Device Variation Effects on Neural Network Inference Accuracy in Analog In-Memory Computing Systems

Qiwen Wang, Yongmo Park, and Wei D. Lu*

In analog in-memory computing systems based on nonvolatile memories such as resistive random-access memory (RRAM), neural network models are often trained offline and then the weights are programmed onto memory devices as conductance values. The programmed weight values inevitably deviate from the target values during the programming process. This effect can be pronounced for emerging memories such as RRAM, PcRAM, and MRAM due to the stochastic nature during programming. Unlike noise, these weight deviations do not change during inference. The performance of neural network models is investigated against this programming variation under realistic system limitations, including limited device on/off ratios, memory array size, analog-to-digital converter (ADC) characteristics, and signed weight representations. Approaches to mitigate such device and circuit nonidealities through architecture-aware training are also evaluated. The effectiveness of variation injection during training to improve the inference robustness, as well as the effects of different neural network training parameters such as learning rate schedule, will be discussed.

1. Introduction

Deep neural networks (DNNs) have achieved unprecedented capabilities in tasks such as image and voice analysis and recognition and have been widely adopted. However, computation requirements and the associated energy consumption of neural network implementations have been growing rapidly.[1] In addition, traditional computing architectures are ineffective for DNN workloads due to the high memory access demands, making it even more challenging to meet these computational requirements. Many systems based on digital CMOS technology have been developed specifically for accelerating DNN workloads, including GPU, FPGA, and more specialized accelerators like DPUs. While these systems have shown significant improvements over traditional CPUs in both computing power and

Q. Wang, Y. Park, W. D. Lu
Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48109, USA
E-mail: wluee@umich.edu

The ORCID identification number(s) for the author(s) of this article can be found under https://doi.org/10.1002/aisy.202100199.

energy efficiency, continued innovation is necessary to meet the growing demand. Particularly, DNN inference workload on edge computing platforms like mobile and internet of things (IoT) has stringent energy efficiency requirements due to limited energy supply, and unconventional approaches like analog computing may prove more advantageous in meeting this requirement.

The most important limiting factor for DNN computing is the transfer of data between processors and off-chip memories due to the limited density of existing on-chip memory technology. In-memory computing (IMC) systems, utilizing the density advantage of emerging memory technologies like RRAM, can potentially store entire DNN models on-chip, thus eliminating off-chip memory access. Analog IMC systems that utilize the device conductance to directly perform vector–matrix multiplication (VMM) operations further allow device-level parallelism that leads to higher performance.[2,3] Meanwhile, neural networks are known for their fault tolerance, making it a feasible workload for analog computing, which is generally unsuitable for traditional arithmetic operations due to its inherently lower precision. Thus, analog IMC systems promise drastic improvement in performance and energy efficiency for DNN applications and have gained much popularity in recent years.[4–9] However, nonidealities in memory devices and peripheral circuits can still cause significant degradation of neural network inference accuracy. In general, for analog computing systems, inference accuracy needs to be ensured before any benefit in energy efficiency can become material.

2. Tiled Analog IMC System and Architecture-Aware Training

2.1. The Necessity of the Tiled Architecture

There are three types of important nonidealities in analog IMC systems for VMM operations, interconnect parasitics, analog-to-digital converter (ADC) limitations, and memory device nonidealities. Because energy efficiency is the most important target, ADC operating frequency is likely to be limited to below ≈100 MHz.[10] At this speed, with more than 10 ns of hold time between input change and ADC sampling, in conjunction with the limited array size, transient effects are generally negligible in

memory arrays. Therefore, line resistance is the primary parasitic effect. In a large array, the effect of line resistance is dependent on data pattern from all cells in the array and the input signals, and thus can only be compensated by performing expensive calculations based on the memory states and input signals, which defeats the purpose of IMC.[11] To address this issue, the array size must be limited to avoid the effect of line resistance, and large-scale neural networks have to be mapped onto multiple arrays.[12] This is one of the reasons for the tiled architecture analyzed in this study. In our proposed system, the array size is 256×64, and assuming line resistance of $\approx 2\,\Omega$ per cell, device LRS resistance of $33\,k\Omega$, line resistance will have a negligible effect on the array operation.

2.2. Cell Defects

In this study, we did not consider stuck at fault cell defects. The defects have been considered by many prior studies, including some of our works.[13] Generally, a small portion of stuck at open devices will not have meaningful impacts. However, a shorted device would saturate the output of an entire column. Columns with shorted cells have to be disabled and replaced with spare ones to deal with shorted devices. Although replacing columns means extra areas are needed for spare columns, once the defective ones are replaced, they will not have an impact on the inference accuracy.

2.3. Error Caused by the Tiled Architecture

Although the tiled architecture avoids the line resistance effects, the additional computing error caused by this implementation needs to be analyzed, including the effects of ADC limitations and device nonidealities. Several prior studies have been published that discuss approaches to implement large DNN models on practical RRAM arrays using a tiled architecture, as shown in **Figure 1**.[12,14,15] The effects of limited RRAM array size, ADC precision, signed weights representation in two RRAM cells, and RRAM cell quantization effects in analog IMC systems have been studied. In such a system, neural network models are trained off-line and programmed onto memory arrays, and large
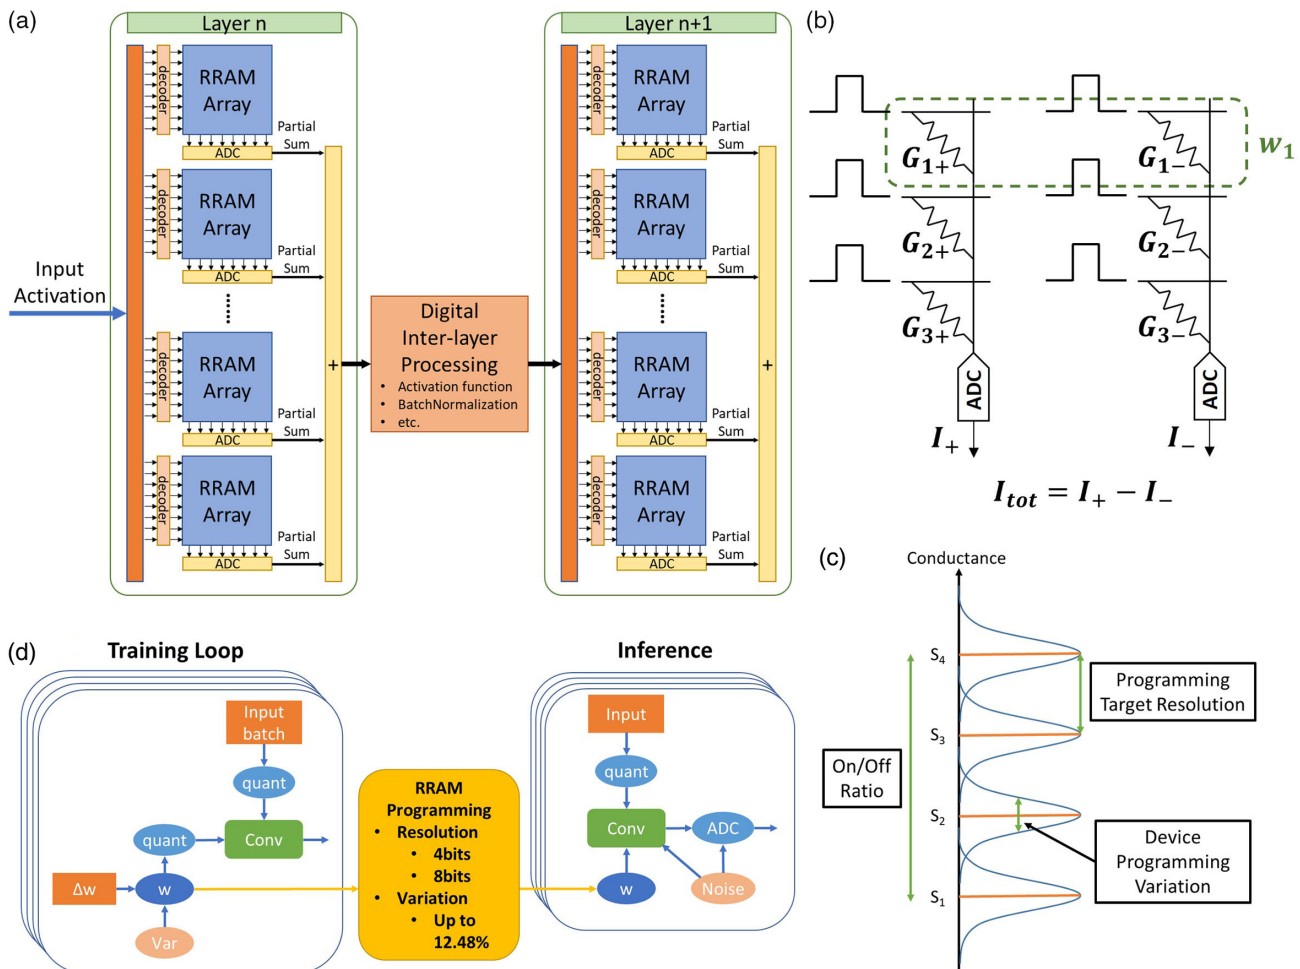
Figure 1. a) Tiled analog IMC systems. Large DNN layers are mapped onto multiple memory arrays. Analog outputs of each array are digitized by ADCs to produce partial sums. The partial sums are then summed in the digital domain to produce the final layer output. b) Signed weights are represented on two memory cells in two different columns. c) Device characteristics considered in this study. d) Neural network models are trained off-line and then programmed onto memory arrays for inference.

