Solid-State Electronics 232 (2026) 109296

Contents lists available at ScienceDirect

Solid State Electronics

journal homepage: www.elsevier.com/locate/sse

Exploring variability and quantization effects in artificial neural networks using the MNIST dataset☆,☆☆

Alan Blumenstein [a,b],*, Eduardo Pérez [c,d], Christian Wenger [c,d], Nadine Dersch [a,b], Alexander Kloes [a], Benjamín Iñíguez [b], Mike Schwarz [a],**

[a] NanoP, THM University of Applied Sciences, Giessen, Germany
[b] DEEEA, Universitat Rovira i Virgili, Tarragona, Spain
[c] IHP - Leibniz Institute for High Performance Microelectronics, Frankfurt (Oder), Germany
[d] BTU Cottbus-Senftenberg, Cottbus, Germany

ARTICLE INFO

Keywords:
Single-Layer Perceptron
Multi-Layer Perceptron
Cycle-to-Cycle variability
Quantization
Stochastic variability
Simulation

ABSTRACT

This paper investigates the impact of introducing variability to trained neural networks and examines the effects of variability and quantization on network accuracy. The study utilizes the MNIST dataset to evaluate various Multi-Layer Perceptron configurations: a baseline model with a Single-Layer Perceptron and an extended model with multiple hidden nodes. The effects of Cycle-to-Cycle variability on network accuracy are explored by varying parameters such as the standard deviation to simulate dynamic changes in network weights. In particular, the performance differences between the Single-Layer Perceptron and the Multi-Layer Perceptron with hidden layers are analyzed, highlighting the network's robustness to stochastic perturbations. These results provide insights into the effects of quantization and network architecture on accuracy under varying levels of variability.

1. Introduction

Artificial Neural Networks (ANN) have become a cornerstone in various machine learning applications due to their ability to model complex patterns and relationships within data [1]. However, understanding how these networks respond to changes in their parameters is useful for developing simulation models and investigation of devices and architectures where the ANN could vary as in the case of **Resistive Random Access Memory** (RRAM) by the filament behavior [2]. Recent state-of-the-art demonstrations of RRAM-based ANN hardware show reliable multi-level programming with **Cycle-to-Cycle** (C2C) variability limited to only a few percent of the conductance range [3]. In particular, the **Single-Layer Perceptron** (SLP), one of the simplest forms of ANN, has been studied at the hardware level to evaluate its feasibility and limitations when implemented with emerging memory technologies, as shown in [4].

This study investigates the effects of introducing stochastic variability to the weights of ANN through a **Variable Neural Network** (VNN) class, which leverages parameters like **Adjustment Rate** (AR), the cumulative change of random synaptic weights in the VNN, and the **Standard Deviation** (σ) to simulate C2C variability that is present on RRAM according to [5]. In this work, **Device-to-Device** (D2D) variability is intentionally excluded from the analysis in order to isolate and evaluate the impact of other factors. Although process-induced D2D variation can be reduced — e.g., by using programming algorithms [6] — it remains inherently present in experimental settings. As a result, only C2C variations are considered, and D2D variation is not included in this study. Prior to quantization, the influence of increasing the number of hidden neurons in the **Multi-Layer Perceptron** (MLP) is considered, as previously explored in [7] to improve precision; however, that work did not address the effects of variability.

Additionally, the study explores the impact of quantization on network accuracy, providing insights into the trade-offs between precision and performance, as shown in [8] but without taking into consideration the AR.

☆ This article is part of a Special issue entitled: 'EuroSOI-ULIS 2025' published in Solid State Electronics.
☆☆ The review of this paper was arranged by Francisco F. Gamiz.

* Corresponding author at: NanoP, THM University of Applied Sciences, Giessen, Germany.
** Corresponding author.
E-mail addresses: alan.luis.blumenstein.leibaschoff@ei.thm.de (A. Blumenstein), mike.schwarz@ei.thm.de (M. Schwarz).

