

Visual Analytics for Process Monitoring: Leveraging Time-Series Imaging for Enhanced Interpretability

Ibrahim Yousef¹, Aditya Tulsyan², Sirish L. Shah³, and R. Bhushan Gopaluni^{1*}

¹Department of Chemical and Biological Engineering, The University of British
Columbia, Vancouver, Canada

²Process Systems Engineering Laboratory, Massachusetts Institute of Technology,
Cambridge, USA

³Department of Chemical and Materials Engineering, University of Alberta,
Edmonton, Canada

*Corresponding author: bhushan.gopaluni@ubc.ca

Abstract

In the era of big data driven by the advent of the Internet of Things (IoT), process industries face the challenge of analyzing massive and complex data to extract relevant information for effective process monitoring. Despite exploring various approaches, scalability and interpretability continue to present practical limitations. To address these limitations, we propose a new framework called *visual analytics*. Visual analytics offers a new perspective on solving process monitoring problems. It involves transforming historical process data into visual clues, thereby converting traditional fault detection problems into image classification problems. This approach allows process experts to visually analyze patterns and textures within the data, making interpretation much easier compared to traditional time domain analysis. Moreover, by treating process data as images, visual analytics can leverage a wide range of computer vision techniques, including convolutional neural networks (CNNs), to accurately classify and detect faults. By integrating human visual perception with advanced computer vision techniques, visual analytics enables the effective analysis of massive and complex process data. To empirically validate the proposed visual analytics approach, we conduct experiments on both the simulated continuous stirred tank heater (CSTH) benchmark and the industrial arc loss benchmark. The experimental results from both benchmarks demonstrate that the proposed visual analytics approach yields competitive performance in predicting process faults while enhancing interpretability by providing meaningful and informative visual representations.

Keywords— Fault detection and diagnosis (FDD), Deep learning, Image analysis, Time-Series imaging, Convolutional neural networks (CNNs), Applications.

1 Introduction

With the increasing automation of industries and advancement in connectivity, modern industries generate large volumes of data continuously [1, 2]. These industrial systems are equipped with a wide range of sensors strategically placed throughout the process, resulting in data being dispersed across multiple dimensions and time [3, 4]. This type of data, referred to as time-series data, encompasses sequential observations or measurements recorded over time, capturing the temporal aspect of system behavior [5, 6]. The classification of these time-series data is crucial in detecting faults, ensuring safety, and upholding high product quality standards [7]. In fact, process monitoring tasks, such as fault detection and diagnosis (FDD), can be configured as time-series classification (TSC) problems.

Traditionally, TSC methods have relied on manually extracting relevant features from the input data [8]. The goal of these methods is to identify important local or global patterns within the time-series that are associated with specific categories or classes [9]. A common feature-based method is the k -nearest neighbors (k -NN) classifier with handcrafted features [10]. In this approach, a set of relevant features is manually extracted from each time-series, such as statistical features (e.g., mean, standard deviation) or frequency domain characteristics (e.g., Fourier coefficients). These features capture key aspects of the time-series data that are assumed to be relevant for classification. The k -NN algorithm then classifies an unseen time-series by measuring the similarity between its extracted features and those of labeled examples in the training set.

However, traditional feature-based approaches have significant drawbacks. Firstly, their performance heavily relies on selecting relevant features, which can be labor-intensive and time-consuming [11]. Different problems may require distinct sets of discriminatory features, limiting the generalizability of these methods. Secondly, feature extraction can result in the loss of information present in the original data [12]. For example, fine-grained temporal changes that are important for understanding the behavior of a dynamic process can be smoothed out or overlooked during feature extraction. Moreover, the computational complexity of these methods evolves according to a power law with respect to the input data size. As industrial data scale up, the computational demands of these methods increase exponentially, rendering them impractical for managing the large-scale datasets prevalent in industrial applications [13].

More recently, deep learning (DL) has sparked numerous breakthroughs across various problem domains, including computer vision and natural language processing [14]. The core strength of DL models lies in their ability to directly learn high-level representations from input data, bypassing the tedious feature engineering process [15]. As a result, DL has garnered significant research interest in addressing TSC problems. For instance, Wang et al. [11] proposed deep neural networks such as multilayer perceptron (MLP) and long short-term memory (LSTM) networks as robust baseline approaches for TSC. These models achieved competitive performance in TSC tasks, effectively learning discriminative representations from raw time-series data, thereby highlighting the effectiveness of DL in TSC. In addition, Zerveas et al. [16] applied the transformer architecture, originally developed for natural language processing tasks, to the realm of multivariate time-series classification (MTSC). By leveraging self-attention mechanisms, transformers directly capture temporal dependencies in time-series data.

TSC presents unique challenges compared to traditional supervised learning for structured data, as algorithms must effectively handle and exploit the temporal information embedded within the signal [9]. Interestingly, there are striking parallels between TSC and computer vision problems like image classification and object localization [17]. In the latter, successful algorithms learn from the spatial information contained in images. Similarly, in TSC, the core problem is fundamentally the same, albeit with one fewer dimension [13]. Inspired by these observations, the opportunity for applying DL-based computer vision algorithms to TSC becomes apparent. One such example is using convolutional neural networks (CNNs), a powerful class

76 of models commonly used for image classification. Although CNNs were originally designed to operate ex-
77 clusively on spatial data, they can be adapted to handle time-series data by treating the temporal dimension
78 as analogous to the spatial dimension in images [18]. Using the hierarchical feature extraction capabilities,
79 CNNs have achieved competitive results in TSC and FDD tasks, often surpassing traditional approaches
80 [19–22].

81 At the same time, many practical limitations remain, especially when dealing with industrial process
82 data. First, process variables in data samples are typically arranged according to their order in the process-
83 ing procedure. Although CNNs can effectively learn correlations between variables within the same local
84 receptive field, they may not fully capture the various correlations among variables that are distant in the
85 physical topological structure (i.e., variables order within the dataset) [20]. The limited scope of the recep-
86 tive field in the first convolutional layers prevents the learning of correlations between process variables that
87 are not in the same local receptive field. Consequently, higher-level representations inadequately capture
88 these correlations. Furthermore, interpretability is crucial for understanding and explaining model predic-
89 tions, particularly for FDD applications [23]. However, CNNs are black-box models, making it challenging
90 to interpret their decisions and restricting their practical applicability in critical industrial settings [24].

91 To address the aforementioned gaps, we introduce a new paradigm for process monitoring called *visual*
92 *analytics*. The main idea behind visual analytics involves the transformation of time-series data (i.e., process
93 data) into visual images, thereby enabling the use of computer vision algorithms (e.g., CNNs) [25]. This
94 approach capitalizes on the strengths of CNNs while acknowledging the inherent dissimilarities between
95 time-series data and images. By representing the data as visual images, process operators can interact with
96 the data more intuitively compared to time-series data analyzed in the time domain. In addition, this visual
97 representation allows operators to develop a deeper understanding and intuition in relating different image
98 patterns to different process operating modes.

99 The remainder of this paper proceeds as follows. In Section 2, we provide the necessary background
100 and discuss related work. Section 3 presents a supervised learning-based visual analytics framework for
101 FDD and describes the network architecture and its main building blocks. In Section 4, we describe the
102 implementation details of our proposed approach with a simulation and industrial case studies. Finally, we
103 conclude with closing remarks in Section 5, highlighting potential future prospects.

104 2 Background & Related Work

105 In this section, we begin by introducing the fundamental concepts and definitions relevant to our study.
106 Following that, we provide an overview of TSC and delve into the foundational aspects of convolution
107 operations, which serve as the primary component in CNNs. Lastly, we discuss two widely known time-
108 series imaging tools.

109 2.1 Definitions

110 The focus of this work revolves around time-series data. A univariate time-series signal, denoted as
111 $S = \{s_1, s_2, \dots, s_L\}$, represents a sequence of L chronologically ordered observations recorded over time.
112 Each observation is associated with a timestamp from the set $T = \{t_1, t_2, \dots, t_L\}$. When multiple time-series
113 signals are recorded simultaneously by a set of p sensors, we refer to the data as a multivariate time-series
114 (MTS) signal denoted as $X = \{S_1, S_2, \dots, S_p\}$, where $S_i \in R^L$. We consider MTS with a fixed and synchro-
115 nized sampling along all dimensions. As a result, we omit the time index from the MTS definition. In this
116 work, we consider an MTS dataset $D = \{(X_1, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)\}$, containing a collection of paired

117 samples (X_i, Y_i) . Each sample consists of an MTS signal X_i with p dimensions and a length of L , and its
118 corresponding label Y_i . The task of TSC involves training a classifier C on dataset D that maps an input X_i
119 to its true label Y_i (i.e., $C : X \rightarrow Y$). In the context of FDD, the labels represent the operating condition
120 of a system, where $Y = 0$ denotes a normal condition and $Y = 1$ denotes a faulty condition. Therefore,
121 the goal of TSC in FDD is to accurately classify the input MTS signals into their corresponding operating
122 conditions.

123 2.2 Time-series classifiers

124 Broadly speaking, time-series classifiers can be grouped based on the algorithmic technique used into three
125 categories: distance-based, feature-based, and deep learning approaches.

126 Distance-based approaches rely on distance measures to evaluate the similarity/ dissimilarity between
127 pairs of time-series [26]. A significant research effort has been dedicated to the development of “elastic”
128 distance measures that compensate for small misalignments between time-series [27]. These measures seek
129 to account for variations in the alignment of time points, allowing for more robust and accurate similarity
130 comparisons. Among these elastic measures, dynamic time warping (DTW) has emerged as one of the most
131 widely used measures [28]. In fact, DTW, in combination with 1-nearest neighbor (DTW+1NN), has gained
132 significant recognition in the field of TSC and has long been considered a benchmark method, often referred
133 to as the gold standard [29].

134 Next, the feature-based category refers to a group of TSC algorithms that rely on extracting relevant
135 features from the time-series data [30]. It can be further divided into two main families: interval-based and
136 dictionary-based approaches. Interval-based approaches use subsequences from the time-series, extracting
137 discriminatory features using statistical measurements. One popular algorithm within the interval-based
138 family is the time-series forest (TSF) [31]. TSF constructs an ensemble of decision trees using randomly
139 selected intervals from the time-series and their corresponding statistical feature values (e.g., mean, slope,
140 and standard deviation).

141 Furthermore, dictionary-based approaches involve discretizing time-series into symbolic sequences, ex-
142 tracting words from these sequences using a sliding window, and quantifying the frequency of each word in
143 a predefined dictionary [32]. The bag of symbolic Fourier approximation symbols (BOSS) algorithm is a
144 prominent example in the dictionary-based family [33]. BOSS represents time-series as bags of words, where
145 each word corresponds to a symbolic Fourier approximation (SFA) coefficient. SFA is a technique that
146 converts time-series data into a compact symbolic representation by approximating their Fourier transforms
147 [34]. BOSS constructs a dictionary of SFA words using a training dataset and maps each time-series into
148 a histogram representation based on the frequency of the SFA words occurring within it. Class labels can
149 then be assigned using similarity or distance measures between the histograms.

150 DL approaches for TSC involve using neural networks that learn hierarchical representations from the
151 input time-series data [35]. A neural network is considered *deep* when it consists of more than one layer
152 between the input and output layers. Specifically, a deep neural network (DNN) is composed of K layers (i.e.,
153 parametric functions), where each layer serves as a representation of the input domain [36]. The simplest
154 architecture within DL models is MLP, also known as a fully connected network (FCN) [11]. In an MLP,
155 each neuron in layer k_i is connected to every neuron in layers k_{i+1} and k_{i-1} , with $i \in [2, K - 1]$. These
156 connections are modeled by the weights within the neural network, enabling the network to capture complex
157 relationships within the data. One impediment to adopting MLPs for TSC is that they do not exhibit
158 any spatial invariance. In other words, each time stamp has its own weight, and the temporal information
159 is lost. To address the unique characteristics of time-series data, recurrent neural networks (RNNs) were

introduced [37, 38]. RNNs maintain a hidden state that can capture information from previous time steps, allowing them to detect patterns in sequential data (e.g., time-series data). However, vanilla RNNs suffer from vanishing and exploding gradients, limiting their ability to capture long-term dependencies [39]. LSTM is a variant of RNNs that overcomes the limitations of vanilla RNNs [40]. Its architecture includes a cell state, input gate, forget gate, and output gate, allowing it to selectively remember or forget information over long sequences while avoiding the vanishing gradient problem [41, 42]. This makes LSTM networks particularly effective for tasks involving sequential data, including TSC and FDD [43, 44].

In this work, we compare and validate our proposed visual analytics framework by benchmarking it against state-of-the-art models in each of the three categories of time-series classifiers: i) distance-based category: DTW+1NN; ii) feature-based category: TSF and BOSS; iii) DL category: LSTM. A conceptual comparison between the aforementioned time-series classifiers is presented in Figure 1, showing the distinct approaches used by each classifier.

