

Detecting and Diagnosing Incipient Building Faults Using Uncertainty Information from Deep Neural Networks

Baihong Jin¹ Dan Li² Seshadhri Srinivasan³ See-Kiong Ng²

Kameshwar Poola^{1,3} Alberto Sangiovanni-Vincentelli^{1,3}

¹Department of EECS, University of California, Berkeley

²Institute of Data Science, National University of Singapore

³The Berkeley Education Alliance for Research in Singapore

Abstract—Early detection of incipient faults is of vital importance to reducing maintenance costs, saving energy, and enhancing occupant comfort in buildings. Popular supervised learning models such as deep neural networks are considered promising due to their ability to directly learn from labeled fault data; however, it is known that the performance of supervised learning approaches highly relies on the availability and quality of labeled training data. In Fault Detection and Diagnosis (FDD) applications, the lack of labeled incipient fault data has posed a major challenge to applying these supervised learning techniques to commercial buildings. To overcome this challenge, this paper proposes using Monte Carlo dropout (MC-dropout) to enhance the supervised learning pipeline, so that the resulting neural network is able to detect and diagnose unseen incipient fault examples. We also examine the proposed MC-dropout method on the RP-1043 dataset to demonstrate its effectiveness in indicating the most likely incipient fault types.

I. INTRODUCTION

Building faults whose impact are less perceivable and/or hinder regular operations are called *soft faults* [23], [36]. These soft faults, especially in their incipient phase, are hard to detect as their signatures are not generally obvious (due to their magnitudes) and are lurking under measurement/system noise or feedback control actions [10], [31]. Nevertheless, they will impact energy consumption, system performance, and maintenance costs adversely in the long-run if left undetected and unattended [14]. In addition, they can lead to costly maintenance and undesirable replacement operations. Therefore, it is an important and challenging task to design methods to detect and diagnose incipient soft faults during their incipient stage for various building systems, such as chillers and Air Handling Units (AHUs).

Fault Detection and Diagnosis (FDD) methods in the literature can be broadly classified into three categories: (i) model-based, (ii) signal-based, and (iii) data-driven [34], [40]. Model-based methods depend on explicit physical models at the device levels and use correlation tests on the input-output data to detect faults [15], [13], [16]. While fault-diagnosis can also be performed with the model used for detection, developing detailed models is a time-consuming and daunting process, especially for complex Cyber-Physical Systems (CPSs) like buildings. Authors in [32] point out that model-based

methods are not as practical as data-driven methods in terms of applying the FDD techniques to real buildings. Signal-based FDD methods find sensor measurement signatures to indicate faults. Signal-based FDD combining wavelet transformation and principal component analysis was presented in [20]. Although the methods achieved good performance, extracting relevant signatures and signals that indicate faulty condition is a daunting task for complex systems such as buildings. In data-driven FDD approaches, when labeled fault data are available, a FDD task are usually modeled as a multiclass classification problem. Then a supervised learning method can be employed to learn a classifier to recognize the faults. Many supervised methods such as multivariate regression models [26], Bayes classifiers [11], [38], [33], neural networks (NN) [6], [41], [4], Fisher Discriminant Analysis (FDA) [5], Gaussian Mixture Models [12], Support Vector Data Description (SVDD) [37], [39], and Support Vector Machines (SVM) [22], [9], [1], [35], [25] have been proposed to classify the faults. Recently, Li et al. proposed a tree-structured learning method [19] that not only recognizes faults but also their severity levels; however, it is hard in practice to obtain such a well-labeled dataset that include incipient faults. Researchers have also proposed unsupervised approaches using Principal Component Analysis (PCA) [21], Statistical Process Control (SPC) [29], and autoencoders [27] for FDD. Depending only on positive (healthy) class data, such unsupervised methods have found their use in detecting anomalies; however, they still lack the ability to diagnose these anomalies.

A review of the literature reveals that data-driven approaches relying on supervised learning are promising methods due to their ability to classify and differentiate data with multiple labels. However, in order to train a well-performing model, large amount of labeled data is typically needed, which is not always easy to obtain. Furthermore, although supervised learning tends to perform well on known (in-distribution) data patterns, the unseen (out-of-distribution) data may lead to unexpected prediction behaviors. In the context of FDD, an incipient fault example not seen in the training phase may fool the classifier into wrong belief, which is certainly not desirable for FDD applications. Although this problem can be conceptually alleviated by using a larger, more comprehensive training dataset, in practice it is technically infeasible to obtain fault data of all different fault types, and of all possible severity

levels, especially for complex building systems such as chillers.

