Computers and Electrical Engineering 101 (2022) 108065

Contents lists available at ScienceDirect

Computers and Electrical Engineering

journal homepage: www.elsevier.com/locate/compeleceng

Two-step residual-error based approach for anomaly detection in engineering systems using variational autoencoders ☆

Ana González-Muñiz [a,*], Ignacio Díaz [a], Abel A. Cuadrado [a], Diego García-Pérez [a], Daniel Pérez [b]

[a] Electrical Engineering Department, University of Oviedo, Edif. Departmental Oeste 2, Campus de Viesques s/n, Gijón 33204, Spain
[b] SUPPRESS Research Group, University of León, Escuela de Ingenierías, Campus de Vegazana, León 24007, Spain

ARTICLE INFO

Keywords:
Anomaly detection
Novelty detection
Deep autoencoder
Variational autoencoder
Engineering systems

ABSTRACT

Anomaly detection is a crucial task in the engineering systems field. However, there is usually little or no information about all possible abnormal modes in systems. Hence, a common approach is to build a model of healthy behaviour, based on normal operation data, so that anomaly detection would depend on how well new data fit this model. According to this idea, we propose a residual-error based approach consisting of: a variational autoencoder, used to model the probability density function of the system's healthy behaviour; and a two-step classification algorithm, which classifies the incoming samples based on their residuals, and reports not only their normal/anomalous nature but also that of their components. We have tested this proposal in three different engineering contexts and we have compared its performance with that of state-of-the-art approaches, demonstrating its capability to successfully detect and characterize anomalies.

1. Introduction

Anomaly detection consists in finding those patterns in data that do not conform to expected normal behaviour. Such patterns, commonly known as anomalies or outliers, represent deviations from normal behaviour, so their detection is not only of great value, but is often critical in a wide variety of real-world applications, like intrusion detection [1], fraud detection [2] or medical anomaly detection [3]. In the context of industrial processes, anomaly detection is also a topic of great importance since both detection and diagnosis of faults are crucial to optimize and guarantee safety in the operation of machines, leading to higher productivity and process efficiency, with benefits such as reduced operating costs, longer operating life or improved operating uptime [4].

Given the relevance of anomaly detection, there has been extensive research on this topic: we can find several approaches in the state of the art, depending on the application domain and even the dataset [5]. In the typical scenario, we deal with datasets containing mostly normal behaviour data and few, if any, examples of anomalous behaviour, which sometimes are not even representative of all possible abnormal modes. This makes traditional multi-class classification schemes unsuitable for anomaly detection [6].

Instead, a common approach in the literature is to construct a model of normal behaviour (based on available normal data), so that anomaly detection would depend on how well new data fit this model [6]. In other words, the aim of this approach is to build an analytical model of the process under study (key concept in the digital twin approach [7]), in such a way that its comparison

☆ This paper is accepted by Associate Editor Dr. Shadi Aljawarneh.

* Corresponding author.
E-mail addresses: gonzalezmunana@uniovi.es (A. González-Muñiz), idiaz@uniovi.es (I. Díaz), cuadrado@isa.uniovi.es (A.A. Cuadrado), diegogarcia@isa.uniovi.es (D. García-Pérez), dperl@unileon.es (D. Pérez).

https://doi.org/10.1016/j.compeleceng.2022.108065
Received 1 July 2021; Received in revised form 28 April 2022; Accepted 1 May 2022
Available online 24 May 2022

with the actual process would bring anomalous behaviours to light. It is therefore a residual-based approach in which instances with large residual errors are considered more likely to be anomalies.

The modelling of normal behaviour has traditionally been addressed by means of model-based methods [8–10], which require expertise and prior knowledge of the system to be modelled. In contrast, there is a growing interest in data-driven methods, especially with the proliferation in recent years of deep learning techniques [11]. In the context of engineering processes, we can find several approaches in the literature that propose the use of deep autoencoders for anomaly detection [12–14], where samples with large residual errors are considered to be anomalous.

Despite the strong interest in deep learning approaches, new techniques have emerged whose potential has yet to be explored. This is the case of the variational autoencoder (VAE), which has proven to be a successful unsupervised learning algorithm [15] showing promising results as generative model, with applications in several fields, such as audio, text, or image generation [16–19]. In contrast to traditional autoencoders, VAEs impose restrictions on the distribution of latent variables and, in doing so, learn the probability density function of training data. Hence, being trained on samples representative of healthy behaviour, VAEs are able to learn the healthy distribution of process data, thus becoming a powerful anomaly detection tool [20–22]. In particular, VAE residuals capture any deviation from process normal behaviour, providing valuable information on the health condition of any incoming sample. In consequence, VAE residuals have started to be used in anomaly detection approaches [23–25], with recent applications as anomaly detection scores in the industrial field [26,27].

These residual-based approaches have proven to be successful in detecting anomalous samples, which, as mentioned above, is a crucial task in many domains. However, it would be of great value to also be able to identify the normal/anomalous nature of each component in the samples; especially in the engineering systems field, where such information, together with the classification of samples, would provide a rich insight into the state of health of the monitored processes. Under such circumstances, we explore in this paper the use of VAE residuals for anomaly detection in the engineering systems field and we propose a novel analysis of the residuals. As discussed in the literature, the closer the model approximates the distribution of training data, the more meaningful its residuals will be [28], so we expect VAE based residual generation to be competitive and of great help in detecting and diagnosing anomalies, thanks to its ability to model the probability distribution of normal conditions as well as to deal with high dimensional and large sample sizes. In line with state-of-the-art approaches [12,14,23,27], we use the residuals of the autoencoder as anomaly indicator, but we propose a two-step classification, which allows us to determine not only the normal/anomalous nature of the samples (as previously done in the literature) but also that of their components. In this way, we aim to provide some valuable additional information about the samples, beyond whether they are anomalous or not.

Our architecture consists of a variational autoencoder, which has been trained to reconstruct samples of healthy behaviour, and a classification algorithm, which classifies the incoming samples based on their reconstruction error. This algorithm classifies each sample in two steps: 1) component-wise classification (classification of each element – component – in the sample); and 2) overall classification (classification of the sample). We have tested this proposal in three different engineering contexts (a rotating machine, a hydraulic system and a body motion system) and we have compared its performance with that of state-of-the-art approaches, both in terms of modelling (variational autoencoder vs. deep autoencoder) and classification (two-step classifier vs. traditional classifiers). The results of the research demonstrate the capability of our proposal to successfully detect anomalies in all three contexts, being the highest accuracy results those obtained with the two-step classifier applied on VAE residuals. Additionally, we present a visual analysis of the results that reveals the potential of the component-wise classification step, providing insight into the anomaly decision and therefore illustrating the explanatory nature (which plays a key role in the interpretable machine learning approach [29]) of our proposal.

In summary, the contributions of this research are as follows: 1) we explore the potential of the VAE for anomaly detection in engineering systems; 2) we propose a system able not only to successfully detect anomalous samples, but also to point out the elements in the samples that do not conform to healthy behaviour, thereby contributing to a better understanding of the process nature.

The rest of this document is organized as follows. In Section 2, we introduce the theoretical foundation of the variational autoencoder. Section 3 describes the architecture of the proposed system. The datasets and the results of the proposal are presented in Section 4. Finally, the conclusions are set out in Section 5.

2. Related literature

Variational autoencoders [15] have become in recent years one of the most popular approaches to unsupervised learning. They emerged as a probability-based extension of deep autoencoders, so we will start this section by introducing the deep autoencoder and then present its variational version.

