

Received September 16, 2019, accepted October 16, 2019, date of publication October 25, 2019, date of current version November 6, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2949740

# Detecting Gas Vapor Leaks Using Uncalibrated Sensors

DIAA BADAWI<sup>1</sup>, TUBA AYHAN<sup>2</sup>, SULE OZEV<sup>3</sup>, CHENGMO YANG<sup>4</sup>,  
ALEX ORAILOGLU<sup>5</sup>, AND AHMET ENIS ÇETİN<sup>1</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, The University of Illinois at Chicago, Chicago, IL, USA

<sup>2</sup>Electronics and Communication Engineering Department, Istanbul Technical University, Istanbul, Turkey

<sup>3</sup>School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ, USA

<sup>4</sup>Department of Electrical and Computer Engineering, University of Delaware, Newark, DE, USA

<sup>5</sup>Department of Computer Science and Engineering, University of California at San Diego, San Diego, CA, USA

Corresponding author: Daa Badawi (dbadaw2@uic.edu)

This work was supported in part by the NSF under Grant 1739396 (UIC), Grant 1739451 (ASU), Grant 1739390 (UD), and Grant 1739684 (UCSD). The work of D. Badawi and A. E. Çetin was supported by the Nvidia.

**ABSTRACT** Chemical and infra-red sensors generate distinct responses under similar conditions because of sensor drift, noise or resolution errors. In this paper, we develop novel machine learning methods for detecting and identifying VOC and Ammonia vapor from time-series data obtained by *uncalibrated* chemical and infrared sensors. We process time-series sensor signals using deep neural networks (DNN). Three neural network algorithms are utilized for this purpose. Additive neural networks (termed AddNet) are based on a multiplication-devoid operator and consequently exhibit energy efficiency compared to regular neural networks. The second algorithm uses generative adversarial neural networks so as to expose the classifying neural network to more realistic data points in order to help the classifier network to deliver improved generalization. Finally, we use conventional convolutional neural networks as a baseline method. Our findings indicate that using raw time-series data obtained from uncalibrated sensors and processing them using deep-learning-based methods yield better results than using hand-crafted feature parameters.

**INDEX TERMS** VOC gas leak detection, sensor drift, additive, convolutional, and generative adversarial (GAN) neural networks, time-series data analysis.

## I. INTRODUCTION

Ammonia and Volatile organic compounds (VOCs) are associated with numerous health problems. Although VOCs and Ammonia are naturally occurring, they can nonetheless cause serious health issues in high concentration. For example, exposure to Ammonia in high concentration causes harm to skin, lungs and eyes. Methane and other VOC compound leaks contribute to global warming. VOC compounds such as benzene and toluene are carcinogenic [1]–[4].

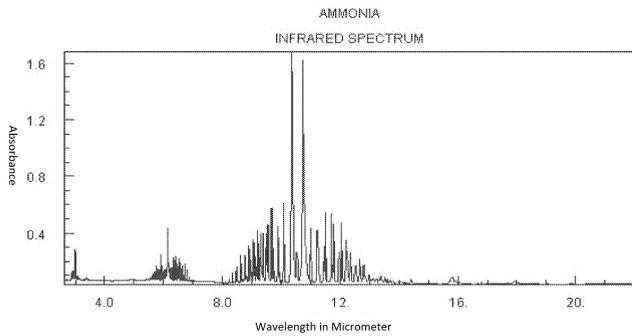
In this paper, we develop novel machine learning methods for uncalibrated VOC and ammonia vapor sensors.<sup>1</sup> We show that there is no need to use either expert-crafted features or thresholds for the purpose of interpreting the time series data that these sensors generate. This is advantageous in that it exempts system designers from extracting discriminative

The associate editor coordinating the review of this manuscript and approving it for publication was Seyedali Mirjalili<sup>id</sup>.

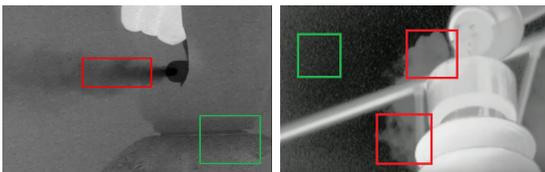
<sup>1</sup>This work was presented in part at the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brighton, UK, May 2019 [5].

and robust features from the raw time-series sensory data, which is a nontrivial task. To the best of our knowledge, this is the first work that uses the raw time-series signal for VOC detection and ammonia vapor sensing using deep learning. We train the deep learning structures using only the time-series data. Our algorithms are applicable to both infrared (IR) and chemical sensor systems. In particular, they can be used for early detection and, thus, prevention of dangerous gas leaks. Mobile infrared and chemical sensors can be part of an open air cyber-physical system (CPS) [5]–[7]. We use the time-series data obtained by the sensors in order to detect accidental and/or deliberate gas vapor leaks. The main contribution of this paper centers on the exploitation of the time-series data that sensors produce rather than conventional reliance on a single or a couple of sensor readings for leak detection.

Some VOC gas vapors such as ethane and ammonia absorb infrared light in the Long Wave Infrared (LWIR) while others such as methane in the Medium Wave Infra-red (MWIR) bands. Absorbance by ammonia of infra-red light at different



**FIGURE 1.** Infrared spectrum of ammonia. The figure is taken from NIST [16].



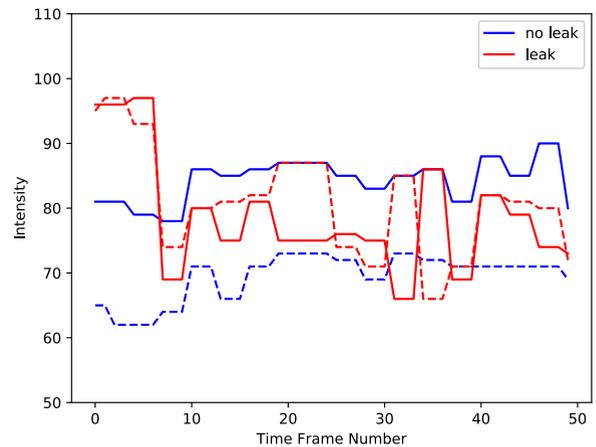
**FIGURE 2.** Two infrared images of VOC gas leaks. Red rectangles contain gas-leak regions. Green rectangles contain leak-free regions. Images are downloaded from FLIR systems and Infrared Cameras Inc. [17], [18], respectively.

wavelengths is shown in Fig 1. We can easily observe the existence of VOC gas vapor using Infra-red (IR) cameras in open air as shown in Fig. 2. In this figure, a dark smoke-like region denotes the image of VOC gas vapor. However, the distance between the sensor and the source, and infrared reflections from the background significantly affect the recorded level [8], [9].

