

Multi-label Prediction in Time Series Data using Deep Neural Networks

Zhang, W.; Jha, D.K.; Laftchiev, E.; Nikovski, D.N.

TR2019-110 January 04, 2020

Abstract

This paper addresses a multi-label predictive fault classification problem for multidimensional time-series data. While fault (event) detection problems have been thoroughly studied in literature, most of the state-of-the-art techniques can't reliably predict faults (events) over a desired future horizon. In the most general setting of these types of problems, one or more samples of data across multiple time series can be assigned several concurrent fault labels from a finite, known set and the task is to predict the possibility of fault occurrence over a desired time horizon. This type of problem is usually accompanied by strong class imbalances where some classes are represented by only a few samples. Importantly, in many applications of the problem such as fault prediction and predictive maintenance, it is exactly these rare classes that are of most interest. To address the problem, this paper proposes a general approach that utilizes a multi-label recurrent neural network with a new cost function that accentuates learning in the imbalanced classes. The proposed algorithm is tested on two public benchmark datasets: an industrial plant dataset from the PHM Society Data Challenge, and a human activity recognition dataset. The results are compared with state-of-the-art techniques for time-series classification and evaluation is performed using the F1-score, precision and recall.

International Journal for Prognostics and Health Management Special Issue on Applications of Deep Learning and Emerging Analytics

© 2020 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Multi-label Prediction in Time Series Data using Deep Neural Networks

Wenyu Zhang¹, Devesh K. Jha², Emil Laftchiev², and Daniel Nikovski²

¹ *Department of Statistics, Cornell University, Ithaca, NY, USA*
wz258@cornell.edu

² *Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA*
{jha, laftchiev, nikovski}@merl.com

ABSTRACT

This paper addresses a multi-label predictive fault classification problem for multidimensional time-series data. While fault (event) detection problems have been thoroughly studied in literature, most of the state-of-the-art techniques can't reliably predict faults (events) over a desired future horizon. In the most general setting of these types of problems, one or more samples of data across multiple time series can be assigned several concurrent fault labels from a finite, known set and the task is to predict the possibility of fault occurrence over a desired time horizon. This type of problem is usually accompanied by strong class imbalances where some classes are represented by only a few samples. Importantly, in many applications of the problem such as fault prediction and predictive maintenance, it is exactly these rare classes that are of most interest. To address the problem, this paper proposes a general approach that utilizes a multi-label recurrent neural network with a new cost function that accentuates learning in the imbalanced classes. The proposed algorithm is tested on two public benchmark datasets: an industrial plant dataset from the PHM Society Data Challenge, and a human activity recognition dataset. The results are compared with state-of-the-art techniques for time-series classification and evaluation is performed using the F1-score, precision and recall.

keywords: Time-series analysis, Fault Detection, Fault Prediction.

1. INTRODUCTION

Time series analysis for rare events such as faults is generally a known and difficult problem (Yamanishi & Takeuchi, 2002). The problem is particularly difficult in the multi-dimensional, or multi-variate setting, where the events may be described by simultaneous occurrences on multiple time

series. A key confounding factor is the number of class labels, or the number of events that must be discovered. In the worst case, the events are described by labels from a known set, but may occur simultaneously. Naturally, this approach to labeling leads to a number of classes that grows combinatorially with the number of individual labels in the original set. The problem in this setting is one of time series analysis on infrequent events with multi-label classes that are characterized by severe class imbalance.

This is a critical problem in the field of prognostics and health management where faults for an equipment occur rarely but often have serious impacts such as service disruption, safety concerns for users and associated costs of repair or replacement. In these fields, the detection of rare events (fault detection, anomaly detection) in a timely fashion is critical to minimizing the detrimental impacts of the events. An example of this type of application was demonstrated in Holst et al. (Holst et al., 2012) where a fault detection tool was used to on a fleet of trains leading to reduced long-term maintenance costs by 5 – 10%.

While detecting faults is important (Chandola, Banerjee, & Kumar, 2009; Romeres, Jha, Yerezunis, Nikovski, & Dau, 2019) and many models have been used including deep learning (Jha, Srivastav, & Ray, 2016; Sarkar, Jha, Lore, Sarkar, & Ray, 2016), today the field of anomaly detection is moving towards predicting when a fault will occur. This capability is commonly known as prognostics, or the ability to predict impending faults. Prognostics is a critical technology in the drive to reduce cost of operation. The key difference is that while anomaly detection leads to corrective maintenance, fault prediction leads to predictive maintenance which can be performed before failures occur or before they develop into major defects. In practice this means that the condition of the equipment is monitored, and the time to failure is continuously estimated. Such an approach is able to reduce the num-

Wenyu Zhang et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.




Figure 1. Notation used in this paper. Example given in the setting of industrial plant fault prediction with two fault types: binary labels denote the presence of fault, observations correspond to sensor measurements, context sequence are control system setpoints.

ber of repairs, allow for optimal scheduling of maintenance workers and improve the safety of human operators.

This paper focuses on the problem of fault prediction in the multidimensional time series case with multi-label classes some of which have few examples. We propose a recurrent neural network-based (RNN) approach that in addition to the observed time series incorporates contextual information for past, present and future time intervals. This contextual information describes the operation of the plant/machine and could contain information about the target set-point, type of control being used for the plant, expected output, speed of plant/machine operation, etc. Using this approach, we are able to predict data labels for future time segments and to localize the faults that could occur in future with better accuracy than several other state-of-the-art algorithms for time-series analysis.

Figure 1 illustrates the problem setting and the paper notation. Here τ is the length of observed data, and $T - \tau$ is the prediction horizon. For a sample i of some monitored system, $z_{1:\tau}^{(i)} = [z_1^{(i)}, \dots, z_\tau^{(i)}]$ of length- τ is a historical sequence of observations where $z_t^{(i)} \in \mathbb{R}^{d_z}$. During the historical period, there are no labels that indicate the state of faults. For the future time segment, fault labels are predicted at two levels of granularity through segment labels and stepwise labels. Given L possible fault labels, $\tilde{y}^{(i)} \in \mathbb{R}^L$ denotes the segment label where each element $\tilde{y}^{(i)}[\ell] \in \{0, 1\}$ indicates the absence or presence of a label ℓ in the next $T - \tau$ steps. For example, if $L = 2$, the set of possible segment labels is $\{[0, 0], [1, 0], [0, 1], [1, 1]\}$. The stepwise label is $\tilde{o}_{\tau+1:T}^{(i)} = [\tilde{o}_{\tau+1}^{(i)}, \dots, \tilde{o}_T^{(i)}]$ where each binary vector $\tilde{o}_t^{(i)} \in \mathbb{R}^L$ indicates the absence or presence of the L labels at step $t \in \{\tau + 1, \dots, T\}$. Here $\tilde{o}_{\tau+1:T}^{(i)}$ provides label localization on top of $\tilde{y}^{(i)}$. Mathematically, $\tilde{y}^{(i)}[\ell] = \mathbb{1}\{\sum_{t=\tau+1}^T \tilde{o}_t^{(i)}[\ell] > 0\}$. There is no restriction on the number of one's in the vectors $\tilde{y}^{(i)}$ and $\tilde{o}_t^{(i)}$, as is consistent with the definition of a multi-label classification problem. In the fault prediction context, this means that multiple types

of faults are allowed to occur simultaneously. The case where multiple components of a system are monitored for failure is also included.