2.1. Deep autoencoders

A deep autoencoder is a feed-forward multi-layer neural network in which the desired output is the input itself. As we see in Fig. 1, the architecture consists of an encoder f_{enc}, which outputs a latent representation $\mathbf{z}$ of the input data $\mathbf{x}$, and a decoder f_{dec} that reconstructs the input data ($\hat{\mathbf{x}}$) from its latent representation $\mathbf{z}$. In between, there is a bottleneck (typically one or more low-dimensional layers), so the identity map is not a possible solution and the model is forced to learn the underlying low-dimensional structure in data. During the learning process, the architecture is trained using the gradient descent method [30] by means of

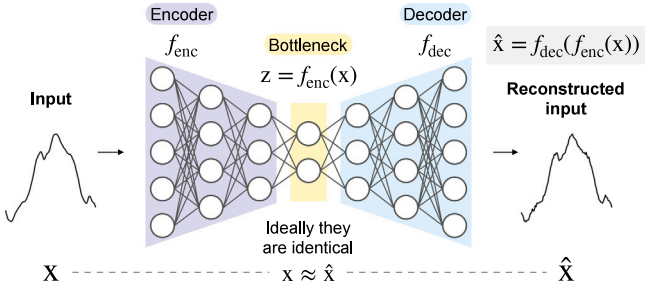

Fig. 1. Deep autoencoder.

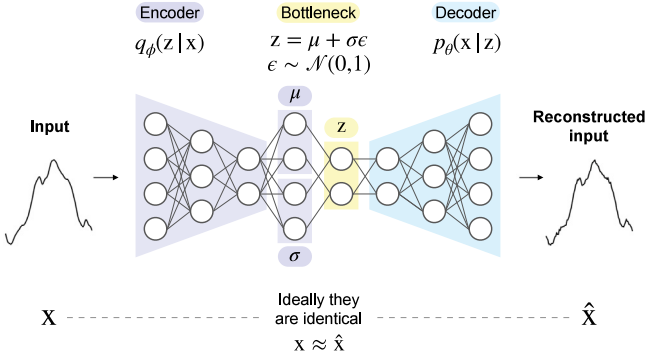

Fig. 2. Variational autoencoder.

backpropagation [31] in order to minimize the difference between $\mathbf{x}$ and $\hat{\mathbf{x}}$. So in short, the autoencoder can be defined as a solution to the optimization problem (1), where we try to minimize a loss function $\mathcal{L}$, being $\|.\|$ usually the l_2-norm.

$$\min_{f_{enc}, f_{dec}} \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) = \min_{f_{enc}, f_{dec}} \|\mathbf{x} - f_{dec}(f_{enc}(\mathbf{x}))\| \tag{1}$$

Although the first versions of autoencoders appeared decades ago [32,33], they have been evolving over the years, especially with the advent of deep learning. Deep learning models have drastically improved the state of the art in several fields, most notably speech recognition, natural language processing and computer vision [34]. They are built by stacking multiple processing layers, which conform compositional hierarchies in which the higher level features are the composition of the lower level ones. This gives them the ability to find the best representations of data, known as *representation learning* [35], thus becoming powerful feature extractors. Consequently, deep autoencoders have the ability to reduce the dimension of the input data in a hierarchical way, leading to high quality reconstructions of data, as shown in the literature [36–38].

However, the quality of the results worsens if we reconstruct samples that do not conform to training data. Considering training data to be representative of normal behaviour, the reconstruction error is used in the state of the art as an anomaly score, so samples with high reconstruction errors are classified as anomalies. Although the details of the anomaly decision depend on each proposal, two common approaches for the classification of the residuals are: a) threshold-based approach, which proposes to classify the residuals by comparison with a predefined anomaly threshold (reconstruction error above which a sample is considered anomalous) [14,39–41]; and b) one-class classifier approach, which proposes to train additional architectures for the classification of the residuals [42–44]. An example of a) is presented in [39], where the maximum reconstruction error of training data is chosen as the anomaly threshold; another proposal can be seen in [41], where the authors choose a percentile of the reconstruction error as the anomaly threshold. In [43] we can see an example of b), where a one-class Support Vector Machine (SVM) is trained to classify the samples based on their residuals.

2.2. Variational autoencoders

Variational autoencoders inherit the architecture of deep autoencoders, but imposing additional constraints on the bottleneck that transform traditional deterministic autoencoders into powerful probabilistic models: while deep autoencoders learn an arbitrary function to encode and decode input data, variational autoencoders learn the parameters of a probability distribution that models the data.

As we can see in Fig. 2, the VAE consists of an encoder $q_\phi(\mathbf{z}|\mathbf{x})$, being an approximate posterior, and a decoder $p_\theta(\mathbf{x}|\mathbf{z})$, being the likelihood of the data $\mathbf{x}$ given the latent variable $\mathbf{z}$. According to this scheme, the encoder becomes a variational inference network,

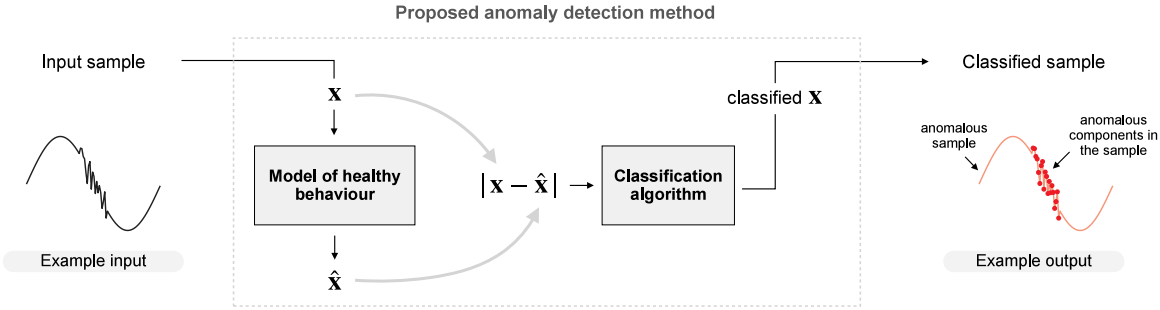

Fig. 3. Proposed method.

mapping input data to (approximate) posterior distributions over the latent space, and the decoder operates as a generative network, mapping arbitrary latent coordinates back to distributions over the original data space. To achieve this, it is assumed that input data can be sampled from a unit Gaussian distribution of latent parameters, so that during the learning process the model is trained by simultaneously optimizing two loss functions: a reconstruction loss $\mathcal{L}$ and the Kullback–Leibler divergence D_{KL} between the learned latent distribution and a prior unit Gaussian. So we can understand the VAE as a deep autoencoder with an additional regularization provided by the D_{KL} term.

The resulting objective function of the VAE is presented in (2) and it is the *variational lowerbound* of the marginal likelihood of data, since the marginal likelihood is intractable. The marginal likelihood is the sum over the marginal likelihood of individual data points $\log p_\theta(\mathbf{x}) = \sum_{i=1}^{n} \log p_\theta(\mathbf{x}^{(i)})$ and it can be rewritten for individual data points $\mathbf{x}^{(i)}$ as follows:

$$\log p_\theta(\mathbf{x}^{(i)}) = D_{\mathrm{KL}}(q_\phi(\mathbf{z}|\mathbf{x}^{(i)})||p_\theta(\mathbf{z}|\mathbf{x}^{(i)})) + \mathcal{L}(\theta, \phi; \mathbf{x}^{(i)}) \tag{2}$$

Considering the D_{KL} term is always bigger than 0, and applying the Bayes rule, the variational lowerbound $\log p_\theta(\mathbf{x}^{(i)})$ results to be:

$$-D_{\mathrm{KL}}(q_\phi(\mathbf{z}|\mathbf{x}^{(i)})||p_\theta(\mathbf{z})) + E_{q_\phi(\mathbf{z}|\mathbf{x}^{(i)})}[\log p_\theta(\mathbf{x}^{(i)}|\mathbf{z})] \tag{3}$$

During the learning process, the VAE is trained in order to maximize (3), using the gradient descent by means of backpropagation. But this entails some difficulties, since the reconstruction error term in (3) requires the Monte Carlo estimate of the expectation, which is not easily differentiable [15]. To overcome this, the VAE includes a *reparameterization trick*, which consists in using a random variable from a standard normal distribution instead of a random variable from the original distribution ($z = \mu + \sigma\epsilon$, with $\epsilon \sim \mathcal{N}(0, 1)$). This trick allows us to propagate the gradient back through the network and apply the backpropagation algorithm to train the VAE.

