A Survey of Ensemble Methods for Mitigating Memristive Neural Network Non-idealities

Muhammad Ahsan Kaleem[1], Jack Cai[1], Amirali Amirsoleimani[2], and Roman Genov[1]
Email: {ahsan.kaleem, jack.cai}@mail.utoronto.ca, amirsol@yorku.ca, roman@eecg.utoronto.ca
[1]Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada
[2]Department of Electrical Engineering and Computer Science, York University, Toronto ON M3J 1P3, Canada

***Abstract*—In this work, ensemble methods are presented and tested as universal ways to improve the performance of Memristive Deep Neural Networks (MDNNs) with non-idealities. The Generalized Ensemble Method and Weighted Voting ensemble methods improve the accuracy of classification on the MNIST dataset by 6.5% and 6.6% respectively, thus showing that they are more effective than basic Ensemble Averaging which has been investigated before, as well as other methods such as Voting. Different weighting schemes for Weighted Voting were tested, and we present Algorithm 1 and 2, which are the theoretically and experimentally optimal weighting schemes respectively. Our work serves as a guideline for choosing ensemble methods for MDNNs.**

***Index Terms*—Ensemble Methods, Memristor, Crossbar, Vector-Matrix Multiplication, Non-idealities, Neural Network.**

I. Introduction

MEMRISTORS are electronic components that can be used to perform operations such as vector matrix multiplication (VMM) entirely in memory, thus overcoming the Von Neumann bottleneck faced by traditional systems [1]. However, they are not perfect and face problems which prevent them from performing at their highest capability [2]. The main problem faced by memristors is non-idealities [3]. These non-idealities include device-to-device variability, random telegraph noise, and memristance drift [4] [5] [6]. Such non-idealities have been shown to negatively affect the performance of memristive devices, and thus it is important to develop methods to combat them [7]. Previously efforts on mitigating non-idealities are mostly hardware-based solutions, and there are many different proposed solutions for different kinds of non-idealities (and even for the same non-idealities, different research has proposed different solutions), and thus such solutions are not feasible. Some examples of these solutions include the ones presented in [8]–[12] which include various approaches based on the specific type of non-ideality tackled; even for the same non-ideality, for example device-to-device variability, [10] and [11] present two completely different solutions, one [10] involving an ultra-thin ALD-TiN buffer layer, and the other [11] proposing a solution involving increasing the roughness of bottom electrodes [13].

However, recently new research suggested a universal method to deal with non-idealities in memristor based neural networks: committee machines [13]; in this paper, committee machines were used to combine networks into multiple committees, resulting in higher accuracy compared to individual networks. Committee machines combine networks such that they perform better in committees consisting of multiple networks compared to the individual networks [14]. The committee machines used in [13] use ensemble averaging to linearly combine the outputs of individual networks. In this paper, we will present some different committee machines (interchangeably referred to as ensemble methods), including some novel algorithms designed specifically for memristive neural network applications, that achieve better results than simply ensemble averaging.

II. Ensemble Methods

A. Ensemble Averaging

In ensemble averaging, the outputs of multiple networks, each trained on the whole dataset, are linearly averaged to produce a final output as follows, where $\mathbf{x}$ is the input, N is the number of networks, and M_i is the i^{th} network (so $M_i(\mathbf{x})$ is the output of M_i for input $\mathbf{x}$): [15]

$$O(\mathbf{x}) = \frac{1}{N}\sum_{i=1}^{N} M_i(\mathbf{x}) \tag{1}$$

In the case of classification, the prediction is the class with the highest probability.

B. Generalized Ensemble Method

Generalized ensemble method (GEM) is an extension of normal ensemble averaging (EA). In this method, each network can be weighed differently, and the weights can be computed using various techniques, such as using the accuracies of the networks. The computation of the output is as follows, where w_i is the weight of the i^{th} network: [15]

$$O(\mathbf{x}) = \frac{\sum_{i=1}^{N} w_i M_i(\mathbf{x})}{\sum_{i=1}^{N} w_i} \tag{2}$$

There are many possible weighting schemes that can be chosen for GEM. Through our testing of various weighting schemes, we found that the most optimal one was one where each network was assigned a weight based on its performance relative to the other networks. The network with the highest accuracy was assigned a weight of N (where N is the number of networks), the network with the second highest accuracy was assigned a weight of N-1, and so on, until the network with the lowest accuracy was assigned a weight of 1. It could be possible that some minor changes to this weighting scheme could result in even better performance, such as using a slightly greater or less weight than just the ”rank” (in terms of accuracies), however the improvement would be negligible.