The aforementioned reasons motivate us to devise a method that can make full use of the available training data. The resulting classifier should not only be good at classifying in-distribution data points, but is also able to give reasonable diagnostic suggestions as well as its prediction uncertainty for out-of-distribution fault examples. The contribution of this paper is two-fold:

- We propose using the uncertainty information given by machine learning models to detect and diagnose unknown incipient faults in building systems.
- To effectively estimate the uncertainty information for this purpose, we propose using MC-dropout networks. The approach requires few modifications to the standard deep learning pipeline, making it attractive for real-world FDD applications. A case study is conducted on the ASHRAE RP-1043 dataset, which demonstrates the effectiveness of our approach.

The rest of the paper is organized as follows. In Sec. II, we give the necessary background about the MC-dropout approach and discuss how it can help us identify faults under development. Next, we describe and analyze the RP-1043 dataset in Sec. III, and then present a case study in Sec. IV with extensive experiment results to show the effectiveness of our proposed approach on the RP-1043 dataset. We later conclude the paper in Sec. V and also discuss potential future steps.

II. MONTE CARLO DROPOUT APPROACH

A. Neural Network Classifiers

The MC-dropout network to be introduced later is based on the classic feed-forward neural network model [18], a.k.a. *multilevel perceptron*. A network usually consists of two or more layers, and can be described as a series of functional transformations on the input vector. We take a simple feedforward neural network with one hidden layer, shown in Fig. 1, as an example. The value of the hidden layer $\mathbf{h}$ is computed from the input vector $\mathbf{x}$ in the following way.

$$\mathbf{h} = g(\mathbf{W}_1 \mathbf{x} + b_1), \quad (1)$$

where $\mathbf{W}_1$ defines an affine transformation to $\mathbf{x}$, and g is an *activation function* that is typically nonlinear and applied element-wise to $\mathbf{W}_1 \mathbf{x} + b_1$. In modern neural networks, the rectified linear unit or ReLU defined by $g(t) = \max\{0, t\}$ is usually used as the activation function [8].

For a multiclass classification problem with C classes, we need the neural network classifier to output a vector $\hat{\mathbf{y}} = (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_C)$ representing a discrete probability distribution. We require that each $\hat{y}_i = P(y = i | \mathbf{x}) \in (0, 1)$, and $\sum_{i=1}^C \hat{y}_i = 1$, i.e. these probabilities sum up to 1. A softmax activation function is usually used at the last layer to obtain the desired $\hat{\mathbf{y}}$. Let $\mathbf{z} = \mathbf{W}_2 \mathbf{h}$ be the *activation* of the last layer¹. Under the softmax transformation, for each class

¹The bias coefficient is not needed for this layer because adding a bias to every element of $\mathbf{z}$ will not change the softmax output.




Fig. 1: An example of a simple feedforward neural network with one hidden layer. The intercept parameter b_1 associated with the hidden layer is omitted for brevity.

$i = 1, 2, \dots, C$, $z_i = \log P(y = i | \mathbf{x})$. And then the softmax probability for class i is given by

$$P(y = i | \mathbf{x}) = \frac{\exp(z_i)}{\sum_{j=0}^C \exp(z_j)}. \quad (2)$$

To train such models for multiclass classification tasks, we usually minimize the cross-entropy loss as below

$$L = -\frac{1}{N} \sum_{j=1}^N \sum_{i=1}^C y_i^{(k)} \log y_i^{(k)}, \quad (3)$$

in order to maximize the log-likelihood of the softmax distribution, over the training samples. We refer the interested readers to the review paper by LeCun et al. [18] and references therein for a more thorough introduction to deep neural network models and techniques.

Here, we illustrate the overconfidence problem of neural networks by using a toy example in two dimensions shown in Fig. 2a. As displayed in the plot, the decision boundary forms a narrow band running across the intermediate states, separating the healthy and severe fault examples; most of the gray points that are not on the decision boundary are either classified as healthy or faulty with high confidence (with network output very close to 0 or 1). By using this decision boundary, the gray points that are closer to the origin will be classified as healthy, while most others will be reckoned faulty; see Fig. 2b. Only the few that reside on the boundary will be considered ambiguous because the fault probabilities are close to 0.5. Although this classification model can recognize some faulty conditions, it is not enough for detecting incipient faults. In addition, the trained model also shows high confidence on the bottom right region where no data are available. It seems that the network has extrapolated its learned pattern to the unseen region. Such extrapolation could be sometimes dangerous—the data distribution might be totally different in that region. It is more desirable for the model to be cautious about what the available data cannot offer.