2100199 (2 of 12)

neural network layers are mapped onto multiple memory arrays where partial sums (Psums) are produced by ADCs at each array and summed in digital domain[14] (Figure 1a). Signed weights are represented in two cells on two different columns that receive the same input activations. Currents from the two columns are quantized by ADCs individually, and then the digital output of the negative column is subtracted from that of the positive column (Figure 1b).

2.4. Architecture-Aware Training

In general, many of the device and circuit nonideality effects can be effectively mitigated through architecture-aware training methods,[14] where hardware details are mimicked in the training process. In architecture-aware training, we developed a simulator based on Google's TensorFlow deep learning framework by modifying the training graph from the standard floating-point pipeline. To compare the impact of different hardware nonidealities, we consider three inference pipelines, level 1 through level 3, and their corresponding training topologies (**Figure 2**). In level 1, only the quantization of weights and activations is considered. In level 2, the effects of signed weights representation on two cells and limited device on/off ratios are introduced. For both training and inference in level 1 and level 2, we used the common scheme for quantization-aware training,[16] where the weights pass through the fakequantization function before calculations are conducted. The fakequantization function does not change the overall range of the weights and instead rounds the weights to a number of fixed values determined by the range and resolution set for the function, and these parameters can be different for each layer. For actual hardware representation of weights where the conductance range is fixed for the whole system, the outputs of each layer need to be multiplied by a high precision scalar to match that of the software model. In level 3, the physical range of memory cells, the multiplier, limited memory array size, and ADC precision limitation are introduced. As described in Figure 2, separate multipliers are assigned to each array and trained during the training process.

By sequentially introducing different levels of architecture details during the training process, the neural network model can potentially account for these architecture and device factors and recover the desired model accuracy.[14] However, high levels of device programming variation, which is indicative of today's analog memory devices, still present challenges in considerable inference accuracy degradation.

3. System Setup

3.1. Networks Used for Benchmarking

In this study, we chose three neural network and dataset combinations of various complexities to investigate the impact of analog IMC accuracy at realistic device nonidealities for different network and dataset complexity (**Table 1**). The first network is a relatively simple VGG-block-based model trained for the CIFAR-10 dataset.[17] This model contains only convolution (Conv) layers, a fully connected (FC) layer, and MaxPool layers. The second network is the Wide ResNet 16-8 model (WRN).[19]

This network uses residual connections and batch normalization in addition to convolution and fully connected layers. We used the WRN 16-8 network for the CIFAR-10 dataset and the more complex CIFAR-100 dataset to test the effects on more challenging tasks.
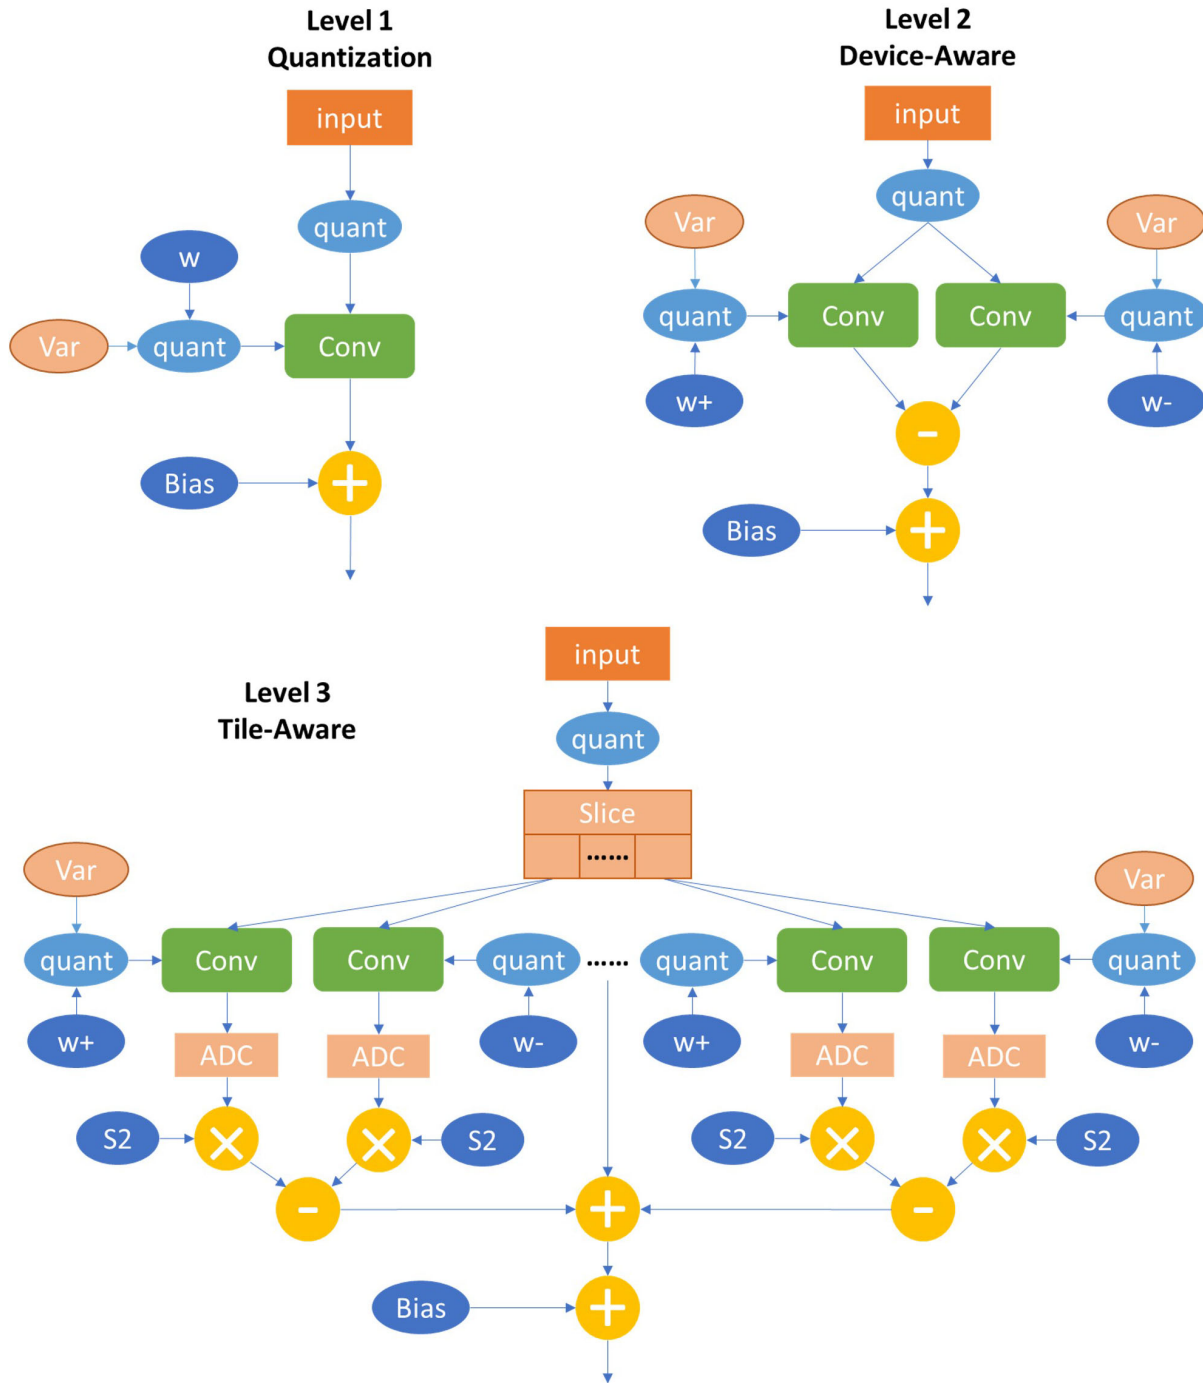
3.2. Hardware Characteristics

We used 8bit ADC in our study because it has been found to offer a good balance between energy efficiency and resolution, as reducing resolution further does not appear to yield a meaningful improvement in energy/sample.[18] RRAM cells with an analog read current range of 0.3–3 µA, ADC input range of 0–45 µA, and array size of 265×64 were considered for the tiled implementation.