In this way, we obtain an autoencoder that works as a generative model. VAEs are therefore a powerful extension of deep autoencoders and they have already shown promise in several fields such as image generation, natural language processing or chemical design [45–47]. In the context of anomaly detection, recent approaches have used the residuals of the VAE as an anomaly score [24,25]. Such works, in line with deep autoencoder approaches, propose to classify the residuals by comparison with an anomaly threshold (as we can see in [23]) or by using a one-class classifier (such as in [48]). The applicability of these ideas in the industrial field has already started to be explored [26,27] and, not surprisingly, VAEs are expected to play a significant role in the future of health monitoring algorithms [49]. Given the potential of VAE approaches, we explore in this paper the use of VAE residuals for anomaly detection in the engineering systems field, and propose a novel threshold-based analysis for the classification of the residuals.

3. Proposed method

In this paper we propose a residual-error based approach to anomaly detection that consists of two elements: 1) a model based on process data, trained to reconstruct healthy samples so that residuals of incoming samples become a measure of their deviation from healthy behaviour; and 2) a classification algorithm, in charge of classifying the incoming samples based on their residuals. Therefore, in this section we present our model of healthy behaviour and the classification algorithm.

An overview of the proposal is shown in Fig. 3. In this figure the model receives a sample $\mathbf{x}$ and returns its reconstruction $\hat{\mathbf{x}}$; then the classifier analyzes the sample residuals $|\mathbf{x} - \hat{\mathbf{x}}|$ (where $|\mathbf{x} - \hat{\mathbf{x}}|$ denotes the component-wise absolute value of $\mathbf{x} - \hat{\mathbf{x}}$) and returns the classification of the sample.

3.1. Model of healthy behaviour

As we can see in Fig. 3, our approach to anomaly detection requires a reconstruction technique, which is in charge of reconstructing the incoming samples, and therefore of providing the residuals with which we feed the classification algorithm.

Table 1
Variational autoencoder architecture for each dataset.

	Number of epochs	Batch size	Number of layers	Number of neurons in the layers		
				Encoder	Bottleneck	Decoder
Rotating machine	800	400	9	(300,60,30)	(2,2,2)	(30,60,300)
Hydraulic system	600	400	11	(17,20,40,20)	(2,2,2)	(20,40,20,17)
Body motion	800	800	11	(23,40,40,40)	(2,2,2)	(40,40,40,23)

Table 2
Deep autoencoder architecture for each dataset.

	Number of epochs	Batch size	Number of layers	Number of neurons in the layers		
				Encoder	Bottleneck	Decoder
Rotating machine	800	100	7	(300,60,30)	(2)	(30,60,300)
Hydraulic system	200	300	9	(17,20,20,20)	(2)	(20,20,20,17)
Body motion	400	600	9	(23,40,40,40)	(2)	(40,40,40,23)

For this purpose, we use a variational autoencoder, which we train to reconstruct samples of normal behaviour, so that residuals of incoming samples become a measure of their deviation from normal behaviour.

We have trained three variational autoencoders, one per dataset, using the minibatch gradient descent [30] and the rmsprop optimizer [50]. The number of epochs and the batch size are indicated in Table 1, along with the architecture of each model (number of layers and number of neurons per layer). Regarding the activation function, the rectified linear unit (ReLU) [51] is a popular choice in the literature. It allows networks to easily obtain sparse representations of the data, helps to alleviate the vanishing gradient problem and accelerates the convergence of learning [52], offering better performance and generalization than other activation functions [53]. Hence, we have used the ReLU function in all model layers (except in the output and bottleneck layers, where we have used sigmoidal and identity functions, respectively). It is worth mentioning that although ReLUs help to alleviate the vanishing gradient problem, they still suffer from it, leading to inactive neurons during the training of the models [54]. To overcome these issues, new variants of ReLU have been proposed in the state of the art (such as Leaky ReLU, Randomized Leaky ReLU or Parametric ReLU [55]). However, the standard ReLU remains to be the preferred option in the literature [56] and is the one we have used in our experiments.

In order to study the potential of the VAE, we have compared its performance with that of the deep autoencoder. To this end, we have trained three deep autoencoders, one per dataset, using the minibatch gradient descent [30] and the Adam optimizer [57]. The number of epochs, the batch size, and the architectures of these models are shown in Table 2, and we have used the ReLU function [51] as activation function in all their layers (except in the output and bottleneck layers, where we have used the identity function).

Regarding the training process, 70% of normal samples have been randomly chosen to compose the training sets. The remaining 30% of normal samples together with the samples representing anomalous behaviour constitute the test sets used to assess the performance of the proposed method (Table 12).

It should also be noted that the tuning of all the parameters mentioned in this section (Tables 1, 2) is based on the testing of different configurations. The chosen configurations are the ones that yielded the best classification results.

3.2. Classification algorithm

Anomaly detection is a critical task in many contexts. Therefore, extensive research has been done on this topic and a wide range of anomaly detection techniques can be found in the literature (such as classification-based, clustering-based or statistical-based techniques [5]). A common approach is to address anomaly detection as a one-class classification problem, in which we aim to differentiate anomalous samples (positive class, represented as 1) from normal samples (negative class, represented as 0). To perform this classification, we propose a threshold-based classification algorithm, which receives as input the residuals of the sample to be classified and returns as output the classification of the sample. More in detail (Fig. 4), the output of the classifier for an incoming sample $\mathbf{x} \in \mathbb{R}^n$ would consist of: a vector $\hat{\mathbf{y}} \in \mathbb{R}^n$, containing the component-wise classification of the sample; and a scalar $\hat{y}_{\text{overall}}$, representing the overall classification of the sample. The two steps required to achieve this classification are described below.

1. **Component-wise classification**
   In this first step we classify each element $\{x_j, j = 1, 2, \dots, n\}$ in the sample $\mathbf{x}$: if the residual of the component x_j exceeds its corresponding anomaly threshold th_j, it is considered anomalous; otherwise, the component is classified as normal.

$$\hat{y}_j = \begin{cases} 1 & \text{if } |x_j - \hat{x}_j| > th_j \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

   The selection of the anomaly threshold th_j is based on the residuals of the training dataset $\{\mathbf{X}_{\text{train}} = \left(x_{\text{train}_{ij}}\right) \in \mathbb{R}^{m \times n}\}$, which contains samples of normal behaviour. Since residuals of anomalous samples are expected to be greater than those of normal

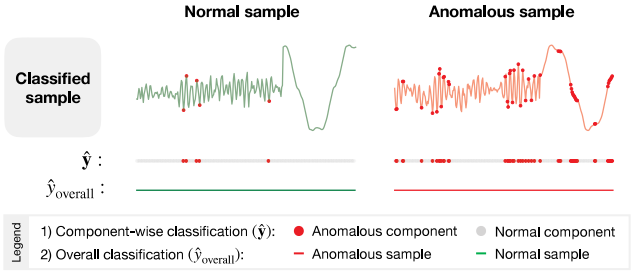

Fig. 4. Results of the classification of two samples (normal and anomalous) of the rotating machine dataset.

samples, we define each component threshold th_j as the 95th percentile of the residual error in training data for component j:

$$th_j = 95\text{th}(|x_{\text{train}_{*j}} - \hat{x}_{\text{train}_{*j}}|) \tag{5}$$

In this way, the anomaly decision depends only on the normal behaviour, of which there are often large amounts of data available; and we choose the 95th percentile as the anomaly threshold, since it is a popular approach [58,59] and provides a representative measure of training data, while being less sensitive to data noise than other choices in the literature (e.g., maximum residual error [39]). It should be noted that this 95th percentile could be replaced by a different percentile depending on the desired trade-off between false positives and false negatives (a lower threshold will increase the number of false positives, while a higher threshold will increase the number of false negatives). In Section 4.2 we will analyse the influence of the chosen percentile on the performance of our proposal.

