rows of matrix Y uniformly at random to create G , which is a $200 \times 299,430$ submatrix of Y . From G , we have created five variants of $G' \in \mathbb{R}^{200 \times 299,430}$ by the addition of Gaussian noise with mean $\mu = 5, 10, 15, 20, 25$ and standard deviation equal to one half of the mean. This corresponds to the peak signal-to-noise ratio (PSNR) of 23.58, 19.97, 16.94, 14.71, and 12.84 dB, respectively, when averaged over the six runs, where PSNR is the ratio between the maximum possible power of a signal and the power of the corrupting noise introduced, that is

$$\text{PSNR}(G, G') = 20 \log_{10}(\max_{ij}(\{G_{ij}\})/\alpha(G, G'))$$

for root mean square error

$$\alpha(G, G') = \|G' - G\|_F = \sqrt{\sum_{i=1}^{200} \sum_{j=1}^{299,430} (G'_{ij} - G_{ij})^2},$$

where G is the $200 \times 299,430$ submatrix of Y that G' is based on. Subsequently, we have worked with a variant Y' of Y , wherein the submatrix G is replaced by G' . Our training data were 1,000 rows of this new matrix Y' chosen uniformly at random and we left the remaining 200 rows as the ground truth for testing. Using the left-out 200 rows, we have evaluated our model with respect to recall, precision, and the so-called F1 score, which is a harmonic mean of the former two measures.

B. THE RESULTS

Figure 3(a) presents the performance of our matrix-factorization component, while Figure 3(b) presents a trade-off between time required for training and reconstruction error in the choice of the rank r . Notice that the reconstruction error is the usual extension of the root mean square error (RMSE) to the matricial setting, i.e., the Frobenius norm of the difference between the matrix Y and the product LR . It is clear that increasing the rank above 10 leads to marginal improvements in the reconstruction error, but increasing it above 40 leads to a sharp increase in the training time. We chose $r = 10$ for our experiments.

Figure 4(a) compares the readings of sensors from the non-event matrix Y with events in G , while omitting zeros. We can observe that G with $\mu = 5$ is hard to distinguish from Y . Figure 4(b) presents the distribution of the values of the samples used for training: the average values of normal samples we




FIGURE 3. Performance of the matrix-factorization component on the instance of Section V-A. Figure 3(a) presents one sample evolution of the reconstruction error over time for $r = 10$. Figure 3(b) displays the reconstruction error and training time (until improvement in the error falls below 10^{-4}), both as functions of rank r . Notice that the approach seems rather robust to the choice of r . Figure 3(c) compares the evolution of reconstruction error for three variants of the method, as described in the text. Notice that online variants seem superior to the offline variant.




FIGURE 4. Properties of the instance of Section V-A: Figure 4(a) shows frequencies of the values reported from sensors, with and without additional Gaussian noise. Figure 4(b) illustrates the mean and standard deviation of the historical flow data at all available sensors (grey), plus the mean values for events (yellow for $\mu = 5$, red for $\mu = 35$) and non-events (green) at the same sensors. Finally, in Figure 4(c), there is a heatmap of a validation matrix, which contains the normal reading (upper half) and event readings with $\mu = 35$ (lower half).

used as input plotted in green and the average values of the samples of events (*i.e.*, with the Gaussian noise for all mean values we used) plotted in red and yellow. We can observe that the supports of the distributions overlap, and especially in the case of $\mu = 5$, PSNR 23.58 dB, the event data seem hard to distinguish from non-event data. In Figure 4(c) we illustrate a heatmap of the validation matrix. The upper half of the matrix contains the normal readings while the lower half contains the event readings when $\mu = 35$ is used for the noise. To evaluate the performance of our subspace proximity tester, we have measured recall, precision, and F1-score using different values of Δ on the 4 matrices G .

Figure 5 presents the evolution of recall, precision, and F1-score as a function of Δ for 5 different values of μ . As can be observed, for small values of Δ , the precision is high, while recall is low, because small values of Δ lead are more likely to lead to infeasibility of the LP, and hence the negative result of the test. As we increase Δ , we observe that our approach identifies more of the input as Normal. On G' matrices with μ ranging from 15 to 25, we can observe that values $\Delta \in [10, 15]$ lead to the perfect performance with F1-Score of 1.0, which should not be too surprising, considering that this regime corresponds to PSNR below 20 dB. We can also observe that for noise of a lesser magnitude ($\mu = 5$), the subspace proximity tester is able to identify samples from G' with maximum F1-score of approximately 0.8 for $\Delta = 8$. By increasing Δ beyond this value, precision falls to 50%, which

is due to the fact that too many input samples are classified as non-event. This behaviour is to be expected, because by increasing the value of Δ , we are “relaxing” the constraints of the linear program, which in turn leads to the Normal outcome being more common.