3.3. Training Process

We first obtain floating-point models using standard practice. For the VGG-block-based[17] models, we trained for 150 epochs using the stochastic gradient descent (SGD) optimizer with a learning rate of 0.001, momentum of 0.9. For WRN models, we follow the parameters described in Zagoruyko and Komodakis.[19] Then, different levels of hardware details are progressively introduced during training, as schematically shown in **Figure 3**, along with the parameters used during the training processes. Specifically, level 1 models are fine-tuned from the floating-point models, level 2 models are fine-tuned level 1 models, and level 3 models are fine-tuned from level 2 models. We found this approach leads to better model inference accuracy compared with training directly the level 2 or level 3 models from scratch with random weights.[14] In fact, we found that level 3 MNIST models trained from random weights reached only 77.78% accuracy (compared to 99.13% for model fine-tuned from level 2 and float model) in previous studies, and level 3 VGG and WRN models produced accuracies of only 10%, which is no more than chance for the CIFAR-10 dataset. In the fine-tuning process, we used a learning rate of 0.001 for the VGG-block-based model and trained for 20 epochs. For the WRN models, we used a learning rate schedule, where the learning rate starts at 0.003, then steps down to 0.001, 0.0005, 0.0002 after 5, 10, 15 epochs and trained for a total of 40 epochs.

4. Effects of Computation Errors in Analog IMC Systems

First, we present the effects of deterministic errors including weight and activation quantization, signed weight representation, limited RRAM array size, ADC precision limitations, and RRAM cell on/off ratios (**Figure 4**). For the three network–dataset combinations we studied, when only quantization (activation and weights quantized to 8bits) and signed weight representation were considered during inference (level 2), there is minimal accuracy drop from just using the quantization-aware trained models[16] (level 1). We do note that the activation quantization range in the inference pipelines must correspond to the input range of activation function used during training (ReLu6,

**ADVANCED
SCIENCE NEWS**
www.advancedsciencenews.com

**ADVANCED
INTELLIGENT
SYSTEMS**
Open Access
www.advintellsyst.com

Figure 2. Architecture-aware training topology. We propose that architecture aware training can be considered in three levels. Level 1 is the standard quantization-aware training method,[16] where high precision weights are passed through a fake-quantization function before computation. At level 2 device-aware training, signed weight representation in memory cells and limited on/off ratio are considered. At level 3 tile-aware training, the limited memory array size and ADC precision limitation are also considered. In addition, for all 3 levels, variation can be injected on per minibatch basis to mimic the effect of programming variation.

etc.), or there is severe degradation in accuracy due to the limited range due to the quantization effects.

However, in the presence of a low device on/off ratio and/or array size and ADC limitations, the quantization-aware trained models cannot produce acceptable accuracies. By introducing the finite on/off ratio properties in the training pipeline, device-aware trained models can successfully mitigate the effect of limited on/off ratio down to 10, along with any effects due to

2100199 (4 of 12)

ADVANCED
SCIENCE NEWS
www.advancedsciencenews.com

ADVANCED
INTELLIGENT
SYSTEMS
Open Access
www.advintellsyst.com

Table 1. Models used for benchmarking. Only CNN and fully connected layers are shown. RRAM array size of 256 × 64 is used.

	Input size	Filter shape	Number of columns	Row vector length	Number of arrays
CIFAR-10 VGG Block					
CNN 1	32 × 32 × 3	3 × 3 × 3 × 32	64	27	1 × 1
CNN 2	32 × 32 × 32	3 × 3 × 32 × 32	64	288	2 × 1
CNN 3	16 × 16 × 32	3 × 3 × 32 × 64	128	288	2 × 2
CNN 4	16 × 16 × 64	3 × 3 × 64 × 64	128	576	3 × 2
CNN 5	8 × 8 × 64	3 × 3 × 64 × 128	256	576	3 × 4
CNN 6	8 × 8 × 128	3 × 3 × 128 × 128	256	1152	5 × 4
FC 1	2048	2048 × 128	256	2048	8 × 4
FC 2	128	128 × 10	10	128	1 × 1
Total					78 arrays
WRN-16-8					
G1_Conv	32 × 32 × 3	3 × 3 × 3 × 16	32	27	1
G2_MB1_Conv1	32 × 32 × 16	3 × 3 × 16 × 128	256	144	1 × 4
G2_MB1_Conv2	32 × 32 × 128	3 × 3 × 128 × 128	256	1152	5 × 4
G2_Res_Conv	32 × 32 × 16	1 × 1 × 16 × 128	256	16	1 × 4
G2_MB2_Conv1	32 × 32 × 128	3 × 3 × 128 × 128	256	1152	5 × 4
G2_MB2_Conv2	32 × 32 × 128	3 × 3 × 128 × 128	256	1152	5 × 4
G3_MB1_Conv1	32 × 32 × 128	3 × 3 × 128 × 256	512	1152	5 × 8
G3_MB1_Conv2	32 × 32 × 256	3 × 3 × 256 × 256	512	2304	10 × 8
G3_Res_Conv	32 × 32 × 128	1 × 1 × 128 × 256	512	128	1 × 8
G3_MB2_Conv1	16 × 16 × 256	3 × 3 × 256 × 256	512	2304	10 × 8
G3_MB2_Conv2	16 × 16 × 256	3 × 3 × 256 × 256	512	2304	10 × 8
G4_MB1_Conv1	16 × 16 × 256	3 × 3 × 256 × 512	1024	2304	10 × 16
G4_MB1_Conv2	8 × 8 × 512	3 × 3 × 512 × 512	1024	4608	19 × 16
G4_Res_Conv	16 × 16 × 256	1 × 1 × 256 × 512	1024	256	1 × 16
G4_MB2_Conv1	8 × 8 × 512	3 × 3 × 512 × 512	1024	4608	19x16
G4_MB2_Conv2	8 × 8 × 512	3 × 3 × 512 × 512	1024	4608	19 × 16
FC	512	512 × 10	20	512	2 × 1
Total					1447

the two-column signed weight representation, as shown in Figure 4. On the other hand, in the presence of array size and ADC limitations, the device-aware training, i.e., level 2 training pipeline, results in poor accuracy for the more complex models or datasets such as WRN. Acceptable results may be produced by level 2 training for simpler models such as VGG-blocks due to the use of only Conv and FC layers which are generally more resilient to errors. As a result, tile-aware training (i.e., level 3 pipeline) must be used for the more complex models or datasets to produce good accuracy, as shown in Figure 4. We believe the more complicated model structure with residue connections and the use of batch normalization layers make the WRN models more sensitive to errors. Particularly, models with batch normalization layers are sensitive to changes in activation distribution, and the quantization of partial sums due to ADC precision and range limitations produce a shift in activation distributions.[20]