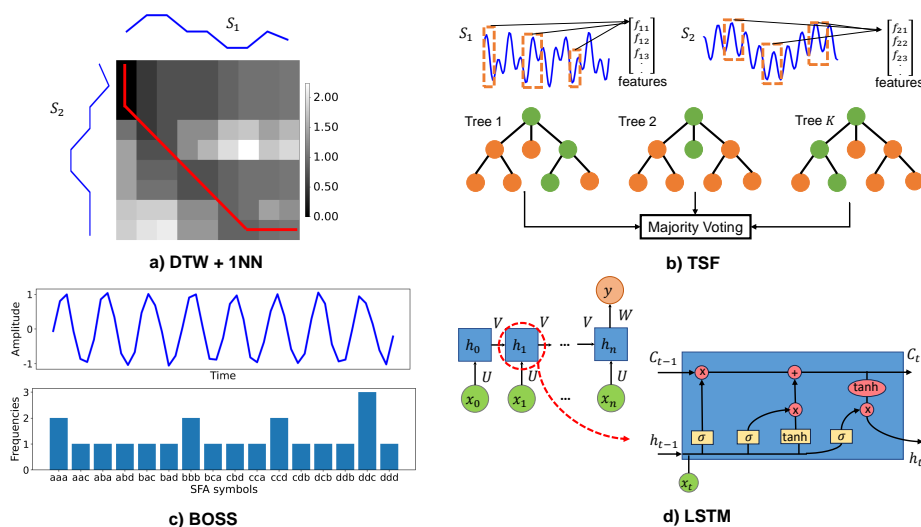


Figure 1: A conceptual illustration of state-of-the-art time-series classifiers: a) DTW+1NN: measures the similarity between time-series using DTW; b) TSF: uses random forest on features extracted from raw time-series data; c) BOSS: constructs a histogram-based representation of time-series using SFA; and d) LSTM: captures temporal dependencies in time-series data through its memory state.

2.3 Convolution operations

Convolution operations are the key functional operations that CNNs use to extract features from the input data [45]. In essence, a convolution operation involves an input signal and a kernel (an operator function). One can think of convolution as a mathematical operation that seeks to transform the input data to uncover and extract relevant features [46]. Convolution can be applied to signals of varying dimensions, such as 1D, 2D, or 3D. Notably, the dimensions of the kernel must align with those of the input. In practical applications, 1D convolution finds common use in processing audio signals [47, 48], while 2D convolution is widely used for image analysis tasks [49, 50]. Similarly, 3D convolution plays a role in video processing applications [51]. For the purpose of this paper, the discussion is narrowed to 1D and 2D convolutions, as our work does not involve video data.

Convolution operations involve sliding a kernel over input data and performing a dot product to extract

183 features, which are referred to as feature maps [52]. The kernel, also known as a filter or a convolutional
 184 filter, is a set of weights that defines the transformation applied to the input. The specific design of the
 185 kernel determines the nature of the extracted features. These features encode various hidden aspects of the
 186 data, such as trends and variations. The kernel is typically smaller in size than the input data, reducing
 187 computational complexity. In 1D convolutions, the kernel is a weight vector, and the resulting feature map
 188 is obtained by adding a bias term to the sliding dot product between the 1D input data and kernel weights
 189 [53]. Figure 2 demonstrates a 1D convolution operation with a gradient kernel on a time-series signal. On
 190 the other hand, 2D convolutions are applied to 2D data (e.g., images), and the kernel takes the form of a
 191 matrix of weights. The feature maps produced by convolutions indicate the degree of similarity between
 192 the input and the kernel pattern. Convolution operations share similarities with feature-based methods, as
 193 both approaches rely on the identification and frequency of occurrence of specific patterns or motifs in the
 194 input data. In convolution operations, multiple kernels with different weights are used to capture a wide
 195 range of patterns and variations within the data. This combination of kernels allows for the detection of
 196 complex and discriminative features. The success of CNNs for TSC and image classification demonstrates
 197 the effectiveness of convolutional kernels as the foundation for extracting informative features from input
 198 data.




Figure 2: An illustration of a 1D convolution operation. The gradient kernel, $[-1, -1, 0, 1, 1]$, is applied to the 1D signal to extract changes in amplitude or slope. The resulting convolved signal highlights regions where the input signal exhibits positive and negative gradients. The convolved signal can provide insights into the overall trend and direction of changes in the original time-series data.

199 2.4 Encoding time-series into images

200 Understanding and analyzing complex systems in the time domain poses a significant challenge in many
 201 scientific and engineering domains. Traditional time-series analysis methods often extract features that
 202 fail to capture the temporal evolution and dynamics of such systems. Consequently, researchers in signal
 203 processing and computer science have been exploring methods to represent temporal characteristics of time-
 204 series signals visually, using 2D images. Such methods provide a visual framework to capture, interpret, and
 205 extract meaningful temporal information from dynamic processes. Additionally, imaging tools transform raw
 206 time-series data into visual representations, facilitating the use of a wide range of image analysis algorithms
 207 (e.g., CNNs). Two prominent imaging techniques are the Gramian Angular Field (GAF) and the Recurrence
 208 Plot (RP).

209 A GAF is a 2D visual representation of a univariate time-series, originally introduced by Wang and
 210 Oates, which captures information about the static behavior of the time-series [54]. Figure 3 illustrates the
 211 step-by-step instructions for encoding a univariate time-series as a GAF image. First, the scaled time-series
 212 $\hat{S} = \{\hat{s}_1, \hat{s}_2, \dots, \hat{s}_L\}$ is transformed from the space coordinate system to polar coordinates. The time step
 213 t_i is encoded as the radius r_i and the scaled value \hat{s}_i of the time-series is encoded as the angular cosine θ_i ,
 214 given by:

$$\begin{aligned} r_i &= \frac{t_i}{L}; & i &\in L \\ \theta_i &= \cos^{-1}(\hat{s}_i); & \hat{s}_i &\in [0, 1] \end{aligned} \quad (1)$$

215 where L represents the time-series length. Once \hat{S} is transformed into polar coordinates, the square GAF
 216 matrix is constructed. In the GAF matrix, each entry represents the pairwise cosine distance between two
 217 angles, i.e.:

$$\text{GAF} [i, j] = \cos(\theta_i + \theta_j); \quad i, j = 1, 2, \dots, L \quad (2)$$

218 GAF offers several key advantages, making it a valuable tool for capturing temporal dynamics. Firstly,
 219 GAF preserves the temporal order of the original time-series, ensuring that the sequential nature of the
 220 data is maintained in the resulting image representation. In addition, GAF is invariant to monotonic
 221 transformations, meaning that it can capture the same underlying patterns regardless of scaling or shifting
 222 of the time-series values.




Figure 3: An illustration of the transformation from a sinusoidal time-series signal into a 2D GAF representation. Left: a scaled 1D time-series $\hat{S} = \{\hat{s}_1, \hat{s}_2, \dots, \hat{s}_{20}\}$ with $L = 20$ time steps. Middle: the first step is to represent \hat{S} in polar coordinates using Equation 1. Right: the GAF image, a square matrix of size 20×20 , is obtained using Equation 2.

223 Furthermore, dynamic nonlinear systems often exhibit recurrent behavior (e.g., periodicities and oscil-
 224 lations) that can be challenging to observe in the time domain. To tackle this challenge, Eckmann et al.
 225 introduced RP, a 2D visual representation of higher dimensional phase space trajectories [55]. RP is a square
 226 matrix that reveals at which points the m -dimensional phase space trajectory revisits a previously visited
 227 state. In this work, we consider the non-binarized version of RP, as proposed by [56], to avoid informa-
 228 tion loss when the matrix is binarized. In practice, one performs two steps to obtain an RP image from a
 229 univariate time-series. First, an embedding dimension m is chosen. The m -dimensional phase space \vec{S} is
 230 constructed from the scaled time-series $\hat{S} = \{\hat{s}_1, \hat{s}_2, \dots, \hat{s}_L\}$ using the time-delay embedding method (i.e.,
 231 $\vec{s}_i = (\hat{s}_i, \hat{s}_{i+1}, \dots, \hat{s}_{i+m-1})$). Next, the RP matrix is calculated as follows:

$$\text{RP} [i, j] = \|\vec{s}_i - \vec{s}_j\|; \quad i, j = 1, 2, \dots, R \quad (3)$$

232 where R is the total number of considered states \vec{s} (i.e., $R = L - m + 1$) and $\|\cdot\|$ is the Euclidean norm.

233 Each pixel in RP denotes the Euclidean distance of two states in the m -dimensional phase space. The full
 234 procedure for encoding a univariate time-series as an RP image is shown in Figure 4. The colors in the RP
 235 image indicate the closeness of the states in the 2D phase space according to the corresponding color bar.
 236 As shown in Figure 5, GAF and RP, both display texture and typology, which provide hints about the *static*
 237 and *recurrent* behaviors of the 1D time-series, respectively.




Figure 4: An illustration of the encoding map from a raw 1D time-series signal to a 2D RP image. Left: a scaled 1D time-series $\hat{S} = \{\hat{s}_1, \hat{s}_2, \dots, \hat{s}_{20}\}$ with $L = 20$ time steps. Middle: the m -dimensional phase space trajectory is constructed from X using the time-delay embedding. In this examples, $m = 3$; hence, the states, represented in dots, $\vec{s}_i = (\hat{s}_i, \hat{s}_{i+1}, \hat{s}_{i+2})$. Right: the RP image, an 18×18 matrix, is obtained using Equation 3.

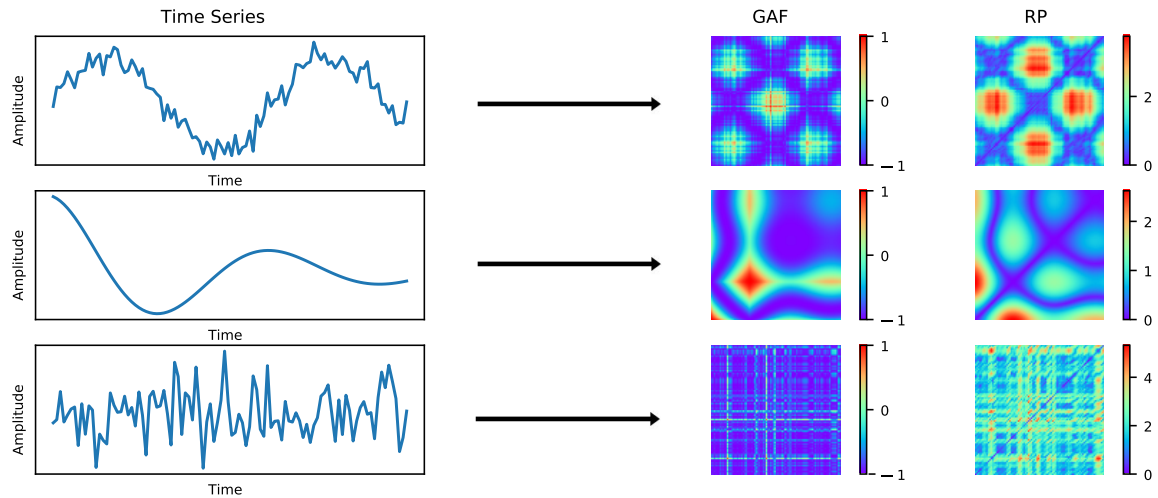


Figure 5: An illustration of the qualitative interpretations of GAF and RP representations for three distinct time-series signals.

238 3 Proposed Approach

239 In this section, we present a new end-to-end visual analytics framework for industrial fault detection using
 240 1D and 2D convolution operations. The proposed approach obtains visual data representations from input
 241 MTS signals in a *supervised* manner, i.e., the proposed model takes an annotated MTS dataset to learn

242 visual representations. The main objective of the proposed network is to maximize the visual distinction
 243 between the visual representation of samples across different classes. The overall architecture of the proposed
 244 visual analytics framework is shown in Figure 6.

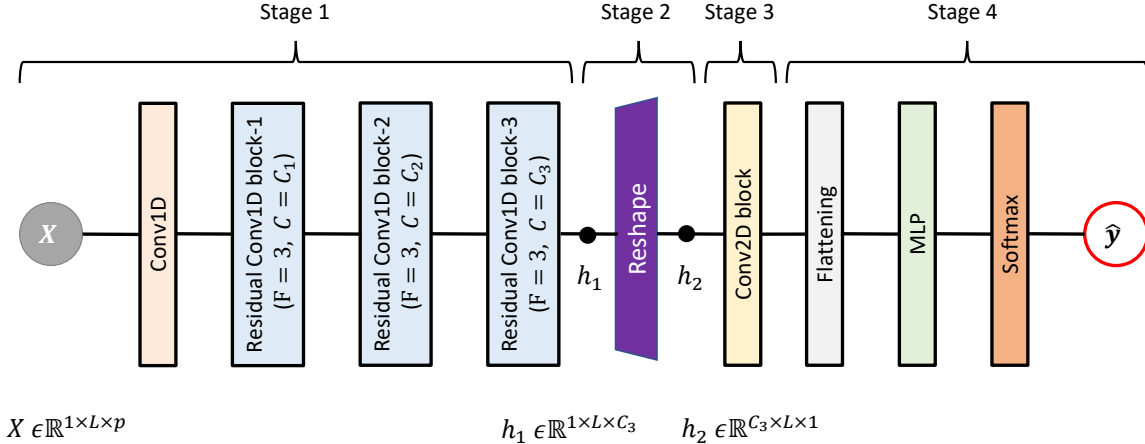


Figure 6: The proposed network architecture. The residual block configuration is presented in Figure 7. Abbreviations: F - kernel size, C - number of kernels, L - the size of the input signal (i.e, the number of time steps), p - number of variables, \hat{y} - predicted class.

245 The proposed network is inspired based on two influential architectures: AlexNet [57] and ResNet [50].
 246 Like AlexNet, the proposed network uses MLP for making predictions. Also, we implement dropout as a
 247 regularization technique to prevent overfitting. By randomly dropping out units during training, we enhance
 248 the network’s ability to generalize and avoid relying too heavily on specific neurons [58]. Furthermore, the
 249 proposed network adopts the concept of residual blocks from ResNet. Using residual connections, we address
 250 the vanishing gradient problem and facilitate the training of deeper networks. The residual blocks enable the
 251 network to capture residual information, making it easier to propagate gradients through the network and
 252 effectively learn both shallow and deep features. This improves network performance, allowing for better
 253 representation learning and enhanced prediction performance.