Conventional optical devices, such as gas chromatographs and MWIR cameras, are generally expensive. A cheaper alternative would be the use of IR sensors and chemical gas sensors. Yet chemical gas sensors incur degradation in their sensitivity over time. Consequently, identically manufactured sensors are likely to yield significantly different responses upon exposure to gas analytes under identical conditions [10]–[14]. This problem is known in the literature as the *sensor drift problem*.

Causes of sensor drift can be summed up by two phenomena, namely, the physical changes in the structure of the sensor and the changes in the operating environment. The former case is known as *first-order sensor drift*. It is caused by sensor aging or by sensor “poisoning”.<sup>2</sup> Unfortunately, neither poisoning nor aging are reversible as the physical structure of the sensor will have been permanently damaged or at the least affected. The latter case is known as *second-order sensor drift* and is caused by external uncontrollable environmental changes, such as temperature and humidity variations. In this regard, the sensor response will be different from that expected from the original settings. Consequently, any decision thresholds that are optimal prior to sensor drift are likely to exhibit sub-optimal sensitivity and/or specificity once the aforementioned changes take place.

<sup>2</sup>A process by which the sensor surface absorbs some compounds irreversibly, thus reducing its resistance sensitivity [15].



**FIGURE 3.** Example infrared sensor time-series data. The red-colored time series are acquired from the vicinity of gas leak, whereas blue-colored samples correspond to time-series with absence of any gas leak.

Similarly, while it is not possible to detect the concentration of the gas using MWIR and LWIR sensors in open air, it is possible to record a time-varying signal and detect the existence of gas leakage using IR sensors as shown in Fig. 3 using a machine learning algorithm such as a neural network. The sensor signal exhibits sudden jumps and fluctuations due to gas vapor leak. Uncalibrated IR sensor intensity measurements suddenly drop from 95 to 70 and fluctuate because of wind as shown in Fig. 3. As one can see from Fig. 3, it is not possible to find a threshold that can isolate leak-free intensity signals from those corresponding to a gas leak. Furthermore, the intensity of leak-free point signals is time-varying. This demonstrates the effect of noise, resolution and lighting factors that in turn lend further complexity to the task of distinguishing between the two classes of signals. The two sets of examples are distinguished by dotting.

In this paper, we analyze the temporal sensor signals using convolutional, additive neural networks and the discriminator of a generative adversarial network (GAN) to detect and classify VOC gas leaks and other dangerous gas emissions. The proposed analysis is applicable to both Chemically-sensitive Field Effect Transistors (ChemFETs) and Electrochemical Impedance Spectroscopy (EIS) and infra-red sensors as they all produce time-varying signals.

The rest of the paper is organized as follows. Section II describes the machine learning algorithms used in this paper. Section III presents experimental results. We use an infrared data set and two publicly available chemical sensor drift data sets obtained at the University of California at San Diego (UCSD) [13] and [19]. The paper finishes by offering a brief set of conclusions in Section IV.

## II. DEEP LEARNING ALGORITHMS FOR IR AND CHEMICAL SENSOR DATA PROCESSING

In this paper, we consider three tasks. Task 1 is infra-red sensor-based gas-leakage detection. In tasks 2 and 3, we identify different types of gas analytes.

Our first network is an energy-efficient network, namely, an additive neural network, which is a neural network that performs no vector-multiplication except in its last layer. Our second neural network is the discriminator of a generative adversarial neural network, to which we refer shortly as DiscGAN.

### A. CONVOLUTIONAL NEURAL NETWORKS

Convolutional neural networks (or ConvNets) have been extensively used in computer vision [20], [21] and time-series data analysis [22]. In ConvNets, convolutions (or local correlations) between the inputs and the filter weights are used to extract local features at different scales in subsequent layers.

### B. ADDITIVE NEURAL NETWORKS (ADDNET)

Despite their ability to learn and recognize images and signals, deep learning algorithms are computationally expensive. This is attributed to the large number of add-multiply operations needed to be implemented in order to realize convolutions. This poses a problem when it comes to using such methods on platforms where energy is limited. As a result, simpler and, thus, more efficient algorithms are generally required to implement computationally expensive deep learning algorithms in such cases.

Nevertheless, there have been attempts to leverage convolutional neural networks across energy-limited devices by means of methods that aim to either implement fewer dot-product operations, or to replace dot-product operations with computationally simpler operations. Binarizing the weights and/or the activations results in replacing real-number multiplication operations with binary logical operations when realizing convolution, as in the case of BinaryConnect [23], XNOR-Net [24] and Binarized Neural Networks [25].

An additive neural network (AddNet) falls under the second category, i.e., replacing real-valued multiplication operations in vector-vector and matrix-vector product operations by special addition operations. The new “product” operation comprises binary sign calculation, unsigned addition and regular addition.

In what follows, we define the scalar version of our binary operation and extend it straightforwardly to its vector operation. In this regard, let  $x$  and  $y \in \mathbf{R}$ , the multiplication-devoid (abbreviated md) operation denoted by  $\oplus$  and defined as follows:

$$x \oplus y := \text{sgn}(x \cdot y)(|x| + |y|) \quad (1)$$

where  $\text{sgn}$  denotes the *signum* function. Alternatively, we can express the  $\oplus$  operation as follows:

$$x \oplus y := \text{sgn}(x)y + \text{sgn}(y)x \quad (2)$$

This is because  $x = \text{sgn}(x)|x|$ . One key property of the md operation is that it preserves the sign of regular multiplication operations [26], [27]. We define the vector version of the md operation as follows. Let  $x$  and  $w$  be two vectors in  $\mathbf{R}^N$ .