Contextual input $c_{1:T}^{(i)} = [c_1^{(i)}, \dots, c_T^{(i)}]$ where $c_t^{(i)} \in \mathbb{R}^{d_c}$ can include any independent variable of the system that is fully known at time τ (ex. pre-specified setpoint control sequences in industrial plants - Figure 1.) The presence of a label ℓ is a function of both input z and c without further assumptions on the structure within the inputs.

In this paper, we would like to estimate $P(\tilde{y}[\ell] = 1 | z_{1:\tau}, c_{1:T})$, the probability of given label in a future segment. To estimate this probability, the usual approach is to first define two functions F and V^ℓ on the input variables $z_{1:\tau}, c_{1:T}$. The function V^ℓ transforms the inputs from the observed space into an embedded space, while the function F maps from the embedded space to the interval $[0, 1]$. A simple map is chosen so that there is a clear and interpretable relationship between embedding and the output. Second, the optimal parameters of these functions are learned such that $F(V^\ell(z_{1:\tau}, c_{1:T}))$ approximates $P(\tilde{y}[\ell] = 1 | z_{1:\tau}, c_{1:T})$ as closely as possible.

This paper represents a new neural network architecture which we call Multi-label Predictive Network (MPN) using RNNs with a loss function that is designed considering the class imbalance. The architecture is fashioned after the sequence-to-sequence (Seq2Seq) model which has been successful in neural machine translation (Sutskever, Vinyals, & Le, 2014). This paper demonstrates that the model can be adapted for a different objective of multi-label prediction for temporal data.

We follow the insight that while prediction of multi-dimensional time-series data is a complex task due to the high-dimensional temporal data and the complex correlations which might be present between the different components of the data. It is, however, easier to learn the dynamics of the faults occurring in the system using a low-dimensional representation of the time-series data in an abstract space.

Neural networks used for sensor data compression are shown to be capable of capturing intramodal and intermodal correlations, and the compressed representations are effective for prediction and missing data imputation tasks (Liu, Zhang, Lin, & Quek, 2017). In prognostics and health management, neural networks can learn useful features to estimate remaining useful life and degradation (Li, Ding, & Sun, 2018).

For accurate prediction of faults over a certain prediction horizon, the faults predicted by the network at every instant of time should be taken into account as the faults may be correlated and thus prediction of a fault at a time-instant ' k ' might affect of the probability of occurrence of the fault at a future instant ' $k + 1$ '. Additionally, this information could also be useful in the multi-fault setting as some faults could always occur together or some faults never occur together and thus,

this can simplify learning of the fault dynamics. As a result of this motivation, the proposed MPN predicts faults one step at a time, and then the predicted fault is fed back to the network for making predictions on the faults occurring in future over a finite window of time.

The loss function used to train MPN is comprised of sub-objectives that predict the true segment labels and the step-wise labels, \tilde{y} and \tilde{o} . Both label types are multi-label, which means that the number of possible labels scales with the number of discrete labels in the label set. Samples can thus be labeled with any number of labels ranging from none (healthy) to all possible labels in the label set. For each label class, the loss is weighted to account for the imbalance in the training data set. Classes that are rare are weighted more heavily. The loss function also includes regularization components to (possibly) prevent overfitting in the network. The proposed ideas for fault prediction using MPN are verified using an industrial plant dataset from the 2015 Prognostics and Health Management Society (PHM Society) Data Challenge which consists a total of 33 plants. The network is also tested on a second dataset which is a human activity recognition (HAR) dataset from the 2011 Opportunity Activity Recognition Challenge, and we aim to predict low-level motions based on sensor readings and high-level activities. Both datasets are publicly available. Comparisons are conducted with several state-of-the-art techniques for time series classification, and MPN is shown to have better performance in general in terms of precision, recall and F1-score. Compared to the best performers in the two data challenges which use hand-engineered features, MPN attains comparable or better performance without the need for extensive feature extraction.

2. RELATED WORK

An important quantity for predictive maintenance is the remaining useful life (RUL). RUL of an equipment is the length from the current time to the end of its useful life, which is usually a unit of time, but can also be other measurements such as revolutions for rotating machines, number of operations for switchgear and load cycles for structural components (Welte & Wang, 2014). If a model of the performance degradation path can be built, RUL can also be found by determining the time at which degradation exceeds the failure threshold (Hu, Zhou, Zhang, & Si, 2015). The rate of degradation is usually not known in real world applications where physical simulation technology or accelerated life testing is not possible. In general, physical and stochastic models may not sufficiently capture the dynamics in the data from a complex system. RUL estimation does not reflect the duration of a fault, while the proposed MPN also attempts to localize the start and end times of faults in a target window.

A conventional data-driven approach for label prediction is to directly perform classification on the historical observations

with standard shallow classifiers (Salfner, Lenk, & Malek, 2010). Ensemble-based fault detection methods frequently use hand-engineered features (Xiao, 2016; Xie, Yang, Huang, & Sun, 2016) that are hard to identify and require domain knowledge. The label power set strategy represents multi-label combinations as distinct classes (Tsoumakas, Katakis, & Vlahavas, 2011) however this approach is limited by unobserved classes and exponential space and computational complexity. The simplifying assumption of independent labels allows for training a separate binary classifier for each label. Classifier chains (Read, Pfahringer, Holmes, & Frank, 2009) can be used to preserve information on label dependence, but chain order can significantly affect performance. A long short term memory (LSTM) RNN on raw time series with replicated targets achieves state-of-the-art performance in medical diagnoses (Lipton, Kale, Elkan, & Wetzel, 2016). In this paper, LSTM models are used in a predictive network to take advantage of their ability to model long-term dependencies.

Alternative distance-based approaches include Shapelet Forests (S-F) (Patri et al., 2014) which extracts subsequences that are the most discriminative for each class and calculates their distances to the time series for classification. K-Nearest Neighbors (KNN) approaches use a majority vote of the neighbors to classify a sample. Dynamic Time Warping (DTW) distance can quantify the qualitative dissimilarity of sequences and is therefore often used for human activity recognition (Seto, Zhang, & Zhou, 2015). Siamese LSTM networks learn a similarity metric from time series by first embedding sample pairs in a new space and then comparing the embedded representations using cosine similarity (Pei, Tax, & van der Maaten, 2016). All KNN methods depend on the diversity and quality of the training samples which is a significant issue in the setting of severe class imbalance.

When the data is fully labeled, Markov models combined with an iterative strategy can be used to do multi-step forward prediction (Read, Martino, & Hollmén, 2017). A common approach is time series prediction followed by classification (Molaei & Keyvanpour, 2015). However such labeling is generally not available in real applications.