5. Programming Variation Effects on Inference Accuracy

Next, we examine the effects of device variations on network inference accuracy. Neural network models are trained off-line and then programmed onto memory arrays for inference, and the weights do not change during the inference process. Combined with analog computation, this means any deviations that occur during the device programming process result in inference to be conducted on models that are effectively different from the trained models, leading to potential accuracy degradation. Different from deterministic errors discussed earlier, the randomness of device variations means each programmed chip maps an essentially different model. Retraining each chip individually may potentially recover the accuracy, but will be very expensive and impractical. In the following, we investigate the impact of device programming variation on large-scale DNN
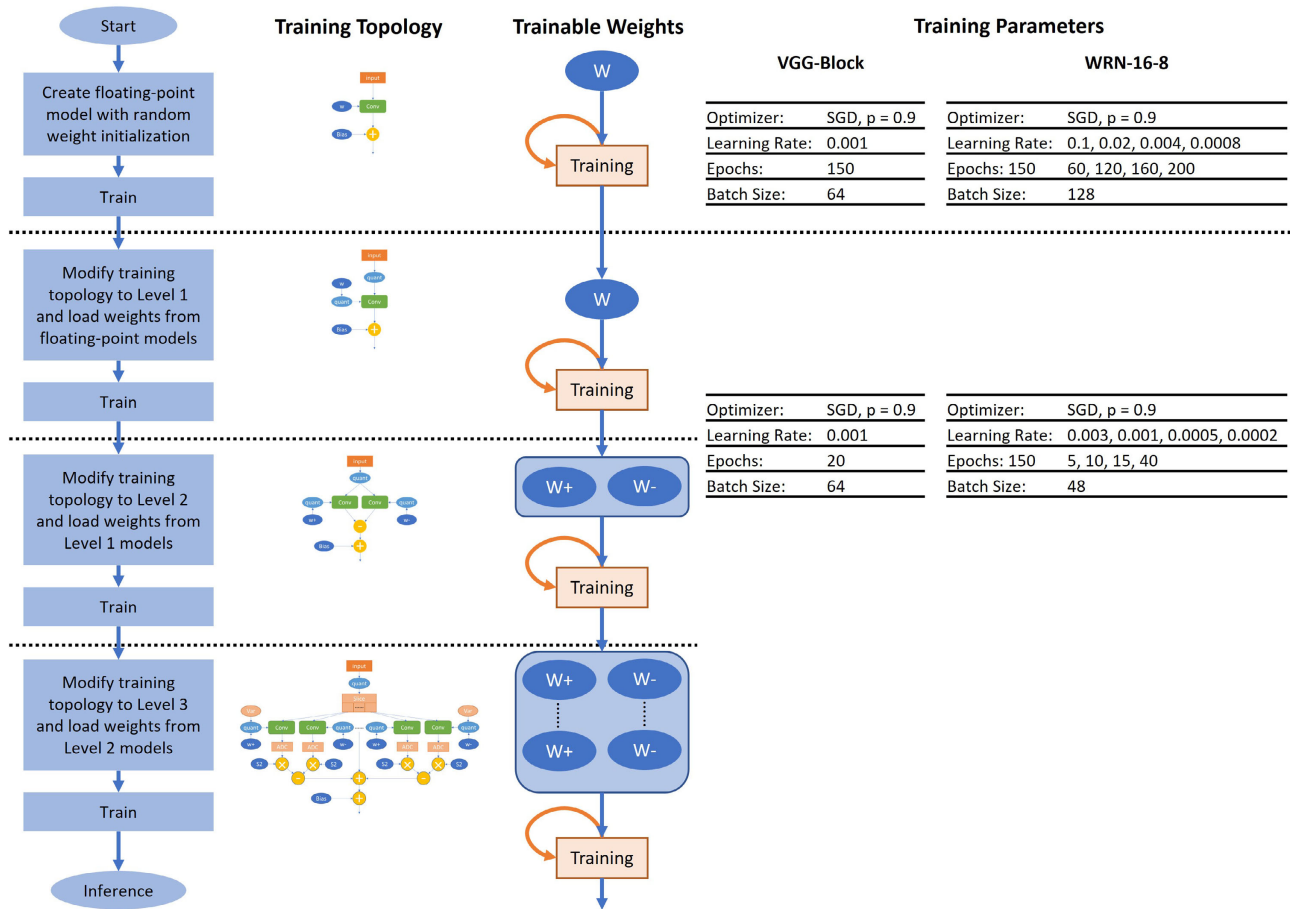
Figure 3. Architecture-aware training process and parameters used during training.
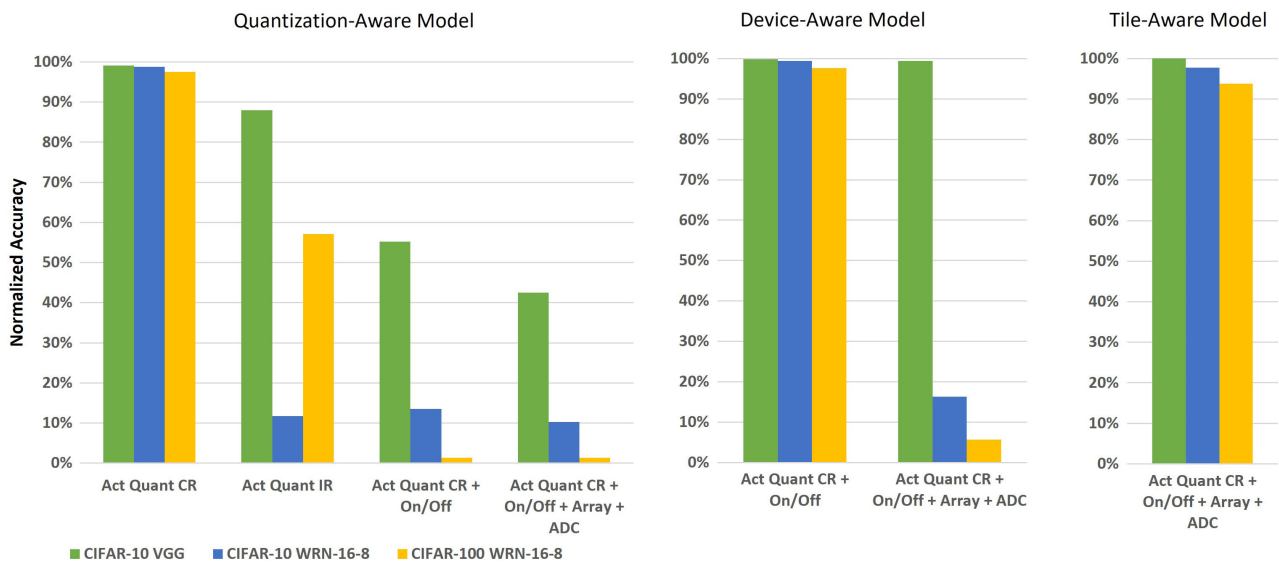

Figure 4. Effect of signed weights represented in two cells, on/off ratio, ADC, and array size limitation. Inference accuracy. Act Quant CR: activation and weight quantization 8bits with activation quantization range corresponding to ReLu6 used during training, on/off ratio 1000. Act Quant IR: activation and weight quantization 8bits with activation quantization range of 0–1 which does not correspond to the ReLu6 range used during training. Act Quant CR + On/Off: low on/off ratio of 10. Act Quant CR + On/Off + Array + ADC: array size 265 × 64, 8bit ADC. Floating-point accuracies for the CIFAR-10 VGG, CIFAR-10 WRN, CIFAR-100 WRN models are 83.73%, 95.11%, and 74.74%.

**ADVANCED
SCIENCE NEWS**

www.advancedsciencenews.com

**ADVANCED
INTELLIGENT
SYSTEMS**
Open Access

www.advintellsyst.com

networks inference accuracy, the effectiveness of mitigations methods, and factors that impact network robustness against device variation under realistic device and circuit conditions.

We examined the effect of weight variations using models trained with level 1, 2, and 3 pipelines, and studied the model accuracy in the corresponding inference conditions (e.g., when only quantization effects, quantization + device on/off, and quantization, on/off and finite array size and ADC precision effects are present during inference, respectively) (**Figure 5**). Previous studies have shown that the VGG-block-based model had minimal accuracy drop even at relatively high variation levels, while more complex models show severe accuracy degradation.[14] In this section, we thus used the more complex WRN-16-8 models for the CIFAR-10 dataset to highlight the effects of device variations.

In the accuracy test, after weight storage, variations were applied additively as Gaussian distributions with a constant standard deviation across all weights (i.e., 4% variation means the standard deviation is 4% of the dynamic range of memory cells). This variation distribution was chosen as a generic example because memory technologies have substantially different characteristics, and it represents a near-worst-case scenario. On one side, many emerging resistive switching devices exhibit state-dependent programming variation, where lower conductance states are associated with lower variations,[5,21] which is less detrimental to inference accuracy. On the other side, programming variations in multibit Flash memories are generally more state-independent while also suffering from additional nonlinear behaviors.[22,23] In level 2 and level 3 inference pipelines, where signed weights are represented in two columns (Figure 1b), the

variations are applied independently to each cell. This is different from variations that are directly applied to the signed weights (level 1) and means the impacts of weight variations are not equivalent between level 1 and the other pipelines. We also note that the signed weight representation we adopted (Figure 1b) is more realistic than the approach where differential cells (cells consisting of two devices) are read out individually[2,5] and more practical to implement in circuits compared to the approach where two rows with positive and negative input voltages signs are used to represent to positive and negative weights.

Because each programming session on each chip results in effectively different models and different inference accuracy, the process was simulated 40 times for each condition. The distribution of inference accuracy is shown in box plots (Figure 5). The models are expected to be programmed onto memory arrays and do not change during inference. Therefore, programming time is less important. Thus, the top 25 percentile in the accuracy distribution is more representative than the average or the median because it can be achieved by attempting programming sessions multiple times. Gray boxes in Figure 5 represent inference accuracies of models trained without any mitigation measures, and, in general, accuracy degradation becomes unacceptable for variations beyond 4%.