The md dot “product” is defined as:

$$\mathbf{w}^T \oplus \mathbf{x} := \sum_{i=1}^N \text{sgn}(x_i \cdot w_i)(|x_i| + |w_i|) \quad (3)$$

It can be seen that the md operation expressed in Eq. 3 requires no real-valued multiplication whatsoever. As such, instead of using add-multiply operations as in an ordinary dot product, we use ordinary addition and addition with sign multiplication in the md vector operation. Furthermore, we can restrict the operands  $x_i$  and  $w_i$  to be 8-bit numbers in order to speed up the vector addition operations. Another property of the md operation is that it induces the  $\ell_1$  norm. This is shown as follows:

$$\mathbf{x}^T \oplus \mathbf{x} = \sum_{i=1}^N \text{sgn}(x_i \cdot x_i)(|x_i| + |x_i|) = 2\|\mathbf{x}\|_1 \quad (4)$$

In the context of neural networks, we use convolution and matrix-vector multiplication operations in convolutional and dense layers, respectively. In AddNet, we replace the aforementioned dot-product operations with the md vector product. The feed-forwarding pass in dense layers in a neural network can be expressed as follows:

$$o_i^l = \phi(\mathbf{w}_i^{lT} \mathbf{o}^{l-1} + b_i^l) \quad (5)$$

where the superscript denotes the layer index, the subscript the neuron index,  $\mathbf{w}_i^l$  the weights connecting the output of the previous layer (the  $(l-1)$ st layer) to the  $i$ th neuron,  $o_i^l$  the output of the  $i$ th neuron, and bold  $\mathbf{o}^{l-1}$  the vector output of the previous layer.  $\phi$  is the non-linearity function applied element-wise and, finally,  $b_i^l$  denotes the bias term added to the pre-activated response  $\mathbf{w}_i^{lT} \mathbf{o}^{l-1}$ . Similarly, we can define AddNet layers by replacing the dot-product  $\mathbf{w}_i^{lT} \mathbf{o}^{l-1}$  by our md operator as follows:

$$o_i^l = \phi(\mathbf{w}_i^{lT} \oplus \mathbf{o}^{l-1} + b_i^l) \quad (6)$$

Since the md operator is additive, it will result in a larger output than ordinary multiplication does when either of the operands is of small magnitude, e.g.  $3 \oplus 0.1 = 3.1 > 3 \times 0.1 = 0.3$ . In the context of neural networks, the layer outputs and the weights are usually small values. As a result, the responses of the md layers will be of larger variance than those of the regular layer. This poses a problem in deep layers, where the dimension of the dot-product is quite large. In other words, if the depth of a convolutional layer is 64 and the kernel size is  $3 \times 3$ , the convolution operations will carry out dot-products between two vectors, each of which  $\in \mathbb{R}^{3 \times 3 \times 64}$ . In the case of the md layer, this will cause the output to exhibit inordinately high magnitudes. In order to overcome this, we introduce a scaling factor  $\alpha$ . As such, the feedforwarding pass in Eq. 6 becomes

$$o_i^l = \phi(\alpha_i^l(\mathbf{w}_i^{lT} \oplus \mathbf{o}^{l-1}) + b_i^l) \quad (7)$$

The scaling factor  $\alpha_i^l$  enables us to control the range of the output prior to applying the activation function  $\phi$  and, thus,

leads to a controlled range of responses in subsequent layers. Note that the scaling by  $\alpha_i^l$  in Eq. 7 implies real-valued multiplication. Nevertheless, it requires only one real-valued multiplication per neuron. Therefore, carrying out scaling is not computationally expensive. Numerous options exist for selecting the scaling factor  $\alpha_i^l$ . One possibility may be the setting of  $\alpha_i^l$  to  $\frac{1}{\|w_i^l\|_1}$ , i.e., the reciprocal of the  $\ell_1$  norm of the associated weights. Another option would be having  $\alpha_i^l$  be trainable by backpropagation. The latter delivers more flexibility for the model.

Nevertheless, batch normalization is a common practice in neural networks and has shown to be quite effective in accelerating the training of deep networks [28]. Therefore, one can simply apply batch normalization to the pre-activation responses in AddNet. Such normalization eliminates the need to carry out scaling by  $\alpha_i^l$  as it will be subsumed by the scaling induced by batch normalization.

The proof of AddNet with linear and/or ReLU activation functions satisfying the *universal approximation property* over the space of Lebesgue integrable functions can be found in [29].

As for training the md layers by backpropagation, it is worth mentioning that the derivative of the *signum* function used in the definition has to be computed. This is because  $\frac{d \operatorname{sgn}(w)}{dw} = 2\delta(w)$ , where  $\delta$  is the Dirac-delta function. In practice, this means that the derivative of the *signum* function is zero almost everywhere except when  $w = 0$ .

The partial derivative of the md operator w.r.t.  $w$  is:

$$\frac{\partial(w \oplus x)}{\partial w} = \operatorname{sgn}(x) + 2x\delta(w) \tag{8}$$

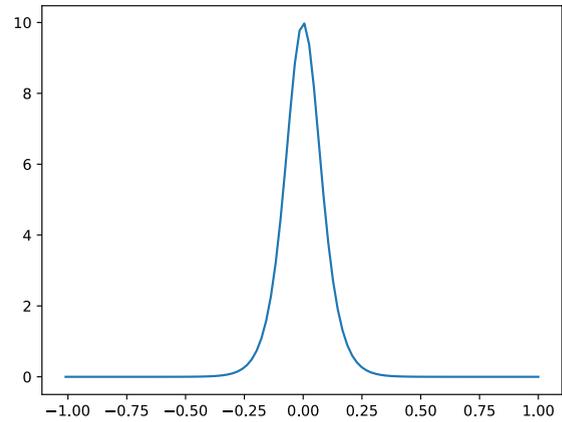
We approximate the derivative of the *signum* operator using the hyperbolic tangent as follows:

$$\frac{d \operatorname{sgn}(w)}{dw} \approx \frac{d \tanh(aw)}{dw} = a \operatorname{sech}^2(aw) \tag{9}$$

where  $\operatorname{sech}(x) = \frac{2}{e^x + e^{-x}}$  is the hyperbolic secant function, and  $a$  is a hyperparameter indicating how sharp the hyperbolic tangent is. The larger the hyperparameter  $a$  is, the closer  $\tanh$  is to the *signum* function. Figure 4 shows the approximate derivative of the *signum* function for  $a = 10$ . As we can see from Fig. 4, the derivative has high magnitude for  $w$  values close to zero, whereas it is effectively zero for large values. This can be seen as allowing small weights to have finer updates than larger weights and thus allowing them to change their sign more often during training. We found empirically that this approximate derivative computation provides satisfactory convergence rates in Google’s Tensorflow software.