B. Estimating Predictive Uncertainty

Incipient faults are often characterized by small deviations from fault-free conditions. As a result, their behaviors often




Fig. 2: A toy example in 2 dimensions. (a) The healthy state data (in blue) are confined in a circle with radius $r = 0.3$ centered at the origin. The severe fault data (in red) reside outside $r = 0.7$. In between the blue and the red points are the intermediate states (gray) that are not observed in the training distribution. (b) The decision boundary (orange) given by a trained non-dropout neural network classifier. Intermediate-level fault data points classified as healthy are shown in light blue, and those classified as faulty by the decision boundary are shown in light red.

resemble both fault-free and fault patterns, which could confuse a classifier. To detect incipient faults, it would be desirable to have a neural network model that can tell its uncertainty in the predictions; however, as previously mentioned, neural networks are usually poor at telling the predictive uncertainties, and often tend to be overly confident in its predictions. The standard and state-of-the-art approach for estimating predictive uncertainty is to use Bayesian neural networks, whose goal is to learn a distribution over weights; however such approaches are typically computationally expensive compared to standard (non-Bayesian) neural neural networks, and do not naturally fit into the standard training pipeline of today's deep learning frameworks, which could limit their use. Details of Bayesian neural networks are beyond the scope of this paper, and we refer interested readers to the paper [17] and references therein for more in-depth discussions.

Besides Bayesian neural networks, there are also approaches that use ensembles of models for estimating predictive uncertainty [17]. The MC-dropout approach to be described next also belongs to these ensemble methods. Such ensemble approaches typically involves some randomization, either in the base learners themselves or in the data used to train each base learner. The idea behind is simple and intuitive: use an ensemble of individual models to obtain multiple predictions, and use the empirical variance of predictions as an approximate measure of uncertainty [17]. For these ensemble methods to work, the individual classifiers must exhibit *diversity* among themselves; the diversity allows individual classifiers to generate different decision boundaries. Due to the randomness of decision boundaries learned by each individual classifier, the ensemble can hopefully give a high predictive uncertainty on out-of-distribution data points, which provides us with a way to detect them. Ensemble learning requires a combination of many diverse base learners in order to build an ensemble classifier. As a result, the idea of ensemble learning is typically adopted in learning schemes where the base learners can be




Fig. 3: A simple feedforward neural network with dropout.

fitted quickly, e.g., random forests, which seems to limit the use of neural networks in an ensemble model because they are time-consuming to train.

C. The MC-dropout Approach

Dropout [28] is a popular and powerful regularization technique to prevent over-fitting neural network parameters. The key idea is to randomly drop units along with their connections from the network during training; see Fig. 3 for an illustration. Each individual hidden node is dropped at a probability of p , i.e. the dropout rate. The training and inference procedure is then run as usual. In effect, the dropout technique provides an inexpensive approximation to training and evaluating an ensemble of exponentially many neural networks.

The dropout mechanism offers a way to incorporate intrinsic randomization into neural network models; recently, Gal and Ghahramani proposed using MC-dropout [7] to estimate a network's predictive uncertainty by using dropout at test time. During testing, we treat a model $\mathcal{M}$ trained using dropout as if we were using it during the training phase. Each time we forward pass a given input x through the network, each hidden node in the network will be dropped at a probability of p and we will obtain a random output $\hat{y}$. By repeating the same process for T times, we will obtain T i.i.d. sampled output vectors $\hat{y}^{(1)}, \hat{y}^{(2)}, \dots, \hat{y}^{(T)}$. Their predictive mean $\mathbb{E}[\hat{y}] = \frac{1}{T} \sum_{k=1}^T \hat{y}^{(k)}$ can be understood as the expected output given input x , and the predictive variance $\frac{1}{T} \sum_{k=1}^T (\hat{y}^{(k)} - \mathbb{E}[\hat{y}])$ can be used to measure the *confidence* of $\mathcal{M}$ in its prediction. The larger the predictive variance is, the more uncertain is the network about its prediction.