When comparing between level 1 and level 2 pipelines, accuracy degradation is more pronounced in level 2 due to the signed weight representation and low on/off ratio of memory devices (Figure 5a,b). When ADC and array size limitation is introduced at level 3, surprisingly, the accuracy degradation at high variation levels (>4%) improved compared to level 1 and level 2. This is
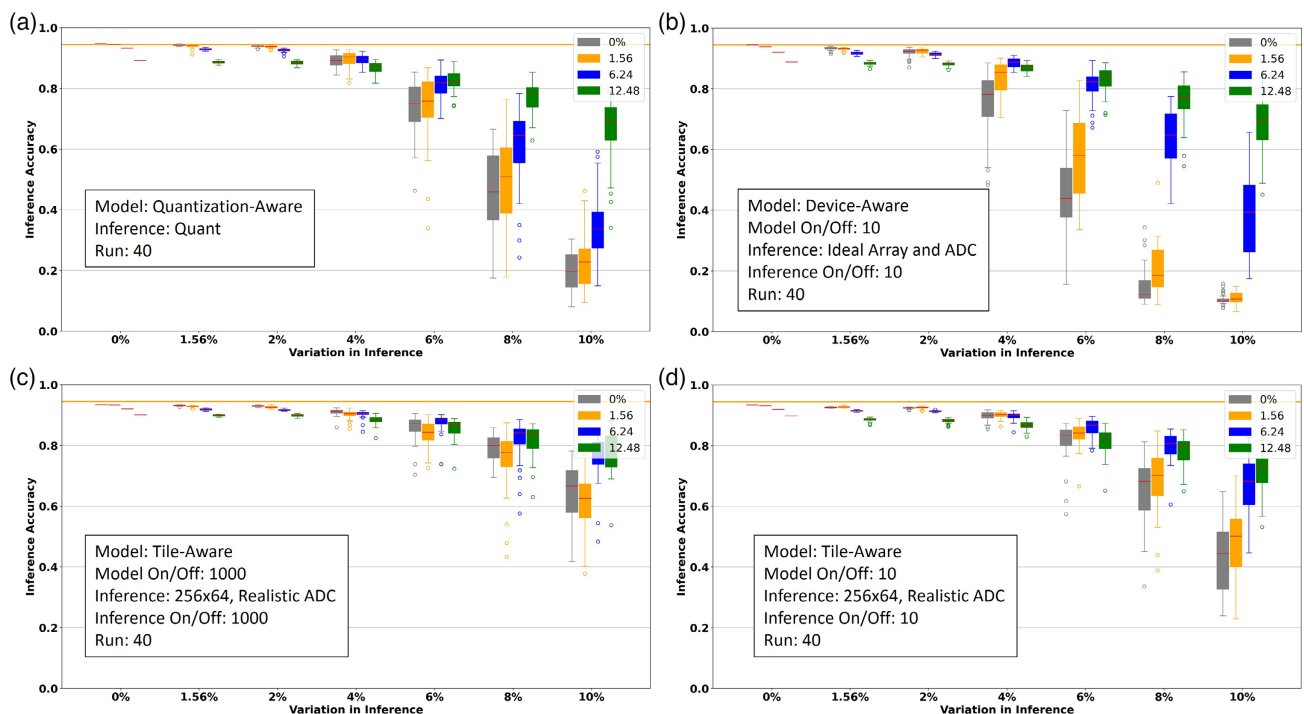
Figure 5. Variation effect under different inference pipelines for WRN-16-8 network on the CIFAR-10 dataset, for models trained with different levels of noise injection. The variation level is defined as standard deviation relative to the dynamic range of the weights. The boxplots show model inference accuracy distribution from 40 runs. Legend: variation injected during training. Orange lines: floating-point baseline. Ideal Array and ADC: no array size limitation, no quantization or range limitation of output. Realistic ADC: 8bit ADC with 0–45 μA as described in Section 3.2, array size 256 × 64.

**ADVANCED
SCIENCE NEWS**
www.advancedsciencenews.com

**ADVANCED
INTELLIGENT
SYSTEMS**
Open Access
www.advintellsyst.com

likely due to the presence of trainable S2 multipliers on a per array basis we implemented in the tiled architecture (Figure 2).[14] At this level, we can also observe the negative effect of a low device on/off ratio (Figure 5c,d). However, in general, accuracy degradations become unacceptable when variations exceed 4% of the dynamic range.

As a natural extension in the architecture-aware training approach, we hypothesize that injecting noise during training may improve inference accuracy. Specifically, we used weight noise injection during training to mimic device programming variations to produce trained DNN models that can produce better inference accuracy in presence of variations. In this implementation, weight noise is added after each minibatch during training, where an error is drawn from a Gaussian distribution for each weight and then added to it. The standard deviation for the Gaussian distribution is defined as relative to the dynamic range of the memory cells. For example, 1.56% noise injection means the Gaussian distribution has a standard deviation that is 1.56% of the dynamic range of memory cells. From a general neural network training perspective, noise injection at inputs, hidden units, and weights during training have long been proposed as methods to improve the generalization ability of neural networks.[24–28] In particular, weight noise injection has been shown mathematically to improve fault tolerance as it produces networks with smoother input–output mapping where the output becomes less sensitive to noise.[26] Recent studies have also applied this method to analog computing systems.[29–32] However, these prior studies are generally limited to small-scale networks or did not consider realistic system limitations like ADC characteristics, device on/off ratios, and especially array size limitations. The improvements in inference accuracy from weight noise injection in training can be observed in Figure 5, and the trend in improvements is consistent across different inference pipelines. In general, higher level noise injection leads to better accuracy recovery. For high device variations, noise injection not only allows the average and the peak accuracy to recover but also reduces the variation in performance between different runs.

The improvements from noise injection can also be observed from model outputs directly. **Figure 6**a shows the error in model outputs caused by device programming variation with CIFAR-10 validation dataset as input. For inference with a programming variation level of 6.24%, the injection of noises of the same level during training significantly reduces the error in the model outputs. For inference with a higher device programming variation of 12.48%, although substantial errors still occur with 12.48% noise injection, the trend in improvement is similar. When random patterns are used as input for models trained on the CIFAR-10, the error caused by programming variation is significantly larger (Figure 5b). This suggests that neural network models are trained for a specific input distribution, and the impact of weight variation can be more pronounced if the inference task is different from that during training.

The robustness of neural network models against weight variation can be characterized as part of the generalization ability. It has been shown that the addition of weight variation exacerbates inference error caused by generalization limitations and roughness of neural network models.[26] This means, in order to obtain acceptable inference accuracy for the same tasks, the models will need to have better generalization ability in the presence of weight variation, and many factors in the training process influence the ability. Indeed, we have found that even when models have similar accuracy with no weight variations, they can have very different robustness against variations. Prior literature has shown ample results for the effects of quantization and learning rate on the generalization ability in standard digital implementation. However, the effects of these factors on neural networks model have not been discussed in the context of analog IMC systems with the presence of realistic hardware limitations. In the next section, we investigate the impact of learning rate, programming target resolution, and different inference pipelines have on model robustness against weight variation.

5.1. Effect of Higher Target Programming Resolution

Although 8bit programming target resolution cannot be reliably represented by devices with a variation of even as low as 1.56%, we found, compared to 4bit programming target, higher target resolution produces models more robust to programming variations (**Figure 7**). Thus, 8bit programming target resolution was used in this work. We believe this is because, although quantization-aware training methods produce models more suited for the specific deterministic error caused by quantization, these models have diminished generalization ability, thus more sensitive to any additional errors like device variation.[16,33] The higher target resolution produces trained models with wider local minima, thus higher robustness against variation in weights.

5.2. Effect of Difference Inference Pipeline on the Same Models

We observed that although the same model can generally achieve very similar accuracies under different inference pipelines, in the presence of weight variations very different behaviors are obtained. For example, level 1 trained models show similar accuracy in the level 1 (Figure 5a) and level 2 (**Figure 8**a) inference pipelines. However, in the presence of relatively high weight variations, the accuracies are consistently lower in the level 2 inference pipeline. This is likely because level 1 trained models are not optimized for the level 2 inference pipeline when additional hardware details are introduced that are not incorporated during training, but the errors may not be large enough to cause accuracy degradation when there is no weight variation. In the presence of weight variations, the effects are amplified and lead to much worse accuracy degradation when training is not matched with the inference conditions.