Figures 6 and 7 illustrate a more challenging scenario: we keep the mean of the Gaussian noise low at $\mu = 5$, but vary the standard deviation $\sigma = 1, 2, \dots, 6$. This corresponds to peak signal-to-noise ratios (PSNR) within 22–24 dB, as detailed in the legend of the figures. Just as above, there is a setting of $\Delta = 10$, where the F1 score approaches 1.0.

Last but not least, we note that in order to classify a new sample in dimension $\mathbb{R}^{(1 \times 299,430)}$, our subspace proximity tester requires approximately ~ 0.009 seconds for subset of cardinality $s = \log r \log(r/e)$ to obtain $e = 0.1$. This runtime has been recorded on a standard PC equipped with an Intel i7-7820X CPU and 64 GB of RAM, where we have executed all the experiments. Having said that, the execution of the sub-space proximity tester is certainly possible in many embedded systems currently available. We note that this does not use the algorithm of [44], and hence can be improved by many orders of magnitude, if needed.

C. BENEFITS OF ONLINE OPTIMIZATION

Next, to demonstrate the benefits of the pursuit in the time-varying setting, we conducted the following experiment.




FIGURE 5. Results of repeated six-fold cross validation as a function of the half-width Δ of the uncertainty set and the PSNR used to generate the synthetic instance of Section V-A: Mean (solid line) and one standard deviation (half-width of the semi-transparent error band around the solid line) of three performance measures (recall, precision and F1 score).




FIGURE 6. Results of repeated six-fold cross validation as a function of the half-width Δ of the uncertainty set and the PSNR used to generate the synthetic instance of Section V-A: Mean (solid line) and one standard deviation (half-width of the semi-transparent error band around the solid line) of three performance measures (recall, precision and F1 score). The variation of PSNR is solely due to varying the standard deviation of the noise used to generate the synthetic instance, while keeping its mean low and constant at $\mu = 5$.

We took 200 rows from our matrix X of the previous section, which corresponded to normal readings from the sensors. Then, for $i = 0 \dots 200$, we sampled a row from X , and added noise, which had mean $\mu = 5$ and zero variance. The 200 rows thus added represented slowly increasing traffic volumes, which we would like the algorithm to adapt to.

We compare three variants of the algorithm. One, which we call “offline”, obtains the estimate of R_{200} and then does not update it further, $R_{200} = R_{200+i}$, $i = 0 \dots 200$. Another “online” variant performs nr updates (11) and rm updates (12) (which is known as 1 epoch) between receiving rows i and $i + 1$, $i = 0 \dots 199$. (We pick $\bar{\tau}$ so as to have the cardinality of $\hat{S}_{row} = \hat{S}_{column}$ equal to the number of hardware threads.) Finally, another online variant performs two epochs between receiving rows i and $i + 1$, $i = 0 \dots 199$. We refer to [32] for a detailed discussion of such online algorithms. Figure 3(c) presents the resulting evolution of the RMSE. As can be expected, using two epochs per update (and hence more CPU cycles) performs better than using one epoch per update, which in turn performs considerably better than the offline version, whose error increases over time.

VI. RELATED WORK

There is much related work in change-point, anomaly, outlier, and event detection, and the related problem of attack detection [56]. Since the work of Lorden [18], [19], there has been much work on change-point detection in univariate time series. See [18], [19], [25] for a book-length history and [26]

for an overview of the latest developments. Within anomaly detection, most statistical approaches have been tested, including hypothesis testing [9], dimension reduction [20], variants of filtering [13], and Gaussian processes [9]. In Computer Science, Complex Event Processing [57]–[59] and deep-learning methods [21]–[24] are popular. Within change-point detection [15], [60], such as cumulative statistics thresholding (CUSUM) or adaptive online thresholding (AOT), there are relatively few papers on the multi-variate problem [61]–[64, e.g.], and fewer still, which allow for missing data [65], [66]. Some of the recent ones [30], [33]–[35], [67], [68] also consider low-rank factorizations, albeit without subsampling. From the methodological works, we differ in our assumptions (uniform, rather than Gaussian noise), focus on efficient algorithms (subsampled subspace proximity testers) for the test, and our PAC guarantees.