254 Stage 1 of the network focuses on the initial processing and feature extraction from the input data
 255 X. This stage comprises a 1D convolutional layer followed by a series of residual blocks. The primary
 256 purpose of the initial 1D convolutional layer is to transform the input data X into a suitable format for
 257 subsequent layers while extracting low-level features from the data. Next, we use a series of “bottleneck” 1D
 258 residual blocks, consisting of three convolutional layers: a 1×1 convolution, a 1×3 convolution, and another
 259 1×1 convolution (refer to Figure 7). This design effectively reduces the number of trainable parameters,
 260 resulting in a computationally efficient network. In addition, we increment the number of 1D kernels after
 261 each residual block to enable the network to capture increasingly higher-level features as the information
 262 propagates through the layers. The number of residual blocks, as well as the number of kernels within each
 263 residual block, are hyperparameters that need to be prespecified based on the complexity of the problem and
 264 data requirements. We show three blocks in Figure 6 for illustration purposes. Note that we omit pooling
 265 layers in the network to preserve valuable spatial information. We also apply “same” padding and a stride
 266 size set to 1. Overall, stage 1 of the network seeks to learn discriminative visual representations via 1D
 267 convolution operations.

268 The output of stage 1 is a combination of C_3 feature maps, corresponding to the number of 1D kernels

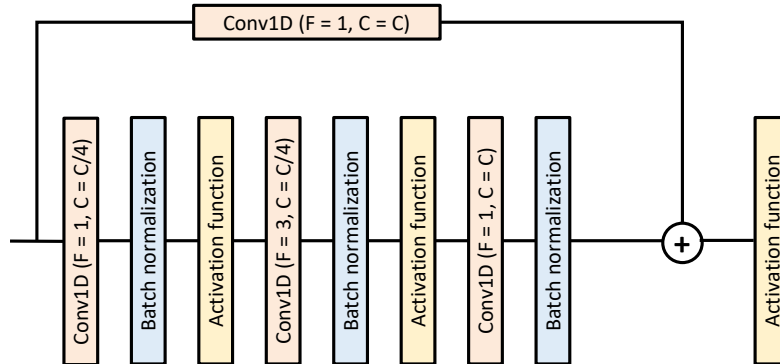


Figure 7: The residual block configuration. Abbreviations: F - kernel size, C - number of kernels.

269 in the last residual block of size L . Each feature map captures the activation level of its corresponding
 270 kernel, indicating the response of that specific kernel pattern across the features of the input X . To visualize
 271 these features and gain a deeper understanding of the network’s representations, we perform a reshaping
 272 step (stage 2). This step involves vertically stacking the 1D feature maps to form a 2D-like matrix with
 273 dimensions $C_3 \times L \times 1$ (height \times width \times channels). This reshaping enables the visualization of the learned
 274 features using color mapping techniques. Each pixel in the resulting image represents the activation level of
 275 a particular feature map at a specific spatial location. Figure 12 presents six samples of the images obtained
 276 in Stage 2.

277 Next, the one-channel image obtained from stage 2 is “visually” recognized, and visual features and
 278 patterns are learned via 2D convolution operations. Stage 3 consists of stacked 2D convolutional blocks.
 279 Each 2D convolution block contains a 2D convolutional layer followed by a batch normalization layer and
 280 a non-linear activation layer. The number of 2D convolution blocks used may vary depending on the
 281 application. For demonstration purposes, we use a single 2D convolution block. The visual features are
 282 extracted using a stride size set to 1, and we use “valid” padding to reduce the dimension of the feature
 283 maps, thus reducing the computational cost. The 2D convolutional blocks enable the network to effectively
 284 capture spatial information and extract discriminative visual features from the input data.

285 In the last stage of the network, the feature maps obtained from the last 2D convolutional layer are
 286 flattened before being fed into an MLP network. Flattening is the process of converting the multidimensional
 287 feature maps into a one-dimensional vector. The MLP network consists of multiple dense layers that map
 288 the extracted visual features into a scalar, representing a class label. The output layer of the MLP network
 289 uses the softmax activation function, which produces a probability distribution over the classes. The number
 290 of neurons in the output layer corresponds to the number of classes. For binary classification problems, the
 291 number of output neurons is two. The class label associated with the highest probability is considered the
 292 predicted class for a given input. During training, the network is trained to minimize the loss function.
 293 Specifically, we use the binary cross-entropy loss for our case (i.e., binary classification problems). Next, the
 294 proposed network is trained by minimizing the loss function using an optimization algorithm. Therefore,
 295 the model learns to assign high probabilities to the correct class and lower probabilities to the other classes.
 296 The optimization algorithm and its arguments are hyperparameters to be specified before training.

4 Applications of the Proposed Approach

In this section, we use two benchmarks to assess the effectiveness of our proposed visual analytics approach. We start with the simulated continuous stirred tank heater (CSTH) benchmark to better demonstrate the applicability and benefits of our approach in a controlled setting. By analyzing the simulated data, we show the interpretability and insights provided by our visual analytics approach. Next, we move on to the industrial arc loss benchmark, which serves as a large-scale scenario to test the scalability and robustness of our proposed approach. The proposed approach will be compared with DTW+1NN, TSF, BOSS, LSTM, GAF followed by CNN, and RP followed by CNN models.

4.1 A simulation case study: the CSTH system

The CSTH system is a dynamic nonlinear system reported in [59]. Figure 8 shows the feedback control system used for the CSTH system. In this system, hot water (HW) and cold water (CW) are mixed, heated by steam flowing through a heating coil, and eventually drained from the tank. To ensure system stability, a closed-loop control system is implemented to regulate the tank’s temperature, level, and CW flow. The input signals for the system consist of steam, HW, and CW valve openings, while the controlled variables are the CW flow, tank level, and temperature. The CSTH model can be classified as a semi-empirical model, combining first principles equations and algebraic equations derived from experimental data. The left-hand column of Figure 9 illustrates measurements of the tank’s level, temperature, and CW flows acquired under normal operating conditions. To formulate a binary fault classification problem, we are considering three scenarios resulting from instrumentation faults and errors. These scenarios are as follows: i) an abrupt pulse change introduced into the level transmitter signals (Figure 9II), ii) a malfunction in the temperature controller characterized by a random change in the controller parameters (Figure 9III), and iii) a random sinusoidal noise introduced into the CW flow controller output signals (Figure 9IV). Table 1 provides a comprehensive overview of the simulated dataset. It is worth noting that this case study specifically addresses a binary classification problem, where discrete outputs are used: normal data inputs are represented by $Y = 0$ and faulty data inputs by $Y = 1$.

Table 1: CSTH data summary

Number of total samples (N)	8000
Training: validation: testing ratio	70:10:20
Number of variables (p)	3
Sampling frequency	1 sec
Signal length (L)	200
Number of classes	2
Class ratio	50:50 (balanced)

4.1.1 Simulation study setup

For model evaluation, we use a hold-out strategy. The entire CSTH dataset is split randomly into three balanced subsets: i) a training set for training the models, ii) a validation set for optimizing the models’ hyperparameters, and iii) a testing set for testing the models on unseen data. We systematically explore different hyperparameter combinations using a random search to identify well-performing model configurations. This involves searching over a manually predefined search space, where various hyperparameter combina-

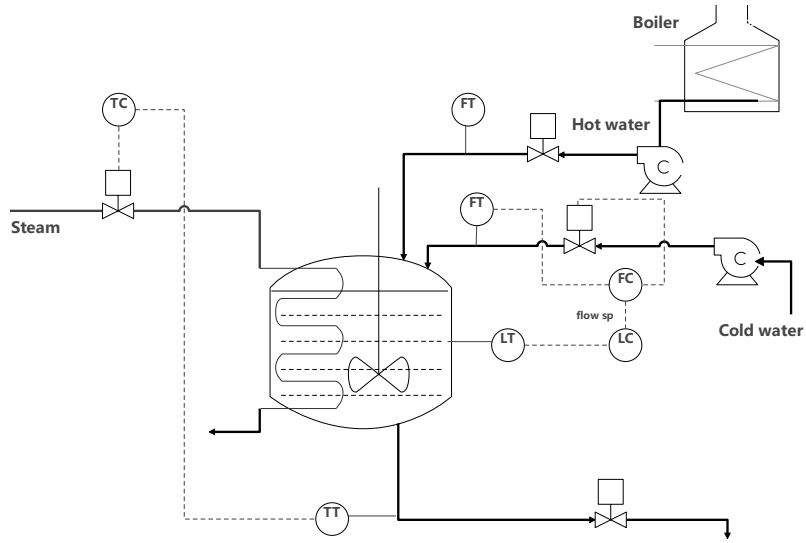


Figure 8: The CSTH feedback control system.

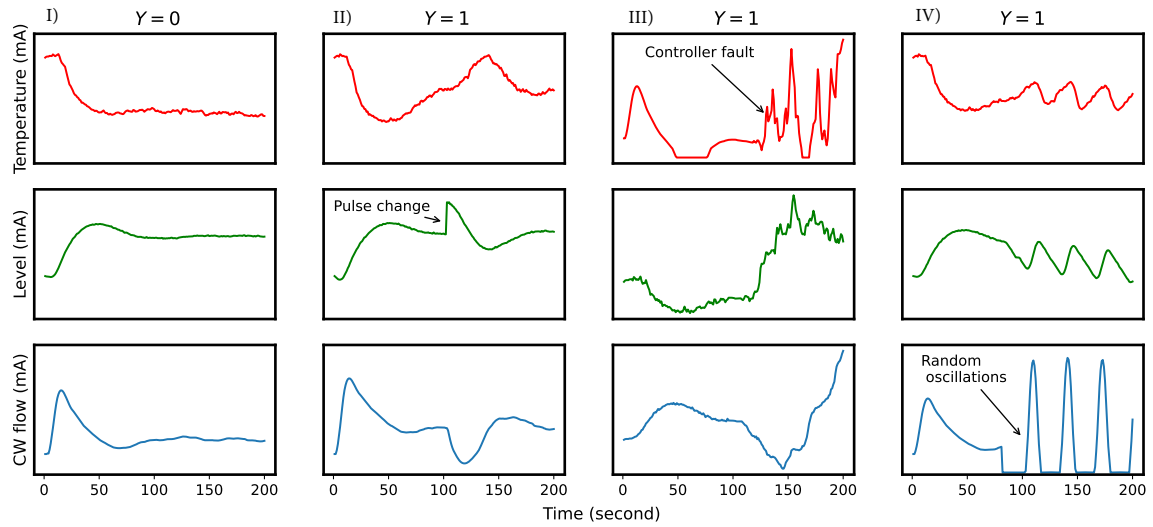


Figure 9: A visual comparison between a smooth operation ($Y = 0$) and a faulty operation ($Y = 1$).

328 tions are sampled. For each sampled set of hyperparameters, we train the models on the training set and
 329 estimate their performance using the validation set. After conducting multiple trials, which are manually
 330 set with consideration for the size of the hyperparameter search space and computational limitations, we
 331 select the models with the hyperparameter configurations that yield the best results on the validation set.
 332 These selected models are then retrained using the entire training set. To ensure effective model training,
 333 we implement an early stopping criterion that prevents overfitting of the data. Specifically, if the validation
 334 loss does not exhibit improvement for a consecutive number of epochs, the training process is halted. A
 335 patience value of 10 epochs is selected for the early stopping criterion. Subsequently, the retrained models
 336 are evaluated on the unseen testing set to provide reliable and unbiased performance metrics.

337 To assess the classification performance of our proposed visual analytics approach, we use carefully tuned
 338 implementations of the following models:

- 339 • **DTW + 1NN**: {implementation = fast DTW, max warping window = 400}
- 340 • **TSF**: {number of classifiers = 3, ensemble method = voting, number of trees = 200, minimum interval
 341 length = 3, number of attributes = 3 (mean, slope, and standard deviation)}
- 342 • **BOSS**: {number of classifiers = 3, ensemble method = voting, word size= 7, number of bins= 20,
 343 window size= 10, window step= 2}
- 344 • **LSTM**: {LSTM layers = 2, LSTM units = 32, recurrent activation function = hard sigmoid, hidden
 345 layers = 1, hidden neurons = 128, hidden activation function = ELU, batch normalization = True,
 346 dropout coefficient = 0.3, L_2 regularization penalty = 0.01, optimizer = RMSprop, learning rate =
 347 0.001, batch size = 64}
- 348 • **GAF + CNN**: {GAF method = summation, convolution layers = 1, convolution kernels = 16, size
 349 of kernels = (3, 3), padding = "same", pooling = max, pool size = (2, 2), convolution activation
 350 function = SELU, hidden layers = 1, hidden neurons = 64, hidden activation function = ELU, batch
 351 normalization = True, L_2 regularization penalty = 0.01, optimizer = AdaGrad, learning rate = 0.001,
 352 batch size = 16}
- 353 • **RP + CNN**: {phase space dimensions $m = 2$, convolution layers = 2, convolution kernels = {8, 16},
 354 size of kernels = (5, 5), padding = "same", pooling = max, pool size = (2, 2), convolution activation
 355 function = tanh, hidden layers = 1, hidden neurons = 16, hidden activation function = ReLU, batch
 356 normalization = True, L_2 regularization penalty = 0.01, optimizer = AdaGrad, learning rate = 0.001,
 357 batch size = 32}
- 358 • **Proposed**: {1D residual blocks = 1, 1D convolution kernels = 32, size of 1D kernels = 3, 2D con-
 359 volution layers = 2, 2D convolution kernels = 16, size of 2D kernels = (3, 3), convolution activation
 360 function = SELU, hidden layers = 3, hidden neurons = 32, hidden activation function = ReLU, batch
 361 normalization = True, dropout coefficient = 0.2, L_2 regularization penalty = 0.0001, optimizer =
 362 SGD, learning rate = 0.001, batch size = 16}

363 4.1.2 Classification performance

364 The experimental results, summarized in Table 2, provide insights into the performance of each model on
 365 the final testing set. Six evaluation metrics are used to comprehensively evaluate the performance. Accuracy
 366 is calculated as the ratio of correctly predicted samples to the total number of testing samples. Precision
 367 measures the percentage of accurately predicted faulty samples out of all faulty predictions, while recall
 368 (i.e., true positive rate (TPR)) quantifies the proportion of correctly predicted faulty samples compared

369 to the total number of faulty samples. The false positive rate (FPR) represents the ratio of false positive
370 predictions to the total number of normal samples. Next, the F_1 score captures the harmonic mean of
371 precision and recall, providing a balanced assessment of model performance. Finally, training time (TT)
372 is used to compare models’ computational complexities. The TT denotes the time required to train the
373 model on the training set. Note that the reported TTs for deep learning models correspond to the time
374 taken to complete 100 training epochs. To ensure a fair comparison, all models considered in this study are
375 trained on an NVIDIA A100 GPU with 51GB of virtual RAM (VRAM). The proposed approach achieved
376 the highest score in five out of the six key metrics, demonstrating its competitive classification performance
377 and computational efficiency.