The Seq2Seq model uses two LSTM networks to generate output sequences for neural machine translation (Sutskever et al., 2014). Extensions of RNNs have benefited from including spatial and temporal information (Liu, Wu, Wang, & Tan, 2016). RNNs have also seen success in RUL estimation, with the benefit that a neural network setup can be configured to include contextual information such as workloads, operating conditions and deterioration modes (Zheng, Ristovski, Farahat, & Gupta, 2017; Heimes, 2008) In this paper, the Seq2Seq model is modified to incorporate temporal contextual input in the forecasting window and to output multi-label predictions.

Moreover, the approach here avoids sequential error sources by directly predicting time series labels.

3. PROPOSED APPROACH

The objective of this work is to predict fault labels in a multidimensional time series on a window of time in the future. This prediction is performed using historical observations of the time series and contextual (application specific) input. Given L labels, the proposed approach has two steps. First, using MPN, we learn an L -dimensional predictive representation of the data in the future time segment. Then second, we use a linear classifier in each dimension of the predictive representation to determine the presence or absence of a given label. Importantly, each dimension ℓ of the representation corresponds to a label ℓ in the total label set and all dimensions are learned jointly. The linear classification is thus performed per dimension and independently of the other dimensions. This independence in the prediction means that the number of dimensions is equal to the number of faults and not the size of the power set which is all possible combinations of faults.

To reduce the complexity of the predictive problem, the MPN predicts each fault separately. This is motivated by the realization that the network learns the coupled, correlated dynamics between different faults during training. This is because the network is trained to predict faults one step at a time and these predictions are then used to predict faults in the future. In such a scenario, the network learns the dependence of faults occurring in time— more specifically, it can possibly learn which faults act as precursors to other faults. As a result of this simplification, the dimension of the predicted vector is reduced to the number of faults occurring in the system instead of all the possible combinations of faults in the system. This simplification also allows to alleviate the sample complexity of the neural network as it has to learn fewer classes.

The predictive representation is modeled using a neural network because neural networks have the advantage of being able to model irregular data shapes and complex system dynamics from data. The network in this paper consists of two connected LSTM networks: the encoder and decoder networks. Each network is turned to capture the time series dynamics in either the historical data (network 1) or the future segment data (network 2). The following equations describe the update step of an LSTM network. Here at each step, given input x_t and network states at the previous time step, the current hidden state h_t is obtained as (*Understanding LSTMs*,

2015):

$$\begin{aligned}
 f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\
 i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\
 \tilde{\xi}_t &= \tanh(W_\xi \cdot [h_{t-1}, x_t] + b_\xi) \\
 \xi_t &= f_t \odot \xi_{t-1} + i_t \odot \tilde{\xi}_t \\
 q_t &= \sigma(W_q \cdot [h_{t-1}, x_t] + b_q) \\
 h_t &= q_t \odot \tanh(\xi_t)
 \end{aligned} \tag{1}$$

where σ is the sigmoid function, \tanh is the hyperbolic tangent and \odot is the Hadamard product. The cell state is ξ_t . The forget gate is denoted as, f_t , the input gate is denoted as i_t , and the output gate is denoted as q_t . The hidden states undergo additional computations to arrive at the final output of the network. These computations are application specific. The weight matrix W and bias vector b corresponding to each gate are learned to minimize the loss function of the network. For $h_t \in \mathbb{R}^L$ and $x_t \in \mathbb{R}^m$, the total number of trainable parameters for an LSTM network is $4(L^2 + Lm + L)$. The MPN encoder is an one-layer LSTM network with L units, x_t is a concatenation of z_t and c_t , and $m = d_z + d_c + L$. The MPN decoder is another one-layer LSTM network with L units, x_t is c_t and $m = d_c + L$.

3.1. Predictive Data Representations

As previously noted, here we are interested in estimating the conditional probability $P(\tilde{y}^{(i)}[\ell] = 1 \mid z_{1:\tau}^{(i)}, c_{1:T}^{(i)})$. This estimate is preformed using MPN which is depicted in Figure 2. In this figure, the blue blocks (square pattern) represent the encoder LSTM network. This network encodes the historical data into the embedded space. The green (stripped pattern) blocks represent the decoder LSTM network which decodes the embedded representation into predictions of the labels on a future segment. The flow of information both within the networks and between the two networks is represented by the arrows. Interpreting the figure from left to right, the encoder network synthesizes historical observations $z_{1:\tau}^{(i)}$ and contextual variables $c_{1:T}^{(i)}$ to output $h_\tau^{(i)} \in \mathbb{R}^L$, an estimated representation of labels present at τ . This representation is input to the decoder network together with the final hidden state of the encoder network. Given this synthesized historical information from the encoder network and the future context $c_{\tau+1:T}^{(i)}$, the decoder network directly learns $h_t^{(i)}$ to predict labels for each future time step.

In the final step, the estimate of $y^{(i)}$ is produced by adding bias to the sum of $h_t^{(i)}$ terms (hidden states) and taking the

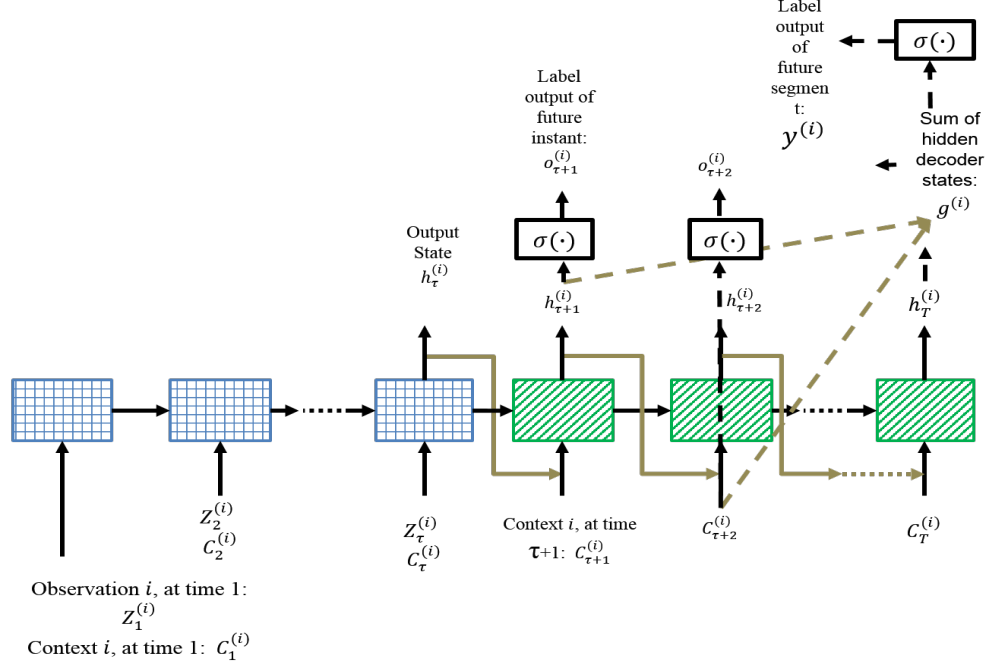


Figure 2. Multi-label Predictive Network (MPN) architecture to embed sample i as $g^{(i)}$. Encoding network (blue square pattern) encodes historical observations $z_{1:\tau}^{(i)}$ and historical context $c_{1:\tau}^{(i)}$, decoding network (green striped pattern) predicts for labels with additional contextual input $c_{\tau+1:T}^{(i)}$. Hidden units are denoted $h_t^{(i)}$. MPN outputs multi-label predictions $y^{(i)}$ for the segment labels and $o_{\tau+1:T}^{(i)}$ for the stepwise labels.