https://doi.org/10.1016/j.sse.2025.109296

Available online 4 December 2025

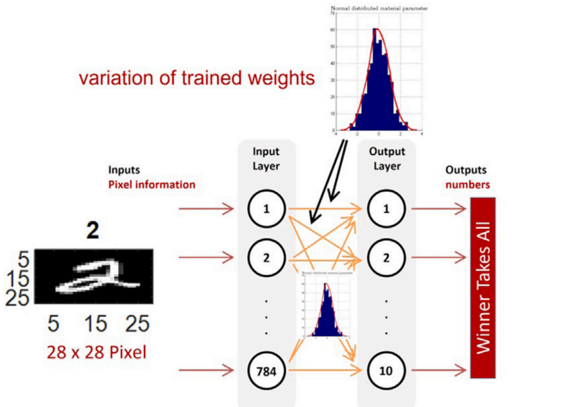

(a) SLP

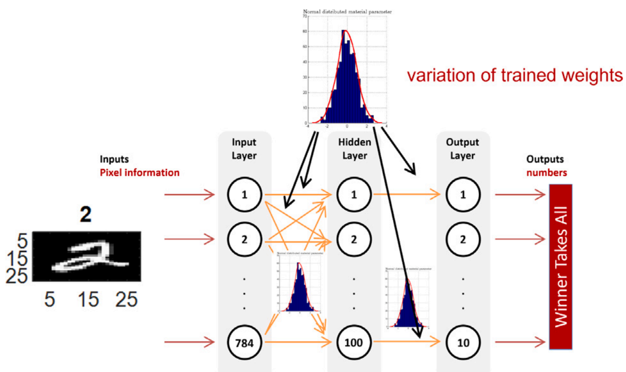

(b) MLP

Fig. 1. (a) SLP is modeled as an MLP without a hidden layer. (b) The MLP includes one hidden layer. In both cases, variability is applied to a proportional number of weights depending on AR. This variability is uniform across all affected weights.

2. Methodology

2.1. Neural network models

This study utilized the MNIST dataset with its standard split of 60,000 training images and 10,000 testing images to train and analyze different neural network configurations [9], focusing on a custom VNN. Two MLP models were used: a baseline MLP without a hidden layer, which is equivalent to a SLP, with an **ANN's Accuracy** (ACC) of 84.49%, based on [4] and shown in Fig. 1(a), an extended MLP with a hidden layer consisting of 100 nodes and a sigmoid activation function, trained for 15 epochs, achieves an initial ACC of 94.59%, as shown in Fig. 1(a).

2.2. VNN model

The VNN, developed in MATLAB, introduces variability into the weights of a trained ANN. It uses two parameters: AR, representing the cumulative synaptic change of random weights in the VNN (e.g., 0.1 = 10%, 1 = 100%, all weights are modified), and σ, which defines the degree of variability. A Monte Carlo simulation applies this stochasticity to random weights, and the ACC is analyzed across parameter levels, with results presented in ACC graphs. σ is identical for each synaptic weight, but the value is chosen randomly (C2C variability). We also tested a level-dependent variability model, where σ scales with each weight's conductance, and observed results similar to the uniform approach, although the shape changes from an exponential trend to an S-shape.

Table 1
Quantization levels.

Set N°	Levels	NL
Set 1	[−1, 1]	2
Set 2	[−1, 0, 1]	3
Set 3	[−1, −0.75, −0.50 …0.50 , 0.75, 1]	9
Set 4	[−1, −0.90, −0.80 …,0.80, 0.90, 1]	21
Set 5	[−1, −0.95, −0.90…,0.90, 0.95, 1]	41

2.3. Quantization approach

Quantization involves mapping ANN weights to the nearest value in a predefined set of levels, reducing the precision of the weights. This can be expressed as:

$$w_q = \arg\min_{l \in \mathcal{L}} |w - l|$$

where w is the original weight, $\mathcal{L}$ is the set of predefined quantization levels, and w_q is the quantized weight.

In this study, quantization was applied after training, using the schemes shown in Table 1. The number of levels (NL) defines the discrete values available for quantization in each scheme. ACC was then measured on the test set to evaluate the effect of each scheme.

2.4. Experimental setup

The experiments were designed to evaluate the impact of variability and quantization on different ANN configurations. The custom VNN implementation was applied to both the SLP and the MLP architectures. For each model, the ACC was measured across different values of AR under varying levels of σ. Additionally, the relationship between ACC and σ was evaluated for fixed AR values.

For the MLP, further experiments were conducted by varying the number of hidden layer nodes. In this case, instead of reporting the absolute ACC, the metric used was the difference between the ACC obtained after the variability simulation and the original ACC, in order to highlight the sensitivity of each configuration to variation. This difference is denoted as ΔACC and calculated as:

$$\Delta\mathrm{ACC} = \mathrm{ACC}_{\mathrm{var}} - \mathrm{ACC}_{\mathrm{orig}}$$

where $\mathrm{ACC}_{\mathrm{var}}$ is the accuracy after the variability simulation and $\mathrm{ACC}_{\mathrm{orig}}$ is the original (software-based) accuracy.

Finally, quantization effects were assessed by measuring the ACC of each model across different quantization levels, without introducing variability. This provided a baseline understanding of how quantization alone impacts network performance.