**C. DISCGAN (DISCRIMINATOR OF GAN AS CLASSIFIER)**

Generative Adversarial Networks (GAN) have become the benchmark in image synthesis [30], [31]. A typical GAN has a generative network, which attempts to generate images (or data) resembling real images from noise input, and a discriminator network, which attempts to discriminate between the



**FIGURE 4.** The derivative of  $\tanh(aw) = a \operatorname{sech}^2(aw)$  as a function of parameter  $w$ , with  $a$  set to 10.

real images and those synthesized by the generator. The generator and the discriminator are optimized in an adversarial scheme, i.e., the generator tries to fool the discriminator by the synthetic data it produces, and, in turn, the discriminator tries to counteract the generator by discriminating between the real data samples and the fake ones.

In this paper, our aim is not to synthesize realistic data but rather to make use of the adversarial nature of GAN training in order to obtain a discriminator network capable of classifying the input with an unbalanced set of training data. As the recordings of gas leak data may fall short of the clean air recordings for this purpose, we have the generator of the GAN to compensate for the data set with a smaller number of data instances by producing “artificial” gas leak data during training.

In this regard, we perform a two-phase training of the GAN. First, we carry out adversarial training of both the discriminator and the generator using the data of one of the classes. In the second phase, we use data from both classes and carry supervised binary-classification training of the discriminator which now acts as a classifier. In this setting, let  $x^i$  represent the *i*th data instance of one of the classes. In this case,  $x^i$ s denote the gas leak recordings (or the anomalous class). Let  $z$  be a random noise vector, e.g. Gaussian noise or uniform noise. Let  $D$  be the discriminator and  $G$  be the generator, with each having a set of parameters  $\theta_D$  and  $\theta_G$ , respectively. In the adversarial-training phase, we seek to optimize the following loss function:

$$\max_{\theta_D} \min_{\theta_G} \sum_i \log(D(x^i)) + \sum_i \log(1 - D(G(z^i))) \tag{10}$$

where  $D(x^i)$  is the soft prediction result of the discriminator corresponding to data point  $x^i$ . From the discriminator perspective, the prediction output  $D(x^i)$  should be close to 1 because  $x^i$  is “real”. The generator  $G$  produces “fake” data signals from noise vector  $z^i$ , that is  $G(z^i)$ , and the prediction  $D(G(z^i))$  should be close to zero because  $G(z^i)$  is an artificial data instance. The generator, on the other hand, will try to produce  $G(z^i)$  that will be assigned the prediction  $D(G(z^i))$

close to 1. Once training the first stage is accomplished, we move on to the second stage of supervised training of the entire training data, in which the cost function we seek to minimize is the regular binary cross entropy function  $CE$  expressed as follows:

$$CE := -\frac{1}{N} \left( \sum_i (1 - t^i) \log(1 - D(x^i)) + t^i \log(D(x^i)) \right) \quad (11)$$

where  $t^i \in \{0, 1\}$  denotes the true class of  $x^i$ .

When there are multiple classes we can still use the discriminator of a GAN with a slight modification of the loss functions. In this regard, let us assume that there are  $N$ -classes. In this case, the one-hot encoded label for each input is an  $N$ -dimensional vector, with all entries equal to zero, except for the  $k$ th entry, where  $k$  is the true class. During training, the discriminator (or classifier) will minimize the cross entropy of the softmax layer applied at the output layer ( $N$  logits). The generator  $G$  will attack the output of the  $k$ th node. Here we consider the output of the  $k$ th node to be the logit of a binary class, i.e., the adversarial loss criterion becomes:

$$\max_{\theta_D} \min_{\theta_G} \log(D_k(G(z))) \quad (12)$$

where  $D_k(G(z))$  is the discriminator *sigmoidal* response of the  $k$ th node, i.e., we apply the sigmoid function to the logits before taking the logarithm in determining the loss.

Note that the loss here is different from the multi-class case, in which we consider multi-class logits, i.e., we use sigmoid normalization instead of the softmax normalization. In practice, since we do a mini-batch update, we take the average of the loss functions and minimize the loss functions based on the mini-batch gradients.

### III. DATASETS AND EXPERIMENTAL RESULTS

#### A. INFRA-RED VOC DATASET

Our first data set consists of infra-red imaging signals of VOC gas leaks in open air and clean air recordings.<sup>3</sup> Specifically, we have two classes of discrete-time signals corresponding to VOC gas leaks and clean air, respectively. Each signal is a time series containing 50 samples corresponding to two seconds of recording with a sampling rate of 25 samples per second. The recorded value varies in open air because of background temperature variations and low resolution error as it can be observed in Fig. 3. Furthermore, the sensors may not be calibrated in practice, so their sensitivity may differ across time. We gathered about 30,000 VOC gas leak and 30,000 clean air data instances. The images are obtained using an MWIR camera produced by FLIR systems and Infrared Cameras Inc. [17], [18]. VOC gas absorbs the infrared light appearing as a white cloud in the black-hot mode infra-red image as shown in Fig. 2. In these videos, a gas leak erupts from the source with the gas spreading out as

<sup>3</sup>The data is available online at the following link: [https://github.com/Diaa0/ir\\_time\\_series\\_data/blob/master/data\\_60k.npz](https://github.com/Diaa0/ir_time_series_data/blob/master/data_60k.npz)

**TABLE 1. Architecture of the convolutional neural network used in classifying the data set of Sec. III-A.**

Layer	Specification
Input Layer	input size: $32 \times 1$
Conv Layer	$16 \times 3 \times 1$ filters, no strides
Max-pooling Layer	down-sampling by 2
Batch-normalization Layer	-
Conv Layer	$32 \times 3 \times 16$ filters, no strides
Max-pooling Layer	down-sampling by 2
Batch-normalization Layer	-
Conv Layer	$64 \times 3 \times 32$ filters, no strides
Max-pooling Layer	down-sampling by 2
Batch-normalization Layer	-
Global Average-pooling Layer	output size: 64
Dense Layer	output size: 64
Batch-normalization Layer	-
Output Linear Layer	output size: 1

time progresses. We manually selected regions of interest and assigned normal event designations to temporal measurements where no gas is present throughout these series, while designating the rest as anomalous events.