To illustrate the approach, we trained a simple feed-forward neural network with MC-dropout on the previously introduced 2D toy example; the results are shown in Fig. 4. It can be observed that the decision boundary in orange (given by points with their predictive mean close to 0.5) looks similar to the boundary given by the non-dropout network (see Fig. 2). Here the predictive uncertainty information will play a crucial role in suggesting potential faults. We can see from the right panel of Fig. 4 that the regions where intermediate states reside, especially in the vicinity of the healthy region, are associated with elevated predictive variance, shown as green shades. In addition, high predictive variance is present in the bottom right region, where no data points have been observed. These

TABLE I: Descriptions of the variables used as features

Sensor	Description	Unit
TEI	Temperature of entering evaporator water	°F
TEO	Temperature of leaving evaporator water	°F
TCI	Temperature of entering condenser water	°F
TCO	Temperature of leaving condenser water	°F
Cond Tons	Calculated Condenser Heat Rejection Rate	Tons
Cooling Tons	Calculated City Water Cooling Rate	Tons
kW	Compressor motor power consumption	kW
FWC	Flow Rate of Condenser Water	gpm
FWE	Flow Rate of Evaporator Water	gpm
PRE	Pressure of refrigerant in evaporator	psig
PRC	Pressure of refrigerant in condenser	psig
TRC	sub Subcooling temperature	°F
T _{suc}	Refrigerant suction temperature	°F
T _{sh_suc}	Refrigerant suction superheat temperature	°F
TR _{dis}	Refrigerant discharge temperature	°F
T _{sh_dis}	Refrigerant discharge superheat temperature	°F

behaviors indicate that the trained MC-dropout network is suspicious about, while still being cautious, potential faults in such regions, although not much prior information is provided by the data distribution for training the network model. Such suspicion is important and can serve as alarms for potential faults, and further inspection and maintenance measures shall be taken if necessary.

The example above demonstrates that MC-dropout technique can provide hints about incipient faults, which is regular neural networks fall short of. But just detecting faults or anomalies is not our sole purpose, nor does it make the MC-dropout approach distinctive from other approaches. Classic anomaly detection methods, such as PCA and autoencoders, can also suggest out-of-distribution inputs; however, these approaches do not possess discriminative ability between fault conditions. Next, we will demonstrate through a case study on the ASHRAE RP-1043 Dataset that the MC-dropout approach can not only detect, but also provide preliminary diagnosis about, underlying health problems of an industrial chiller system.

III. CHILLER SYSTEM AND DATASET

We used the ASHRAE RP-1043 Dataset [3] to test out the proposed MC-dropout approach for incipient fault detection. In RP-1043, sensor measurements of a typical cooling system—a 90-ton centrifugal water-cooled chiller—were recorded under both fault-free and various fault conditions. The 90-ton chiller is representative of chillers used in larger installations [24], and consists of the following parts: evaporator, compressor, condenser, economizer, motor, pumps, fans, and distribution pipes etc. with multiple sensor mounted in the system. Fig. 5 depicts the cooling system with sensors mounted in both evaporation and condensing circuits.

In RP-1043 experimental data, eight different types of process faults were injected into the chiller, and each fault was




Fig. 4: The spatial distribution of the predictive mean (in orange shades) and the predictive variance (in green shades) of an MC-dropout model trained on the aforementioned toy example data. For the predictive mean plot, the color intensity signifies the proximity between the predictive mean at a point and 0.5. In the predictive variance plot, the more intense the color, the higher predictive variance at a given point.




Fig. 5: Schematic of the cooling system test facility and sensors mounted in the related water circuits.

introduced at four levels of severity (SL1 - SL4, from slightest to severest). In this study, we only included the six faults shown in Table II, because an earlier study by Reddy [30] found certain limitations with the excess oil and faulty TXV operation data.

The condenser fouling (CF) fault was emulated by plugging tubes into condenser. The reduced condenser water flow rate (FWC) fault and reduced evaporator water flow rate (FWE) fault were emulated directly by reducing water flow rate in the condenser and evaporator. The refrigerant overcharge (RO) fault and refrigerant leakage (RL) fault were emulated by reducing or increasing the refrigerant charge respectively. The excess oil (EO) fault was emulated by charging more oil than nominal. And the non-condensable in refrigerant (NC) fault was emulated by adding Nitrogen to the refrigerant.

All faults were tested at 27 different operating conditions with varying chiller thermal load, chilled water outlet and inlet temperature settings. The data were collected at ten-second intervals, not only when the system has reached steady state, but also at transient states in between. We focused on only the steady-state data in this study. A more detailed review and discussion on the dataset can be found in [30]. The sixteen




Fig. 6: A visualization of the RP-1043 dataset, including the six faults being studied as well as the fault-free data. To visualize these high-dimensional data points, we employed LDA for dimensionality reduction. The color intensity of a data point signifies the severity of the corresponding fault. In (a), the SL1-SL4 data form clusters that are easily distinguishable. In (b) and (c) faults of different severity levels are harder to separate.