3. Results and discussion

3.1. C2C simulation in SLP

The results are presented through a series of comparative graphs. Fig. 2 illustrates the performance of the baseline SLP by showing ACC as a function of σ and AR. These graphs show the effect of applying variability at different levels.

The test is conducted as described in Section 2.4, the different values of σ are shown in Fig. 2(a), with the labels indicating the varying levels, and the different values of AR are shown in Fig. 2(b), where the levels are labeled accordingly. In both cases, an increase in either σ or AR results in a decrease in ACC, highlighting the sensitivity of the model to variability.

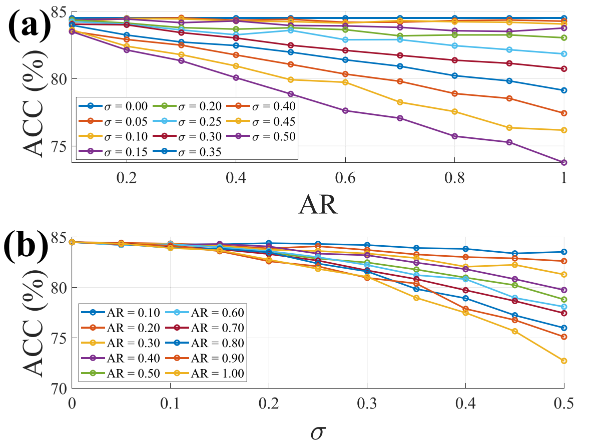

Fig. 2. (a) ACC variation for different σ values vs. ARs, and (b) ACC variation for various ARs vs. σ values for a ANN values for a SLP.

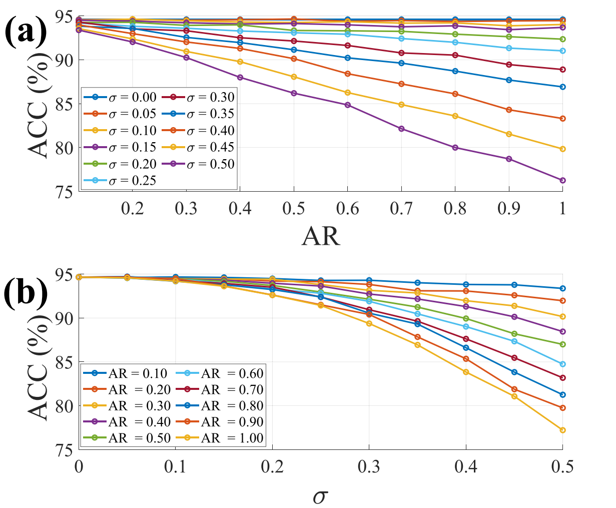

Fig. 3. (a) ACC variation for different σ values vs. ARs, and (b) ACC variation for various ARs vs. σ values for a MLP with 100 hidden nodes.

3.2. C2C simulation in MLP

Fig. 3 shows the results for the MLP. As illustrated in Figs. 3a and 3b, the accuracy decreases progressively with increasing values of AR and σ. The curves indicate a clear trend where higher levels of variability lead to performance degradation. A similar observation is reported in [8].

As described in Section 2.4, the same levels of σ and AR used for the baseline SLP model were applied to the MLP, allowing for a direct comparison between the two models. These levels of σ and AR can also be seen in Figs. 3(a) and 3(b), where the specific values used in the tests are labeled accordingly.

Finally, the SLP and MLP are compared in the maximal variation case, where AR is set to 1 and $\sigma = 0.5$, in order to evaluate the performance of both ANNs under extreme conditions. As shown in Fig. 4, subfigures (a) and (b) show that the $-\Delta$ACC of the MLP is generally higher than that of the SLP, which can be attributed to the greater number of weights in the MLP resulting from its additional layer.

3.3. C2C simulation in MLP with different number of hidden nodes

In this section, the C2C simulation is extended to the MLP with varying numbers of hidden nodes, each trained for 15 epochs. To

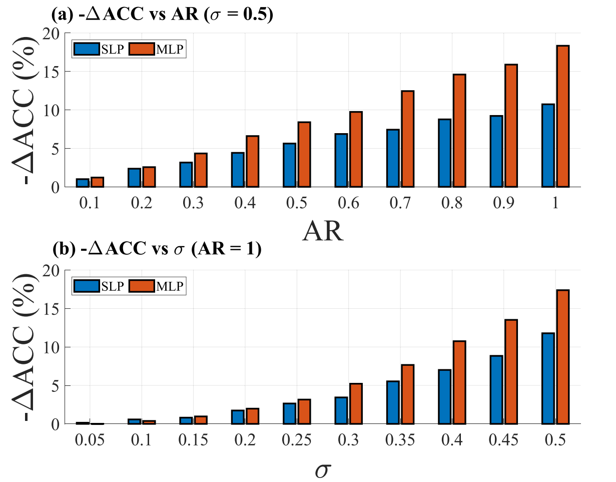

Fig. 4. Comparison between SLP and MLP: (a) $-\Delta$ACC versus AR for $\sigma = 0.5$; (b) $-\Delta$ACC versus σ for AR = 0.1.