sigmoid transformation on the total sum.

$$g^{(i)} = \sum_{t=\tau+1}^T h_t^{(i)} + b_g \quad (2)$$

$$y^{(i)} = \sigma(g^{(i)}) \quad (3)$$

The final output, $y^{(i)}[\ell]$, is the estimate of the desired conditional probability $\mathbb{P}(\tilde{y}^{(i)}[\ell] = 1 \mid z_{1:\tau}^{(i)}, c_{1:T}^{(i)})$. That is, using notations in Section 1, the function V^ℓ is the first portion of the proposed MPN that maps to \mathbb{R} to produce the embedding g , and the function F is the sigmoid function.

3.1.1. Designing the Loss Function

The network described in the last section are trained using a loss function considering two main objectives. Specifically, the loss function has two components: the component which penalizes errors in labeling the whole future time segment, and a component which penalizes errors in labeling each step in the future segment.

The first component is the component which focuses on learning the labels on the whole future segment. A variable in this component of the objective function is p_ℓ . p_ℓ is the probability of label occurrence in the training set, and it is used

as a weight in the cross-entropy loss, $w_\ell = -\log(p_\ell)$ (Jain, Prabhu, & Varma, 2016). Using this weighting is important because it accentuates rare classes and forces the network to learn classes which have very few examples. This weighting is known to increase the false positive rate in the rare classes. However, in this application this increase is acceptable as compared to the relative importance of true positive and the importance of preserving label diversity in the predictions.

We note that there are many different weight schemes (Jain et al., 2016). The choice in this paper is motivated by the empirical experiments which show that this choice optimizes the performance of MPN. Edge cases where $p_\ell = 0$ are assigned a weight of 1 to indicate no preference in predicting label ℓ as either present or absent in the test set. The label prediction loss is then expressed as follows with true labels denoted by \tilde{y} ,

$$L_Y^{(i)} = - \sum_{\ell=1}^L \left[w_\ell \tilde{y}^{(i)}[\ell] \log(y^{(i)}[\ell]) + (1 - \tilde{y}^{(i)}[\ell]) \log(1 - y^{(i)}[\ell]) \right] \quad (4)$$

The second component which penalizes errors in labeling each step in the future segment. This component of the objective function uses $o_t^{(i)} = \sigma\left(h_t^{(i)}\right) \sigma\left(g^{(i)}\right) \in (0, 1)$, the step wise label vector which annotates the occurrence of labels in each of future segment time steps. At each time instant in the future segment, $o_t^{(i)}[\ell]$ approximates the probability $P\left(\tilde{o}_t^{(i)}[\ell] = 1 \mid z_{1:\tau}^{(i)}, c_{1:T}^{(i)}\right)$ that the label is present or absent at t for sample i .

When $\tilde{o}_t^{(i)}[\ell] = 1$, both $\sigma\left(h_t^{(i)}\right)$ and $\sigma\left(g^{(i)}\right)$ are ≈ 1 . When a given label is present at t and $h_t^{(i)}[\ell]$ is negative, this function penalizes the result. Similarly, negative value of $g^{(i)}[\ell]$ are penalized when label ℓ is present in the forecasting window. Therefore $o_t^{(i)}[\ell]$ tends to be positive for a present label and close to zero for an absent label. $o_t^{(i)}[\ell]$ is used in a squared loss function that ensures training stability.

$$L_O^{(i)} = \frac{1}{L(T-\tau)} \sum_{t=\tau+1}^T \left[\tilde{o}_t^{(i)} \cdot \left(1 - o_t^{(i)}\right)^2 + \left(1 - \tilde{o}_t^{(i)}\right) \cdot \left(o_t^{(i)}\right)^2 \right] \quad (5)$$

Note that the summand is similar to $\left\| \tilde{o}_t^{(i)} - o_t^{(i)} \right\|_2^2$. The true step wise label $\tilde{o}_t^{(i)}[\ell] \in \{0, 1\}$. When $\tilde{o}_t^{(i)}[\ell] = 1$, the loss incurred is $\left(1 - o_t^{(i)}[\ell]\right)^2$, and when $\tilde{o}_t^{(i)}[\ell] = 0$, the loss incurred is $\left(0 - o_t^{(i)}[\ell]\right)^2$. We write $L_O^{(i)}$ as such so that it has similar form as $L_Y^{(i)}$.

The objective function further has an ℓ_2 regularization on the weights of the network, collectively denoted as W , to prevent overfitting. Letting B denote the set of samples in a batch during training, the overall batch loss functions with hyperparameter λ are:

$$L_{base}^B = \frac{1}{|B|} \sum_{i \in B} L_Y^{(i)} + \frac{\lambda}{2} \|W\|_F^2 \quad (6)$$

$$L_{localize}^B = \frac{1}{|B|} \sum_{i \in B} \left[L_Y^{(i)} + L_O^{(i)} \right] + \frac{\lambda}{2} \|W\|_F^2 \quad (7)$$

The loss functions L_{base}^B and $L_{localize}^B$ correspond to network configurations MPN-base and MPN-localize, respectively. MPN-base does not make use of stepwise labels and hence does not incorporate the L_O loss. It is suitable for situations where stepwise labels are unavailable or label localization is not required.

3.2. Multi-label Decision

Figure 3c shows the t-SNE representation of the original samples of the PHM dataset for two fault types, which demonstrates that the samples are not readily separable at their original state. Figure 3a shows the ideal distribution of $g^{(i)}$ across i learned for two labels. In this representation, samples with identical label vectors are distinctly clustered together, and a threshold at zero would clearly separate samples that disagree on a given label ℓ . Figure 3b plots the actual embeddings produced by MPN, which are evidently more linearly separable than the original samples.

Cluster radius is typically governed by the data. Two observations can be made. First, the more data points in a cluster, the larger the cluster's radius. Second, the larger the noise component observed in the data, the larger the cluster radius. Both cases may lead to cluster overlap as shown in Figure 3b. Thus in practice it may be difficult to immediately note the correct classification boundary. However, standard classification algorithms can be used to find this boundary. An example here is a linear one-vs-all SVM classifier which was found to be robust to input data during our experiment. Label localization, or prediction of labels per step in the future interval, can be achieved using a linear SVM classifier.