Our approach builds upon a rich history of research in low-rank matrix completion methodologically. There, [69] suggested to replace the rank with the nuclear norm in the objective. The corresponding use of semidefinite programming (SDP) has been very successful in theory [70], while augmented-Lagrangian methods [71]–[74] and alternating least-squares (ALS) algorithms [75], [76] have been widely used in practice [75]–[79]. As it turns out [80], [81], they also allow for perfect recovery in some settings. The inequality-constrained variant of matrix completion, which we employ, has been introduced by [27] and extended towards on-line applications in Computer Vision by [32].

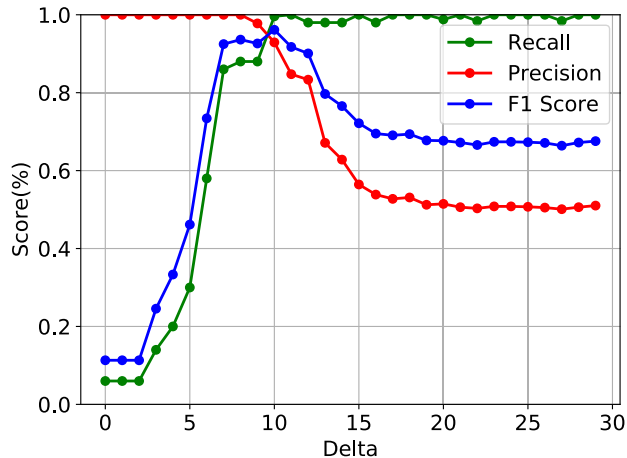


FIGURE 7. A further illustration of the results of repeated six-fold cross validation: mean of three performance measures (recall, precision and F1 score) plotted as a function of the half-width Δ of the uncertainty set. Here, we have used a low constant mean $\mu = 5$ and high standard deviation $\sigma = 6.0$, which correspond to PSNR = 22.26 dB.

IoT applications of anomaly detection are numerous and varied [15], [60], [82], mirroring much of the development in change-point, anomaly, outlier, and event detection at large. As suggested in the introduction, notable examples include transportation applications [2]–[4], [6], [33], [34], [67], power systems [1], [8]–[10], manufacturing [12], [13] and environmental applications [14]. These are, clearly, only some sample references in a much larger field.

In particular, the related work to our motivating application of IoT in Urban Traffic Management goes back at least to [83]. More recently, [2] proposes a method for detecting traffic events that have an impact on the road traffic conditions by extending the Bayesian Robust Principal Component Analysis. They create a sparse structure composed of multiple traffic data streams (e.g., traffic flow and road occupancy) and use it to localize traffic events in space and time. The data streams are subsequently processed so that with little computational cost they are able to detect events in an on-line and real-time fashion. [3] analyze road traffic accidents based on their severity using a space-time multivariate Bayesian model. They include both multivariate spatially structured and unstructured effects, as well a temporal component to capture the dependencies between the severity and time effects within a Bayesian hierarchical formulation.

Beyond the Internet of Things, Computer Vision studies a large number of related problems within “background modelling”, where the aim is to distinguish moving objects from stationary or dynamic backgrounds in a video feed. These are closely related to event detection, although typically focus on a single video feed, uniformly sampled, with no missing data. We refer to the recent handbook [84] and to the August 2018 special issue of the Proceedings of the IEEE [85] for up-to-date surveys. Compared to the work in Computer Vision, we develop both subsampled subspace proximity testers (point-to-subspace distance queries), and focus on the needs of applications in IoT, where there is more variety of less reliable data sources.

VII. CONCLUSION

Within a framework for representing what is an event and what is a non-event considering heterogeneous data, which are possibly not sampled uniformly, with missing values and measurement errors, we have presented a novel randomized event detection technique, implemented via a point-to-subspace distance query, with guarantees within probably approximately correct (PAC) learning. This is the first time such guarantees have been provided for any subsampling in matrix completion. The proofs use elaborate techniques from computational geometry (a bound on the VC dimension). We have also presented an experimental evaluation on data from a traffic-control system in Dublin, Ireland, which shows that it is possible to process data collected from thousands of sensors over the course of one year within minutes, to answer point-to-subspace distance queries in milliseconds and thus detect even hard-to-detect events. We envision that this approach may have wide-ranging applications, wherever asynchronous high-dimensional data streams need to be monitored.