Table 2: Simulation results summary

	Accuracy	Precision	Recall	1-FPR	F_1	TT (min)
DTW+1NN	0.9569	0.9816	0.9315	0.9824	0.9559	112.15*
TSF	0.9775	0.9773	0.9773	0.9777	0.9773	2.80
BOSS	0.9413	0.9917	0.8904	0.9925	0.9383	0.56
LSTM	0.9338	0.9861	0.8804	0.9875	0.9303	370.11
GAF+CNN	0.9825	0.9936	0.9714	0.9937	0.9824	5.65**
RP+CNN	0.9788	0.9949	0.96264	0.9950	0.9785	9.12**
Proposed	0.9881	0.9987	0.9776	0.9987	0.9880	7.06

* TT corresponds to the time required to compute the DTW similarity matrix.

** TT includes the time required to encode time-series data into images.

378 4.1.3 Visual representations comparative analysis

379 In this subsection, we illustrate how our visual analytics approach reconciles performance and visual inter-
380 pretability. We demonstrate how our approach enables process operators to visually validate and analyze
381 the correlation between the model’s predictions and the underlying process. While visual interpretability
382 cannot be quantified by a metric, we use a qualitative approach to analyze the visual representations derived
383 from GAF encoding, RP encoding, and our proposed approach.

384 Figure 10 displays two negative MTS samples extracted from the Csth dataset and their corresponding
385 GAF and RP representations. The GAF and RP encodings are visualized as RGB images, with red, green,
386 and blue channels denoting temperature, level, and CW flow measurements, respectively. Analyzing the
387 GAF representations, we identify distinct “L-shaped” patterns, which serve as indicators of changes in the
388 system’s set points. Moreover, we observe homogeneous patterns, represented by uniform boxes, in the
389 upper-right section of the GAF images. These consistent patterns symbolize steady-state behavior, where
390 the system maintains a stable operational regime over a certain period. The extent of this uniformity
391 directly corresponds to the duration of stationarity. The presence of “L-shaped” transitions followed by
392 homogeneous patterns in the GAF images reveals a distinct sequence in the underlying process: a change in
393 set point followed by a period of stationarity. This sequence is a hallmark of normal operating regimes within
394 the Csth system. Contrastingly, the RP images appear mostly blank, indicating the absence of recurrent
395 behaviors. In other words, the system lacks repeating patterns or periodic behaviors that RP is designed
396 to highlight. Next, Figure 11 provides a different perspective as it showcases two positive MTS samples
397 alongside their corresponding GAF and RP images. The presence of cyclicities within the system is reflected
398 in the RP and GAF images, which exhibit distinctive periodic patterns. These patterns offer insights into
399 the cyclic behavior of the underlying process, with the time distance between the patterns representing the

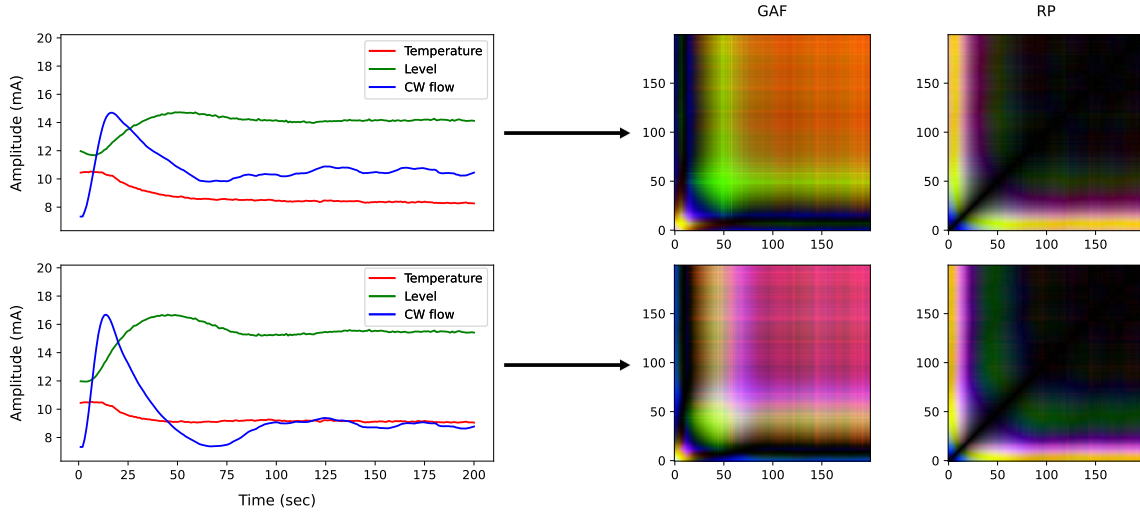


Figure 10: Two normal CSTH samples ($Y = 0$) alongside their corresponding GAF and RP encodings, revealing meaningful insights: i) A greenish-yellow (green + red) “L-shaped” pattern emerged around the 52-second mark in the upper GAF implies shifts in both level and temperature set points. Given the greater shift in the level set point (green) compared to the temperature set point (red), the resulting color leans more towards green. ii) The appearance of a reddish-white (red + green + blue) homogeneous box in the top GAF image represents stationarity in temperature, level, and CW flow measurement. Because temperature measurements exhibit greater stationarity, the red color is more dominant. iii) The bottom GAF image reveals a reddish-purple (red + blue) homogeneous box, indicating stable temperature and CW flow measurements. iv) Blank RP images denote the absence of recurrent patterns in the underlying system.

400 period of this recurring behavior. It is important to note that the construction of the simulated CSTH
 401 dataset emphasizes that periodic patterns are a strong indication of faulty operations. Thus, the periodicity
 402 observed in the RP and GAF images serves as a valuable indicator for identifying faulty operating regimes.

403 Figure 12 offers a visual comparison of the visual representations of three normal and three faulty
 404 MTS signals obtained using the proposed visual analytics framework. Unlike RP and GAF encodings, our
 405 proposed approach generates a single-channel matrix that can be visualized via color mapping. The color
 406 mapping assigns colors to matrix values, following a selected color map, which facilitates enhanced visual
 407 interpretation. The visual representations of normal MTS samples are visually distinct from those of faulty
 408 samples. Specifically, the visual representations for normal samples exhibit an even texture, whereas the
 409 visual representations for faulty samples display an irregular texture. This clear visual distinction is achieved
 410 through the application of 1D convolution operations. By applying 1D convolutional kernels, our approach
 411 effectively captures discriminative patterns present in the raw time-series data. As a result, normal and faulty
 412 MTS signals respond differently when convolved with these kernels. This differential response contributes
 413 to the observed visual distinctions between normal and faulty MTS samples in the visual representations
 414 produced by our framework.

4.1.4 Fault magnitude sensitivity analysis

416 While previous subsections have empirically demonstrated the effectiveness of the proposed approach in
 417 distinguishing normal and faulty MTS signals using distinctive image representations, a pivotal question

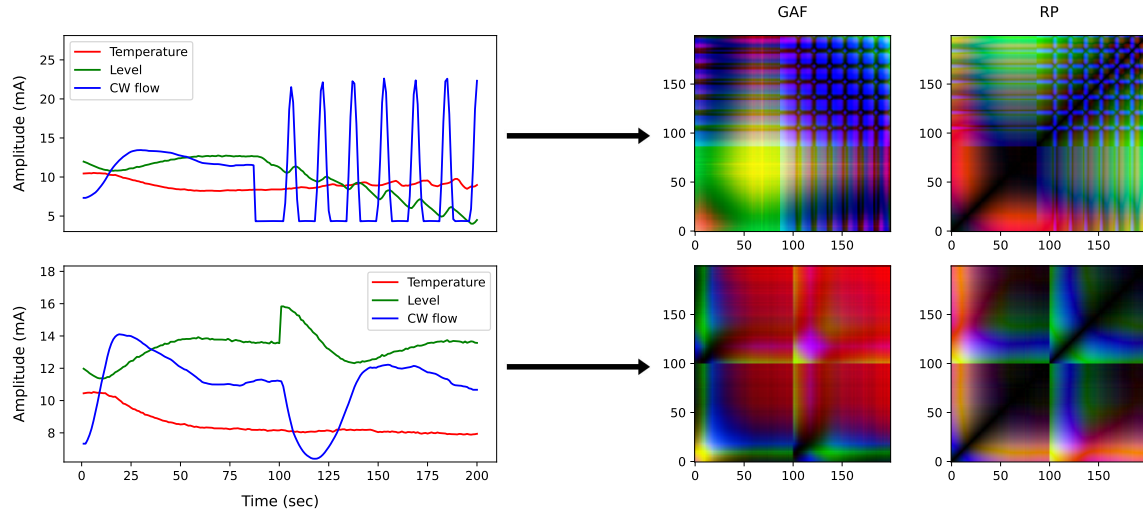


Figure 11: Two faulty CSTH signals ($Y = 1$) and their corresponding GAF and RP representations. Several observations and their qualitative interpretations can be made: i) The presence of cyclic behavior in the system results in distinct periodic patterns within the GAF and RP images. The temporal interval between these patterns corresponds directly to the oscillation period of the system. ii) Bright corners in the GAF images (e.g., the bright red upper corner in the lower GAF image) indicate shifts in trends—either decreasing or increasing.

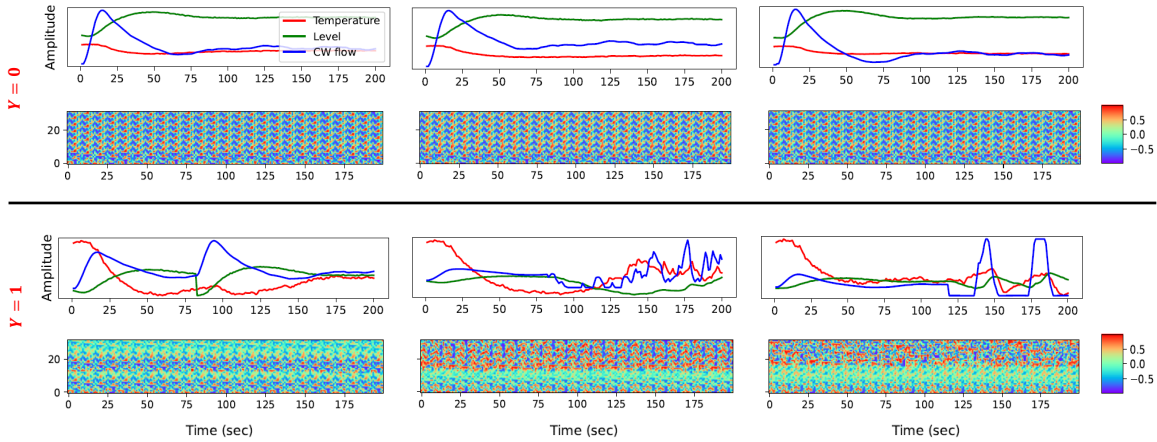


Figure 12: A visual comparison of the representations of three normal (top) and three faulty MTS signals (bottom) obtained using our proposed visual analytics framework. Homogenous texture indicates normal operating conditions, whereas uneven texture represents faulty operations. Normal and faulty signals respond to the learned 1D convolutional kernels differently. This dissimilarity in response leads to visually distinct representations between normal and faulty signals.

418 arises: Does the proposed approach have the requisite sensitivity to detect variations in fault magnitudes,
 419 specifically those of small ones? In this subsection, we explore the detection capabilities of the proposed
 420 approach across various fault sizes. The objective is to demonstrate that the proposed approach is not only
 421 able to detect substantial deviations (i.e., faults with a great magnitude) but also to capture minor faults
 422 within its image representations. Similar to subsection 4.1.3, we compare the proposed approach against the
 423 GAF and RP tools.