We used min-max normalization in order to scale signal data points between 0 and 1. The normalized signal  $\hat{x}$  is obtained as follows

$$\hat{x}[n] = \frac{x[n] - \min(x)}{\max(x) - \min(x)}, \quad n = 0, 1, \dots, 49 \quad (13)$$

where  $\max(x)$  and  $\min(x)$  represent the maximum and minimum values of a given infrared signal  $x$ , respectively.

We used convolutional neural networks with the architecture specified in Table 1. In order to obtain more temporal data points, and in order to ensure the network is translation-invariant to the gas eruption location, we chose to randomly crop the input data into temporal signals of size 32 each. We divided our data set into three disjoint sets. The training data consists of 8,000 recordings of each class. Another set of 8,000 recordings of each class is used as the validation data set. The rest of the data was reserved for testing. We trained our networks using the RMSProp optimizer algorithm [32]. We tested the hypothesis of whether dropout helps achieve better results [33] by using a dropout rate of 50%. As for the GAN approach, we used a generator which is a multi-layer perceptron (MLP) with one hidden layer of size 256.

The regular convolutional neural network and AddNet exhibit comparable results. We obtained an accuracy of 99.8% for no-gas data and 99.7% for gas-leak data for a regular ConvNet. AddNet attained a recognition rate of 98.9% for no-gas data and 99.3% for gas-leak data.

In the second set of experiments, we assumed that we have an unbalanced data set. In practice, we may not have VOC or ammonia gas leak recordings as clean air. We trained the models with only 50 recordings of gas leak signals against 8,000 recordings of clean air recordings. The test data set contains 14,000 recording instances of VOC gas leaks and

**TABLE 2.** Accuracy results for infra-red VOC data. Classifiers are trained with only 50 VOC gas leak recordings vs 8000 clean air recordings.

Model	No-gas Accuracy (specificity)	Gas-leak Accuracy (sensitivity)	Total Accuracy
ConvNet (dropout 50%)	98.3%	95.8%	97.1%
ConvNet no dropout	98.0%	94.2%	96.1%
AddNet (dropout 50%)	98.2%	96.0%	97.1%
AddNet (no dropout)	99.1%	97.3%	98.2%
DiscGAN	99.0%	97.1%	98.1%

**TABLE 3.** Confusion matrix for the best achieving neural network (AddNet with no dropout) over the testing data. The true positive rate (sensitivity) is 97.3%. The true negative rate (specificity) is 99.1%.

Actual Class	Predicted Class		Total Count
	Leak (positive)	No Leak (negative)	
Leak (positive)	13,622	378	14,000
No Leak (negative)	126	13,874	14,000

clean air recordings. Classification results are also summarized in Table 2. AddNet produces the best results but the discriminator of the GAN Network is also quite close to AddNet. The confusion matrix of the results of the best model is given in Table 3.

We also investigated pruning the weights in both AddNet and ConvNet during inference. In this regard, we discard the magnitudes of the smallest magnitude weights while retaining their sign information. We keep the bias coefficients and the coefficients of the last layer intact. Results of various pruning rates are shown in Table 4. Apparently, in AddNet, we can discard the magnitude information of the weights up to a high rate (67.4%) without severely degrading performance. On the other hand, the magnitude information is quite critical in the case of a regular ConvNet. These results clearly show the advantages of AddNet, which requires reduced memory space in a mobile device and consumes less energy as it performs much fewer arithmetic operations during inference.

### B. GAS SENSOR ARRAY RECORDINGS UNDER DYNAMIC GAS MIXTURES

We consider a gas type identification problem, in which we have three types of gases to identify, namely, CO, Ethylene and Methane. We used the data set obtained by Fonollosa et al. [19]. The data set consists of time-series measurements of a sensor array of 16 chemical metal-oxide sensors under exposure to two different kinds of gas mixtures,

**TABLE 4.** Effect of compressing weights of AddNet and ConvNet by discarding the smallest  $K\%$  magnitude while keeping the sign information. ConvNet fails to produce reasonable results when the compression rate exceeds 16.1%. The compression rate is estimated by allocating 32 bits to intact weight values and 1 bit for every binarized weight factor.

Model Accuracy	Weight Compression Rate (smallest $K\%$ )					
	0	16.1	19.7	67.4	76.8	86.6
AddNet	98.9	97.2	97.9	98.0	97.1	61.3
ConvNet	99.8	67.4	—	—	—	—

ethylene and methane in air, and ethylene and CO in air. Sensors were exposed to volatile organic compounds at different concentration levels under tightly-controlled operating conditions during the experiment. The data is obtained at a sampling frequency of 100 Hz. The 16 chemical sensors are of four different types, with each having four identical sensors.<sup>4</sup> Furthermore, switching between different mixtures of VOCs may occur too fast making it challenging if not impossible for the sensors to reach steady state. This makes identifying the gas analytes difficult using a machine learning method.

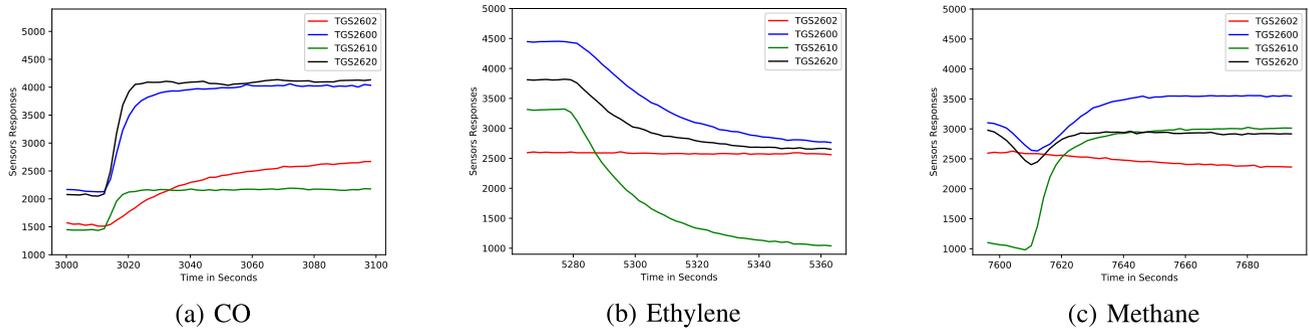
The recorded sensor data is deposited to the UC Irvine Machine Learning Repository online in the form of two long time-series. We extracted portions of the time series such that the sensor array is exposed to one type of analyte at a given time. Each recording corresponds to 100 seconds of data. We observed that it is enough to sample the sensor response every 2 seconds.