REFERENCES

- [1] X. Lu, B. Chen, C. Chen, and J. Wang, “Coupled cyber and physical systems: Embracing smart cities with multistream data flow,” *IEEE Electrific. Mag.*, vol. 6, no. 2, pp. 73–83, Jun. 2018.
- [2] S. Yang, K. Kalpakis, and A. Biem, “Detecting road traffic events by coupling multiple timeseries with a nonparametric Bayesian method,” *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 5, pp. 1936–1946, Mar. 2014.
- [3] A. Boulieri, S. Liverani, K. de Hoogh, and M. Blangiardo, “A space-time multivariate Bayesian model to analyse road traffic accidents by severity,” *J. Roy. Stat. Soc., A, Statist. Soc.*, vol. 180, no. 1, pp. 119–139, Jan. 2017.
- [4] X. J. Kong, X. Song, F. Xia, H. Guo, J. Wang, and A. Tolba, “LoTAD: Long-term traffic anomaly detection based on crowdsourced bus trajectory data,” *World Wide Web*, vol. 21, no. 3, pp. 825–847, 2018.
- [5] P. Ta-Shma, A. Akbar, G. Gerson-Golan, G. Hadash, F. Carrez, and A. Moessner, “An ingestion and analytics architecture for IoT applied to smart city use cases,” *IEEE Internet Things J.*, vol. 5, no. 2, pp. 765–774, Apr. 2018.
- [6] H. Lu, Y. Li, S. Mu, D. Wang, H. Kim, and S. Serikawa, “Motor anomaly detection for unmanned aerial vehicles using reinforcement learning,” *IEEE Internet Things J.*, vol. 5, no. 4, pp. 2315–2322, Aug. 2017.
- [7] J. Clausen, A. Larsen, J. Larsen, and N. J. Rezanova, “Disruption management in the airline industry—Concepts, models and methods,” *Comput. Oper. Res.*, vol. 37, no. 5, pp. 809–821, May 2010.
- [8] M. Farajollahi, A. Shahsavari, and H. Mohsenian-Rad, “Location identification of distribution network events using synchrophasor data,” in *Proc. North Amer. Power Symp. (NAPS)*, Sep. 2017, pp. 1–6.
- [9] W. B. Wu, M. X. Cheng, and B. Gou, “A hypothesis testing approach for topology error detection in power grids,” *IEEE Internet Things J.*, vol. 3, no. 6, pp. 979–985, Dec. 2016.
- [10] R. Moghaddass and J. Wang, “A hierarchical framework for smart grid anomaly detection using large-scale smart meter data,” *IEEE Trans. Smart Grid*, vol. 9, no. 6, pp. 5820–5830, Nov. 2018.
- [11] B. Zhou, V. Lau, and X. Wang, “Machine-learning-based leakage-event identification for smart water supply systems,” *IEEE Internet Things J.*, vol. 7, no. 3, pp. 2277–2292, Mar. 2020.
- [12] A. Kanawaday and A. Sane, “Machine learning for predictive maintenance of industrial machines using IoT sensor data,” in *Proc. 8th IEEE Int. Conf. Softw. Eng. Service Sci. (ICSESS)*, Nov. 2017, pp. 87–90.
- [13] Y. Liu, T. Dillon, W. Yu, W. Rahayu, and F. Mostafa, “Noise removal in the presence of significant anomalies for industrial IoT sensor data in manufacturing,” *IEEE Internet Things J.*, vol. 7, no. 8, pp. 7084–7096, Aug. 2020.
- [14] L.-J. Chen, Y.-H. Ho, H.-H. Hsieh, S.-T. Huang, H.-C. Lee, and S. Mahajan, “ADF: An anomaly detection framework for large-scale PM_{2.5} sensing systems,” *IEEE Internet Things J.*, vol. 5, no. 2, pp. 559–570, Apr. 2018.