424 In pursuit of our objective, we introduce faults of varying magnitudes (small and large) into a normal
 425 CSTH sample and then assess how the proposed approach, GAF, and RP representations respond. Figure 13
 426 displays the progression of GAF and RP encodings for different fault magnitudes alongside the underlying
 427 MTS in the time domain. The encodings are presented for a normal MTS, a faulty MTS with a small
 428 fault magnitude (equivalent to 5% of steady-state value), and a faulty MTS with a large fault magnitude
 429 (equivalent to 30% of steady-state value). Moreover, Figure 14 shows the visual representations of the same
 430 aforementioned MTS obtained using the proposed approach. The visual encodings generated for the normal
 431 MTS and the faulty MTS of large magnitude (30%) using GAF, RP, and the proposed approach exhibit
 432 perceptible disparities. However, the proposed approach better distinguishes between the normal MTS and
 433 the faulty MTS of a small magnitude (5%). Notably, for the proposed approach, the extent of textural
 434 irregularity in the visual representations directly corresponds to the magnitude of the fault.

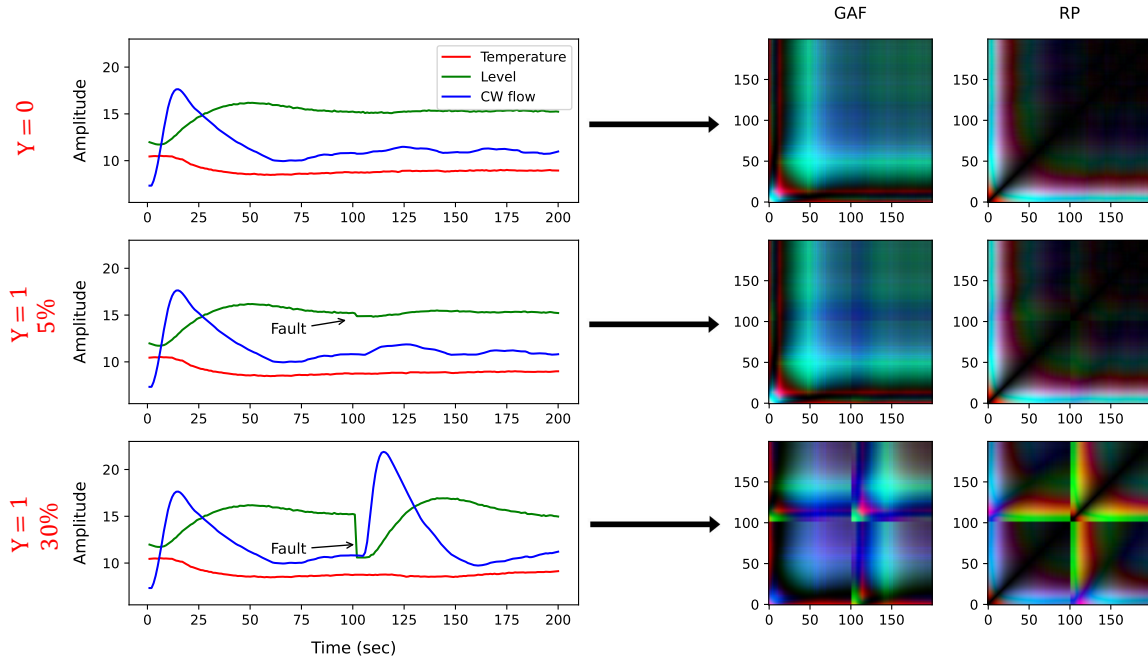


Figure 13: Responses of GAF and RP encodings across different fault magnitudes. Left: Original time-series plots. Middle: GAF encodings. Right: RP encodings. The visual encodings are shown for three scenarios: the top row shows a baseline normal CSTH sample, the middle row shows a faulty sample with a minor fault magnitude (5% deviation from steady-state value), and the bottom row represents a faulty sample with a large fault magnitude (30% deviation from steady-state value).

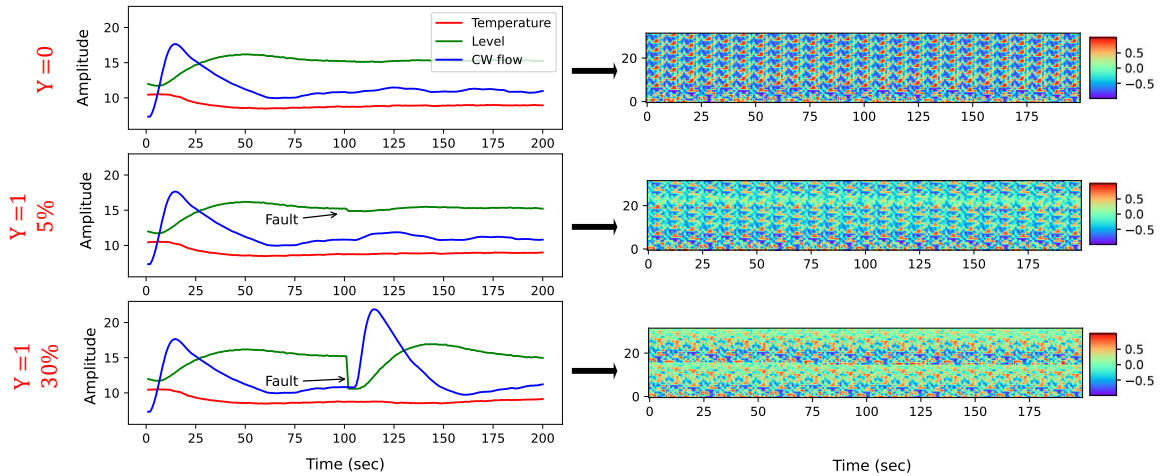


Figure 14: A visual comparison of the representations of a normal CSTH sample (top), a faulty CSTH sample of size 5% (middle), and a faulty CSTH sample of size 30% (bottom) obtained using our proposed visual analytics framework.

4.2 An industrial case study: the arc loss benchmark

The arc loss benchmark [60] was proposed as a real-world data benchmark for developing and validating data-driven process monitoring workflows. It provides a large-scale dataset obtained from an industrial pyrometallurgical smelting process. In this process, high-grade oxidized ore deposits are converted into refined base metals. Figure 15 illustrates the high-level pyrometallurgical processing. First, ore deposits extracted from open pits undergo grinding and drying operations using hammer mill flash dryers, resulting in fine ores with low moisture content. The dried ores are then dehydrated and deoxidized through a series of flash calciners and fluidized bed reducers. Subsequently, the processed ores are fed into a direct current electric arc furnace (DC EAF) unit to obtain base metals. The base metals are later processed by shotting and packaging units before being shipped to customers. An in-depth description of the process and its schematics can be found in Yousef et al. [60].




Figure 15: The broader pyrometallurgical processes.

The DC EAF unit, shown in Figure 16, is a high-temperature furnace that converts the electrical energy attained from the DC power supply into thermal energy by means of two plasma arcs. The two plasma arcs serve as the main heating component in the DC EAF unit, facilitating the separation of the base metals from other components of the ore (i.e., slag). Ensuring stable DC EAF operation is crucial for maximizing production efficiency and stability. However, the DC EAF unit faces a persistent fault known as *the arc loss fault*. The arc loss fault refers to the intermittent or sudden disruptions of the plasma arcs within the furnace.

452 These disruptions cause undesirable fluctuations in temperature and hinder the smelting efficiency, leading
 453 to potential production loss and variations in product quality. In this case study, the arc loss benchmark
 454 dataset is used to validate the proposed visual analytics framework in predicting the onset of arc loss in the
 455 DC EAF unit.




Figure 16: An illustration of the DC EAF unit.

456 The arc loss benchmark dataset comprises one year of high-frequency operating data collected from
 457 various pyrometallurgical process units. However, the raw data contains problematic artifacts, including
 458 outliers, irrelevant data, missing data, and class imbalance. To address the aforementioned issues, we
 459 processed the raw data, resulting in a balanced, clean dataset containing 3226 MTS samples. Each sample
 460 corresponds to 55 consecutive minutes of 96 different process measurements taken during either a smooth
 461 or faulty operating period (i.e., arc loss period). In the interest of brevity, we refrain from detailing the
 462 data preprocessing steps in this paper. Instead, we refer readers to [61] for comprehensive insights into
 463 the preprocessing techniques performed. Table 3 provides an overview of the processed arc loss data. The
 464 primary objective of this case study is to validate and compare the performance of our visual analytics
 465 workflow in accurately predicting whether a given MTS signal belongs to a smooth ($Y = 0$) or a faulty
 466 ($Y = 1$) operation.

Table 3: The processed arc loss data summary

Number of total samples (N)	3226
Training: validation: testing ratio	70:10:20
Number of variables (p)	96
Sampling frequency	3 sec
Signal length (L)	1100
Number of classes	2
Class ratio	50:50 (balanced)

467 4.2.1 Classification performance

468 We adopt a parallel approach to that used in the simulation case study, using consistent strategies for
 469 model evaluation and hyperparameter optimization. Notably, in the context of the industrial case study,
 470 we encounter a substantial difference in data scale compared to the simulated scenario. This difference
 471 requires a tailored preprocessing step to reduce data dimensionality, mitigating computational complexity.
 472 Specifically, the computational demands of imaging tools (i.e., GAF and RP) are directly proportional to

473 the length of the time-series data, following a power-law relationship. Given the limitations of our hardware
 474 in terms of memory and processing capacity, imaging a full-length MTS sample extracted from the arc
 475 loss data becomes challenging. To tackle this challenge, we incorporate techniques like piecewise aggregate
 476 approximation (PAA) [62] prior to encoding MTS data into GAF and RP images. PAA segments the time-
 477 series into non-overlapping windows and computes the mean value for each window, reducing the temporal
 478 complexity.

479 We consider the following models, accompanied by their respective tuned implementations:

- 480 • **DTW + 1NN**: {implementation = fast DTW, max warping window = 1000}
- 481 • **TSF**: {number of classifiers = 96, ensemble method = voting, number of trees = 300, minimum
 482 interval length = 3, number of attributes = 3 (mean, slope, and standard deviation)}
- 483 • **BOSS**: {number of classifiers = 96, ensemble method = voting, word size= 5, number of bins= 20,
 484 window size= 10, window step= 3}
- 485 • **LSTM**: {LSTM layers = 3, LSTM units = 16, recurrent activation function = hard sigmoid, hidden
 486 layers = 2, hidden neurons = 16, hidden activation function = SELU, batch normalization = True,
 487 dropout coefficient = 0.1, L_2 regularization penalty = 0.01, optimizer = Adam, learning rate = 0.001,
 488 batch size = 32}
- 489 • **GAF + CNN**: {PAA window size = 10, GAF method = summation, convolution layers = 2, con-
 490 volution kernels = 32, size of kernels = (3, 3), padding = "same", pooling = max, pool size = (2,
 491 2), convolution activation function = ELU, hidden layers = 2, hidden neurons = 16, hidden activa-
 492 tion function = ReLU, batch normalization = True, L_2 regularization penalty = 0.001, optimizer =
 493 AdaMax, learning rate = 0.001, batch size = 32}
- 494 • **RP + CNN**: {PAA window size = 10, RP phase space dimensions $m = 2$, convolution layers =
 495 3, convolution kernels = 64, size of kernels = (5, 5), padding = "same", pooling = max, pool size
 496 = (2, 2), convolution activation function = tanh, hidden layers = 1, hidden neurons = 32, hidden
 497 activation function = ReLU, batch normalization = True, L_2 regularization penalty = 0.1, optimizer
 498 = RMSprop, learning rate = 0.001, batch size = 32}
- 499 • **Proposed**: {1D residual blocks = 3, 1D convolution kernels = {64, 128, 256}, size of 1D kernels
 500 = 3, 2D convolution layers = 2, 2D convolution kernels = {16, 32}, size of 2D kernels = (3, 3),
 501 convolution activation function = ReLU, hidden layers = 3, hidden neurons = 16, hidden activation
 502 function = SELU, batch normalization = True, dropout coefficient = 0.3, L_2 regularization penalty
 503 = 0.1, optimizer = Adam, learning rate = 0.0001, batch size = 16}

504 Table 4 summarizes the classification performance of each model on the held-out testing set. The
 505 proposed approach achieved the highest score in five out of six key performance metrics, whereas GAF +
 506 CNN was the best configuration with respect to recall. Notably, non-deep learning models, such as DTW +
 507 1NN, TSF, and BOSS, exhibit significantly longer training times. This prolonged training duration can be
 508 attributed to the high dimensionality of the data and the intricate computations required by these models.
 509 Additionally, LSTM, owing to its recurrent nature and the long signal length (i.e., 1100-time steps for each
 510 signal), also incurs a substantial training time. On the other hand, the performance of GAF + CNN and RP
 511 + CNN models has been adversely impacted by the use of PAA for data preprocessing. PAA introduces a loss
 512 of temporal information, which is essential for these convolutional models to capture meaningful patterns.
 513 As a result, their classification performance, while still competitive, is somewhat compromised. In contrast,
 514 our proposed approach scales efficiently, with relatively short training times. This efficiency is attributed

515 to the use of 1D convolutions, which are computationally efficient and well-suited for processing large-scale
 516 industrial time-series data (e.g., the arc loss benchmark dataset). This scalability makes the proposed
 517 approach an appealing choice for real-world applications where computational resources are limited.

Table 4: Industrial case study results summary

	Accuracy	Precision	Recall	1-FPR	F_1	TT (min)
DTW+1NN	0.6388	0.6486	0.6227	0.6552	0.6354	1372.27*
TSF	0.6977	0.8499	0.6787	0.7389	0.7547	79.62
BOSS	0.6682	0.6474	0.7546	0.5799	0.6969	15.68
LSTM	0.7519	0.7280	0.8129	0.6897	0.7681	203.33
GAF+CNN	0.6915	0.6377	0.9018	0.4765	0.7471	8.05**
RP+CNN	0.7147	0.6830	0.8129	0.6144	0.7423	9.12**
Proposed	0.7721	0.8650	0.7325	0.8308	0.7932	3.33

* TT corresponds to the time required to compute the DTW similarity matrix.