Example sensor response signals to CO, ethylene and methane gas vapor exposures are shown in Fig. 5. Each sub-figure contains four different sensor responses.

We gathered a total of 215 instances from the raw recordings, in which we have 49 CO, 116 ethylene and 50 methane time-series signals. Each instance has 50 time measurements for each sensor. Thus, a total of  $50 \times 16$  measurements per instance are used. Since the number of instances in the data set is small, we employed cross validation with holdout method, where our validation set consists of 35 examples, with the experiment repeated 4 times. Thus, we validated our results over 140 examples. Furthermore, since the number of instances is small compared to the input dimension  $50 \times 16$ , we opted to randomly crop data points during training into smaller time series of size  $40 \times 16$ . This allows the classifier to be invariant to the exact time where the exposure takes place. Furthermore, it increases the number of data points during training.

Since the sensors are of different types, and since even the sensors of the same type produce different temporal responses, we process the temporal sensor data using 1-D convolutional networks. The input to each neural network is a matrix of size  $40 \times 16$ , for 40 time instances and 16 sensors.

<sup>4</sup>For more details, the reader may refer to the original paper [19].



**FIGURE 5.** Time-series data generated by four different sensors under exposure to different type of gases (50 time samples for each sensor).

**TABLE 5.** Architecture of the convolutional neural network used in classifying the data set of Sec. III-B.

Layer	Specification
Input Layer	input size: $40 \times 16$
Conv Layer	$64 \times 5 \times 16$ filters, no strides
Max-pooling Layer	pooling size: 4, stride size: 3
Batch-normalization Layer	-
Conv Layer	$128 \times 5 \times 64$ filters, no strides
Max-pooling Layer	pooling size: 4, stride size: 3
Batch-normalization Layer	-
Conv Layer	$256 \times 5 \times 128$ filters, no strides
Max-pooling Layer	pooling size: 4, stride size: 3
Batch-normalization Layer	-
Global Average-pooling Layer	output size: 256
Dense Layer	output size: 256
Batch-normalization Layer	-
Output Linear Layer	output size: 3

**TABLE 6.** Recognition rates for the two neural networks over the test data set.

	Gas Type-based Accuracy			Average Accuracy
	CO	Ethylene	Methane	
ConvNet	91.1%	98.6%	100%	96.6%
AddNet	91.1%	97.2%	100%	96.1%

We used ReLU non-linearity between layers. Our loss function is the cross-entropy with the softmax operator. We used the RMSProp optimizer to carry out the parameter updates during training. We trained a regular ConvNet and an AddNet of the same architecture as in Table 5. Our classification accuracy results over the testing data are shown in Table 6. The confusion matrix of the results over the validation data set obtained by AddNet is given in Table 7.

It can be observed in Table 6 that the recognition capabilities of both AddNet and ConvNet are on par with one another. It is worth emphasizing the computational frugality of the scheme as use of the regular dot-product is confined solely to the last layer in AddNet.

**TABLE 7.** Confusion matrix for AddNet over the validation data sets for repeated trials.

True Class	Predicted Class			Total Count
	CO	Ethylene	Methane	
CO	31	3	0	34
Ethylene	2	70	0	72
Methane	0	0	34	34

### C. CHEMICAL GAS SENSOR ARRAY DRIFT DATASET

The third data set is the publicly available chemical VOC gas sensor drift data set compiled by Vergara *et al.* at UCSD [13]. The data set was obtained by exposing an array of 16 distinct chemical sensors to 6 types of gas mixtures (ammonia, acetone, ethylene, ethanol, toluene and acetaldehyde) at a variety of concentration levels. Each data record is a vector time series. Vectors contain 8 feature parameters extracted from the sensor time series signals during a gas release experiment, conducted over a period of three years at UCSD. The feature parameters include the steady state resistance value and the normalized resistance change. The remaining 6 parameter features are the maxima and minima of the exponential moving average ( $\text{ema}_\alpha$ ) transform governed by the following input-output relation:

$$y[k] = (1 - \alpha)y[k - 1] + \alpha(r[k] - r[k - 1]) \quad (14)$$

where  $r[k]$  is the resistance value at time step  $k$ , and  $y[k]$  is the transformed value after applying the ema filter. The maxima and minima features are reported for  $\alpha$  values equal to 0.1, 0.01 and 0.001 over an entire experiment. These ema features have distinct time constants for different  $\alpha$  values, as they contain temporal information.

Unfortunately, the raw time-domain sensor signals are not available in this data set.

Since there are 16 sensors, a total of  $16 \times 8 = 128$  feature values are recorded per experiment.

The data set is divided into 10 batches ordered chronologically. Full details about the experiment and the data set can be found in [13].

We carried out our classification tasks by training neural networks for  $N = 2$  batches and testing on successive

**TABLE 8.** Comparative accuracy (in %) results of the various models when training on batches 1 and 2 and testing on batches 3-10. Bold-text numbers correspond to the best accuracy results obtained amongst the different algorithms for each batch.

Batch ID	SVM Classifier Ensemble [13]	MLP	AddNet-MLP	DiscGAN	AddNet-DiscGAN	Domain adaptation
Batch 3	87.8	<b>98.6</b>	<b>98.6</b>	98.3	97.8	98.7
Batch 4	<b>90.6</b>	83.8	75.1	71.4	69.6	73.3
Batch 5	72.1	<b>99.5</b>	99.4	98.4	98.9	99.5
Batch 6	44.5	74.9	<b>75.9</b>	72.3	73.9	76.4
Batch 7	42.5	59.8	57.4	61.5	<b>66.3</b>	59.2
Batch 8	29.9	34.0	34.0	<b>62.3</b>	58.8	39.1
Batch 9	59.8	31.6	38.9	63.2	<b>63.8</b>	52.3
Batch 10	39.7	47.3	<b>54.3</b>	43.8	44.5	46.1

batches. This is identical to the sensor drift estimation approach given in [13].

Because feature values have huge variances, we opted to apply the signed square root function element-wise to control the ranges of the reported values. The modification delivered improved results in our experiments, especially for later batches.

We trained an MLP model with two hidden layers, each with 512 output units, and an output layer. Furthermore, we trained the network for 100 epochs using the RMSProp optimizer [32]. We applied a dropout rate of 20% and used a batch size of 128 in order to prevent complex co-adaptation. To augment the data, we added a zero-mean Gaussian noise with standard deviation of 0.1.