Table 2
ΔACC at different AR values for fixed $\sigma = 0.1$.

Hidden nodes	ΔACC at AR = 0.1	ΔACC at AR = 1
10	−0.29	−1.71
20	−0.09	−0.53
50	−0.06	−0.46
100	−0.22	−0.61
200	−0.06	−1.16

Table 3
ΔACC at different σ values for fixed AR = 1.

Hidden nodes	ΔACC at $\sigma = 0.05$	ΔACC at $\sigma = 0.5$
10	−0.09	−3.50
20	0.06	−2.33
50	0.03	−2.98
100	0	−3.77
200	−0.04	−7.59

analyze the impact of AR, a fixed σ value of 0.1 is used. This isolates the effect of increasing adjustment rates on the accuracy of networks with different hidden layer sizes. Conversely, to evaluate the influence of σ, AR is held constant at 1, allowing for an assessment of how increased variability affects each model's robustness.

The results are presented in Fig. 5, where (a) shows the variation of ΔACC with respect to AR, and (b) presents the variation of ΔACC with respect to σ. For clarity, the numerical results are also summarized in Tables 2 and 3. Table 2 presents the change in accuracy (ΔACC) as a function of AR for different numbers of hidden nodes. Similarly, Table 3 reports ΔACC as a function of σ.

To further validate this observation, additional tests were conducted without randomness to ensure these irregularities were a result of stochastic variation. These tests, performed with more epochs and fixed seeds, reduced the observed differences but are not included in the paper.

The results indicate that increasing the number of hidden nodes tends to reduce the sensitivity to changes in AR at low variability ($\sigma = 0.1$), with ΔACC generally smaller for larger networks. However, under higher variability conditions ($\sigma = 0.5$), the degradation becomes more pronounced, especially in networks with a large number of hidden nodes, such as 200, where ΔACC reaches as much as −7.59%. These findings are visually supported by the trends shown in Fig. 5.

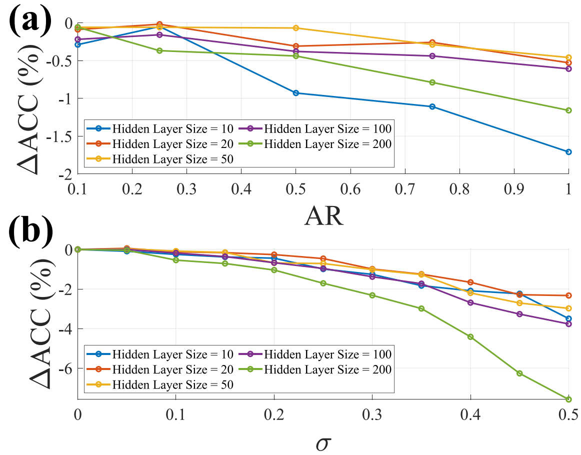

Fig. 5. Performance of the MLP with multiple hidden nodes under C2C simulation. (a) ACC versus AR for a fixed σ of 0.1. (b) ACC versus σ for a fixed AR of 1.

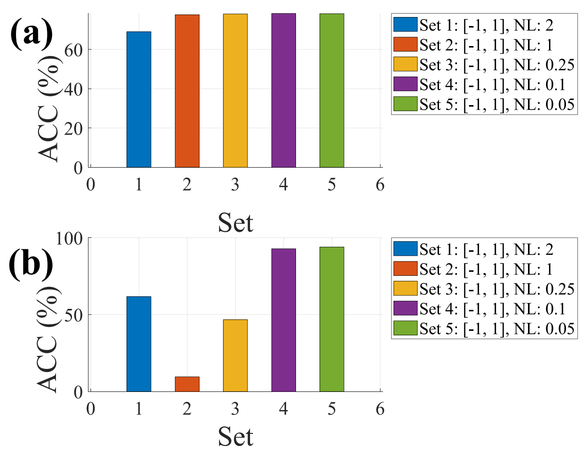

Fig. 6. The variation in ACC for different quantization ranges is shown in (a) for SLP, and in (b) for MLP one hidden layer with 100 hidden nodes.