- [15] F. Li, A. Shinde, Y. Shi, J. Ye, X.-Y. Li, and W.-Z. Song, "System statistics learning-based IoT security: Feasibility and suitability," *IEEE Internet Things J.*, vol. 6, no. 4, pp. 6396–6403, Aug. 2019.
- [16] M. Eskandari, Z. H. Janjua, M. Vecchio, and F. Antonelli, "Passban IDS: An intelligent anomaly-based intrusion detection system for IoT edge devices," *IEEE Internet Things J.*, vol. 7, no. 8, pp. 6882–6897, Aug. 2020.
- [17] Y. An, F. R. Yu, J. Li, J. Chen, and V. C. M. Leung, "Edge intelligence (EI)-enabled HTTP anomaly detection framework for the Internet of Things (IoT)," *IEEE Internet Things J.*, vol. 8, no. 5, pp. 3554–3566, Mar. 2021.
- [18] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, no. 6, pp. 1897–1908, Dec. 1971.
- [19] V. Barnett and T. Lewis, *Outliers in Statistical Data* (Wiley Series in Probability and Mathematical Statistics. Applied Probability and Statistics). Chichester, U.K.: Wiley, 1978.
- [20] T. Yu, X. Wang, and A. Shami, "Recursive principal component analysis-based data outlier detection and sensor data aggregation in IoT systems," *IEEE Internet Things J.*, vol. 4, no. 6, pp. 2207–2216, Dec. 2017.
- [21] Y. Liu, S. Garg, J. Nie, Y. Zhang, Z. Xiong, J. Kang, and M. S. Hossain, "Deep anomaly detection for time-series data in industrial IoT: A communication-efficient on-device federated learning approach," *IEEE Internet Things J.*, vol. 8, no. 8, pp. 6348–6358, Apr. 2021.
- [22] S. Dou, K. Yang, and H. V. Poor, "PC2A: Predicting collective contextual anomalies via LSTM with deep generative model," *IEEE Internet Things J.*, vol. 6, no. 6, pp. 9645–9655, Dec. 2019.
- [23] M. Razian, M. Fathian, H. Wu, A. Akbari, and R. Buyya, "SAIoT: Scalable anomaly-aware services composition in CloudIoT environments," *IEEE Internet Things J.*, vol. 8, no. 5, pp. 3665–3677, Mar. 2021.
- [24] Y. Cheng, Y. Xu, H. Zhong, and Y. Liu, "Leveraging semi-supervised hierarchical stacking temporal convolutional network for anomaly detection in IoT communication," *IEEE Internet Things J.*, vol. 8, no. 1, pp. 144–155, Jan. 2021.
- [25] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*, vol. 104. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [26] A. A. Cook, G. Misirlı, and Z. Fan, "Anomaly detection for IoT time-series data: A survey," *IEEE Internet Things J.*, vol. 7, no. 7, pp. 6481–6494, Jul. 2020.
- [27] J. Mareček, P. Richtárik, and M. Takáč, "Matrix completion under interval uncertainty," *Eur. J. Oper. Res.*, vol. 256, no. 1, pp. 35–43, Jan. 2017.
- [28] D. Haussler, "Probably approximately correct learning," in *Proc. 8th Nat. Conf. Artif. Intell.*, 1990, pp. 1101–1108.
- [29] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge Univ. Press, 2004.
- [30] J. Feng, H. Xu, S. Mannor, and S. Yan, "Online PCA for contaminated data," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 26, 2013, pp. 764–772.
- [31] H. Guo, C. Qiu, and N. Vaswani, "An online algorithm for separating sparse and low-dimensional signal sequences from their sum," *IEEE Trans. Signal Process.*, vol. 62, no. 16, pp. 4284–4297, Aug. 2014.
- [32] A. Akhriev, J. Marecek, and A. Simonetto, "Pursuit of low-rank models of time-varying matrices robust to sparse and measurement noise," in *Proc. AAAI*, 2020, pp. 3171–3178.
- [33] M. Mardani, G. Mateos, and G. B. Giannakis, "Dynamic anomalography: Tracking network anomalies via sparsity and low rank," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 1, pp. 50–66, Feb. 2013.
- [34] M. Mardani, G. Mateos, and G. B. Giannakis, "Recovery of low-rank plus compressed sparse matrices with application to unveiling traffic anomalies," *IEEE Trans. Inf. Theory*, vol. 59, no. 8, pp. 5186–5205, Aug. 2013.
- [35] F. Li, R. Xie, B. Yang, L. Guo, P. Ma, J. Shi, J. Ye, and W. Song, "Detection and identification of cyber and physical attacks on distribution power grids with PVs: An online high-dimensional data-driven approach," *IEEE J. Emerg. Sel. Topics Power Electron.*, vol. 10, no. 1, pp. 1282–1291, Feb. 2022.
- [36] J. Zhan, B. Lois, H. Guo, and N. Vaswani, "Online (and offline) robust PCA: Novel algorithms and performance guarantees," in *Proc. Artif. Intell. Statist.*, 2016, pp. 1488–1496.
- [37] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Feb. 1995.
- [38] N. J. A. Harvey, D. R. Karger, and S. Yekhanin, "The complexity of matrix completion," in *Proc. 17th Annu. ACM-SIAM Symp. Discrete Algorithm (SODA)*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2006, pp. 1103–1111.