4. EXPERIMENTS USING REAL DATA

The experiments for this study compare the performance of MPN to six state-of-the-art multi-label classifiers in two applications: An ensemble of classifiers on hand-engineered features and observations (Feature); a KNN classifier using majority vote and the DTW distance (DTW); Shapelet Forests (S-F) (Patri et al., 2014) which require learning on a powerset of labels; a multi-label LSTM network with replicated targets (ML-LSTM) (Lipton et al., 2016); a Siamese LSTM network learned metric (McClure, 2017) used in a KNN classifier (Siamese-LSTM); and a SVM using the RBF kernel which does not explicitly model the temporal input (SVM-rbf). Details of these methods can be found in Section 2. As noted before, the inputs to these classifiers are $z'_{1:\tau}$ and $c_{1:T}$, where $z'_{1:\tau}$ is $z_{1:\tau}$ padded to the length of c with the feature mean. This is a common approach taken to remedy the case of missing data from real life data collection. Because these classifiers do not learn the stepwise labels \tilde{o} they are compared with MPN-base which does not have the loss component L_O .

In addition, different features and ensemble of classifiers are used for each of the two applications due to differences of the two domains. Details of the datasets are in Sections 4.1 and 4.2. The two methods are named Feature (PHM) and Feature (HAR) according to the two domains. Feature (PHM) includes all features used in the winning entry (Xiao, 2016) of the 2015 PHM Society Data Challenge, namely month, week-day, hour, day, time, and historical observations. The ensemble




Figure 3. Embedded representations g and original distribution for two labels. $[n_1 n_2]$ denotes the segment label to predict.

ble comprises gradient boosting machine, random forest and penalized logistic regression, which are the best-performing methods for the tasks in Xiao’s paper (Xiao, 2016). The ensemble in Feature (HAR) comprises three best-performing methods in the 2011 Opportunity Activity Recognition Challenge, namely SVM, 1-Nearest Neighbor and decision tree (Chavarriaga, Sagha, Roggen, & Ferscha, 2011). The corresponding features (mean, variance, maximum, minimum and historical observations) are included for Feature (HAR). Feature engineering is a common preprocessing technique for human activity recognition (Figo, Diniz, Ferreira, & Cardoso, 2010). More recent methods utilize deep learning and are moving away from being feature-based (Ordóñez & Roggen, 2016), but these methods are mainly focused on segmentation and classification and not prediction.

Detailed experiments for this paper also studied the effects of different multi-label decision classifiers and auxiliary losses in MPN. The performance of three classifiers is tested. The default classifier on the embedded representations is the linear SVM. Two other classifiers tested are hard thresholding at zero and assuming the nearest class as measured by distance from the empirical class mean. MPN-base and MPN-localize are network configurations with loss functions L_{base}^B and $L_{localize}^B$, respectively. Siamese MPN-localize is a Siamese version of MPN-localize with an auxiliary Siamese objective measuring the similarity $\tilde{s}^{(i,j)}[\ell]$ of label ℓ for sample pair (i, j) through $s^{(i,j)}[\ell] = \exp(-|g^{(i)}[\ell] - g^{(j)}[\ell]|)$. This version is included since the Siamese structure has shown success in classification tasks with class imbalance (Pei et al., 2016; Mueller & Thyagarajan, 2016; Neculoiu, Versteegh, & Rotaru, 2016). Siamese MPN-localize is trained with loss function:

$$L_{siamese}^B = \frac{1}{|pair(B)|} \sum_{(i,j) \in pair(B)} \left[\beta \left(L_Y^{(i)} + L_Y^{(j)} + L_O^{(i)} + L_O^{(j)} \right) + (1 - \beta) L_S^{(i,j)} \right] + \frac{\lambda}{2} \|W\|_F^2 \quad (8)$$

where $pair(B)$ is a set of unique pairs of samples in B , and $L_S^{(i,j)}$ is the squared loss for $s^{(i,j)}$:

$$L_S^{(i,j)} = \frac{1}{L} \left[\tilde{s}^{(i,j)} \cdot \left(1 - s^{(i,j)} \right)^2 + \left(1 - \tilde{s}^{(i,j)} \right) \cdot \left(s^{(i,j)} \right)^2 \right] \quad (9)$$

For the experiments, we train all three network configurations using $L_{siamese}^B$ with the appropriate loss components turned off, so that all configurations have the same number of trainable parameters and are comparable.

Performance is evaluated using the F1-score, precision and recall. Micro and macro averages are presented because metric choice depends on the importance of the rare labels to the user. Macro metrics average the performance of each label type, and micro metrics average the performance across each sample.

For all methods, a hyperparameter grid search is performed on a held out validation set to maximize the sum of micro and macro F1-scores. The hyperparameters tuned and the search values are presented in Table 1. In all experiments, there are 500, 100, and 400 training, validation, and test samples, respectively.

4.1. Fault Prediction for Industrial Plants

The PHM dataset contains plant data from 33 plants with labeled data for 6 types of failure events (2015 PHM Data Challenge, 2015). The sampling interval for all plants is 15 minutes. The dataset contains time series samples for several years. For this reason fixed-length samples (for each plant) are chosen using random sampling. Sample length for historical observations is 30 (7.5 hours), and the forecast length for future events is 10 (2.5 hours). These durations are chosen because the dataset description stipulates that faults are predictable at most three hours in advance. Observation input z contains sensor measurements and contextual input c contains control system setpoints, and their combined dimension is 20

Method	Hyperparameter
Feature (PHM)	Regularization parameter (C)
Feature (HAR)	RBF kernel parameter (γ), regularization parameter (C)
DTW	number of nearest neighbors (k), warping window (w)
SVMrbf	RBF kernel parameter (γ), regularization parameter (C)
S-F	None
ML-LSTM	learning rate (η), intermediary target parameter (α), dropout probability (p)
Siamese-LSTM	number of nearest neighbors (k), learning rate (η), margin parameter (m), dropout probability (p)
MPN-base	learning rate (η), regularization parameter (λ)
MPN-localize	learning rate (η), regularization parameter (λ)
Siamese MPN-localize	learning rate (ℓ), weighing parameter (β), regularization parameter (λ)

(a) Hyperparameters tuned

Hyperparameter	Values
k	1, 3, 5, 7, 9
w	0, 1, 2
γ, C	0.001, 0.01, 0.1, 1, 10
η	0.001, 0.01, 0.1
α, β, p, m	0.1, 0.3, 0.5, 0.7, 0.9
λ	0.01, 0.1, 1

(b) Values tuned

Table 1. Hyperparameters tuned for each method

to 116 across plants. Because each plant has a unique sensor set, testing and training is performed separately for each plant.

Each plant has a particular number of monitored components, and each component corresponds to four sensors and four control references. The components are further grouped by zones, and there are two sensors providing environmental readings per zone. The true component-to-zone mapping is not provided in the dataset, hence no mapping is explicitly enforced through model input. Fault instances are given in terms of a start and end time, and a code specifying the fault type. This information is formatted according to our multi-label objective as \tilde{y} and $\tilde{o}_{\tau+1:T}$ for training and testing the models. An example experiment setup for Plant 1 is shown in Figure 4. Plant 1 has six components split over three zones.