C. Voting

Voting is an ensemble method that is only applicable for classification systems, such as image classification on the MNIST dataset which we used for testing. The idea of voting is very simple; each individual network makes a prediction

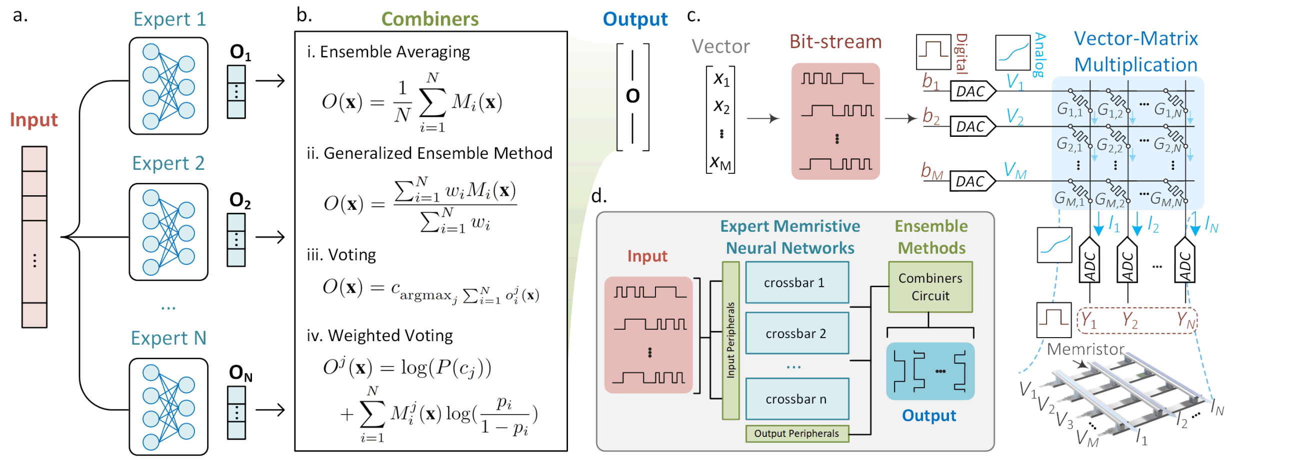

Fig. 1: (a) Input sent to multiple experts, instead of a single expert. (b) Four different ways of combining the outputs of multiple experts to produce one final output. (c) Memristive vector-matrix multiplication scheme. (d) Overview of the ensemble methods pipeline.

based on its output, and the most common prediction is chosen. This specific method is known as plurality voting [15]. If the classes are represented as c_k and the output of network M_i for class c_j for an input $\mathbf{x}$ is represented as $o_i^j(\mathbf{x})$, which takes the value 1 if the prediction of M_i is class c_j and 0 otherwise, then the final prediction with voting will be the following (in this case O refers to the predicted class itself, rather than the output vectors as in the case of the averaging ensemble methods):

$$O(\mathbf{x}) = c_{\text{argmax}_j \sum_{i=1}^{N} o_i^j(\mathbf{x})} \tag{3}$$

D. Weighted Voting

Weighted voting is an extension of voting in which each network can be weighed differently, similar to GEM; again, the question is what should the optimal weighting scheme be. In this case, with some assumptions we can mathematically derive what the optimal weighting scheme might be [15]: Let $\mathbf{l} = (l_1, ..., l_N)^\intercal$ be the outputs of the N individual networks; l_i is the class (a digit in the case of classification on the MNIST dataset) predicted by network M_i for some input $\mathbf{x}$. Let a_i be the accuracy of network M_i. Then, there is the following Bayesian optimal discriminant function for the combined output on class c_j:

$$O^j(\mathbf{x}) = \log(P(c_j)P(\mathbf{l}|c_j)) \tag{4}$$

We can simplify this expression as follows:

$$O^j(\mathbf{x}) = \log(P(c_j)) + \log(P(\mathbf{l}|c_j)) \tag{5}$$

Making the assumption that the outputs of each of the N networks are conditionally independent, $P(\mathbf{l}|c_j) = \prod_i P(l_i|c_j)$, (5) can be further simplified as follows:

$$\begin{aligned} O^j(\mathbf{x}) &= \log(P(c_j)) + \log(\prod_{i=1}^{N} P(l_i|c_j)) \\ &= \log(P(c_j)) + \log(\prod_{i=1,l_i=c_j}^{N} P(l_i|c_j) \prod_{i=1,l_i\neq c_j}^{N} P(l_i|c_j)) \\ &= \log(P(c_j)) + \log(\prod_{i=1,l_i=c_j}^{N} a_i \prod_{i=1,l_i\neq c_j}^{N} (1-a_i)) \\ &= \log(P(c_j)) + \sum_{i=1,l_i=c_j}^{N} \log(\frac{a_i}{1-a_i}) + \sum_{i=1}^{N} \log(1-a_i) \end{aligned} \tag{6}$$

The last term in this expression does not depend on the class c_j and $l_i = c_j$ can be written as the result of $M_i^j(\mathbf{x})$, so the expression can be further simplified to the following:

$$O^j(\mathbf{x}) = \log(P(c_j)) + \sum_{i=1}^{N} M_i^j(\mathbf{x}) \log(\frac{a_i}{1-a_i}) \tag{7}$$

Therefore, since $\log(P(c_j))$ does not depend on the individual networks M_i, the optimal weighting scheme would have weights computed using the following expression:

$$w_i \propto \log(\frac{a_i}{1-a_i}) \tag{8}$$

With this weighting scheme, the algorithm for weighted voting is as follows (for the computation of weights a proportionality constant of 1 was used as that was found to give the best results in comparison to other values):

Algorithm 1 Weighted Voting

1. For each MDNN M_i with accuracy a_i, assign weights as follows: $w_i = \log(\frac{a_i}{1-a_i})$
2. Compute prediction $p_i(\mathbf{x})$ for each M_i for some input $\mathbf{x}$; let $\mathbf{p(x)} = [p_1(\mathbf{x}), ..., p_N(\mathbf{x})]$
3. For each class c_j compute the following sum, where $o_i^j(\mathbf{x})$ is equal to 1 if $p_i(\mathbf{x}) = c_j$ and 0 otherwise: $\sum_{i=1}^{N} w_i o_i^j(\mathbf{x})$
4. The output prediction is the class c_j which maximises this sum, that is, $p(\mathbf{x}) = c_{\text{argmax}_j \sum_{i=1}^{N} w_i o_i^j(\mathbf{x})}$

As we will see later in this paper, this weighting scheme does not give the best results and instead an alternative scheme gives better results. In this weighting scheme, the 2 most accurate of the 5 MDNNs used for testing are given an extra vote (meaning their vote is equivalent to 2 votes rather than 1), because in our case only 5 MDNNs were used so increasing the weight of the most accurate MDNNs too much would not be beneficial as that would result in the output of weighted voting simply being equivalent to the output of those most accurate MDNNs. By a similar argument, increasing the weight by a smaller amount would almost be equivalent to an unweighted vote. This weighted scheme is generalized for an arbitrary amount of MDNNs in the following algorithm:

Algorithm 2 Weighted Voting

1. Select all MDNNs whose accuracy is at least 0.5 standard deviations greater than the mean accuracy of the MDNNs; let X be the set containing these MDNNs

2. Set weights as follows: $w_i = \begin{cases} 2, & \text{if } M_i \in X \\ 1, & \text{otherwise} \end{cases}$

3. Compute prediction $p_i(\mathbf{x})$ for each M_i for some input $\mathbf{x}$; let $\mathbf{p(x)} = [p_1(\mathbf{x}), ..., p_N(\mathbf{x})]$

4. **for** i = *1 to N* **do**

    if $w_i = 2$

        $\mathbf{p(x)}$.append($p_i(\mathbf{x})$)