- [39] J. Tanner and K. Wei, "Normalized iterative hard thresholding for matrix completion," *SIAM J. Sci. Comput.*, vol. 35, no. 5, pp. S104–S125, Jan. 2013.
- [40] V. Vapnik and A. Y. Chervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," *Theory Probab. Appl.*, vol. 16, no. 2, p. 264, 1971.
- [41] D. Haussler and E. Welzl, "ε-nets and simplex range queries," *Discrete Comput. Geometry*, vol. 2, no. 2, pp. 127–151, Jun. 1987.
- [42] D. Haussler and E. Welzl, "Epsilon-nets and simplex range queries," in *Proc. 2nd Annu. Symp. Comput. Geometry (SCG)*, New York, NY, USA, 1986, pp. 61–71.
- [43] K. L. Clarkson, "Further applications of random sampling to computational geometry," in *Proc. 18th Annu. ACM Symp. Theory Comput. (STOC)*, New York, NY, USA, 1986, pp. 414–423.
- [44] K. L. Clarkson, "New applications of random sampling in computational geometry," *Discrete Comput. Geometry*, vol. 2, no. 2, pp. 195–222, Jun. 1987.
- [45] B. Gärtner and M. Hoffmann, *Computational Geometry*. Zürich, Switzerland: ETH, Jan. 2013.
- [46] J. Gondzio, "Interior point methods 25 years later," *Eur. J. Oper. Res.*, vol. 218, no. 3, pp. 587–601, 2012.
- [47] M. Kojima, N. Megiddo, and S. Mizuno, "A general framework of continuation methods for complementarity problems," *Math. Oper. Res.*, vol. 18, no. 4, pp. 945–963, 1993.
- [48] Y. Ye, M. J. Todd, and S. Mizuno, "An $O(\sqrt{nL})$ -iteration homogeneous and self-dual linear programming algorithm," *Math. Oper. Res.*, vol. 19, no. 1, pp. 53–67, Feb. 1994.
- [49] J. Mattingley and S. Boyd, "Real-time convex optimization in signal processing," *IEEE Signal Process. Mag.*, vol. 27, no. 3, pp. 50–61, May 2010.
- [50] J. Mattingley and S. Boyd, "CVXGEN: A code generator for embedded convex optimization," *Optim. Eng.*, vol. 13, no. 1, pp. 1–27, 2012.
- [51] M. De Berg, M. Van Kreveld, M. Overmars, and O. C. Schwarzkopf, *Computational Geometry*. Berlin, Germany: Springer, 2000.
- [52] R. P. Stanley, *An Introduction to Hyperplane Arrangements*. Park City, UT, USA: IAS/Park City Mathematics Institute, 2004.
- [53] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, "Spark: Cluster computing with working sets," in *Proc. 2nd USENIX Conf. Hot Topics Cloud Comput. (HotCloud)*. Berkeley, CA, USA: USENIX Association, 2010, p. 10.
- [54] S. van der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," *Comput. Sci. Eng.*, vol. 13, no. 2, pp. 22–30, 2011.
- [55] R. Wu and E. J. Keogh, "Current time series anomaly detection benchmarks are flawed and are creating the illusion of progress," 2020, *arXiv:2009.13807*.
- [56] F. Pasqualetti, F. Dörfler, and F. Bullo, "Attack detection and identification in cyber-physical systems," *IEEE Trans. Autom. Control*, vol. 58, no. 11, pp. 2715–2729, Nov. 2013.
- [57] A. Artakis, M. Weidlich, F. Schnitzler, I. Boutsis, T. Liebig, N. Piatkowski, C. Bockermann, K. Morik, V. Kalogeraki, J. Marecek, A. Gal, S. Mannor, D. Kinane, and D. Gunopulos, "Heterogeneous stream processing and crowdsourcing for urban traffic management," in *Proc. EDBT*, vol. 14, 2014, pp. 712–723.
- [58] M. Ma, P. Wang, C. H. Chu, and L. Liu, "Efficient multipattern event processing over high-speed train data streams," *IEEE Internet Things J.*, vol. 2, no. 4, pp. 295–309, Aug. 2015.
- [59] R. Mayer, B. Koldehofe, and K. Rothermel, "Predictable low-latency event detection with parallel complex event processing," *IEEE Internet Things J.*, vol. 2, no. 4, pp. 274–286, Aug. 2015.
- [60] Y. Gu, Y. Jiao, X. Xu, and Q. Yu, "Request–response and censoring-based energy-efficient decentralized change-point detection with IoT applications," *IEEE Internet Things J.*, vol. 8, no. 8, pp. 6771–6788, Apr. 2021.
- [61] J. A. D. Aston and C. Kirch, "High dimensional efficiency with applications to change point tests," *Electron. J. Statist.*, vol. 12, no. 1, pp. 1901–1947, Jan. 2018.
- [62] H. Cho and P. Fryzlewicz, "Multiple-change-point detection for high dimensional time series via sparsified binary segmentation," *J. Roy. Stat. Soc., B, Stat. Methodol.*, vol. 77, no. 2, pp. 475–507, Mar. 2015.
- [63] C. Zou, Z. Wang, X. Zi, and W. Jiang, "An efficient online monitoring method for high-dimensional data streams," *Technometrics*, vol. 57, no. 3, pp. 374–387, Jul. 2015.
- [64] G. Wang, C. Zou, and G. Yin, "Change-point detection in multinomial data with a large number of categories," *Ann. Statist.*, vol. 46, no. 5, Oct. 2018.
- [65] Y. Xie, J. Huang, and R. Willett, "Change-point detection for high-dimensional time series with missing data," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 1, pp. 12–27, Feb. 2013.