2. **Overall classification**

In the second step, we assume that the nature of a sample depends on the number of normal/anomalous components present in it. Hence, we classify the sample $\mathbf{x}$ in terms of its component-wise classification: if the number of anomalous components in the sample exceeds the anomaly threshold th_{overall}, the sample is considered anomalous; otherwise, the sample is classified as normal.

$$\hat{y}_{\text{overall}} = \begin{cases} 1 & \text{if } \sum_{j=1}^{n} \hat{y}_j > th_{\text{overall}} \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

The overall threshold th_{overall} represents the 95th percentile of the number of anomalous components per sample in training data:

$$th_{\text{overall}} = 95\text{th}\left(\sum_{j=1}^{n} \hat{y}_{\text{train}_{*j}}\right) \tag{7}$$

The datasets used in this research provide us with the overall classification of the samples, but their component-wise classification is not available. That is why in (7) we operate on the basis of the estimated classification ($\hat{y}_{\text{train}}$), obtained according to (4).

These two steps result in a classification algorithm that allows us to detect anomalous samples, but also informs us of the presence of anomalous components in the samples, thus providing an explainable diagnostic containing valuable information about the nature of the process under study. An overview of this algorithm is presented in Fig. 5.

Finally, in order to study the potential of the proposed classifier, we have compared its performance with that of two different state-of-the-art approaches: threshold-based classifier [27,41] and one-class classifier [43,48]. In line with the first approach, we have used a threshold-based classifier, which classifies incoming samples as anomalous if the l_2-norm of their residuals exceeds an anomaly threshold (for reasons of comparison with the two-step classifier, we have chosen the 95th percentile of the l_2-norm of training data residuals to be the anomaly threshold). In line with the second approach, we have used a one-class SVM with non-linear kernel (RBF), trained with training data residuals and different configurations depending on the dataset (rotating machine: $\nu = 0.01, \gamma = 0.1$; hydraulic system: $\nu = 0.01, \gamma = 0.1$; body motion system: $\nu = 0.03, \gamma = 0.1$). It should be noted that the chosen configurations are the ones that yielded the best classification results (having considered possible values of ν and γ in the range 0.001–10).

In the next section we will present the overall classification results of these classifiers, and thus determine whether the novelty of our proposal, which lies in the component-wise classification step, leads to an improvement of performance in the overall classification of the samples. We will also analyse the contributions of the classification to the understanding of process nature, which is a unique feature of our proposal.

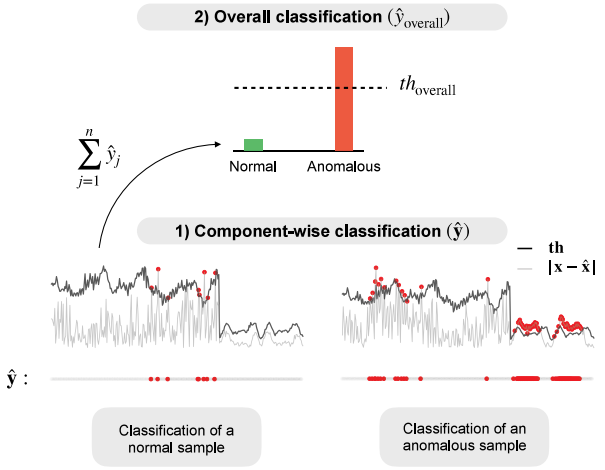

Fig. 5. Classification of the residuals of two samples (normal and anomalous) of the rotating machine dataset.

Fig. 6. Testing machine.

4. Results

We have proposed a residual-error based approach to anomaly detection consisting of a variational autoencoder and a two-step classification algorithm, and we have tested its performance in three different engineering contexts, which are described below. We also present in this section the results of the overall classification and the contribution of the component-wise classification to the understanding of process nature.

4.1. Datasets

The proposed method has been tested on three different datasets, which have been chosen to illustrate the contribution of our proposal in a wide range of engineering contexts – covering health monitoring in machines (rotating machine dataset) and industrial systems (hydraulic system dataset), as well as in the biomedical field (body motion dataset) – where anomaly detection is not only of great value, but a crucial task. In this section we describe the corresponding datasets and their preprocessing.

The states that we have selected as representative of normal/anomalous behaviour in each engineering context are also indicated below. This selection of states is a problem-dependent decision that we have made in a convenient way to better illustrate the performance of our method, choosing as normal/anomalous states those that we have considered more appropriate depending on the nature of each dataset: whether it actually contains data for states that are usually considered as faulty or anomalous for that particular kind of system or not, whether the amount of available data for normal states is larger than for faulty or anomalous states (as often occurs in practice) or not, etc. See Tables 4, 7 and 10 for the actual descriptions of the states present in each dataset and Tables 5, 8 and 11 for the normal/anomalous labelling chosen in our experiments.

4.1.1. Rotating machine

Our testing machine is shown in Fig. 6. It is a 4 kW induction motor that rotates at 1500 rpm with a supply frequency of 50 Hz. We have subjected this machine to seven different tests (Table 4), for each of which three operating variables have been measured (Table 3), thus obtaining a dataset containing vibration and current data of the rotating machine [60]. In order for the three variables to have the same range, we normalize this dataset by means of min–max scaling [61] with range [0,1]; and then we apply the windowing technique, using windows of 100 samples and a 50% overlap. Finally, we concatenate the three resulting windows (one per variable), thus obtaining samples of 300 elements containing temporal information of the three variables under study. It should also be noted that tests of mechanical nature (T1, T2) have been considered anomalous and the rest (the normal test – T3 – as well as the tests of electrical nature – T4, T5, T6, T7 –) have been considered representative of normal behaviour (Table 5).

Table 3
Variables in the rotating machine dataset.

Variable	Description
a_x	Horizontal vibration acceleration
a_y	Vertical vibration acceleration
i_r	Phase R current

Table 4
Tests in the rotating machine dataset.

Test	Machine condition
T1	Mechanical fault (eccentric mass on pulley)
T2	Combined electrical and mechanical fault
T3	Normal operation
T4	Electrical fault (5 Ω resistor in phase R)
T5	Electrical fault (10 Ω resistor in phase R)
T6	Electrical fault (15 Ω resistor in phase R)
T7	Electrical fault (20 Ω resistor in phase R)

Table 5
Rotating machine dataset expressed in terms of normal and anomalous behaviour (data size indicates the number of samples and the number of elements in the samples).

	Normal behaviour (T3,T4,T5,T6,T7)	Anomalous behaviour (T1,T2)
Size:	1990 × 300	796 × 300

Table 6
Variables in the hydraulic system dataset.

Variable	Description
$PS_1, PS_2, PS_3, PS_4, PS_5, PS_6$	Pressure
EPS_1	Motor power
FS_1, FS_2	Volume flow
TS_1, TS_2, TS_3, TS_4	Temperature
VS_1	Vibration
CE	Cooling efficiency
CP	Cooling power
SE	Efficiency factor

Table 7
Tests in the hydraulic system dataset.