We also tried combining AddNet with the GAN approach, in which case, the discriminator is an AddNet and the generator is a regular MLP. The architecture of the network is the same as that of the GAN we use. Furthermore, we tried utilizing the other batches by passing them to the classifier and carrying out backpropagation according to their guessed labels. This is done in order for the network to utilize the correctly guessed labels so that it could be helpful in improving the classification accuracy for the mis-classified data point. A numerical comparison of the proposed methods to the SVM-classifier ensemble used in [13] is given in Table 8. In general, the AddNet-MLP, the MLP and the multi-class GAN discriminator produce better sensor-drift compensated results than does the SVM based method.

As we can see from Table 8, using DiscGAN (with a regular discriminator or AddNet), we were able to obtain better recognition rates for later batches (batches 7, 8 and 9). This could be attributed to the fact that the generator did expose the discriminator to novel unseen points in the data space during training. Therefore, the discriminator would have been able to learn additional meaningful features. As for AddNet, it can perform as well as the regular MLP, either in conventional binary classification or in the case of DiscGAN. It is also worth noting that the domain adaptation scheme we employed did not yield any significant improvements. We believe that improved classification results would have been

attained, if the entire temporal sensor signal set were at our disposal as input to our algorithms.

#### IV. CONCLUSION

In this paper, we have introduced a variety of deep-learning based algorithms and applied them to VOC gas and ammonia vapor leak detection and gas type identification problems. This is the first work that uses the raw time-series signal for VOC detection and ammonia vapor sensing using deep learning. The first algorithm is based on AddNet. In AddNet, we replace the computationally expensive dot-product operations in deep neural networks with a modified addition operation that retains the sign of multiplication. Its computational efficiency enables AddNet to be used in embedded and mobile systems, in which we envision a smart gas leakage monitoring and detection CPS being reliably used.

The second algorithm is called DiscGAN, which uses the discriminator of a generative adversarial neural network as a classifier in a bid to enhance the recognition capabilities of the system. The generator part helps in exposing the discriminator to realistic synthetic data points that can be helpful in classification tasks.

We considered three detection and classification tasks. The first task is to detect VOC gas leakage from temporal IR data. Our proposed algorithms achieved accuracy rates of 97 – 98%. The second task we considered is to identify gas types using temporal data of sensor arrays. We were able to attain recognition rates of 96.1 – 96.5%. Our third task was to identify gas types using non-temporal data, where the readings are obtained for the same sensor array over a period of 36 months. The sensor measurements suffer from degradation due to sensor drift.

Although our gas identification accuracy results for the early batches in the last data set were quite high, the degradation incurred in later batches resulted in significant identification accuracy drop. We believe that the non-temporal global features reported for the experiments are highly affected by sensor drift. As a result, the features are not sufficiently expressive of the sensor responses for different gas analyte types. Based on our high recognition rates for two temporal

data sets considered in this work, we conclude that using sensor measurements in their temporal presentation, and feeding these recordings into deep neural network algorithms, achieves better performance as these algorithms learn discriminative features by themselves with no need to hand-craft features that could be sensitive to error as in the case of the sensor drift problem.