- [66] Y. S. Soh and V. Chandrasekaran, "High-dimensional change-point estimation: Combining filtering with convex optimization," *Appl. Comput. Harmon. Anal.*, vol. 43, no. 1, pp. 122–147, Jul. 2017.
- [67] Q. Zhang and T. Chu, "Structure regularized traffic monitoring for traffic matrix estimation and anomaly detection by link-load measurements," *IEEE Trans. Instrum. Meas.*, vol. 65, no. 12, pp. 2797–2807, Dec. 2016.
- [68] L. Balzano, Y. Chi, and Y. M. Lu, "Streaming PCA and subspace tracking: The missing data case," *Proc. IEEE*, vol. 106, no. 8, pp. 1293–1310, Aug. 2018.
- [69] M. Fazel, "Matrix rank minimization with applications," Ph.D. dissertation, Dept. Elect. Eng., Stanford Univ., Stanford, CA, USA, 2002.
- [70] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Found. Comput. Math.*, vol. 9, no. 6, pp. 717–772, Apr. 2009.
- [71] M. Jaggi and M. Y. Sulovský, "A simple algorithm for nuclear norm regularized problems," in *Proc. 27th Int. Conf. Mach. Learn. (ICML)*, New York, NY, USA, 2010, pp. 471–478.
- [72] K. Lee and Y. Bresler, "ADMIRA: Atomic decomposition for minimum rank approximation," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4402–4416, Sep. 2010.
- [73] S. Shalev-Shwartz, A. Gonen, and O. Shamir, "Large-scale convex minimization with a low-rank constraint," in *Proc. 28th Int. Conf. Mach. Learn. (ICML)*, New York, NY, USA, 2011, pp. 329–336.
- [74] Z. Wang, M.-J. Lai, Z. Lu, W. Fan, H. Davulcu, and J. Ye, "Rank-one matrix pursuit for matrix completion," in *Proc. 31st Int. Conf. Mach. Learn. (ICML)*, New York, NY, USA, 2014, pp. 91–99.
- [75] N. Srebro, J. Rennie, and T. S. Jaakkola, "Maximum-margin matrix factorization," in *Advances in Neural Information Processing Systems*. Red Hook, NY, USA: Curran Associates, 2004, pp. 1329–1336.
- [76] J. D. M. Rennie and N. Srebro, "Fast maximum margin matrix factorization for collaborative prediction," in *Proc. 31st Int. Conf. Mach. Learn. (ICML)*, New York, NY, USA, 2005, pp. 713–719.
- [77] A. Mnih and R. Salakhutdinov, "Probabilistic matrix factorization," in *Advances in Neural Information Processing Systems*. Red Hook, NY, USA: Curran Associates, 2007, pp. 1257–1264.
- [78] R. Bell and Y. Koren, "Scalable collaborative filtering with jointly derived neighborhood interpolation weights," in *Proc. IEEE 7th Int. Conf. Data Mining*, Piscataway, NJ, USA, Oct. 2007, pp. 43–52.
- [79] J. P. Haldar and D. Hernando, "Rank-constrained solutions to linear matrix equations using powerfactorization," *IEEE Signal Process. Lett.*, vol. 16, no. 7, pp. 584–587, Jul. 2009.
- [80] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2980–2998, Jun. 2010.
- [81] P. Jain, P. Netrapalli, and S. Sanghavi, "Low-rank matrix completion using alternating minimization," in *Proc. 45th Annu. ACM Symp. Symp. Theory Comput. (STOC)*, New York, NY, USA, 2013, pp. 665–674.
- [82] F. Li, R. Xie, Z. Wang, L. Guo, J. Ye, P. Ma, and W. Song, "Online distributed IoT security monitoring with multidimensional streaming big data," *IEEE Internet Things J.*, vol. 7, no. 5, pp. 4387–4394, May 2020.
- [83] J. West, "Proposed real-time surveillance and control system for Los Angeles," California Dept. Transp. Freeway Oper. Branch, Los Angeles, CA, USA, Tech. Rep., 1969. [Online]. Available: <https://trid.trb.org/view/489112>
- [84] T. Bouwmans, N. S. Aybat, and E.-H. Zahzah, *Handbook of Robust Low-Rank and Sparse Matrix Decomposition: Applications in Image and Video Processing*. London, U.K.: Chapman & Hall, 2016.
- [85] N. Vaswani, Y. Chi, and T. Bouwmans, "Rethinking PCA for modern data sets: Theory, algorithms, and applications [scanning the issue]," *Proc. IEEE*, vol. 106, no. 8, pp. 1274–1276, Aug. 2018.