** TT includes the time required to encode time-series data into images.

518 4.2.2 Visual representations comparative analysis

519 Similar to the simulation case study, in this subsection, we analyze the visual patterns within the visual
 520 representations obtained using GAF, RP, and our proposed approach. However, this industrial case study
 521 offers a distinctive perspective as it focuses on how the visual interpretability of these representations is
 522 influenced by the substantial differences in data dimensionality compared to the simulation case.

523 Firstly, both GAF and RP encodings operate on a single dimension, specifically designed for univariate
 524 time-series data. Consequently, when applied to a multivariate system like the arc loss benchmark, with 96
 525 process variables, these encoding techniques generate visual representations comprising 96 channels. The
 526 enormity of this dimensionality poses a significant challenge for visualization, as conventional computers
 527 struggle to render images with more than three channels effectively. To overcome this limitation, we care-
 528 fully select three pivotal variables for imaging: total power (TP), furnace feed (FF), and furnace off-gas
 529 temperature (FOGET). This selection is based on process knowledge, as these variables exhibit strong cor-
 530 relations that are important in predicting arc loss faults. Figure 17 shows the GAF and RP encodings as
 531 RGB images for FOGET, FF, and TP measurements recorded during two distinct normal operating peri-
 532 ods. As previously discussed, we use the PAA method to reduce the raw time-series size, mitigating the
 533 computational complexity associated with these imaging tools. Specifically, we use a PAA window size of
 534 10, reducing the length of the time-series by a factor of 10. Figure 18 further illustrates the application of
 535 GAF and RP encodings as RGB images, this time focusing on FOGET, FF, and TP measurements during
 536 two distinct faulty operating periods.

537 While these encodings reveal valuable visual patterns, it is important to acknowledge their limitations,
 538 particularly when applied to large-scale industrial datasets. One major challenge arises from the process of
 539 selecting the optimal set of three variables for imaging. This task can be time-consuming and labor-intensive,
 540 particularly in high-dimensional systems where process knowledge is limited, such as in industrial settings.
 541 Additionally, the use of the PAA method, although essential for managing the computational demands of
 542 these imaging techniques, introduces its own set of limitations. PAA inherently smooths the original time-
 543 series data, which can have unintended consequences. For instance, singular abnormal outlier measurements
 544 may get averaged out in the reduced time-series. Consequently, the resulting RP or GAF images may not
 545 capture these anomalous data points, potentially affecting the accuracy of FDD. Furthermore, it is important

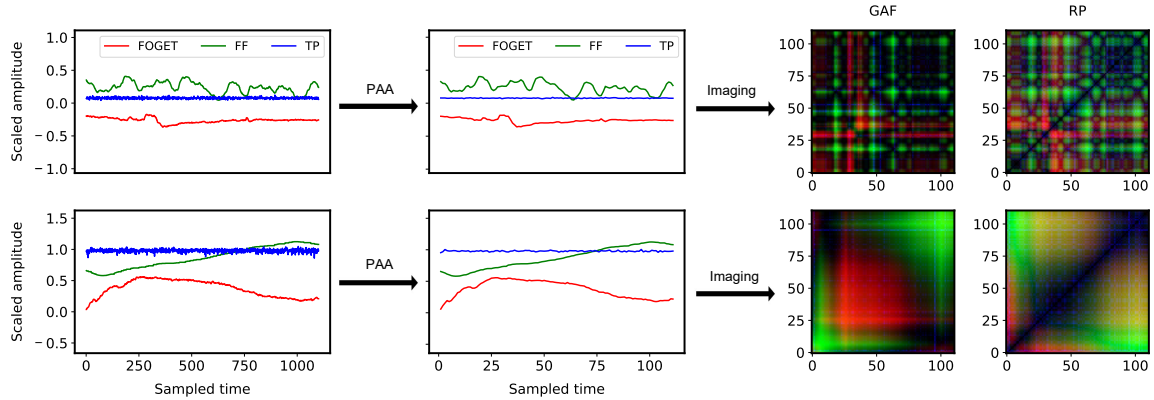


Figure 17: An illustration of two normal arc loss samples ($Y = 0$) and their step-by-step encoding into GAF and RP images. Firstly, a set of three variables (FOGET, TP, and FF) was carefully chosen for imaging. Subsequently, the PAA method is used to reduce the dimensionality of the raw signals, using a window size of 10. Finally, the RP and GAF encodings are applied to the reduced time-series. The resulting GAF and RP representations are displayed as RGB images, where the red, green, and blue channels represent FOGET, FF, and TP measurements, respectively. Qualitative interpretations are as follows: In the first sample, uniform patterns indicate stationary measurements, while red vertical/horizontal lines mark the time length in which a disturbance in FOGET measurements is introduced. In the second sample, vibrant green and red corners highlight increasing and decreasing trends in FF and FOGET measurements, respectively.

546 to note that the time distance represented in the GAF and RP encodings corresponds to the reduced time-
 547 series length resulting from PAA. In other words, each pixel in the GAF and RP representations represents
 548 the average of two-time windows of size 10 in the underlying system. This temporal aggregation may mask
 549 certain transient dynamics that occur within shorter time frames, which can be critical for understanding
 550 the dynamics of industrial processes, especially during fault conditions.

551 Next, Figure 19 compares two normal arc loss samples ($Y = 0$) along with their visual representations
 552 obtained using the proposed approach versus two faulty arc loss samples ($Y = 1$) and their visual representa-
 553 tions. The size of the visual representation is $256 \times 1100 \times 1$. In other words, the width of the one-channelled
 554 image corresponds to the length of the input signals (i.e., 1100-time steps), while the height represents
 555 the number of representation dimensions, which equals the number of 1D kernels in the last 1D residual
 556 block. Although the global texture is differentiable between the normal and faulty visual representations,
 557 the low-level details are not as pronounced as in RP and GAF representations due to the high resolution
 558 of the images. As shown in Figure 19, the normal visual representations exhibit regular patterns, while the
 559 faulty visual representations display irregular textures.

560 The proposed approach offers several distinctive advantages when compared to existing imaging tools like
 561 RP and GAF. Firstly, the proposed method operates directly on MTS signals and produces single-channel
 562 visual representations, simplifying the interpretation process. This simplicity in representation can enhance
 563 the ease of understanding and analysis. Furthermore, the proposed approach handles MTS signals of any
 564 size directly without the need for dimensionality reduction techniques such as PAA. This advantage stems
 565 from the computational efficiency of 1D convolution operations, which scale linearly with the input size. As a
 566 result, the proposed approach can efficiently handle large-scale industrial datasets without the preprocessing
 567 steps that might introduce data loss and compromise temporal resolution.

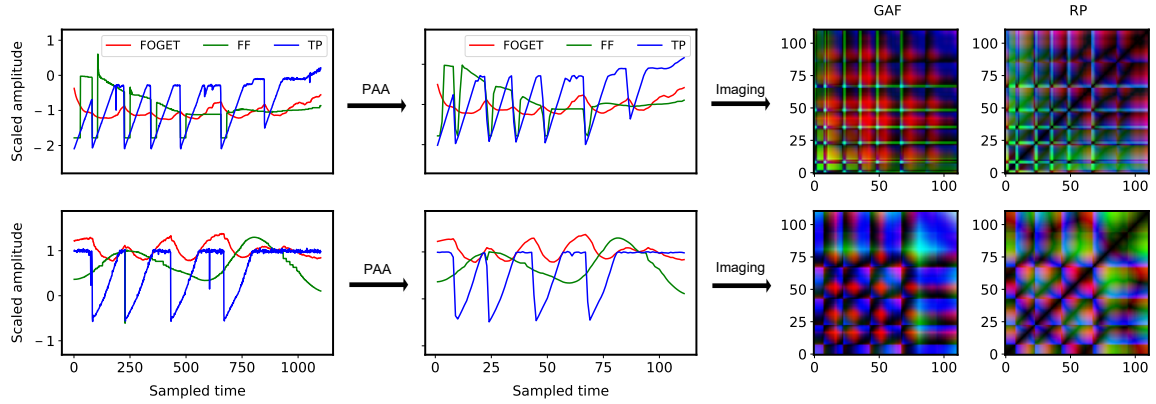


Figure 18: GAF and RP encodings presented as RGB images for FOGET, FF, and TP measurements recorded during two different faulty operating periods ($Y = 1$). The presence of a periodic checkerboard structure in both GAF and RP images suggests underlying fluctuations in FOGET, TP, and FF measurements. Unexpected power drops, leading to pronounced variations in FF and FOGET measurements, define an arc loss event.

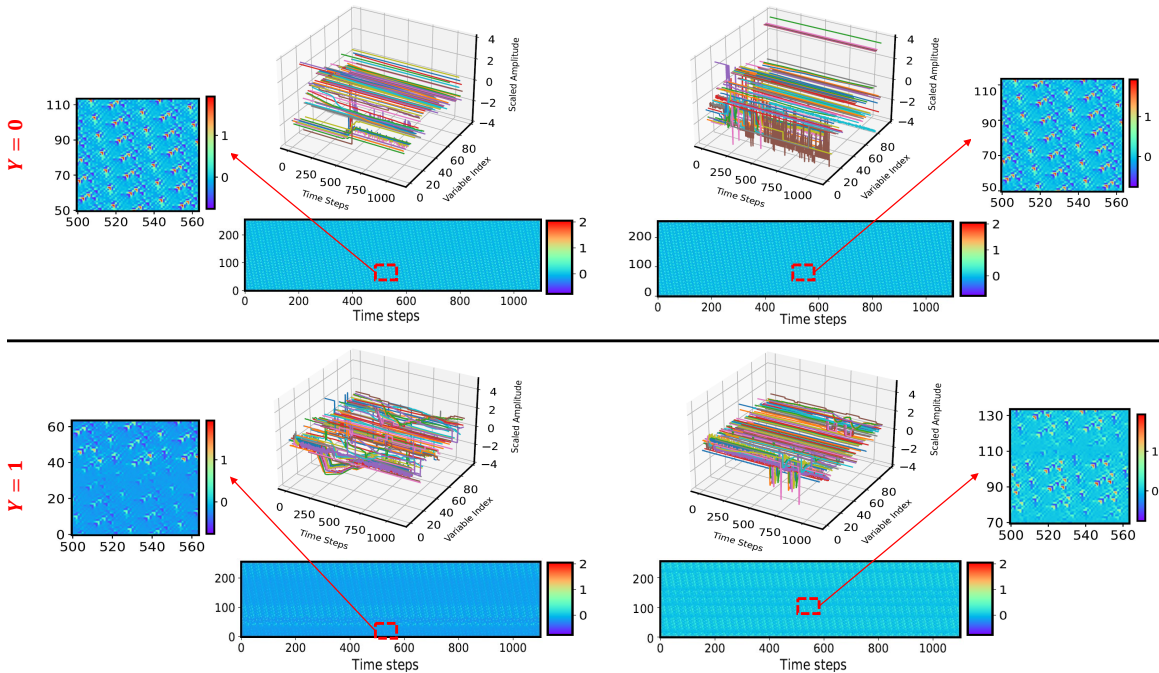


Figure 19: A visual comparison of two normal ($Y = 0$) and two faulty ($Y = 1$) arc loss samples using our proposed visual analytics framework. Normal samples exhibit homogeneous textures, while faulty samples display unsmooth textures. Zoomed-in patches highlight local patterns.

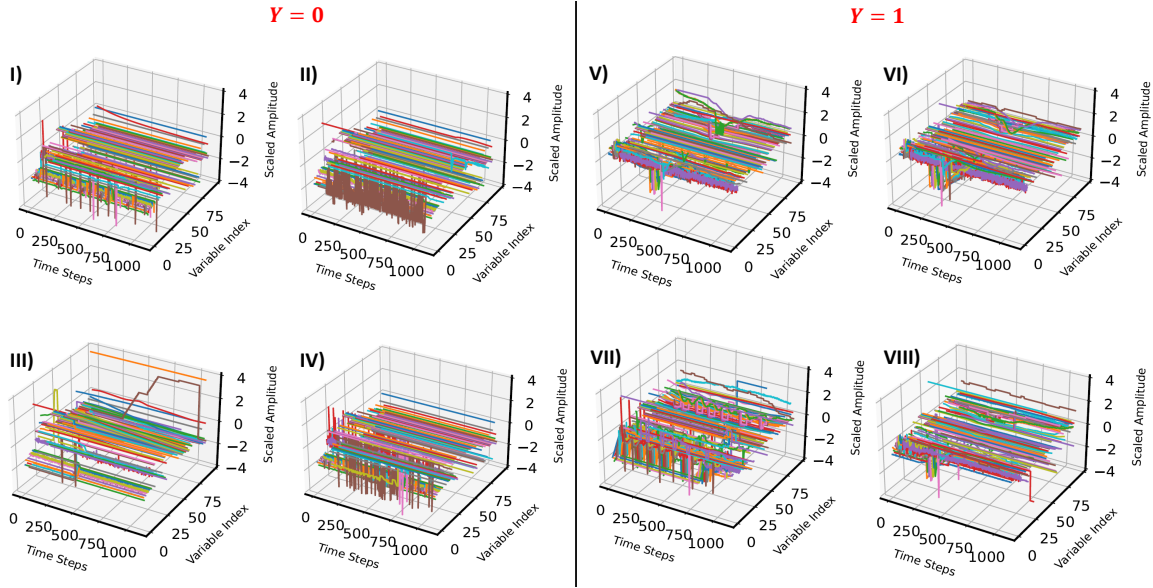


Figure 20: time domain data comparison: normal ($Y = 0$) vs. faulty ($Y = 1$) operating conditions. Readers are referred to Figure 21 for a more detailed view of some of the most important features, as identified by process experts.