TABLE II: The six chiller faults in our study

Fault	Normal Operation
1 Reduced Condenser Water Flow (FWC)	270 gpm
2 Reduced Evaporator Water Flow (FWE)	216 gpm
3 Refrigerant Leak (RL)	300 lb
4 Refrigerant Overcharge (RO)	300 lb
5 Condenser Fouling (CF)	164 tubes
6 Non-condensables in System (NC)	No nitrogen

key features identified by RP-1043, as listed in Table I, were selected to train our neural network models.

To give the readers an intuition about the distribution of RP-1043 data, we employed LDA to reduce the data into 2 dimensions, and visualized the dimension-reduced data in Fig. 6. As can be seen in the plots, FWE, FWC and NC faults are further away from normal than RL, RO and CL faults are. We can also see a general trend for a data point to deviate further away from the normal when the corresponding fault develops into a more severe level. It can also be seen from the plots that, SL1 data points are often closer to the fault-free region than to their corresponding SL4 regions. In other words, these incipient fault data may look more like the fault-free data than their high-severity versions, which cast a challenge in using supervised learning algorithms (e.g., neural networks) to detect and diagnose these slight incipient faults.

IV. EXPERIMENTAL RESULTS

A. Experimental Setup

We conducted a case study to evaluate the performance of our MC-dropout approach on the RP-1043 dataset. We implemented the neural networks in Python using Keras [2]. The MC-dropout network $\mathcal{M}$ has four fully-connected hidden layers, each containing 20 nodes and with dropout layers interleaved. The output layer is a softmax layer with 7 output nodes, each representing an output class label (fault-free and the six faults in our study). Subscript p indicates the dropout rate being used. By setting $p = 0$, we will obtain a non-dropout network $\mathcal{M}_0$, which is used as a baseline to show the effect of MC-dropout on FDD performance.

B. Dropout Rate Selection

As with many other machine learning models, the selection of hyperparameters of a MC-dropout network can have a

significant impact on its prediction performance. In our case, we need to find a suitable dropout rate p , such that the resulting MC-dropout network can 1) accurately classify in-distribution examples with high confidence, and 2) identify ambiguous out-of-distribution examples by indicating high predictive uncertainty. The first requirement can be checked by using the usual cross-validation process; however, cross-validation will not help with the second requirement since we do not have access to the out-of-distribution data at training time.

In the experiment, we tested a number of dropout rates ranging from 0 to 0.5. Due to space limit, we only display in Fig. 7 the prediction results under five typical dropout rate settings ($p = 0, 0.03, 0.1, 0.2, 0.5$). Each network was trained for 30 epochs, with categorical crossentropy used as the loss function. For each network with a non-zero dropout rate, $N = 100$ Monte Carlo samples were drawn for estimating the predictive mean and variance.

With p set to zero, our MC-dropout network will degenerate to a regular non-dropout neural network—no output variance will exist given the same input because the network does not have any inherent randomness. The network is thus not able to provide predictive uncertainty information. When the dropout rate is high ($p = 0.5$), the excessive dropout randomness as can be expected will undermine the model's predictive capability even on in-distribution examples as shown in Fig. 7. In order to achieve a balance between the two aforementioned requirements, we used the following method to select a p value empirically. We gradually increase p until the predictive means (variances) start to drop (increase) fast. In this way, the trained model will still have good performance on in-distribution data, and can likely indicate out-of-distribution data via its predictive uncertainty. We chose $p = 0.1$ for our subsequent experiments. Next, we are going to analyze the model's performance on out-of-distribution data, and explain how the model's outputs can be used to indicate potential faults.

C. Comparison between Dropout and Non-dropout Methods

We show in Fig. 8 a comparison between the non-dropout network $\mathcal{M}_0$ and the MC-dropout network $\mathcal{M}_{0.1}$ with the chosen dropout rate $p = 0.1$. Their prediction results on examples of all four severity levels are displayed. Again, for in-distribution (SL0 & SL4) examples, both networks demonstrate




Fig. 7: The prediction results on in-distribution data, under five different dropout rate settings $p = 0, 0.03, 0.1, 0.2, 0.5$. In the displayed heatmaps, each row corresponds to a chiller state (NM or a SL4 fault condition), and the values in each row are the respective predictive mean/variance at each output node.

good classification ability; in addition, the results given by $\mathcal{M}_{0.1}$ shows little variance on in-distribution examples (SL0 & SL4), which indicates its high confidence on these decisions.