5.3. Impact of Learning Rate

When a simple learning rate of 0.0002 is used in the fine-tuning process instead of the learning rate schedule described in Section 2.2, models achieved similar accuracies with no weight variation. When weight variation is introduced, device-aware trained models (level 2) showed no clear pattern between the two different learning rates, while tile-aware trained (level 3) models with a learning rate of 0.0002 are more sensitive to weight variations compared to ones obtained with the learning rate schedule (**Figure 9**). It is well known that selecting suitable learning rates is critical in the training process for neural network models to converge to optimal
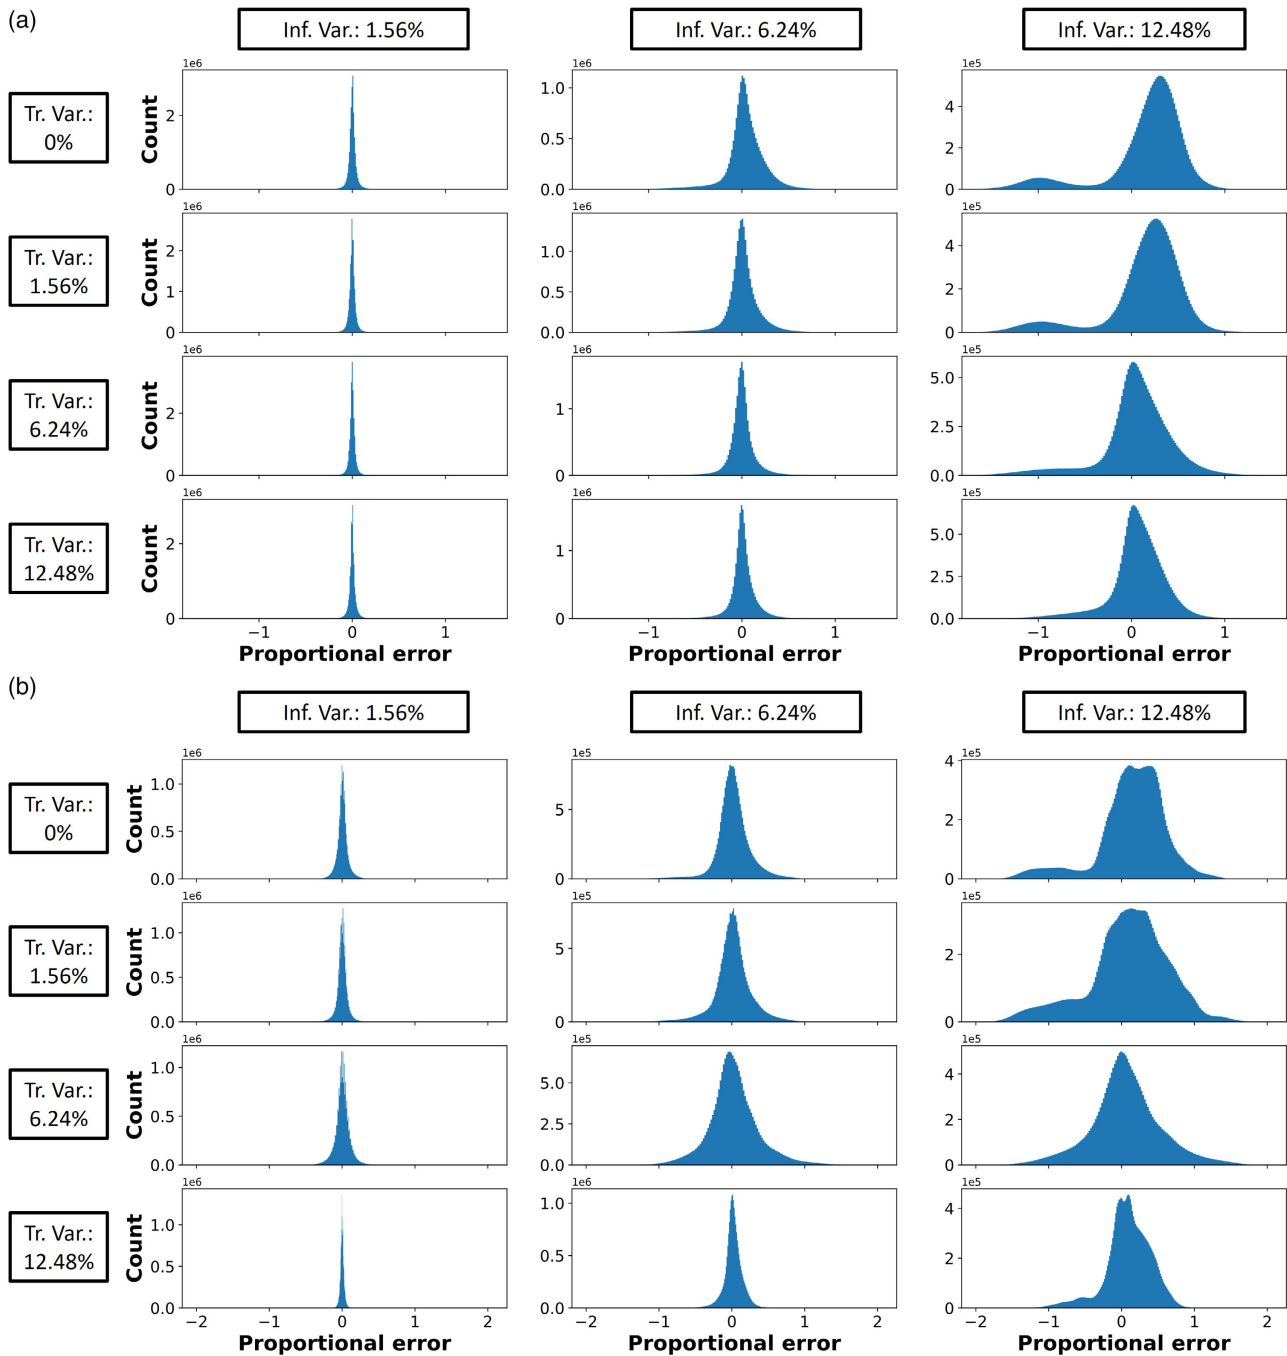
Figure 6. The proportional error of network model inference output with weight programming variation compared to inference output without variation, level 3 model in level 3 inference pipeline with on/off ratio of 10. Programming variations are simulated 40 times, and results are aggregated. Tr. Var.: variation injected during training. Inf. Var: variations experienced in the device programming process. a) Images from the validation dataset as input to the network. b) Random pattern as input to the network.

states, and learning rate schedules are often superior to constant learning rates.[34] In this particular case, the learning rate schedule showed an advantage in level 3 training but not in leve 2 training. We believe the addition of array size and ADC limitations at level 3, analogous to weight quantization, results in models with less smooth input-mapping, thus illuminating the difference between models produced by the different learning rates.

6. CIFAR-100 Results and Discussions

We also considered the WRN-16-8 model for the more complex CIFAR-100 dataset. The results showed a similar general trend to the results for the CIFAR-10 dataset, while models are much more sensitive to variations across the board with significant accuracy degradation at as low as 2% variation (**Figure 10**).
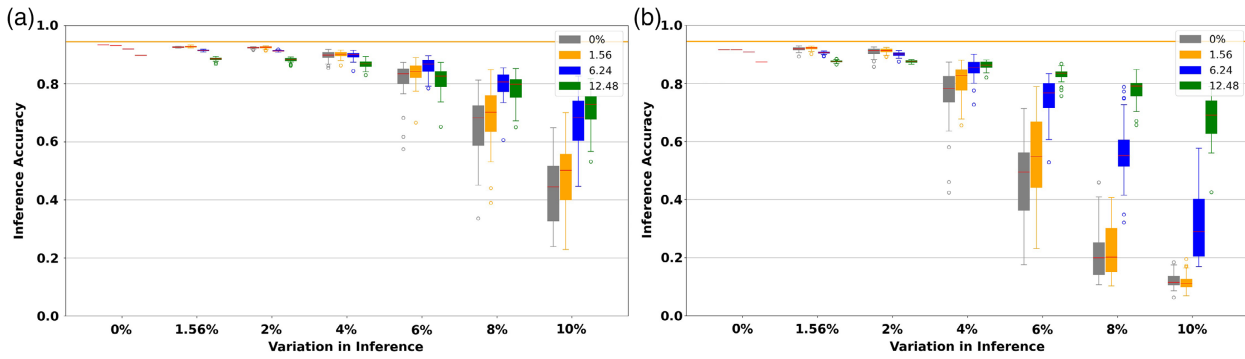
Figure 7. Sensitivity to variation of tile-aware trained WRN-16-8 CIFAR-10 model in tiled inference pipeline with ADC limitation, array size of 265 × 64, on/off ratio of 10 for both training and inference, accuracy evaluations were run 40 times. a) Models trained with 8bit quantized weights and 8bit programming target resolution during inference. b) Models trained with 4bit quantized weights and 4bit programming target resolution during inference.
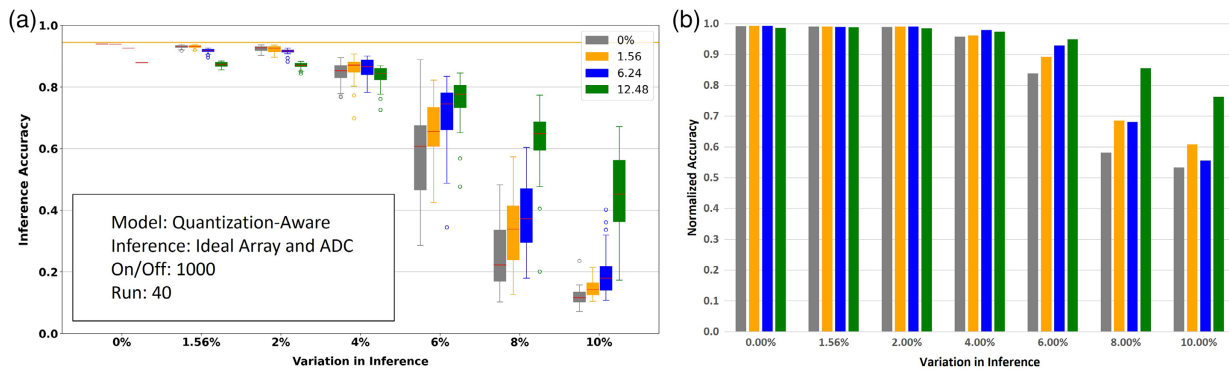
Figure 8. WRN-16-8 models for the CIFAR-10 dataset. a) Accuracy of level 1 trained model in level 2 inference pipeline. b) 75th percentile accuracies of level 1 trained models in level 2 inference pipeline, normalized to level 1 inference pipeline accuracies of corresponding training and inference variation.