568 4.2.3 Why visual analytics for industrial process monitoring?

569 In the realm of process monitoring, visual analytics involves transforming historical process data in the form
 570 of MTS into visually meaningful representations. The visual nature of the representations allows process
 571 experts to manually analyze patterns and textures within the images, facilitating easier interpretation and
 572 identification of significant insights. This is done by enabling process experts to develop intuition in relating
 573 different patterns in the visual representations to distinct process operating modes. In this work, we propose
 574 a visual analytics framework based on supervised learning, which converts MTS into pixelated images with
 575 patterns directly associated with various operating conditions.

576 Once the proposed network is trained on the training set, the trained stages 1 and 2 of the network can
 577 be used to convert input MTS signals into two-dimensional images. Figure 20 shows a visual comparison
 578 of four normal ($Y = 0$) and four faulty ($Y = 1$) arc loss samples in the time domain, while Figure 22
 579 illustrates their corresponding visual representations obtained using the proposed network. In the time
 580 domain, the data from both normal and faulty operations may appear as intricate, overlapping patterns,
 581 making it difficult for process operators to discern meaningful insights. However, when using the visual
 582 analytics approach, the corresponding images reveal a stark contrast, revealing the underlying nature of the
 583 operating conditions. During normal operation, these images exhibit regular and smooth patterns, reflecting
 584 the consistent and expected behavior of the process. In contrast, under faulty conditions, the images display
 585 a distinctly unsmooth texture. This intuitive visual distinction allows process operators to quickly identify
 586 and understand the operating condition of the data without the need for extensive data analysis. The
 587 integration of human visual perception and computer vision algorithms in visual analytics offers a powerful
 588 tool for enhancing process monitoring and responding effectively to the challenges posed by the era of big
 589 data in process industries.

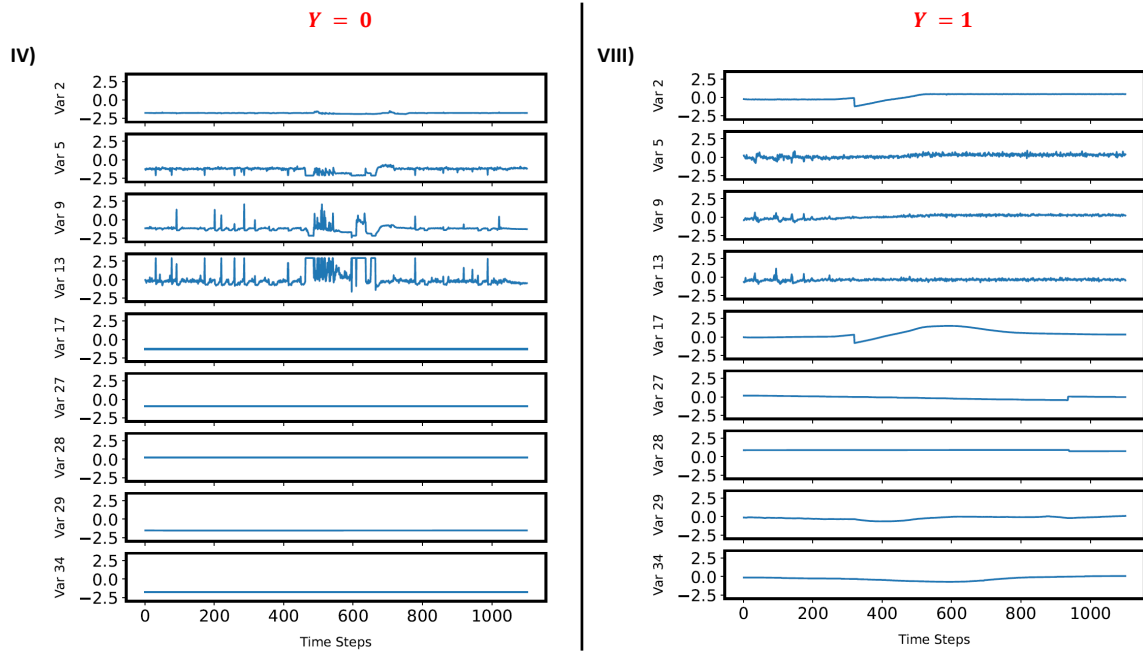


Figure 21: A visual demonstration of nine process variable measurements for normal ($Y = 0$) and faulty ($Y = 1$) operating conditions. Discriminative patterns are non-evident in the time domain. Latin numbers match Figure 20.

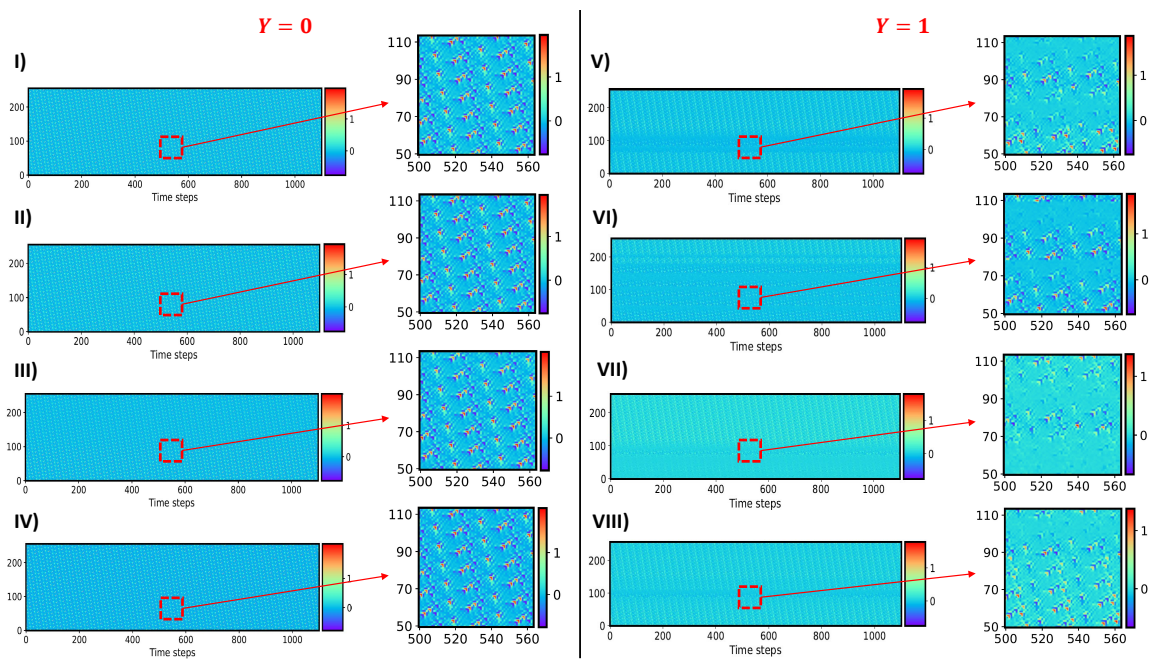


Figure 22: Visual representations, obtained using the proposed approach, comparison: normal ($Y = 0$) vs. faulty ($Y = 1$) operating conditions. Latin numbers match Figure 20.

5 Conclusion

This paper introduces a new paradigm for process monitoring called visual analytics. Visual analytics framework provides a powerful solution for industrial process monitoring by integrating the human visual system and computer vision algorithms. In this context, we propose a new end-to-end visual analytics pipeline for industrial fault detection, using both 1D and 2D convolution operations. Our approach begins with a series of 1D convolutions to capture relevant temporal information from the input MTS data. Subsequently, the extracted features are transformed into a 2D matrix, facilitating analysis and interpretation in the image domain, which proves to be more intuitive than the traditional time domain. By leveraging 2D convolution operations, our framework enables the network to visually recognize and classify these extracted features. To validate the effectiveness and interpretability of our approach, we conduct a simulated case study using the CSTH benchmark and an industrial case study using the arc loss benchmark. Experimental results demonstrate the superiority of our proposed visual analytics approach over state-of-the-art algorithms, providing not only improved performance but also meaningful and informative visual representations that enhance interpretability. Future works include examining the application of visual analytics in multiple fault scenarios. Furthermore, we plan to investigate the potential of visual analytics to determine the time of fault occurrence, which is an important aspect of process monitoring and FDD.

Acknowledgement

We gratefully acknowledge the financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- [1] V. Alcácer and V. Cruz-Machado, “Scanning the industry 4.0: A literature review on technologies for manufacturing systems,” *Engineering Science and Technology, an International Journal*, vol. 22, no. 3, pp. 899–919, 2019, ISSN: 2215-0986.
- [2] F. Yang and S. Gu, “Industry 4.0, a revolution that requires technology and national strategies,” *Complex & Intelligent Systems*, vol. 7, pp. 1311–1325, 2021.
- [3] C. Li, Y. Chen, and Y. Shang, “A review of industrial big data for decision making in intelligent manufacturing,” *Engineering Science and Technology, an International Journal*, vol. 29, p. 101 021, 2022, ISSN: 2215-0986. DOI: <https://doi.org/10.1016/j.jestch.2021.06.001>.
- [4] R. B. Gopaluni *et al.*, “Modern machine learning tools for monitoring and control of industrial processes: A survey,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 218–229, 2020, 21st IFAC World Congress, ISSN: 2405-8963. DOI: <https://doi.org/10.1016/j.ifacol.2020.12.126>.
- [5] H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller, “Deep learning for time series classification: A review,” *Data Mining and Knowledge Discovery*, vol. 33, no. 4, pp. 917–963, Mar. 2019. DOI: [10.1007/s10618-019-00619-1](https://doi.org/10.1007/s10618-019-00619-1).
- [6] P. Esling and C. Agon, “Time-series data mining,” *ACM Computing Surveys*, vol. 45, no. 1, p. 12, 2012. DOI: [10.1145/2379776.2379788](https://doi.org/10.1145/2379776.2379788).

- 626 [7] Y.-J. Park, S.-K. S. Fan, and C.-Y. Hsu, “A review on fault detection and process diagnostics
627 in industrial processes,” *Processes*, vol. 8, no. 9, p. 1123, Sep. 2020, ISSN: 2227-9717. DOI:
628 10.3390/pr8091123.
- 629 [8] J. Faouzi, “Time Series Classification: A review of Algorithms and Implementations,” in *Ma-
630 chine Learning (Emerging Trends and Applications)*, K. Kotecha, Ed., Proud Pen, 2022.
- 631 [9] A. Bagnall, J. Lines, A. Bostrom, J. Large, and E. Keogh, “The great time series classification
632 bake off: A review and experimental evaluation of recent algorithmic advances,” *Data Mining
633 and Knowledge Discovery*, vol. 31, pp. 606–660, May 2017. DOI: 10.1007/s10618-016-0483-9.
- 634 [10] Y.-H. Lee, C.-P. Wei, T.-H. Cheng, and C.-T. Yang, “Nearest-neighbor-based approach to
635 time-series classification,” *Decision Support Systems*, vol. 53, no. 1, pp. 207–217, 2012, ISSN:
636 0167-9236. DOI: <https://doi.org/10.1016/j.dss.2011.12.014>.
- 637 [11] Z. Wang, W. Yan, and T. Oates, “Time series classification from scratch with deep neural
638 networks: A strong baseline,” May 2017, pp. 1578–1585. DOI: 10.1109/IJCNN.2017.7966039.
- 639 [12] Y. Zheng, Q. Liu, E. Chen, Y. Ge, and J. L. Zhao, “Time series classification using multi-
640 channels deep convolutional neural networks,” in *Web-Age Information Management*, Springer
641 International Publishing, 2014, pp. 298–310, ISBN: 978-3-319-08010-9.
- 642 [13] A. Dempster, F. Petitjean, and G. I. Webb, “ROCKET: Exceptionally fast and accurate time
643 series classification using random convolutional kernels,” *Data Mining and Knowledge Discov-
644 ery*, vol. 34, no. 5, pp. 1454–1495, Jul. 2020. DOI: 10.1007/s10618-020-00701-z.
- 645 [14] I. Sarker, “Deep learning: A comprehensive overview on techniques, taxonomy, applications and
646 research directions,” *SN Computer Science*, vol. 2, p. 420, Aug. 2021. DOI: 10.1007/s42979-
647 021-00815-1.
- 648 [15] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.
- 649 [16] G. Zerveas, S. Jayaraman, D. Patel, A. Bhamidipaty, and C. Eickhoff, “A transformer-based
650 framework for multivariate time series representation learning,” in *Proceedings of the 27th ACM
651 SIGKDD Conference on Knowledge Discovery Data Mining*, ser. KDD ’21, Virtual Event,
652 Singapore: Association for Computing Machinery, 2021, pp. 2114–2124, ISBN: 9781450383325.
653 DOI: 10.1145/3447548.3467401.
- 654 [17] H. Ismail Fawaz *et al.*, “Inceptiontime: Finding alexnet for time series classification,” *Data
655 Mining and Knowledge Discovery*, vol. 34, pp. 1–27, Nov. 2020. DOI: 10.1007/s10618-020-
656 00710-y.
- 657 [18] S. Jiang and V. M. Zavala, “Convolutional neural nets in chemical engineering: Foundations,
658 computations, and applications,” *AIChE Journal*, vol. 67, no. 9, e17282, 2021. DOI: <https://doi.org/10.1002/aic.17282>.
- 659 [19] Y. Zheng, Q. Liu, E. Chen, Y. Ge, and J. Zhao, “Time series classification using multi-channels
660 deep convolutional neural networks,” *WAIM 2014. LNCS*, vol. 8485, pp. 298–310, Jan. 2014.
661