This dataset has extreme label imbalance with important events that happen infrequently. This is demonstrated in the label distribution table in Table 2, where fault label 4 is observed only 8 times in the 600 training and validation samples of Plant 1. Furthermore, multiple concurrent faults can create unique system states. Figure 5 shows two samples of measurements from the first sensor and first control reference across the six components monitored in Plant 1. Figure 5a comes from a healthy sample with no faults. The sensor read-




Figure 4. PHM: Monitoring setup for Plant 1 with components (yellow circles) split over zones. True component-to-zone mapping is not provided, and included only for illustration purposes.

Fault Code	Occurrences
Code 1	58
Code 2	43
Code 3	31
Code 4	8
Code 5	20
Code 6	155

Table 2. PHM: Number of Plant 1 training samples per fault label.

ings demonstrate an upward trend and fluctuations, possibly due to frequent changes of the setpoints in the control reference. This indicates the importance of including information of future context for prediction. Figure 5b comes from a sample with co-occurring fault labels 1 and 6. Fault 6 starts in the middle and persists till the end of the observed time frame, and Fault 1 starts and ends in the second half of the time frame. Multi-label objectives are required due to the variety of label combinations.

4.1.1. Results and Discussion

Figure 6 shows the multi-label prediction test performance, where the performance of MPN-base for each plant is subtracted from that of competing methods for the corresponding plant. The illustration allows direct comparison of the methods for each of the 33 plants. A negative value after subtraction indicates a superior performance of MPN-base, which is mostly the case at both micro and macro levels for precision, recall and F1-scores. The hyperparameters most frequently selected for each model across the 33 plants are presented in Table 3.

S-F performed the worst, because the limited number of samples corresponding to each label combination does not allow sufficiently diverse shapelets to be extracted for classification. The other five methods performed similarly with SVMrbf being slightly worse overall since it does not attempt to capture the dynamics of time series. Feature (PHM) has slightly higher micro-F1 scores but lower macro-F1 scores, which implies that it has worse prediction performance for rare labels. In general, the competing methods did not perform as well as MPN-base, possibly because they are not constructed to make

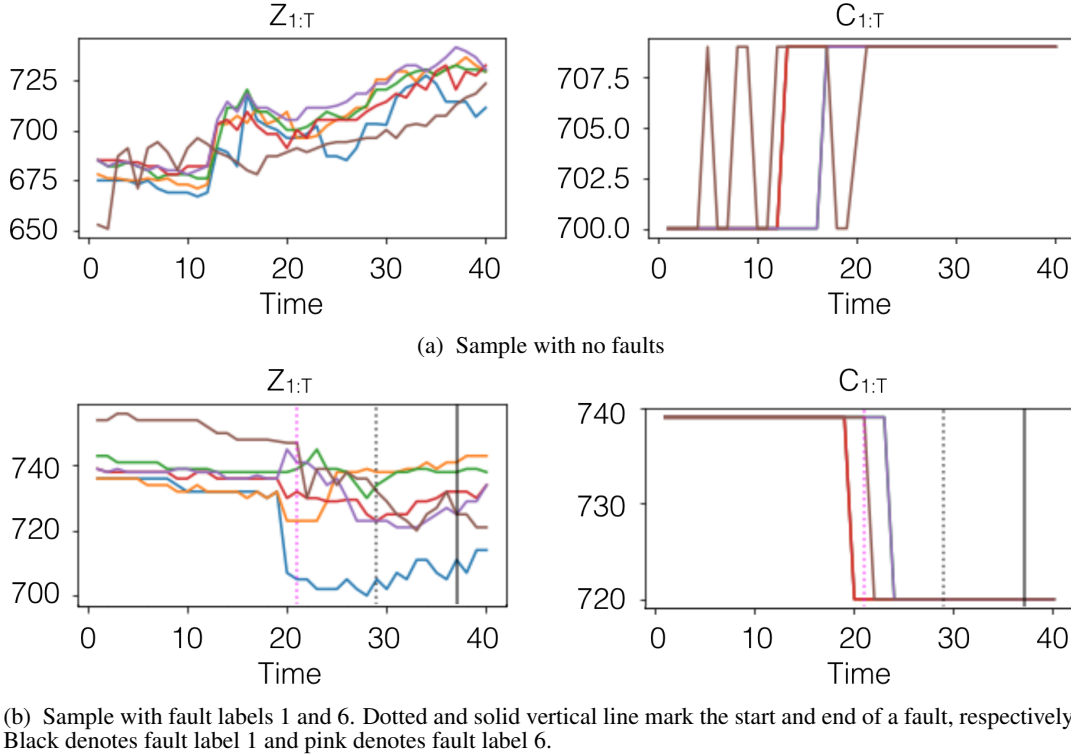


Figure 5. PHM: Plant 1 sample measurements of the first sensor and first control reference across six components.

Method	Selected hyperparameter
Feature (PHM)	$C = 0.1$
DTW	$k = 1, w = 2$
SVMrbf	$\gamma = 0.001, C = 10$
S-F	None
ML-LSTM	$\eta = 0.001, \alpha = 0.1, p = 0.9$
Siamese-LSTM	$k = 1, \eta = 0.001, m = 0.1, p = 0.9$
MPN-base	$\eta = 0.001, \lambda = 1$
MPN-localize	$\eta = 0.01, \lambda = 0.1$
Siamese MPN-localize	$\eta = 0.01, \beta = 0.3, \lambda = 0.1$

Table 3. PHM: Hyperparameters most frequently selected across 33 plants.

effective use of future context. The additional future context results in an irregular input data shape. Padding of the samples, a common remedy for missing data, allows the inclusion of the extra information in these models, but padding with a constant can cause a bias in learning since the missing data is not random and the data is temporal in nature.

Although some training labels are rare as the Plant 1 example shows in Table 2, the proposed MPN-base is able to learn the full variety of labels. Figure 7 shows the distribution of the true Plant 1 test labels versus the labels predicted by MPN-base. The counts are similar across all labels, and MPN-base

Classifier	Macro-F1		Micro-F1	
	Mean	Var	Mean	Var
Linear SVM	0.610	0.017	0.774	0.007
Threshold at 0	0.603	0.018	0.765	0.007
Nearest class	0.599	0.019	0.590	0.044
Network	Macro-F1		Micro-F1	
	Mean	Var	Mean	Var
MPN-base	0.610	0.017	0.774	0.007
MPN-localize	0.605	0.021	0.771	0.007
Siamese MPN-localize	0.603	0.020	0.773	0.008

Table 4. PHM: Performance of MPN-base network given choice of classifier, and performance of three versions of the MPN network. The mean values are averaged across 33 plants.

is capable of predicting label 4 despite it being underrepresented in the training set.