This, again, illustrates that larger scale networks, as well as more complex tasks, are more sensitive to weight variations during inference, and as neural networks and tasks become more complicated, further improvements in DNN model robustness and device performance may be required.
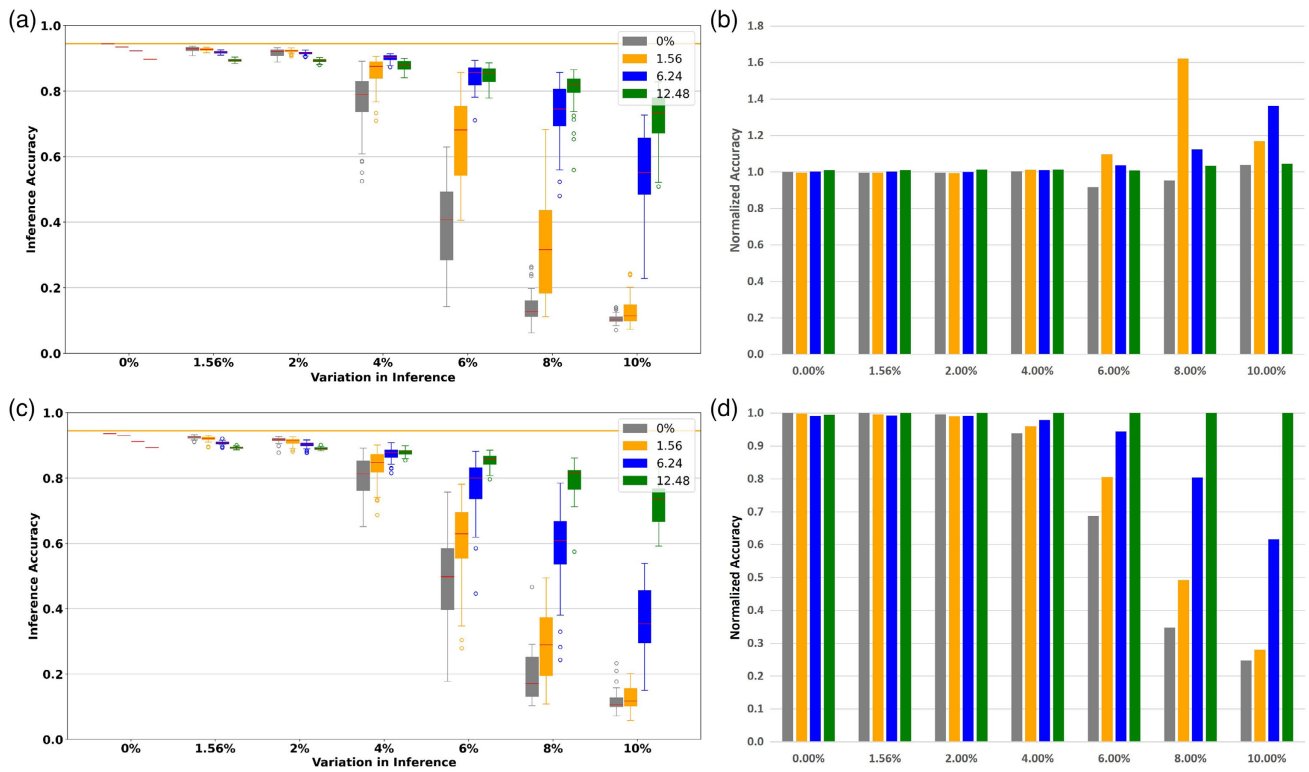
Other than improving the intrinsic precision of memory devices, two main methods have been proposed to improve weight storage precision. The first is using multiple cells to encode different bits of one weight,[21,35] and the second is to use closed-loop, write-verify programming schemes.[36,37] Using multiple cells drastically decreases memory density and incurs additional peripheral circuit overhead, thus resulting in decreases in computing efficiency in terms of both area and energy. Closed-loop programming processes have already been widely implemented, but have not been able to yield programming precision high enough for 8bit or even 4bit weights. As shown in Figure 10, even 2% variation can lead to significant accuracy degradation for complex tasks.

Here, we provide another possibility for future device and programming algorithm design. Neural networks are generally sparse, where weights close to zero constitute a large portion of all weights. This means programming variations at low

conductance states have much higher impacts, and nonuniform device programming variation characteristics can be engineered to minimize the effect at low conductance value. In particular, most resistive switching memory technologies have a limited analog dynamic range where conductance can be changed continuously, compared to the entire dynamic range. For weights close to zero, the corresponding memory devices can be hard reset, where conductance is set to the lowest possible value beyond the analog range. This would greatly reduce weight variation for weights close to zero because the absolute variation at the lowest conductance state is generally much lower compared to that inside the analog range. Although hard reset can have an impact on endurance in some memory technology, programming is expected to be infrequent, and endurance is unlikely to be the primary limiting factor.

7. Conclusion

In this work, we took a systematic look at the weight variation effect caused by memory device programming in analog IMC systems, which appears to be the most difficult error source

ADVANCED
SCIENCE NEWS

www.advancedsciencenews.com

ADVANCED
INTELLIGENT
SYSTEMS
Open Access

www.advintellsyst.com

Figure 9. Effects of learning rates. a) Device-aware models trained with a learning rate of 0.0002, evaluated in tiled inference pipeline. b) 75th percentile accuracies in (a) normalized to that of models trained with learning rate schedule (Figure 5b). c) 75th percentile tile-aware models evaluated in the tiled pipeline with an on/off ratio of 10. d) Accuracies in (c) normalized to that of models trained with learning rate schedule (Figure 5d).
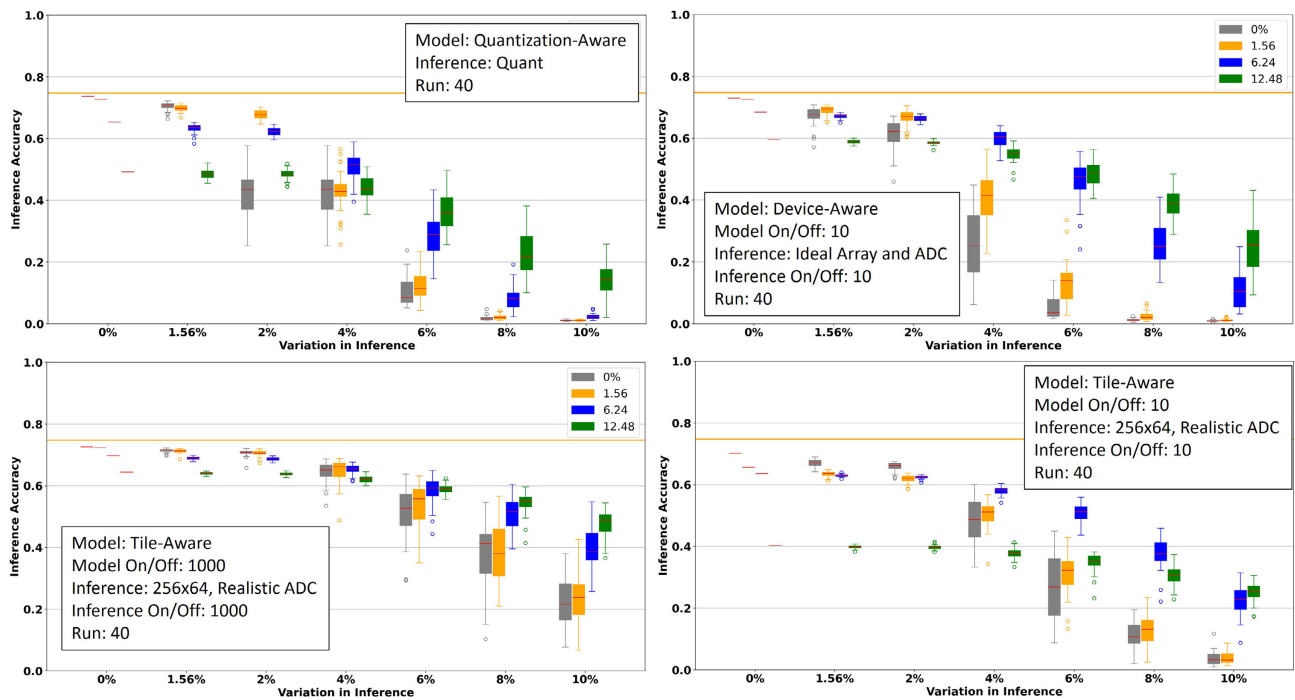
Figure 10. Variation effect under different inference pipeline for WRN-16-8 network on CIFAR-100 dataset.