JAKUB MARECEK (Member, IEEE) received the Ph.D. degree in computer science from the University of Nottingham, Nottingham, U.K., in 2012.

He has been a faculty member at the Czech Technical University in Prague, the Czech Republic, since 2020. Previously, he had also worked in two start-ups, at ARM Ltd., at the University of Edinburgh, at the University of Toronto, at IBM Research – Ireland, and at the University of California, Los Angeles. His research interests

include the design and analysis of algorithms for optimisation and control problems across a range of application domains. He has coauthored more than 17 journal and 40 conference papers. He is a co-inventor on 9 families of awarded U.S. patents and several pending patent applications.



STATHIS MAROULIS received the bachelor's degree in computer science from the University of Athens and the M.Sc. degree from the Athens University of Economics and Business, where he is currently pursuing the Ph.D. degree with the Department of Computer Science, under the supervision of Vana Kalogeraki. His work has appeared in the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS and the IEEE International Conference on Data Engineering (ICDE), among other venues.



VANA KALOGERAKI received the Ph.D. degree from the University of California at Santa Barbara, Santa Barbara, CA, USA, in 2000. Previously, she has held positions as an associate professor and assistant professor with the Department of Computer Science, University of California at Riverside, Riverside, CA, USA, and a research Scientist at Hewlett-Packard Labs, Palo Alto, CA, USA. She is currently an Associate Professor with the Department of Informatics, Athens University

of Economics and Business, Athens, Greece, where she is leading the Computer Systems and Communications Laboratory. She has been working in the field of distributed and real-time systems, distributed sensor systems, peer-to-peer systems, resource management, and fault-tolerance for more than 20 years. She has published more than 150 journals and conference papers and contributions to books. She has coauthored the OMG CORBA Dynamic Scheduling Standard. She was invited to give keynote talks and deliver tutorials and seminars. She has served as the general chair and the program chair on many conferences in the area of distributed systems and program committee member on more than 200 conferences.



DIMITRIOS GUNOPOULOS received the Ph.D. degree from Princeton University, Princeton, NJ, USA, in 1995. He was a Researcher at the IBM Almaden Research Center and a Professor with the University of California at Riverside, Riverside, CA, USA. He is currently a Professor with the Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Greece. He has developed clustering, classification, and indexing techniques

for high-dimensional and trajectory data, data analysis tools for streams and sensor networks, and on web and document analysis problems. He has coauthored more than 50 journals and 100 conference papers, seven patents and one patent application, 20 book chapters, and a book. His research interests include data mining, web mining, data science, databases, sensor and peer-to-peer systems, and algorithms.

• • •