Test	Cooler condition
T1	Full efficiency
T2	Reduced efficiency
T3	Close to total failure

Table 8
Hydraulic system dataset expressed in terms of normal and anomalous behaviour (data size indicates the number of samples and the number of elements in the samples).

	Normal behaviour (T1,T2)	Anomalous behaviour (T3)
Size:	88 380 × 17	43 920 × 17

4.1.2. Hydraulic system dataset

It is a dataset containing operating data of a hydraulic test rig [62], which consists of a primary working and a secondary cooling-filtration circuit. The system cyclically repeats constant load cycles and measures 17 process values (Table 6) while the condition of the cooler is quantitatively varied (Table 7). We have normalized the dataset by means of min–max scaling [61] with range [0,1] and we have considered the total failure of the cooling circuit (T3) to be representative of anomalous behaviour (Table 8).

4.1.3. Body motion dataset

The MHEALTH dataset [63] provides 23 process values corresponding to body motion and vital signs recordings (Table 9) for 10 volunteers of diverse profile while performing 12 different physical activities. Discarding the unlabelled samples from the original data, we have built our body motion dataset, consisting of 12 tests (Table 10), where each test (T1, T2, . . . ,Tn) contains

Table 9
Variables in the body motion dataset.

Variable	Description
a_{cx}, a_{cy}, a_{cz}	Chest acceleration (axis: x,y,z)
ecg_1, ecg_2	Electrocardiogram signal (lead: 1,2)
a_{ax}, a_{ay}, a_{az}	Ankle acceleration (axis: x,y,z)
g_{ax}, g_{ay}, g_{az}	Ankle angular speed (axis: x,y,z)
m_{ax}, m_{ay}, m_{az}	Ankle magnetic field (axis: x,y,z)
$a_{lax}, a_{lay}, a_{laz}$	Lower arm acceleration (axis: x,y,z)
$g_{lax}, g_{lay}, g_{laz}$	Lower arm angular speed (axis: x,y,z)
$m_{lax}, m_{lay}, m_{laz}$	Lower arm magnetic field (axis: x,y,z)

Table 10
Tests in the body motion dataset.

Test	Activity
T1	Standing still
T2	Sitting and relaxing
T3	Lying down
T4	Walking
T5	Climbing stairs
T6	Waist bends forward
T7	Frontal elevation of arms
T8	Knees bending (crouching)
T9	Cycling
T10	Jogging
T11	Running
T12	Jump front and back

Table 11
Body motion dataset expressed in terms of normal and anomalous behaviour (data size indicates the number of samples and the number of elements in the samples).

	Normal behaviour (T1,T2,T3,T4)	Anomalous behaviour (T5,T6,T7,T8,T9,T10,T11,T12)
Size:	122 880 × 23	220 315 × 23

Table 12
Overall classification results in terms of f1-score (%) for the three test sets, using random cross validation executed 5 times (we report the mean and standard deviation of all executions).

Model of healthy behaviour	Classification algorithm		Dataset		
			Rotating machine	Body motion	Hydraulic system
Variational autoencoder	Two-step classifier		**97.75** (σ = **0.23**)	**95.85** (σ = **0.34**)	97.92 (σ = 0.44)
	State-of-the-art classifiers	Threshold-based	93.83 (σ = 2.02)	95.59 (σ = 0.37)	98.42 (σ = 0.08)
		One-class SVM	94.33 (σ = 2.33)	94.38 (σ = 0.25)	**99.68** (σ = **0.02**)
Deep autoencoder	Two-step classifier		93.26 (σ = 2.70)	93.74 (σ = 1.10)	96.45 (σ = 1.56)
	State-of-the-art classifiers	Threshold-based	92.28 (σ = 2.89)	92.91 (σ = 1.00)	95.57 (σ = 2.78)
		One-class SVM	92.46 (σ = 2.65)	90.69 (σ = 0.49)	91.61 (σ = 5.44)

all the samples of the 10 volunteers while performing the activity n. We have normalized this dataset by means of min–max scaling [61] with range [0,1] and we have considered low-intensity activities (T1, T2, T3, T4) to be representative of normal behaviour (Table 11).

4.2. Overall classification results

In order to assess the performance of our proposal, we have classified the test set of each engineering system (rotating machine, hydraulic system, body motion system) using different reconstruction techniques (variational autoencoder, deep autoencoder) and classifiers (two-step classifier, traditional classifiers). The results of the classification are shown in Table 12.

According to Table 12, the VAE residuals lead to better classification results than the deep autoencoder residuals, and this can be seen in the results of the three datasets. With regard to the classification algorithm, the two-step classifier performs better than state-of-the-art classifiers in all cases except in the hydraulic system dataset, where applied on VAE residuals both approaches achieve high accuracy results and traditional classifiers win by a small gap.

In addition, we have visualized the influence of the anomaly threshold on the results of the proposed method. As mentioned before, we have chosen the 95th percentile of training data residuals as the anomaly threshold, which means that our proposal

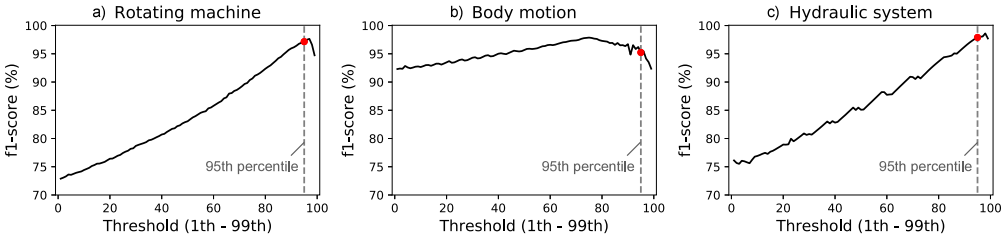

Fig. 7. Overall classification results of the proposal (VAE and two-step classifier) in terms of f1-score (%) for the three test sets, using random cross validation executed 5 times (we report the mean of all executions).

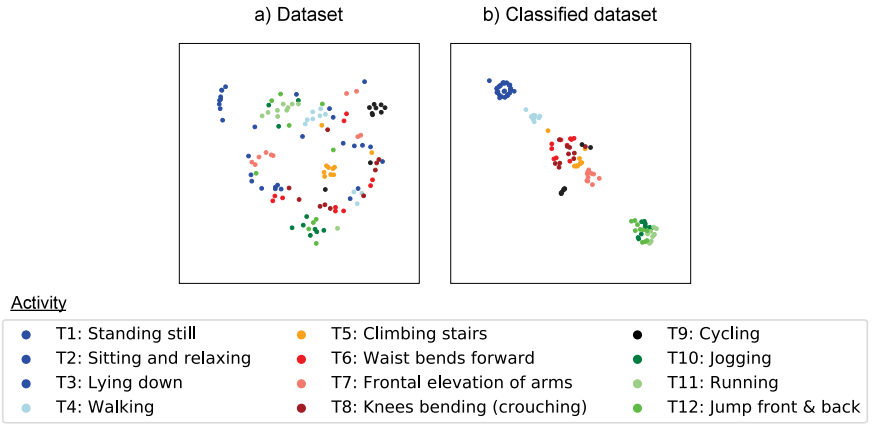

Fig. 8. Dimensionality reduction of the body motion dataset (a) and its component-wise classification (b), both obtained using a t-SNE with perplexity 30.

depends only on normal behaviour and does not need any anomaly information to perform the classification. As we can see in Fig. 7, this threshold is not optimal for any of the three datasets, and yet the two-step classifier stands out from traditional classifiers.

In summary, these results demonstrate the potential of the variational autoencoder and also the contribution to the overall classification of the samples of the component-wise classification step.