- 662 [20] X. Yuan, S. Qi, Y. A. Shardt, Y. Wang, C. Yang, and W. Gui, "Soft sensor model for dynamic
663 processes based on multichannel convolutional neural network," *Chemometrics and Intelligent*
664 *Laboratory Systems*, vol. 203, p. 104050, 2020, ISSN: 0169-7439. DOI: [https://doi.org/10.](https://doi.org/10.1016/j.chemolab.2020.104050)
665 [1016/j.chemolab.2020.104050](https://doi.org/10.1016/j.chemolab.2020.104050).
- 666 [21] B. Zhao, H. Lu, S. Chen, J. Liu, and D. Wu, "Convolutional neural networks for time series
667 classification," *Journal of Systems Engineering and Electronics*, vol. 28, no. 1, pp. 162–169,
668 2017. DOI: [10.21629/JSEE.2017.01.18](https://doi.org/10.21629/JSEE.2017.01.18).
- 669 [22] K. Fauvel, T. Lin, V. Masson, É. Fromont, and A. Termier, "Xcm: An explainable convolutional
670 neural network for multivariate time series classification," *Mathematics*, vol. 9, no. 23, p. 3137,
671 Dec. 2021, ISSN: 2227-7390. DOI: [10.3390/math9233137](https://doi.org/10.3390/math9233137).
- 672 [23] L. Brito, G. A. Susto, J. Brito, and M. Duarte, "An explainable artificial intelligence approach
673 for unsupervised fault detection and diagnosis in rotating machinery," *Mechanical Systems*
674 *and Signal Processing*, vol. 163, p. 108105, Jan. 2022. DOI: [10.1016/j.ymsp.2021.108105](https://doi.org/10.1016/j.ymsp.2021.108105).
- 675 [24] J. Grezmak, P. Wang, C. Sun, and R. X. Gao, "Explainable convolutional neural network for
676 gearbox fault diagnosis," *Procedia CIRP*, vol. 80, pp. 476–481, 2019, 26th CIRP Conference
677 on Life Cycle Engineering (LCE) Purdue University, West Lafayette, IN, USA May 7-9, 2019,
678 ISSN: 2212-8271. DOI: <https://doi.org/10.1016/j.procir.2018.12.008>.
- 679 [25] I. Yousef, S. L. Shah, and R. B. Gopaluni, "Visual analytics: A new paradigm for process
680 monitoring," *IFAC-PapersOnLine*, vol. 55, no. 7, pp. 376–383, 2022, 13th IFAC Symposium
681 on Dynamics and Control of Process Systems, including Biosystems DYCOPS 2022, ISSN:
682 [2405-8963](https://doi.org/10.1016/j.ifacol.2022.07.473). DOI: <https://doi.org/10.1016/j.ifacol.2022.07.473>.
- 683 [26] A. Abanda, U. Mori, and J. Lozano, "A review on distance based time series classification,"
684 *Data Mining and Knowledge Discovery*, vol. 33, pp. 378–412, Mar. 2019. DOI: [10.1007/
685 s10618-018-0596-4](https://doi.org/10.1007/s10618-018-0596-4).
- 686 [27] J. Lines, S. Taylor, and A. Bagnall, "Hive-cote: The hierarchical vote collective of transformation-
687 based ensembles for time series classification," in *2016 IEEE 16th International Conference*
688 *on Data Mining (ICDM)*, 2016, pp. 1041–1046. DOI: [10.1109/ICDM.2016.0133](https://doi.org/10.1109/ICDM.2016.0133).
- 689 [28] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recog-
690 nition," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 26, no. 1, pp. 43–
691 49, 1978. DOI: [10.1109/TASSP.1978.1163055](https://doi.org/10.1109/TASSP.1978.1163055).
- 692 [29] J. Lines and A. Bagnall, "Time series classification with ensembles of elastic distance mea-
693 sures," English, *Data Mining and Knowledge Discovery*, vol. 29, no. 3, pp. 565–592, May 2015,
694 ISSN: 1384-5810. DOI: [10.1007/s10618-014-0361-2](https://doi.org/10.1007/s10618-014-0361-2).
- 695 [30] Y. Lei and Z. Wu, "Time series classification based on statistical features," *EURASIP Journal*
696 *on Wireless Communications and Networking*, vol. 2020, p. 46, Feb. 2020. DOI: [10.1186/
697 s13638-020-1661-4](https://doi.org/10.1186/s13638-020-1661-4).
- 698 [31] H. Deng, G. Runger, E. Tuv, and M. Vladimir, "A time series forest for classification and
699 feature extraction," *Information Sciences*, vol. 239, pp. 142–153, 2013, ISSN: 0020-0255. DOI:
700 <https://doi.org/10.1016/j.ins.2013.02.030>.

- 701 [32] J. Large, A. Bagnall, S. Malinowski, and R. Tavenard, “On time series classification with
702 dictionary-based classifiers,” *Intelligent Data Analysis*, vol. 23, pp. 1073–1089, Oct. 2019. DOI:
703 10.3233/IDA-184333.
- 704 [33] P. Schäfer, “The boss is concerned with time series classification in the presence of noise,” *Data
705 Mining and Knowledge Discovery*, vol. 29, pp. 1505–1530, Nov. 2015. DOI: 10.1007/s10618-
706 014-0377-7.
- 707 [34] P. Schäfer and M. Höggqvist, “Sfa: A symbolic fourier approximation and index for similarity
708 search in high dimensional datasets,” in *Proceedings of the 15th International Conference on
709 Extending Database Technology*, ser. EDBT ’12, Berlin, Germany: Association for Computing
710 Machinery, 2012, pp. 516–527, ISBN: 9781450307901. DOI: 10.1145/2247596.2247656.
- 711 [35] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural Networks*, vol. 61,
712 pp. 85–117, 2015, ISSN: 0893-6080. DOI: <https://doi.org/10.1016/j.neunet.2014.09.003>.
- 713 [36] N. Papernot and P. McDaniel, *Deep k-nearest neighbors: Towards confident, interpretable and
714 robust deep learning*, 2018. arXiv: 1803.04765 [cs.LG].
- 715 [37] A. Sherstinsky, “Fundamentals of recurrent neural network (RNN) and long short-term mem-
716 ory (LSTM) network,” *Physica D: Nonlinear Phenomena*, vol. 404, p. 132 306, Mar. 2020. DOI:
717 10.1016/j.physd.2019.132306.
- 718 [38] R. M. Schmidt, *Recurrent neural networks (rnns): A gentle introduction and overview*, 2019.
719 arXiv: 1912.05911 [cs.LG].
- 720 [39] R. Pascanu, T. Mikolov, and Y. Bengio, *On the difficulty of training recurrent neural networks*,
721 2013. arXiv: 1211.5063 [cs.LG].
- 722 [40] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation*, vol. 9,
723 no. 8, pp. 1735–1780, Nov. 1997, ISSN: 0899-7667. DOI: 10.1162/neco.1997.9.8.1735.
- 724 [41] G. Van Houdt, C. Mosquera, and G. Nápoles, “A review on the long short-term memory
725 model,” *Artif. Intell. Rev.*, vol. 53, no. 8, pp. 5929–5955, Dec. 2020, ISSN: 0269-2821. DOI:
726 10.1007/s10462-020-09838-1.
- 727 [42] Y. Yu, X. Si, C. Hu, and J. Zhang, “A review of recurrent neural networks: Lstm cells and
728 network architectures,” *Neural Computation*, vol. 31, no. 7, pp. 1235–1270, 2019. DOI: 10 .
729 1162/neco_a_01199.
- 730 [43] R. Yang, M. Huang, Q. Lu, and M. Zhong, “Rotating machinery fault diagnosis using long-
731 short-term memory recurrent neural network,” *IFAC-PapersOnLine*, vol. 51, no. 24, pp. 228–
732 232, 2018, 10th IFAC Symposium on Fault Detection, Supervision and Safety for Technical
733 Processes SAFEPROCESS 2018, ISSN: 2405-8963. DOI: [https://doi.org/10.1016/j .
734 ifacol.2018.09.582](https://doi.org/10.1016/j.ifacol.2018.09.582).
- 735 [44] S. Belagoune, N. Bali, A. Bakdi, B. Baadji, and K. Atif, “Deep learning through lstm classifi-
736 cation and regression for transmission line fault detection, diagnosis and location in large-scale
737 multi-machine power systems,” *Measurement*, vol. 177, p. 109 330, 2021, ISSN: 0263-2241. DOI:
738 <https://doi.org/10.1016/j.measurement.2021.109330>.

- 739 [45] R. Yamashita, M. Nishio, R. Do, and K. Togashi, “Convolutional neural networks: An overview
740 and application in radiology,” *Insights into Imaging*, vol. 9, pp. 611–629, Jun. 2018. DOI:
741 10.1007/s13244-018-0639-9.
- 742 [46] A. Ajit, K. Acharya, and A. Samanta, “A review of convolutional neural networks,” in *2020*
743 *International Conference on Emerging Trends in Information Technology and Engineering*
744 *(ic-ETITE)*, 2020, pp. 1–5. DOI: 10.1109/ic-ETITE47903.2020.049.
- 745 [47] F. Li *et al.*, “Feature extraction and classification of heart sound using 1D convolutional neural
746 networks,” *EURASIP Journal on Applied Signal Processing*, vol. 2019, no. 1, 59, p. 59, 2019.
747 DOI: 10.1186/s13634-019-0651-3.
- 748 [48] S. Abdoli, P. Cardinal, and A. Lameiras Koerich, “End-to-end environmental sound classifi-
749 cation using a 1d convolutional neural network,” *Expert Systems with Applications*, vol. 136,
750 pp. 252–263, 2019, ISSN: 0957-4174. DOI: <https://doi.org/10.1016/j.eswa.2019.06.040>.
- 751 [49] C. Szegedy *et al.*, “Going deeper with convolutions,” in *2015 IEEE Conference on Computer*
752 *Vision and Pattern Recognition (CVPR)*, 2015, pp. 1–9. DOI: 10.1109/CVPR.2015.7298594.
- 753 [50] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016*
754 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
755 DOI: 10.1109/CVPR.2016.90.
- 756 [51] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, “Large-scale
757 video classification with convolutional neural networks,” in *2014 IEEE Conference on Com-*
758 *puter Vision and Pattern Recognition*, 2014, pp. 1725–1732. DOI: 10.1109/CVPR.2014.223.
- 759 [52] K. O’Shea and R. Nash, *An introduction to convolutional neural networks*, 2015. arXiv: 1511.
760 08458 [cs.NE].
- 761 [53] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, “1d convolu-
762 tional neural networks and applications: A survey,” *Mechanical Systems and Signal Processing*,
763 vol. 151, p. 107398, 2021, ISSN: 0888-3270. DOI: [https://doi.org/10.1016/j.ymsp.2020.](https://doi.org/10.1016/j.ymsp.2020.107398)
764 107398.
- 765 [54] Z. Wang and T. Oates, *Imaging time-series to improve classification and imputation*, 2015.
766 arXiv: 1506.00327 [cs.LG].
- 767 [55] J.-P. Eckmann, S. O. Kamphorst, and D. Ruelle, “Recurrence plots of dynamical systems,”
768 *Europhysics Letters (EPL)*, vol. 4, no. 9, pp. 973–977, Nov. 1987.
- 769 [56] V. M. Souza, D. F. Silva, and G. E. Batista, “Extracting texture features for time series
770 classification,” in *2014 22nd International Conference on Pattern Recognition*, 2014, pp. 1425–
771 1430.
- 772 [57] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolu-
773 tional neural networks,” in *Advances in Neural Information Processing Systems*, F. Pereira,
774 C. Burges, L. Bottou, and K. Weinberger, Eds., vol. 25, Curran Associates, Inc., 2012.

- 775 [58] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple
776 way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*,
777 vol. 15, no. 56, pp. 1929–1958, 2014.
- 778 [59] N. F. Thornhill, S. C. Patwardhan, and S. L. Shah, “A continuous stirred tank heater simula-
779 tion model with applications,” *Journal of Process Control*, vol. 18, no. 3, pp. 347–360, 2008,
780 Festschrift honouring Professor Dale Seborg, ISSN: 0959-1524.
- 781 [60] I. Yousef, L. D. Rippon, C. Prévost, S. L. Shah, and R. B. Gopaluni, “The arc loss challenge:
782 A novel industrial benchmark for process analytics and machine learning,” *Journal of Process*
783 *Control*, vol. 128, p. 103 023, 2023, ISSN: 0959-1524. DOI: <https://doi.org/10.1016/j.jprocont.2023.103023>.
784
- 785 [61] L. Rippon *et al.*, “Representation learning and predictive classification: Application with an
786 electric arc furnace,” *Computers Chemical Engineering*, vol. 150, p. 107 304, 2021, ISSN: 0098-
787 1354. DOI: <https://doi.org/10.1016/j.compchemeng.2021.107304>.
- 788 [62] E. Keogh, K. Chakrabarti, M. Pazzani, and S. Mehrotra, “Dimensionality reduction for fast
789 similarity search in large time series databases,” *Knowledge and Information Systems*, vol. 3,
790 no. 3, pp. 263–286, Aug. 2001, ISSN: 0219-1377. DOI: 10.1007/PL00011669.