SL3 faults do not appear in the training distribution, and are less severe than the SL4 faults used for training the models. As can be seen from Fig. 8, both networks can still correctly recognize FWE-SL3, FWC-SL3 and RO-SL3 faults with high confidence; however, they also demonstrate some uncertainty about RL-SL3, CF-SL3 and NC-SL3 cases as can be seen from their predictive means. When the underlying fault is RL-SL3 or CF-SL3, the classifier is uncertain whether the chiller is in NM state, or the respective fault state. When the underlying fault is NC-SL3, the classifier is unsure about the underlying state being RO or NC. This phenomenon can be roughly explained using the LDA analysis in Fig. 6. From the plots we can see that RL, RO, CF faults all reside very closely to the NM data points, which is one cause of the confusion. We can also see the presence of elevated predictive variance in blocks under suspicion; the predictive variance given by the MC-dropout network serves as another metric for indicating predictive uncertainty in such scenarios.

Our MC-dropout approach further demonstrates its usefulness in detecting and diagnosing slighter faults (SL1 & SL2). It can be seen from the SL1 panel of Fig. 8a and Fig. 8b that both classifiers make wrong predictions under all fault conditions, often with high confidence. SL1 fault conditions are classified as fault-free (as in RL-SL1, RO-SL1 and CF-SL1), or as other fault types (as in FWE-SL1, FWC-SL1 and NC-SL1). Similar problems are seen in SL2 cases as well. These phenomena can be partially understood using the visualization in Fig. 6. These low-severity fault data points often reside in the proximity of the data points of NM or other fault types,

which results in the misclassification. Despite the ambiguity, the MC-dropout network $\mathcal{M}_{0.1}$ casts its skepticism through the predictive uncertainty information given by Monte Carlo sampling. For example, when the chiller is under RO-SL1 fault, the predictive mean given by $\mathcal{M}_{0.1}$ shows high confidence in believing the underlying state is NM, whereas the predictive variance shows high uncertainty in both NM and RO. The high predictive uncertainties indicate the possibility of a potential RO fault.

We show the diagnostic results for all SL1 and SL2 fault conditions in Table III. For the non-dropout network $\mathcal{M}_0$, we select class labels with softmax probability above 20% as possible diagnoses. For the MC-dropout network $\mathcal{M}_{0.1}$, we also include class labels with high predictive variance as possible diagnoses, in addition to those with softmax probability above 0.2. A class label is considered to be with high predictive variance, if the ratio between the standard deviation of this particular class label and that of all class labels is above 10%. It can be seen from the table that $\mathcal{M}_0$ cannot correctly diagnose any of the SL1 faults, and is only able to correctly identify FWE-SL2 and FWC-SL2 among all SL2 faults. In comparison, the diagnoses given by the MC-dropout network $\mathcal{M}_{0.1}$ contain the correct labels in all 12 cases listed in Table III. It can also be observed that $\mathcal{M}_{0.1}$ is more certain about prediction results on SL2 cases than on SL1 cases. In FWE-SL1, RL-SL1, CF-SL1 and NC-SL1 cases, $\mathcal{M}_{0.1}$ suggests three or more possible states, while in all SL2 cases the network only suggests two. This is understandable, because the fault signatures presented in SL1 cases are presumably less obvious than those in SL2 cases, thus creating less confusion for the neural network. Although the MC-dropout network is unable to give a definite diagnosis due to lack of information in such cases, it indicates some

TABLE III: Diagnosis Results

Actual Chiller State	Diagnosis from Non-dropout Network	Diagnosis from MC-dropout Network
<i>FWE-SL1</i>	NM	NM, FWE , RL
<i>FWC-SL1</i>	CF	NM, FWC , CF
<i>RL-SL1</i>	NM	NM, RL , RO, CF
<i>RO-SL1</i>	NM	NM, RO
<i>CF-SL1</i>	NM	NM, RL, RO, CF
<i>NC-SL1</i>	RO	NM, RO, CF, NC
<i>FWE-SL2</i>	NM, FWE	FWE , RL
<i>FWC-SL2</i>	FWC, CF	FWC , CF
<i>RL-SL2</i>	NM	NM, RL
<i>RO-SL2</i>	NM	NM, RO
<i>CF-SL2</i>	NM	NM, CF
<i>NC-SL2</i>	RO	RO, NC

potential fault and narrows down the possible causes, which is valuable for further maintenance decisions to uncover the true status of the chiller system.

From the above analyses, we can see that the uncertainty information given by MC-dropout networks is useful in cases where there is a lack of information about the characteristics of incipient faults. Under such situations, it is better and more reasonable for a classifier to show the uncertainty about its decisions, rather than just giving definite but often incorrect predictions. We can also see that the softmax probabilities alone have limited capability in telling the predictive uncertainty of a neural network, which limits its use in detecting unseen, incipient fault conditions. With MC-dropout, the network has gained another way to indicate its uncertainty through the predictive variance. Our experimental results on RP-1043 have demonstrated its usefulness in detecting and diagnosing incipient faults.