3.4. Quantization levels

Fig. 6 highlights the effects of quantization on the ACC. For the SLP, ACC is lower with Set 1 (69.70%) compared to the other schemes, which range from 77.69% to 78.21% (Fig. 6a). This suggests that Set 1 introduces greater quantization-induced variability, a result consistent with findings in [4].

For the MLP, the performance shows greater variation between quantization sets (Fig. 6b). Set 1 yields the lowest ACC at 61.72%, followed by Set 3 at 46.69% and Set 2 at just 9.58%. In contrast, Sets 4 and 5 lead to significantly better results, with ACC values of 92.83% and 93.36%, respectively. This demonstrates the significant influence of quantization schemes on performance, highlighting the importance of range selection in optimizing the ACC of deeper networks.

4. Conclusion and future work

The study demonstrates that increasing AR and σ values consistently leads to reduced ACC, reflecting the detrimental effect of variability on network performance. In the SLP, accuracy drops from 84.49% to 72.71% as σ increases from 0 to 0.5 at AR = 1, and similarly decreases from 84.36% to 73.76% with increasing AR under fixed σ. The MLP shows higher robustness, maintaining ACC above 94% at low variability, but still degrades under high AR and σ.

Quantization analysis reveals that for the baseline SLP, Set 1 yields the lowest ACC (69.70%), while other schemes produce more stable results (77.69–78.21%). In contrast, the MLP is more sensitive to quantization scheme choice. Set 1 performs poorly (61.72%), and Set 2 even worse (9.58%), while Sets 4 and 5 provide the best performance (92.83% and 93.36%). These results highlight that deeper networks not only tolerate variability better but also require carefully selected quantization strategies.

Overall, the findings emphasize the critical role of architectural design in managing the effects of AR, σ, and quantization schemes to optimize ANN performance. Future work will investigate applying distinct σ values per synaptic weight (D2D variability) to better reflect RRAM stochasticity in hardware implementations.

CRediT authorship contribution statement

Alan Blumenstein: Writing – original draft, Software, Methodology, Investigation, Conceptualization. **Eduardo Pérez:** Writing – review & editing, Resources, Funding acquisition. **Christian Wenger:** Writing – review & editing. **Nadine Dersch:** Writing – review & editing, Methodology, Conceptualization. **Alexander Kloes:** Writing – review & editing. **Benjamín Iñíguez:** Writing – review & editing, Supervision. **Mike Schwarz:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to thank the German Research Foundation (DFG) for funding this work under grant 546680029.

Data availability

Data will be made available on request.

References

[1] Raman R, Kumar V, Pillai BG, Verma A, Rastogi S, Meenakshi R. Advances in multi-spectral image categorization through convolutional neural networks with an optimized activation function. In: 2024 international conference on data science and network security. IEEE; 2024, p. 1–6.
[2] Perez E, Mahadevaiah MK, Perez-Bosch Quesada E, Wenger C. In-depth characterization of switching dynamics in amorphous HfO2 memristive arrays for the implementation of synaptic updating rules. Japan J Appl Phys 2022;61(SM):SM1007.
[3] Pazos S, Zhu K, Villena MA, Alharbi O, Zheng W, Shen Y, Yuan Y, Ping Y, Lanza M. Synaptic and neural behaviours in a standard silicon transistor. Nature 2025;640(8057):69–76.
[4] Dersch N, Perez-Bosch Quesada E, Perez E, Wenger C, Roemer C, Schwarz M, Kloes A. Efficient circuit simulation of a memristive crossbar array with synaptic weight variability. Solid-State Electron 2023;209:108760.
[5] Kloes A, Bischoff C, Leise J, Perez-Bosch Quesada E, Wenger C, Perez E. Stochastic switching of memristors and consideration in circuit simulation. Solid-State Electron 2023;201:108606.
[6] Perez E, Grossi A, Zambelli C, Olivo P, Roelofs R, Wenger C. Reduction of the cell-to-cell variability in Hf1-xAlxOyBased RRAM arrays by using program algorithms. IEEE Electron Device Lett 2017;38(2):175–8.
[7] Haryono KB, Purnomo HC, Ferdinand R, Lucky H, Suhartono D. Investigating the influence of layers towards speed and accuracy of neural networks. In: 2022 international conference on informatics electrical and electronics. IEEE; 2022, p. 1–6.
[8] Chang C-C, Li S-T, Pan T-L, Tsai C-M, Wang I-T, Chang T-S, Hou T-H. Device quantization policy in variation-aware in-memory computing design. Sci Rep 2022;12(1).
[9] Deng L. The MNIST database of handwritten digit images for machine learning research [Best of the Web]. IEEE Signal Process Mag 2012;29(6):141–2.