4.3. Contributions to the understanding of process nature

The proposed method reports not only the overall classification of the samples, but also their component-wise classification, thus providing valuable additional information about the samples, beyond whether they are anomalous or not. Hence, we propose to explore the contribution of this classification to the understanding of process nature.

To this end, we have analyzed the component-wise classification results of each engineering system (rotating machine, hydraulic system, body motion system) using the proposed method (two-step classifier applied on VAE residuals) to perform the classification.

4.3.1. Body motion dataset: Classes underlying normal/anomalous behaviour

In this section, we explore the contribution of our proposal through data visualization: we have used a dimensionality reduction technique (*t-stochastic neighbour embedding*, t-SNE [64]) to generate a two-dimensional (2D) representation of both the dataset (Fig. 8a) and its component-wise classification (Fig. 8b), so that we can easily notice the contribution of the classification to the understanding of the body motion system. In order to simplify the visualization of the results and favour interpretability, the 343 195 samples of the dataset and their resulting component-wise classifications were averaged by subject and activity before projection, resulting in 120 points that represent the average results for the 12 activities and the 10 subjects.

As we can see in Fig. 8b, since our proposal eliminates the variability associated with normal behaviour (T1, T2, T3, T4) and thus enhances the abnormal modes present in the samples, the component-wise classification brings to light the activities performed by the volunteers, apparently clustered according to their level of intensity. Meanwhile, it is difficult to extract any information from the representation of the original dataset (Fig. 8a).

Therefore, the proposed method not only informs us about the normal/anomalous behaviour of the samples but also gives us qualitative information about the classes (in this context, physical activities) underlying such behaviours.

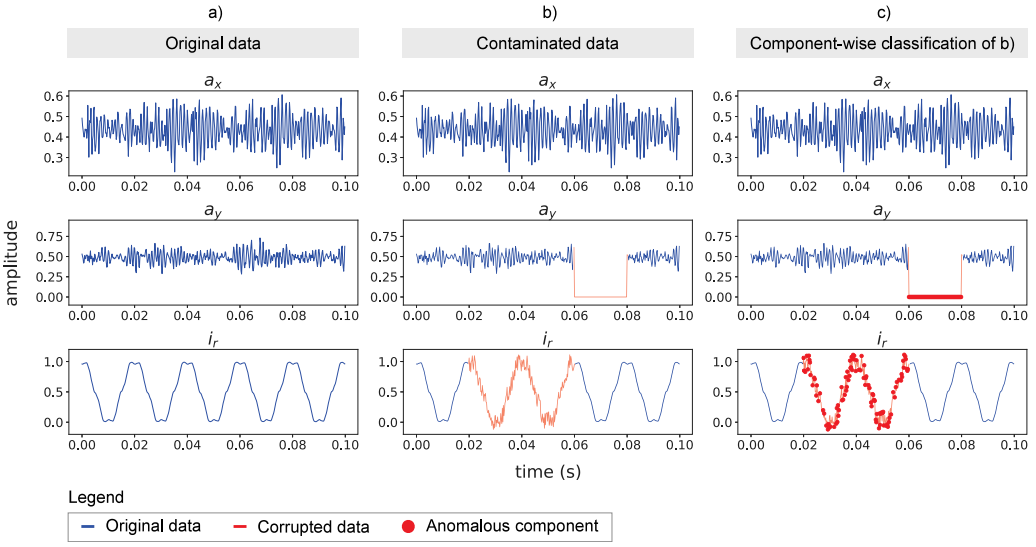

Fig. 9. Component-wise classification of contaminated data coming from the rotating machine.

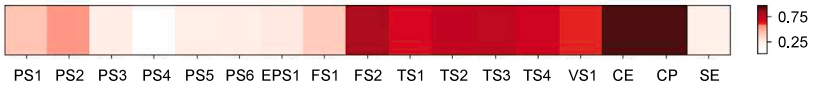

Fig. 10. Average component-wise classification results for all the anomalous samples in the hydraulic system dataset.

4.3.2. *Rotating machine dataset: Anomalous components in the samples*

Sometimes not only the nature of a sample is relevant, but also that of its elements. In such cases, the information provided by the component-wise classification step becomes of great value, as illustrated in Fig. 9.

In this figure we can see a 100 ms slice of normal behaviour data (Fig. 9a) that has been contaminated (Fig. 9b) with corrupted data in variables i_r and a_y: we have added white noise to i_r and simulated missing data, sensor fault, etc. in a_y. The component-wise classification of the contaminated data is shown in Fig. 9c.

The results of the classification shown in Fig. 9c reveal the presence of anomalous elements in the samples, matching the corrupted data. Therefore, this figure illustrates the value of the information provided by the component-wise classification step, which in this case allows us to successfully detect in which variables and instants of time the machine has deviated from its normal behaviour.

In view of these results, it can be said that the proposed method contributes to a better understanding of the process under study, since it reveals the normal/anomalous nature of each element in the samples, thus providing an explanatory picture of the normal/anomalous behaviour of the rotating machine.

4.3.3. *Hydraulic system dataset: Underlying causes of anomalous behaviour*

The component-wise classification step may also help us to identify the underlying causes of anomalous behaviour. For this purpose we have analyzed the results of the classification, not for one sample, but for all the anomalous samples in the hydraulic system dataset.

Fig. 10 shows the average component-wise classification of the anomalous samples, where we can see the contribution of each component to the anomalous behaviour of the system, being the most relevant: FS_2, $TS_1 - TS_4$, VS_1, CE and CP (volume flow, temperatures, vibration, cooling efficiency and cooling power, respectively).

Those components are either related to the temperature of the hydraulic system or present in the cooling circuit (Fig. 11). As mentioned in Section 4.1.2, the anomalous behaviour of the system comes from the total failure of the cooling circuit, so the insight provided by the component-wise classification is consistent with the nature of the process.

Therefore, the component-wise classification of the samples has proven again to be helpful for a better understanding of the processes, revealing in this case the components involved in the anomalous behaviour of the hydraulic system.

5. Conclusion

In this paper, we have proposed a residual-error based approach to anomaly detection consisting of: 1) a variational autoencoder, which is trained to reconstruct samples of healthy behaviour based on learning their underlying probability density function (pdf),

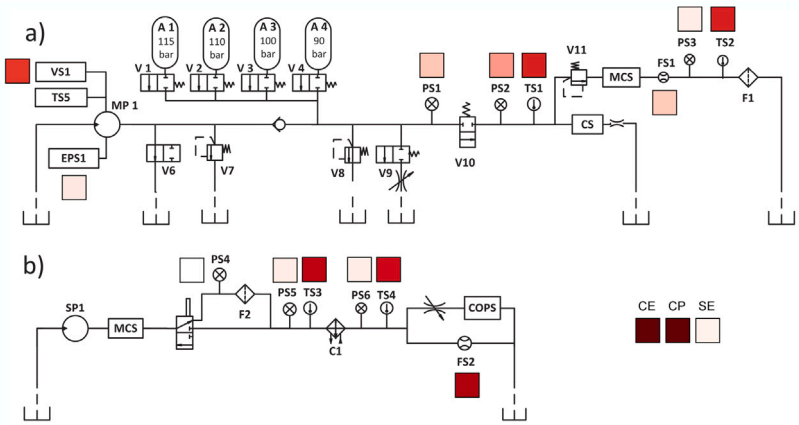

Fig. 11. Hydraulic system scheme [62] (consisting of a primary working – Fig. 11a – and a secondary cooling-filtration circuit – Fig. 11b –) including the contribution of each component to the anomalous behaviour of the system (total failure of the cooling circuit).

so that residuals of incoming samples become a measure of their deviation from healthy behaviour; and 2) a two-step classification algorithm, which classifies the incoming samples based on their residuals, reporting the normal/anomalous nature of the samples (overall classification step) and also that of their elements (component-wise classification step).