V. CONCLUSIONS

In this paper, we proposed using MC-dropout, a method for estimating a deep learning model's uncertainty in its decisions, to detect incipient or unknown faults in modern CPSs with a neural network trained with limited fault data. By presenting a case study on ASHRAE RP-1043 dataset, we have shown the effectiveness of MC-dropout in detecting and diagnosing chiller faults. As part of our future work, we plan to conduct a more theoretical analysis on the proposed approach to gain a better understanding of it.

ACKNOWLEDGMENT

This work is supported in part by the National Research Foundation of Singapore through a grant to the Berkeley Education Alliance for Research in Singapore (BEARS) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) program, and by the National Science Foundation under Grant No. 1645964. BEARS has been established by the University of California, Berkeley as a center for intellectual excellence in research and education in Singapore.

REFERENCES

[1] Kai-Ying Chen, Long-Sheng Chen, Mu-Chen Chen, and Chia-Lung Lee. Using SVM based method for equipment fault detection in a thermal power plant. *Computers in industry*, 62(1):42–50, 2011.

[2] François Chollet et al. Keras. <https://keras.io>, 2015.

[3] Matthew C Comstock, James E Braun, and Robert Bernhard. *Development of analysis tools for the evaluation of fault detection and diagnostics in chillers*. Purdue University, 1999.

[4] Zhimin Du, Bo Fan, Xinqiao Jin, and Jinlei Chi. Fault detection and diagnosis for buildings and HVAC systems using combined neural networks and subtractive clustering analysis. *Building and Environment*, 73:1–11, 2014.

[5] Zhimin Du and Xinqiao Jin. Multiple faults diagnosis for sensors in air handling unit using fisher discriminant analysis. *Energy Conversion and Management*, 49(12):3654–3665, 2008.

[6] Bo Fan, Zhimin Du, Xinqiao Jin, Xuebin Yang, and Yibo Guo. A hybrid FDD strategy for local system of AHU based on artificial neural network and wavelet analysis. *Building and environment*, 45(12):2698–2708, 2010.

[7] Yarín Gal. Uncertainty in deep learning. *University of Cambridge*, 2016.

[8] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.

[9] Hua Han, Zhikun Cao, Bo Gu, and Neng Ren. PCA-SVM-based automated fault detection and diagnosis (AFDD) for vapor-compression refrigeration systems. *HVAC & R Research*, 16(3):295–313, 2010.

[10] Zhangming He, Yuri AW Shardt, Dayi Wang, Bowen Hou, Haiyin Zhou, and Jiongqi Wang. An incipient fault detection approach via detrending and denoising. *Control Engineering Practice*, 74:1–12, 2018.

[11] David J Hill, Barbara S Minsker, and Eyal Amir. Real-time bayesian anomaly detection for environmental sensor data. In *Proceedings of the Congress-International Association for Hydraulic Research*, volume 32, page 503. Citeseer, 2007.

[12] Padmini Jaikumar, Aca Gacic, Burton Andrews, and Michael Dambier. Detection of anomalous events from unlabeled sensor data in smart building environments. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2268–2271. IEEE, 2011.

[13] Saththasivam Jayaprakash and Ng Kim Choon. Predictive and diagnostic methods for centrifugal chillers. *ASHRAE Trans*, 114(1):282–287, 2008.

[14] R. Jia, B. Jin, M. Jin, Y. Zhou, I. C. Konstantakopoulos, H. Zou, J. Kim, D. Li, W. Gu, R. Arghandeh, P. Nuzzo, S. Schiavon, A. L. Sangiovanni-Vincentelli, and C. J. Spanos. Design automation for smart building systems. *Proceedings of the IEEE*, 106(9):1680–1699, Sept 2018.

[15] Qin Jianying and Wang Shengwei. A fault detection and diagnosis strategy of VAV air-conditioning systems for improved energy and control performances. *Energy and Buildings*, 37(10):1035–1048, 2005.

[16] Ru Jifeng and Li X Rong. Variable-structure multiple-model approach to fault detection, identification, and estimation. *IEEE Transactions on Control Systems Technology*, 16(5):1029–1038, 2008.

[17] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pages 6402–6413, 2017.

[18] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.

[19] Dan Li, Yuxun Zhou, Guoqiang Hu, and Costas J Spanos. Fault detection and diagnosis for building cooling system with a tree-structured learning method. *Energy and Buildings*, 127:540–551, 2016.

[20] Shun Li and Jin Wen. A model-based fault detection and diagnostic methodology based on pca method and wavelet transform. *Energy and Buildings*, 68:63–71, 2014.