Effects of model design: Table 4 shows three different methods of classification on the embedded representations: linear SVM; hard thresholding at zero; and assuming the nearest class as measured by distance from the empirical class mean. The SVM shows the best overall performance. Note the poor performance of the third approach is because the variance of the representations is different across the classes.

Table 4 shows a comparison of MPN-base, MPN-localize and Siamese MPN-localize. All three versions have similar F1-scores, showing that the auxiliary Siamese targets added in




Figure 6. PHM: Label prediction test performance of methods *subtracted by MPN-base* for each of 33 plants. Negative value indicates that MPN-base outperformed in the metric for the plant.




Figure 7. PHM: Number of true and estimated Plant 1 test samples per fault label.

this experiment do not improve learning. Separation of the embedded representations across different labels is achieved sufficiently by the loss functions L_{base}^B and $L_{localize}^B$.

Fault localization: For testing the label localization performance of MPN-localize, the baseline is the multi-label prediction output without temporal localization, i.e. a fault label is either present/absent at all forecast time steps. Figure 8 shows the localization performance subtracted from the baseline performance for each of the 33 plants. As expected, the baseline has higher recall but also more false positives. Overall, localization improves accuracy and the F1-score.




Figure 8. PHM: Fault localization test performance of multi-label prediction baseline subtracted by label localization for each of 33 plants.

Method	Selected hyperparameter
Feature (HAR)	$\gamma = 0.001, C = 0.001$
DTW	$k = 1, w = 2$
SVMrbf	$\gamma = 10, C = 0.1$
S-F	None
ML-LSTM	$\eta = 0.001, \alpha = 0.1, p = 0.9$
Siamese-LSTM	$k = 1, \eta = 0.01, m = 0.1, p = 0.7$
MPN-base	$\eta = 0.001, \lambda = 1$
MPN-localize	$\eta = 0.01, \lambda = 1$
Siamese MPN-localize	$\ell = 0.01, \beta = 0.3, \lambda = 0.1$

Table 5. HAR: Hyperparameters selected.

4.2. Label Prediction for Human Activity Recognition

This application of human activity prediction is included to further demonstrate the effectiveness of the proposed MPN in multi-label prediction, and also its usage outside of PHM domains.

The Opportunity Dataset contains time series data used for HAR tasks (Roggen et al., 2010). Each of 4 subjects perform an activity of daily living (ADL) 5 times. Sensor readings are collected from sensing systems deployed on the body of the subjects, on objects and in the environment. The ADL consists of a sequence of high-level activities, namely early morning, coffee time, sandwich time, cleanup and relaxing. On a finer level, each action is described by the motion of the arms and the objects the arms interact with. We are interested in predicting the fine-level motions based on the sensor readings and high-level activities.

In this experiment, fixed-length samples are uniformly sampled with observed and forecast lengths as 75 and 25, respectively. The observation input z is the 243-dimensional measurements from GPS, motion sensors and other sensing units, and the contextual input c is the indicator matrix for high-level activities which empirically can be inferred from the time of day and the subject’s pattern of past schedules. Labels to be predicted are the low-level activities of the right arm. This is a combination of its motion in 13 categories including open and unlock, and the object the arm is interacting with which consists 23 options such as bottle and dishwasher. That is, $L = 36$, Label occurrences range from 1 to 180 times. This is a challenging task due to label sparsity and the inherent difficulty of predicting human activities (Ryoo, 2011).

4.2.1. Results and Discussion

Figure 9 shows the test performance of each competing method subtracted by that of MPN-base. The proposed method has the best overall performance. This is similar to the results in the PHM application in Section 4.1.1.

The hyperparameters selected for each model are presented in Table 5.

Classifier	Macro-F1	Micro-F1
Linear SVM	0.361	0.417
Threshold at 0	0.362	0.413
Nearest class	0.343	0.298
Network	Macro-F1	Micro-F1
MPN-base	0.361	0.417
MPN-localize	0.411	0.426
Siamese MPN-localize	0.344	0.420

Table 6. HAR: Performance of MPN-base network given choice of classifier, and performance of three versions of the MPN network.

Effects of model design: Table 6 shows similar performance for SVM and thresholding, indicating that the learned representations can be classified by the zero cutoff as designed. MPN-localize has the best F1-score, showing that MPN benefits from the additional stepwise label targets in this application.

5. CONCLUSION

This paper proposes a recurrent neural network-based approach for multi-label prediction problem in high-dimensional time series data with severe class imbalance. The proposed network, which we refer to as multi-label predictive network (MPN), embeds time series samples into a target space where the data is linearly separable. The loss function of MPN allows a variable number of labels for each sample and appropriately weighs each class to compensate for the rarity of some labels. Co-occurring labels and rare labels are common characteristics in fault prediction applications. The proposed algorithm for fault prediction using MPN performed well against state-of-the-art techniques on two different benchmark datasets, and is applicable to prediction tasks beyond PHM domains as well.

REFERENCES

- 2015 *PHM Data Challenge*. (2015). <https://www.phmsociety.org/events/conference/phm/15/data-challenge>. (Accessed: 2018-09-20)
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3), 15.
- Chavarriga, R., Sagha, H., Roggen, D., & Ferscha, A. (2011). *Opportunity Activity Recognition Challenge: Results and Conclusions*. http://www.opportunity-project.eu/system/files/docs/Workshop/2011_SMC_Challenge.pdf. (Accessed: 2019-08-27)
- Figo, D., Diniz, P. C., Ferreira, D. R., & Cardoso, J. M. P. (2010, Oct 01). Preprocessing techniques for context recognition from accelerometer data. *Personal and Ubiquitous Computing*, 14(7), 645–662.




Figure 9. HAR: Label prediction test performance of methods subtracted by MPN-base. Negative value indicates that MPN-base outperformed in the metric.

- Heimes, F. O. (2008, Oct). Recurrent neural networks for remaining useful life estimation. In *2008 International Conference on Prognostics and Health Management* (p. 1-6).
- Holst, A., Bohlin, M., Ekman, J., Sellin, O., Lindström, B., & Larsen, S. (2012). Statistical anomaly detection for train fleets. *AI Magazine*, 34, 33-42.
- Hu, C., Zhou, Z., Zhang, J., & Si, X. (2015). A survey on life prediction of equipment. *Chinese Journal of Aeronautics*, 28(1), 25 - 33.
- Jain, H., Prabhu, Y., & Varma, M. (2016). Extreme multi-label loss functions for recommendation, tagging, ranking, other missing label applications. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 935–944). New York, NY, USA: ACM.
- Jha, D. K., Srivastav, A., & Ray, A. (2016). Temporal learning in video data using deep learning and gaussian processes. In *Workshop on Machine Learning for Prognostics and Health Management at 2016 KDD, San Francisco, CA*.
- Li, X., Ding, Q., & Sun, J.-Q. (2018). Remaining useful life estimation in prognostics using deep convolution neural networks. *Reliability Engineering System Safety*, 172, 1 - 11.
- Lipton, Z. C., Kale, D. C., Elkan, C., & Wetzel, R. C. (2016). Learning to diagnose with LSTM recurrent neural networks. *International Conference on Learning Representations*.
- Liu, Q., Wu, S., Wang, L., & Tan, T. (2016). Predicting the next location: A recurrent model with spatial and temporal contexts. In *AAAI Conference on Artificial Intelligence*.
- Liu, Z., Zhang, W., Lin, S., & Quek, T. Q. S. (2017, April). Heterogeneous sensor data fusion by deep multimodal encoding. *IEEE Journal of Selected Topics in Signal Processing*, 11(3), 479-491.
- McClure, N. (2017). *Tensorflow machine learning cookbook*. Packt Publishing.
- Molaei, S. M., & Keyvanpour, M. R. (2015, March). An analytical review for event prediction system on time series. In *2015 2nd International Conference on Pattern Recognition and Image Analysis (IPRIA)* (p. 1-6).
- Mueller, J., & Thyagarajan, A. (2016). Siamese recurrent architectures for learning sentence similarity. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence* (pp. 2786–2792). AAAI Press.
- Neculoiu, P., Versteegh, M., & Rotaru, M. (2016). Learning text similarity with Siamese recurrent networks. In *Rep4NLP@ACL*.