We have tested this proposal in three different engineering contexts (rotating machine, hydraulic system, body motion system), considering different reconstruction techniques (variational autoencoder, deep autoencoder) and classifiers (two-step classifier, traditional classifiers). The results of the research have proven the ability of our proposal to successfully detect anomalies in all three contexts, being the highest accuracy results those obtained with the two-step classifier applied on VAE residuals. This suggests that the ability of the VAE to learn the pdf of healthy states gives it a superior performance over the deep autoencoder. Indeed, while the deep autoencoder models the geometry of the data in the input space, it does not model its density, which makes it accurate in reconstructing healthy states, but also some non-healthy states; the VAE, in turn, constrains the reconstruction to the support of the pdf of healthy states; this seems to make deep autoencoders less discriminative for an anomaly detection approach, and thereby, less efficient. In addition, we have presented a visual analysis of the component-wise classification results that illustrates the contribution of the proposed classifier to the explainability of the anomaly decision and to the better understanding of the processes under study.

In summary, the proposed method has proven to be useful for the detection of anomalies in engineering systems and also to provide valuable and explanatory additional information about the samples, beyond whether they are anomalous or not.

CRediT authorship contribution statement

Ana González-Muñiz: Conceptualization, Investigation, Software, Writing – original draft, Writing – review & editing. **Ignacio Díaz:** Supervision, Writing – review & editing. **Abel A. Cuadrado:** Writing – Review & Editing. **Diego García-Pérez:** Writing – review & editing. **Daniel Pérez:** Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work has been financed by the Spanish National Research Agency (under grant number PID2020-115401GB-I00/AEI/10.13039/501100011033). The authors would also like to thank the financial support provided by the Principado de Asturias government, Spain through the predoctoral grant “Severo Ochoa”.

References

[1] Alrawashdeh K, Purdy C. Toward an online anomaly intrusion detection system based on deep learning. In: 2016 15th IEEE international conference on machine learning and applications (ICMLA). IEEE; 2016, p. 195–200.
[2] Ahmed M, Mahmood AN, Islam MR. A survey of anomaly detection techniques in financial domain. Future Gener Comput Syst 2016;55:278–88.
[3] Schlegl T, Seeböck P, Waldstein SM, Schmidt-Erfurth U, Langs G. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In: International conference on information processing in medical imaging. Springer; 2017, p. 146–57.
[4] Mobley RK. An introduction to predictive maintenance. Elsevier; 2002.
[5] Chandola V, Banerjee A, Kumar V. Anomaly detection: A survey. ACM Comput Surv 2009;41(3). http://dx.doi.org/10.1145/1541880.1541882.

[6] Pimentel MA, Clifton DA, Clifton L, Tarassenko L. A review of novelty detection. Signal Process 2014;99:215–49. http://dx.doi.org/10.1016/j.sigpro.2013.12.026.
[7] Qi Q, Tao F. Digital twin and big data towards smart manufacturing and industry 4.0: 360 degree comparison. IEEE Access 2018;6:3585–93.
[8] Isermann R. Model-based fault-detection and diagnosis–status and applications. Annu Rev Control 2005;29(1):71–85.
[9] Balasubramanian A, Muthu R. Model based fault detection and diagnosis of doubly fed induction generators–a review. Energy Procedia 2017;117:935–42.
[10] Nazir M, Khan AQ, Mustafa G, Abid M. Robust fault detection for wind turbines using reference model-based approach. J King Saud Univ-Eng Sci 2017;29(3):244–52.
[11] Chalapathy R, Chawla S. Deep learning for anomaly detection: A survey. 2019, arXiv preprint arXiv:1901.03407.
[12] Martinelli M, Tronci E, Dipoppa G, Balducelli C. Electric power system anomaly detection using neural networks. In: International conference on knowledge-based and intelligent information and engineering systems. Springer; 2004, p. 1242–8.
[13] Sakurada M, Yairi T. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In: Proceedings of the MLSDA 2014 2nd workshop on machine learning for sensory data analysis. 2014, p. 4–11.
[14] Oh D, Yun I. Residual error based anomaly detection using auto-encoder in smd machine sound. Sensors 2018;18(5):1308.
[15] Kingma DP, Welling M. Auto-encoding variational bayes. 2013, arXiv preprint arXiv:1312.6114.
[16] Roberts A, Engel J, Raffel C, Hawthorne C, Eck D. A hierarchical latent vector model for learning long-term structure in music. 2018, arXiv preprint arXiv:1803.05428.
[17] Kusner MJ, Paige B, Hernández-Lobato JM. Grammar variational autoencoder. 2017, arXiv preprint arXiv:1703.01925.
[18] Pu Y, Gan Z, Henao R, Yuan X, Li C, Stevens A, Carin L. Variational autoencoder for deep learning of images, labels and captions. In: Advances in neural information processing systems. 2016, p. 2352–60.
[19] Hou X, Shen L, Sun K, Qiu G. Deep feature consistent variational autoencoder. In: 2017 IEEE winter conference on applications of computer vision (WACV). IEEE; 2017, p. 1133–41.
[20] An J, Cho S. Variational autoencoder based anomaly detection using reconstruction probability. Spec Lect IE 2015;2(1).
[21] Xu H, Chen W, Zhao N, Li Z, Bu J, Li Z, Liu Y, Zhao Y, Pei D, Feng Y, et al. Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications. In: Proceedings of the 2018 world wide web conference. 2018, p. 187–96.
[22] Fan Y, Wen G, Li D, Qiu S, Levine MD, Xiao F. Video anomaly detection and localization via Gaussian mixture fully convolutional variational autoencoder. Comput Vis Image Underst 2020;102920.
[23] Zhou J, Komuro T. Recognizing fall actions from videos using reconstruction error of variational autoencoder. In: 2019 IEEE international conference on image processing (ICIP). IEEE; 2019, p. 3372–6.
[24] Gao Y, Shi B, Dong B, Chen Y, Mi L, Huang Z, Shi Y. RVAE-ABFA: Robust anomaly detection for highdimensional data using variational autoencoder. In: 2020 IEEE 44th annual computers, software, and applications conference (COMPSAC). IEEE; 2020, p. 334–9.
[25] Memarzadeh M, Matthews B, Avrekh I. Unsupervised anomaly detection in flight data using convolutional variational auto-encoder. Aerospace 2020;7(8):115.
[26] Chadha GS, Rabbani A, Schwung A. Comparison of semi-supervised deep neural networks for anomaly detection in industrial processes. In: 2019 IEEE 17th international conference on industrial informatics (INDIN), Vol. 1. IEEE; 2019, p. 214–9.
[27] Lin S, Clark R, Birke R, Schönborn S, Trigoni N, Roberts S. Anomaly detection for time series using VAE-LSTM hybrid model. In: ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing (ICASSP). IEEE; 2020, p. 4322–6.
[28] Diaz I, Hollmen J. Residual generation and visualization for understanding novel process conditions. In: Proceedings of the 2002 international joint conference on neural networks. IJCNN'02 (Cat. No. 02CH37290), Vol. 3. IEEE; 2002, p. 2070–5.
[29] Doshi-Velez F, Kim B. Towards a rigorous science of interpretable machine learning. 2017, arXiv preprint arXiv:1702.08608.
[30] Ruder S. An overview of gradient descent optimization algorithms. 2016, arXiv preprint arXiv:1609.04747.
[31] Rumelhart DE, Hinton GE, Williams RJ, et al. Learning representations by back-propagating errors. Cogn Model 1988;5(3):1.
[32] Bourlard H, Kamp Y. Auto-association by multilayer perceptrons and singular value decomposition. Biol Cybernet 1988;59(4–5):291–4.
[33] Hinton GE, Zemel RS. Autoencoders, minimum description length and Helmholtz free energy. In: Advances in neural information processing systems. 1994, p. 3–10.
[34] LeCun Y, Bengio Y, Hinton G. Deep learning. Nature 2015;521(7553):436.
[35] Bengio Y, Courville A, Vincent P. Representation learning: A review and new perspectives. IEEE Trans Pattern Anal Mach Intell 2013;35(8):1798–828.
[36] Gao S, Zhang Y, Jia K, Lu J, Zhang Y. Single sample face recognition via learning deep supervised autoencoders. IEEE Trans Inf Forensics Secur 2015;10(10):2108–18.
[37] Tewari A, Zollhofer M, Kim H, Garrido P, Bernard F, Perez P, Theobalt C. Mofa: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In: Proceedings of the IEEE international conference on computer vision. 2017, p. 1274–83.
[38] Zhao F, Feng J, Zhao J, Yang W, Yan S. Robust lstm-autoencoders for face de-occlusion in the wild. IEEE Trans Image Process 2017;27(2):778–90.
[39] Dau HA, Ciesielski V, Song A. Anomaly detection using replicator neural networks trained on examples of one class. In: Asia-Pacific conference on simulated evolution and learning. Springer; 2014, p. 311–22.
[40] Zhou C, Paffenroth RC. Anomaly detection with robust deep autoencoders. In: Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining. ACM; 2017, p. 665–74.
[41] Nguimbous YN, Ksantini R, Bouhoula A. Anomaly-based intrusion detection using auto-encoder. In: 2019 international conference on software, telecommunications and computer networks (SoftCOM). IEEE; 2019, p. 1–5.
[42] Andresini G, Appice A, Di Mauro N, Loglisci C, Malerba D. Exploiting the auto-encoder residual error for intrusion detection. In: 2019 IEEE European symposium on security and privacy workshops (EuroS&PW). IEEE; 2019, p. 281–90.
[43] Harrou F, Dairi A, Taghezouit B, Sun Y. An unsupervised monitoring procedure for detecting anomalies in photovoltaic systems using a one-class Support Vector Machine. Sol Energy 2019;179:48–58.
[44] Lee J, Lee YC, Kim JT. Fault detection based on one-class deep learning for manufacturing applications limited to an imbalanced database. J Manuf Syst 2020;57:357–66.
[45] White T. Sampling generative networks. 2016, arXiv preprint arXiv:1609.04468.
[46] Bowman SR, Vilnis L, Vinyals O, Dai AM, Jozefowicz R, Bengio S. Generating sentences from a continuous space. 2015, arXiv preprint arXiv:1511.06349.
[47] Gómez-Bombarelli R, Wei JN, Duvenaud D, Hernández-Lobato JM, Sánchez-Lengeling B, Sheberla D, Aguilera-Iparraguirre J, Hirzel TD, Adams RP, Aspuru-Guzik A. Automatic chemical design using a data-driven continuous representation of molecules. ACS Cent Sci 2018;4(2):268–76.
[48] Luo P, Wang B, Li T, Tian J. ADS-B anomaly data detection model based on VAE-SVDD. Comput Secur 2021;104:102213.
[49] Khan S, Yairi T. A review on the application of deep learning in system health management. Mech Syst Signal Process 2018. http://dx.doi.org/10.1016/j.ymssp.2017.11.024.
[50] Tieleman T, Hinton G. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Netw Mach Learn 2012;4(2):26–31.
[51] Hahnloser RH, Sarpeshkar R, Mahowald MA, Douglas RJ, Seung HS. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. Nature 2000;405(6789):947.