[21] Shun Li and Jin Wen. A model-based fault detection and diagnostic methodology based on PCA method and wavelet transform. *Energy and Buildings*, 68:63–71, 2014.

[22] Jian Liang and Ruxu Du. Model-based fault detection and diagnosis of HVAC systems using support vector machine method. *International Journal of refrigeration*, 30(6):1104–1114, 2007.

[23] J. Braun M. Comstock and E. Groll. The sensitivity of chiller performance to common faults. *HVAC & R Res*, 7(3):263–279, 2001.

[24] Comstock MC and Braun JE. Fault detection and diagnostic (FDD) requirements and evaluation tools for chillers. *ASHRAE*, 2002.

[25] Timothy Mulumba, Afshin Afshari, Ke Yan, Wen Shen, and Leslie K Norford. Robust model-based fault diagnosis for air handling units. *Energy and Buildings*, 86:698–707, 2015.

[26] G. Mustafaraj, J. Chen, and G. Lowry. Development of room temperature and relative humidity linear parametric models for an open office using BMS data. *Energy and Buildings*, 42:348–356, August 2010.




Fig. 8: Output comparison of the non-dropout network $\mathcal{M}_0$ and the MC-dropout network $\mathcal{M}_{0.1}$, on faults of all four severity levels.

- [27] Mayu Sakurada and Takehisa Yairi. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis*, page 4. ACM, 2014.
- [28] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [29] Biao Sun, Peter B Luh, Qing-Shan Jia, Zheng O'Neill, and Fangting Song. Building energy doctors: An SPC and kalman filter-based method for system-level fault detection in HVAC systems. *IEEE Transactions on Automation Science and Engineering*, 11(1):215–229, 2014.
- [30] PE T Agami Reddy PhD. Development and evaluation of a simple model-based automated fault detection and diagnosis (FDD) method suitable for process faults of large chillers/discussion. *ASHRAE Transactions*, 113:27, 2007.
- [31] Kajiyo Watanabe, Ichiro Matsuura, Masahiro Abe, Makoto Kubota, and DM Himmelblau. Incipient fault diagnosis of chemical processes via artificial neural networks. *AIChE journal*, 35(11):1803–1812, 1989.
- [32] Dumidu Wijayasekara, Ondrej Linda, Milos Manic, and Craig Rieger. Mining building energy management system data using fuzzy anomaly detection and linguistic descriptions. *Industrial Informatics, IEEE Transactions on*, 10(3):1829–1840, 2014.
- [33] Fu Xiao, Yang Zhao, Jin Wen, and Shengwei Wang. Bayesian network based FDD strategy for variable air volume terminals. *Automation in Construction*, 41:106–118, 2014.
- [34] Dai Xuewu and Gao Zhiwei. From model, signal to knowledge: A data-driven perspective of fault detection and diagnosis. *IEEE Transactions on Industrial Informatics*, 9(4):2226–2238, 2013.
- [35] Ke Yan, Wen Shen, Timothy Mulumba, and Afshin Afshari. ARX model based fault detection and diagnosis for chillers using support vector machines. *Energy and Buildings*, 81:287–295, 2014.
- [36] Wen Shen Yan Ke and Timothy Mulumba. ARX model based fault detection and diagnosis for chillers using support vector machines. *Energy and Buildings*, 81:287–295, 2014.
- [37] Yang Zhao, Shengwei Wang, and Fu Xiao. Pattern recognition-based chillers fault detection method using support vector data description (SVDD). *Applied Energy*, 112:1041–1048, 2013.
- [38] Yang Zhao, Fu Xiao, and Shengwei Wang. An intelligent chiller fault detection and diagnosis methodology using bayesian belief network. *Energy and Buildings*, 57:278–288, 2013.
- [39] Yang Zhao, Fu Xiao, Jin Wen, Yuehong Lu, and Shengwei Wang. A robust pattern recognition-based fault detection and diagnosis (FDD) method for chillers. *HVAC&R Research*, 20(7):798–809, 2014.
- [40] Gao Zhiwei, Carlo Cecati, and Ding Steven X. A survey of fault diagnosis and fault-tolerant techniques-part i: fault diagnosis with model-based and signal-based approaches. *IEEE Transactions on Industrial Electronics*, 62(6):3757–3767, 2015.
- [41] Yonghua Zhu, Xinqiao Jin, and Zhimin Du. Fault diagnosis for sensors in air handling unit based on neural network pre-processed by wavelet and fractal. *Energy and buildings*, 44:7–16, 2012.