- Ordóñez, F. J., & Roggen, D. (2016). Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1).
- Patri, O. P., Sharma, A. B., Chen, H., Jiang, G., Panangadan, A. V., & Prasanna, V. K. (2014, Oct). Extracting discriminative shapelets from heterogeneous sensor data. In *2014 IEEE International Conference on Big Data* (p. 1095-1104).
- Pei, W., Tax, D. M. J., & van der Maaten, L. (2016). Modeling time series similarity with Siamese recurrent networks. *CoRR*, abs/1603.04713.
- Read, J., Martino, L., & Hollmén, J. (2017). Multi-label methods for prediction with sequential data. *Pattern Recognition*, 63, 45 - 55.
- Read, J., Pfahringer, B., Holmes, G., & Frank, E. (2009). Classifier chains for multi-label classification. In *Proceedings of the 2009th European Conference on Machine Learning and Knowledge Discovery in Databases - Volume Part II* (pp. 254–269). Berlin, Heidelberg: Springer-Verlag.
- Roggen, D., Calatroni, A., Rossi, M., Holleczeck, T., Frster, K., Trster, G., ... d. R. Milln, J. (2010, June). Collecting complex activity datasets in highly rich networked sensor environments. In *2010 Seventh International Conference on Networked Sensing Systems (INSS)* (p. 233-240).
- Romeres, D., Jha, D. K., Yerazunis, W., Nikovski, D., & Dau, H. A. (2019, June). Anomaly detection for insertion tasks in robotic assembly using Gaussian process models. In *2019 18th European Control Conference (ECC)* (p. 1017-1022). doi: 10.23919/ECC.2019.8795698
- Ryoo, M. S. (2011, Nov). Human activity prediction: Early recognition of ongoing activities from streaming videos. In *2011 International Conference on Computer Vision* (p. 1036-1043).
- Salfner, F., Lenk, M., & Malek, M. (2010, March). A survey of online failure prediction methods. *ACM Comput. Surv.*, 42(3), 10:1–10:42.
- Sarkar, S., Jha, D. K., Lore, K. G., Sarkar, S., & Ray, A. (2016). Multimodal spatiotemporal information fusion using neural-symbolic modeling for early detection of combustion instabilities. In *2016 American Control Conference (ACC)* (pp. 4918–4923).
- Seto, S., Zhang, W., & Zhou, Y. (2015, Dec). Multivariate time series classification using dynamic time warping template selection for human activity recognition. In *2015 IEEE Symposium Series on Computational Intelligence* (p. 1399-1406).
- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2* (pp. 3104–3112). Cambridge, MA, USA: MIT Press.
- Tsoumakas, G., Katakis, I., & Vlahavas, I. (2011, July). Random k-labelsets for multilabel classification. *IEEE Transactions on Knowledge and Data Engineering*, 23(7), 1079-1089.
- Understanding LSTMs*. (2015). <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>. (Accessed: 2018-09-20)
- Welte, T. M., & Wang, K. (2014, Mar 01). Models for lifetime estimation: An overview with focus on applications to wind turbines. *Advances in Manufacturing*, 2(1), 79–87.
- Xiao, W. (2016, 03). A probabilistic machine learning approach to detect industrial plant faults. *International Journal of Prognostics and Health Management*, 7, 11.
- Xie, C., Yang, D., Huang, Y., & Sun, D. (2016). Feature extraction and ensemble decision tree classifier in plant failure detection. *International Journal of Prognostics and Health Management*, 7.
- Yamanishi, K., & Takeuchi, J.-i. (2002). A unifying framework for detecting outliers and change points from non-stationary time series data. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 676–681).
- Zheng, S., Ristovski, K., Farahat, A., & Gupta, C. (2017, June). Long Short-Term Memory network for remaining useful life estimation. In *2017 IEEE International Conference on Prognostics and Health Management (ICPHM)* (p. 88-95).

BIOGRAPHIES



Wenyu Zhang Wenyu received the B.A. degree with double major in Applied Mathematics and Statistics from the University of California, Berkeley, Berkeley, CA, USA, in 2013. She is currently working toward the Ph.D. degree in Statistics and Data Science at Cornell University, Ithaca, NY, USA, where she works on anomaly and change detection methods for time series. Her research interests are in developing statistical and machine learning methodology for sequential data.



Devesh K. Jha Devesh received PhD in Mechanical Engineering from Penn State in December 2016. He received M.S. degrees in Mechanical Engineering and Mathematics also from Penn State. He is currently a Research Scientist at MERL. His research interests are in the areas of Machine Learning, Robotics and Time Series Analytics. He was a recipient of the best student paper award at the 1st ACM SIGKDD workshop on Machine Learning for Prognostics and Health Management at KDD 2016, San Francisco. He also received the Kalman Best Paper Award 2019 from the American Society of Mechanical Engineers (ASME) Dynamic Society and Control Divisions.



Emil Laftchiev Emil I. Laftchiev received the B.S. and Ph.D. degrees in Electrical Engineering from the Pennsylvania State University in 2007 and 2015, respectively. In 2015 Emil joined MERL (Mitsubishi Electric Research Labs) in Cambridge, MA, USA as a research staff member. He has a wide range of research interests including

machine learning, the internet of things, time series analysis, and signal processing. Within these fields Emil is particularly interested in developing subfields of machine learning such as generative modeling and reinforcement learning. Emil is a member of IEEE and SIGKDD. His work has appeared in IEEE Conferences, Journal and Workshops.



Daniel Nikovski Daniel Nikovski received his PhD degree in robotics from Carnegie Mellon University in Pittsburgh, USA, in 2002, and is currently the group manager of the Data Analytics group at Mitsubishi Electric Research Laboratories in Cambridge, Massachusetts. His research interests include artificial intelligence, robotics,

machine learning, optimization and control, and numerical methods for the analysis of complex industrial systems.