to mitigate. We show proper noise injection can improve model robustness against weight variations. However, in the presence of moderate to high variations and for complex tasks and models, these methods may not be able to fully recover the accuracy drop. Thus, further developments in algorithms to produce neural networks that are more robust against weight variations could be critical for practical deployment for analog IMC systems for neural network workload.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

[1] E. Strubell, A. Ganesh, A. McCallum, in *AAAI 2020—34th AAAI Conf. Artif. Intell.*, New York **2020**, pp. 1393–13696.

[2] S. Yu, H. Jiang, S. Huang, X. Peng, A. Lu, *IEEE Circuits Syst. Mag.* **2021**, *21*, 31.

[3] M. A. Zidan, W. D. Lu, in *Memristive Devices for Brain-Inspired Computing*, Elsevier, San Diego **2020**, pp. 221–254.

[4] X. Peng, S. Huang, Y. Luo, X. Sun, S. Yu, *Technical Digest - Int. Electron Devices meeting, IEDM 2019*, San Francisco **2019**, pp. 32.5.1–32.5.4.

[5] V. Joshi, M. Le Gallo, S. Haefeli, I. Boybat, S. R. Nandakumar, C. Piveteau, M. Dazzi, B. Rajendran, A. Sebastian, E. Eleftheriou, *Nat. Commun.* **2020**, *11*, 2473.

[6] F. Cai, J. M. Correll, S. H. Lee, Y. Lim, V. Bothra, Z. Zhang, M. P. Flynn, W. D. Lu, *Nat. Electron.* **2019**, *2*, 290.

[7] P. Yao, H. Wu, B. Gao, J. Tang, Q. Zhang, W. Zhang, J. J. Yang, H. Qian, *Nature* **2020**, *577*, 641.

[8] O. Krestinskaya, B. Choubey, A. P. James, *Sci. Rep.* **2020**, *10*, 5838.

[9] O. Krestinskaya, A. Irmanova, A. P. James, in *2019 IEEE Int. Symp. Circuits Syst.*, Sapporo, **2019**, pp. 1–5.

[10] B. Murmann, Boris Murmann: ADC Performance Survey, https://web.stanford.edu/~murmann/adcsurvey.html, **2021**.

[11] Y. Jeong, M. A. Zidan, W. D. Lu, *IEEE Trans. Nanotechnol.* **2018**, *17*, 184.

[12] Q. Wang, X. Wang, S. H. Lee, F.-H. Meng, W. D. Lu, in *Technical Digest - Int. Electron Devices Meeting, IEDM 2019*, San Francisco **2019**, pp. 14.4.1–14.4.4.

[13] W. Ma, F. Cai, C. Du, Y. Jeong, M. Zidan, W. D. Lu, in *Technical Digest—Int. Electron Devices Meeting IEDM 2017*, San Francisco **2017**, p. 16.7.1.

[14] Q. Wang, Y. Park, W. D. Lu, in *2021 IEEE Int. Symp. Circuits Systems*, IEEE, Piscataway, NJ **2021** pp. 1–5.

[15] X. Wang, Q. Wang, F. H. Meng, S. H. Lee, W. D. Lu, in *Proc.—2020 IEEE Int. Conf. Artificial Intelligence Circuits and Systems AICAS 2020*, Genoa **2020**, pp. 141–144.

[16] R. Krishnamoorthi, (Preprint) arXiv:1806.08342, v1, unpublished **2018**.

[17] K. Simonyan, A. Zisserman, in *3rd Int. Conf. on Learning Representations, ICLR 2015*, San Diego **2015**.

[18] B. E. Jonsson, in *IMEKO TC4 Int. Workshop ADC Model. Test. Data Convert. Anal. Des. 2011, IWADC 2011*, Orvieto **2011**, pp. 132–137.

[19] S. Zagoruyko, N. Komodakis, in *Proc. Br. Mach. Vis. Conf. 2016*, York **2016**, pp. 87.1–87.12.

[20] L.-H. Tsai, S.-C. Chang, Y.-T. Chen, J.-Y. Pan, W. Wei, D.-C. Juan, (Preprint) arXiv:2007.03230, v2, unpublished **2020**.

[21] M. A. Zidan, Y. Jeong, J. Lee, B. Chen, S. Huang, M. J. Kushner, W. D. Lu, *Nat. Electron.* **2018**, *1*, 411.

[22] X. Guo, F. M. Bayat, M. Prezioso, Y. Chen, B. Nguyen, N. Do, D. B. Strukov, in *2017 IEEE Custom Integrated Circuits Conf.*, Austin **2017**, pp. 1–4.

[23] Y. Cai, E. F. Haratsch, O. Mutlu, K. Mai, in *Design, Automation & Test in Europe Conf. and Exhibition (DATE), 2013*, New Jersey **2013**, pp. 1285–1290.

[24] A. F. Murray, P. J. Edwards, *IEEE Trans. Neural Netw.* **1993**, *4*, 722.

[25] R. D. Clay, C. H. Sequin, in *JCNN Int. Jt. Conf. Neural Networks*, Baltimore **1992**, pp. 769–774.

[26] G. An, *Neural Comput.* **1996**, *8*, 643.

[27] D. P. Kingma, T. Salimans, M. Welling, in *Proc. 28th Int. Conf. Neural Inf. Process. Syst., NIPS 2015*, Montréal **2015**, pp. 2575–2583.

[28] H. Noh, T. You, J. Mun, B. Han, in *Advances in Neural Information Processing Systems*, **2017**, pp. 5110–5119.

[29] S. Moon, K. Shin, D. Jeon, *IEEE Trans. Very Large Scale Integr. Syst.* **2019**, *27*, 1455.

[30] D. Miyashita, S. Kousai, T. Suzuki, J. Deguchi, *IEEE J. Solid-State Circuits* **2017**, *52*, 2679.

[31] M. Klachko, M. R. Mahmoodi, D. Strukov, in *Proc. Int. Joint Conf. Neural Networks IJCNN'19*, Budapest **2019**, pp. 1–8.

[32] A. S. Rekhi, B. Zimmer, N. Nedovic, N. Liu, R. Venkatesan, M. Wang, B. Khailany, W. J. Dally, C. T. Gray, in *Proc. 56th Annual Design Automation Conf. DAC 2019, Las Vegas*, **2019**, pp. 1–6.

[33] Y. Boo, S. Shin, W. Sung, in *2020 IEEE Work. Signal Process. Syst.*, Coimbra **2020**, pp. 1–6.

[34] L. N. Smith, in *Proc.—2017 IEEE Winter Conf. Appl. Comput. Vision, WACV 2017*, Santa Rosa **2017**, pp. 464–472.

[35] A. P. James, L. O. Chua, *IEEE Trans. Circuits Syst. I. Regul. Pap.* **2021**, *68*, 4470.

[36] S.-S. Sheu, M.-F. Chang, K.-F. Lin, C.-W. Wu, Y.-S. Chen, P.-F. Chiu, C.-C. Kuo, Y.-S. Yang, P.-C. Chiang, W.-P. Lin, C.-H. Lin, H.-Y. Lee, P.-Y. Gu, S.-M. Wang, F. T. Chen, K.-L. Su, C.-H. Lien, K.-H. Cheng, H.-T. Wu, T.-K. Ku, M.-J. Kao, M.-J. Tsai, in *2011 IEEE Int. Solid-State Circuits Conf., ISSCC 2011*, San Francisco **2011**, pp. 200–202.

[37] N. Papandreou, H. Pozidis, A. Pantazi, A. Sebastian, M. Breitwisch, C. Lam, E. Eleftheriou, in *2011 IEEE Int. Symp. Circuits Systems*, Rio de Janeiro **2011**, pp. 329–332.