[52] Zeiler MD, Ranzato M, Monga R, Mao M, Yang K, Le QV, Nguyen P, Senior A, Vanhoucke V, Dean J, et al. On rectified linear units for speech processing. In: 2013 IEEE international conference on acoustics, speech and signal processing. IEEE; 2013, p. 3517–21.
[53] Ding B, Qian H, Zhou J. Activation functions and their characteristics in deep neural networks. In: 2018 Chinese control and decision conference (CCDC). IEEE; 2018, p. 1836–41.
[54] Lu L, Shin Y, Su Y, Karniadakis GE. Dying relu and initialization: Theory and numerical examples. 2019, arXiv preprint arXiv:1903.06733.
[55] Xu B, Wang N, Chen T, Li M. Empirical evaluation of rectified activations in convolutional network. 2015, arXiv preprint arXiv:1505.00853.
[56] Nwankpa C, Ijomah W, Gachagan A, Marshall S. Activation functions: Comparison of trends in practice and research for deep learning. 2018, arXiv preprint arXiv:1811.03378.
[57] Kingma DP, Ba J. Adam: A method for stochastic optimization. 2014, arXiv preprint arXiv:1412.6980.
[58] Daniluk P, Goździewski M, Kapka S, Kośmider M. Ensemble of auto-encoder based and wavenet like systems for unsupervised anomaly detection. In: Challenge on detection and classification of acoustic scenes and events (DCASE 2020 Challenge). Tech. Rep, 2020.
[59] Ásgrímsson DS, González I, Salvi G, Karoumi R. Bayesian deep learning for vibration-based bridge damage detection. In: Structural health monitoring based on data science techniques. Springer; 2022, p. 27–43.
[60] Díaz Blanco I, Cuadrado Vega AA, González Muñiz A, García Pérez D. Dataicann: datos de vibración y corriente de un motor de inducción. 2019, http://hdl.handle.net/10651/53461.
[61] Patterson J, Gibson A. Deep learning: A practitioner's approach. " O'Reilly Media, Inc."; 2017.
[62] Helwig N, Pignanelli E, Schütze A. Condition monitoring of a complex hydraulic system using multivariate statistics. In: 2015 IEEE international instrumentation and measurement technology conference (I2MTC) proceedings. IEEE; 2015, p. 210–5.
[63] Banos O, Garcia R, Holgado-Terriza JA, Damas M, Pomares H, Rojas I, Saez A, Villalonga C. mHealthDroid: a novel framework for agile development of mobile health applications. In: International workshop on ambient assisted living. Springer; 2014, p. 91–8.
[64] Maaten Lvd, Hinton G. Visualizing data using t-SNE. J Mach Learn Res 2008;9(Nov):2579–605.

Ana González-Muñiz received the M.S. degree in automation and industrial engineering in 2018 from the University of Oviedo (Spain), where she is currently pursuing the Ph.D. degree. Her research focuses on the application of deep learning techniques for process analysis and monitoring.

Ignacio Díaz received the MEng and PhD from the University of Oviedo in 1995 and 2000, where he is Associate Professor at the Electrical Engineering department since 2004. He has published more than 80 papers in journals and conferences and led R&D national and EU projects. He researchs in the application of visual analytics to complex problems in engineering and biomedicine.

Abel A. Cuadrado received the M.S. and Ph.D. degrees in electronic and control engineering from the University of Oviedo, Spain, in 1998 and 2003, respectively. There he is currently an Associate Professor with the Department of Electrical, Electronic, Communications, and Systems Engineering, teaching control theory and automation. His research interests include mostly the supervision of complex industrial processes.

Diego García-Pérez is an Assistant Professor of Electrical Engineering Department at the University of Oviedo since 2021. He received a M.Eng. in 2017 and his Ph.D. in industrial engineering in 2021. His research interests are the application of deep learning and visual analytics techniques to energy data in order to improve the energy efficiency in large buildings.

Daniel Pérez received the M.S. and Ph.D degree in industrial engineering from the University of Oviedo in 2007 and 2015, respectively. From 2006 to 2017, he worked in the University of Oviedo. He is currently an Assistant Professor with the University of León. His research interests include information visualization and machine learning techniques applied to industrial environments.