## REFERENCES

- [1] A. P. Jones, "Indoor air quality and health," *Atmos. Environ.*, vol. 33, no. 28, pp. 4535–4564, 1999.
- [2] S. K. Brown, M. R. Sim, M. J. Abramson, and C. N. Gray, "Concentrations of volatile organic compounds in indoor air—A review," *Indoor Air*, vol. 4, no. 2, pp. 123–134, 1994.
- [3] C. Maltoni, A. Ciliberti, G. Cotti, B. Conti, and F. Belpoggi, "Benzene, an experimental multipotential carcinogen: Results of the long-term bioassays performed at the Bologna Institute of Oncology," *Environ. Health Perspect.*, vol. 82, pp. 109–124, Jul. 1989.
- [4] J. Kang, J. Liu, and J. Pei, "The indoor volatile organic compound (VOC) characteristics and source identification in a new university campus in Tianjin, China," *J. Air Waste Manage. Assoc.*, vol. 67, no. 6, pp. 725–737, 2017.
- [5] D. Badawi, S. Ozev, J. B. Christen, C. Yang, A. Orailoglu, and A. E. Çetin, "Detecting gas vapor leaks through uncalibrated sensor based CPS," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 8296–8300.
- [6] J. Wu, S. Guo, H. Huang, W. Liu, and Y. Xiang, "Information and communications technologies for sustainable development goals: State-of-the-art, needs and perspectives," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 3, pp. 2389–2406, 3rd Quart., 2018.
- [7] R. Atat, L. Liu, J. Wu, G. Li, C. Ye, and Y. Yang, "Big data meet cyber-physical systems: A panoramic survey," *IEEE Access*, vol. 6, pp. 73603–73636, 2018.
- [8] F. Erden, E. B. Soyer, B. U. Toreyin, and A. E. Çetin, "VOC gas leak detection using pyro-electric infrared sensors," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2010, pp. 1682–1685.
- [9] A. E. Cetin and B. U. Toreyin, "Method, device and system for determining the presence of volatile organic compounds (VOC) in video," U.S. Patent 8432451 B2, Apr. 30 2013.
- [10] W. Göpel and K.-D. Schierbaum, "Definitions and typical examples," in *Sensors: Chemical and Biochemical Sensors-Part I*, vol. 2, 1991, pp. 1–27.
- [11] F. Davide, C. Di Natale, M. Holmberg, and F. Winquist, "Frequency analysis of drift in chemical sensors," in *Proc. 1st Italian Conf. Sensors Microsyst.*, Rome, Italy, 1996, pp. 150–154.
- [12] M. Zuppa, C. Distanto, P. Siciliano, and K. C. Persaud, "Drift counteraction with multiple self-organising maps for an electronic nose," *Sens. Actuators B, Chem.*, vol. 98, nos. 2–3, pp. 305–317, Mar. 2004.
- [13] A. Vergara, S. Vembu, T. Ayhan, M. A. Ryan, M. L. Homer, and R. Huerta, "Chemical gas sensor drift compensation using classifier ensembles," *Sens. Actuators B, Chem.*, vol. 166, pp. 320–329, May 2012.
- [14] T. Artursson, T. Eklöv, I. Lundström, P. Mårtensson, M. Sjöström, and M. Holmberg, "Drift correction for gas sensors using multivariate methods," *J. Chemometrics*, vol. 14, nos. 5–6, pp. 711–723, 2000.
- [15] D. E. Williams, G. S. Henshaw, and K. F. Pratt, "Detection of sensor poisoning using self-diagnostic gas sensors," *J. Chem. Soc., Faraday Trans.*, vol. 91, no. 18, pp. 3307–3308, 1995.
- [16] Ammonia. *National Institute of Standards and Technology*. Accessed: Jun. 16, 2019. [Online]. Available: <https://webbook.nist.gov/cgi/cbook.cgi?ID=C7664417&Type=IR-SPEC&Index=1>
- [17] *FLIR Systems Australia*. Accessed: Jun. 16 2019. [Online]. Available: <http://www.ferret.com.au/latest-videos>
- [18] *Infrared Cameras*. Accessed: Jun. 6, 2019. [Online]. Available: <https://infraredcameras.com/thermal-infrared-media/videos/>
- [19] J. Fonollosa, S. Sheik, R. Huerta, and S. Marco, "Reservoir computing compensates slow response of chemosensor arrays exposed to fast varying gas concentrations in continuous monitoring," *Sens. Actuators B, Chem.*, vol. 215, pp. 618–629, Aug. 2015.
- [20] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [21] Y. LeCun and Y. Bengio, "Convolutional networks for images, speech, and time series," in *The Handbook of Brain Theory and Neural Networks*, vol. 3361, 1995, no. 10.
- [22] M. Längkvist, L. Karlsson, and A. Loutfi, "A review of unsupervised feature learning and deep learning for time-series modeling," *Pattern Recognit. Lett.*, vol. 42, pp. 11–24, Jun. 2014.
- [23] M. Courbariaux, Y. Bengio, and J.-P. David, "Binaryconnect: Training deep neural networks with binary weights during propagations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 3123–3131.
- [24] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "XNOR-Net: ImageNet classification using binary convolutional neural networks," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2016, pp. 525–542.
- [25] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1," Feb. 2016, *arXiv:1602.02830*. [Online]. Available: <https://arxiv.org/abs/1602.02830>
- [26] D. Badawi, E. Akhan, M. Mallah, A. Üner, R. Çetin-Atalay, and A. E. Çetin, "Multiplication free neural network for cancer stem cell detection in H-and-E stained liver images," *Proc. SPIE*, vol. 10211, May 2017, Art. no. 102110C.
- [27] A. Afrasiyabi, D. Badawi, B. Nasir, O. Yıldız, F. T. Y. Vural, and A. E. Çetin, "Non-euclidean vector product for neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Apr. 2018, pp. 6862–6866.
- [28] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," Feb. 2015, *arXiv:1502.03167*. [Online]. Available: <https://arxiv.org/abs/1502.03167>
- [29] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Math. Control, Signals Syst.*, vol. 2, no. 4, pp. 303–314, 1989.
- [30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [31] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," Nov. 2015, *arXiv:1511.06434*. [Online]. Available: <https://arxiv.org/abs/1511.06434>
- [32] T. Tieleman and G. Hinton, "Lecture 6.5-RMSPROP: Divide the gradient by a running average of its recent magnitude," *COURSERA, Neural Netw. Mach. Learn.*, vol. 4, no. 2, pp. 26–31, 2012.
- [33] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.

**DIAA BADAWI** received the M.Sc. degree in electrical and electronics engineering from Bilkent University, Ankara, Turkey, in 2016. He is currently pursuing the Ph.D. degree in electrical engineering with The University of Illinois at Chicago. His research interests include intelligent monitoring systems, energy-efficient implementations of neural networks, and machine learning algorithms.

**TUBA AYHAN** received the Ph.D. degree in electrical engineering from KU Leuven, Belgium, in 2016. She is currently an Assistant Professor in electrical engineering with MEF University. She also continues the research with the ECC Laboratory, Istanbul Technical University. Her research interests include embedded system design and power-efficient digital circuit implementation.

**SULE OZEV** received the Ph.D. degree in computer science and engineering from the University of California, San Diego, in 2002. She is currently an Associate Professor with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe. Her research interests include self-test and self-calibration for wireless transceivers, analysis and mitigation of process variations for mixed signal and digital circuits, fault-tolerant and reconfigurable heterogeneous systems, and mixed signal circuit testing.

**CHENGMO YANG** received the B.S. degree from Peking University, China, and the M.S. and Ph.D. degrees from the University of California at San Diego, USA. She is currently an Associate Professor with the Department of Electrical and Computer Engineering, University of Delaware. She has authored over 70 articles in peer-reviewed international journals and conferences. Her research interests include broad areas of computer architecture, embedded systems, and design automation, with a particular focus on improving security, reliability, non-volatility, and energy-efficiency of embedded systems, cyber physical systems (CPS), and the Internet-of-Things (IoT). She has served as a member for program/organization committees in many prestigious international conferences and workshops. She has received the National Science Foundation Career Award, in 2013. She received two best paper awards and four best paper nominations. She is serving as an Associate Editor for the IEEE TRANSACTIONS ON COMPUTER AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS (TCAD) and an Associate Editor of *ACM Transactions on Design Automation of Electronic Systems* (TODAES).

**ALEX ORAILOGLU** received the Ph.D. degree in computer science from the University of Illinois, Urbana, in 1983. He has been a Professor with the Computer Science and Engineering Department, University of California, since 1987. He leads the Reliable System Synthesis Group. He has authored or coauthored over 80 publications. His research interests include electronic design automation, VLSI test, and synthesis of fault-tolerant ICs.

**AHMET ENIS ÇETIN** received the B.Sc. degree and the M.S.E. and Ph.D. degrees in systems engineering from the Moore School of Electrical Engineering, University of Pennsylvania. He studied electrical engineering with the Orta Dogu Teknik Universitesi (METU). From 1987 to 1989, he was an Assistant Professor in electrical engineering with the University of Toronto, Canada. He is on leave from Bilkent University, Ankara, Turkey. He was a Visiting Professor with the UCSD, from 2016 to 2017. He is currently a Research Professor with The University of Illinois at Chicago. He is one of the founders of GrandEye that produces omnidirectional cameras.

• • •