5. Output prediction $p(\mathbf{x}) = mode(\mathbf{p(x)})$

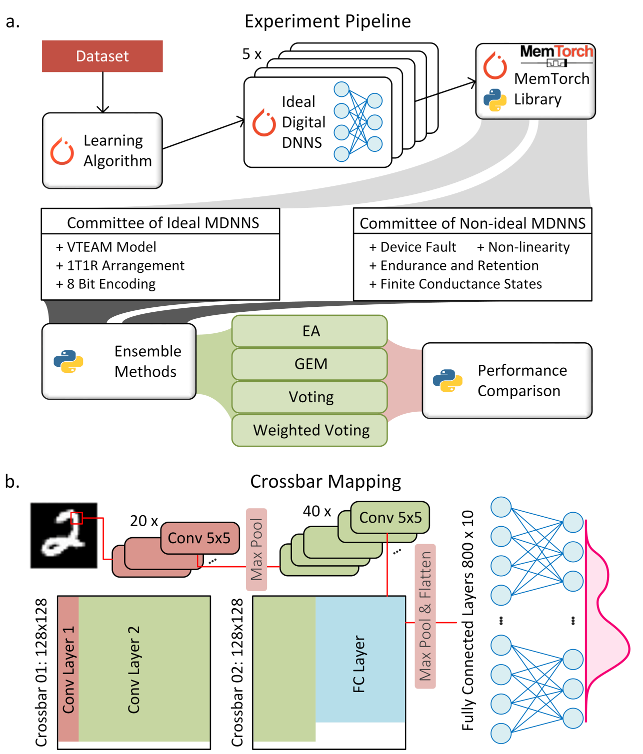

Fig. 2: (a) Detailed overview of the experiment pipeline. (b) Mapping of the neural network layers onto the memristor crossbar.

III. Methodology

In order to simulate the memristive networks, the MemTorch library was used [16]. Firstly, conventional Deep Neural Networks (DNNs) were trained to perform image classification on the MNIST dataset. These DNNs were not state of the art so they achieved relatively standard results, with average test accuracies of around 97%. Stochastic gradient descent was used for optimization, and in order for the ensemble methods to perform more effectively, parameters such as the learning rate and momentum were varied for the 5 DNNs used for testing. The DNNs were then converted to Memristive Deep Neural Networks (MDNNs) using the MemTorch library; the VTEAM model for memristors was used with a time series resolution of 1e-10 [17]. The Conv2d and Linear layers from the DNNs were converted to equivalent memristive layers, with the weights mapped to conductances using the naive mapping scheme - using two crossbars, each of size 128x128 and in 1T1R arrangements, to represent positive and negative weights. Inputs to each layer were encoded between +0.3V and -0.3V, and 8-bit Analogue to Digital Converters (ADCs) with an initial overflow rate of 0 were emulated. These memristive layers were then tuned through linear regression. Finally, non-idealities were introduced in these memristive networks. Non-idealities that were simulated include device faults, device endurance and retention, and finite number of conductance states. Methods presented in this paper do not take into account any specific types of non-idealities and can be used universally to deal with any type of non-ideality, in contrast with hardware based solutions which are only able to deal with the specific type of non-ideality that they are designed for.

TABLE I: Comparison of Ensemble Methods

Ensemble Method	Improvement in Performance (%)
Basic EA	6.3
GEM	6.5
Weighted Voting Algorithm 1	6.4
Weighted Voting Algorithm 2	6.6

To obtain more accurate results from the simulations, each ensemble method was tested 5 times, and the accuracies were averaged over each of the 5 trials. When evaluating the results of the different ensemble methods, the accuracy of the combined ensemble of 5 MDNNs is compared to the average accuracy of the 5 individual MDNNs (this average accuracy of the 5 individual MDNNs represents a single MDNN with non-idealities).

IV. Non-Idealities

Device Faults is a non-ideality in which individual memristors become stuck in high or low resistance states, thus effectively producing open and short circuits respectively (as these states are the high and low extremes) [18]. Since the weights of the neural networks are mapped to conductance values, these faults lead to non-ideal behavior because the memristors are stuck at low and high conductance values and will thus be stuck at low and high weights.

Endurance is a non-ideality that refers to the non-ideal behavior exhibited over time due to the usage of the memristors; since memristors are hardware components, just like any other hardware they eventually will decay with more usage.

Retention is a non-ideality that refers to the non-ideal behavior that arises due to the limited retention time of information in memristors [19]. Essentially, this retention time is determined by how stable the low and high resistance states are (LRS and HRS). If these states are unstable, meaning that over time there is a degradation of these states because of the LRS and HRS going out of their normal ranges, non-ideal behavior is produced because it causes the memristor to be stuck in a specific state, effectively producing a device fault non-ideality. Typically, for retention time, the HRS decreases while the LRS increases, thus making it impossible to switch between these two states.

Finite Conductance States is a non-ideality that arises due to the fact that physical memristors only have quantized, switchable conductance/resistance states, which introduces non-ideal behavior because the resolution of MDNN weights are reduced when mapped [20].

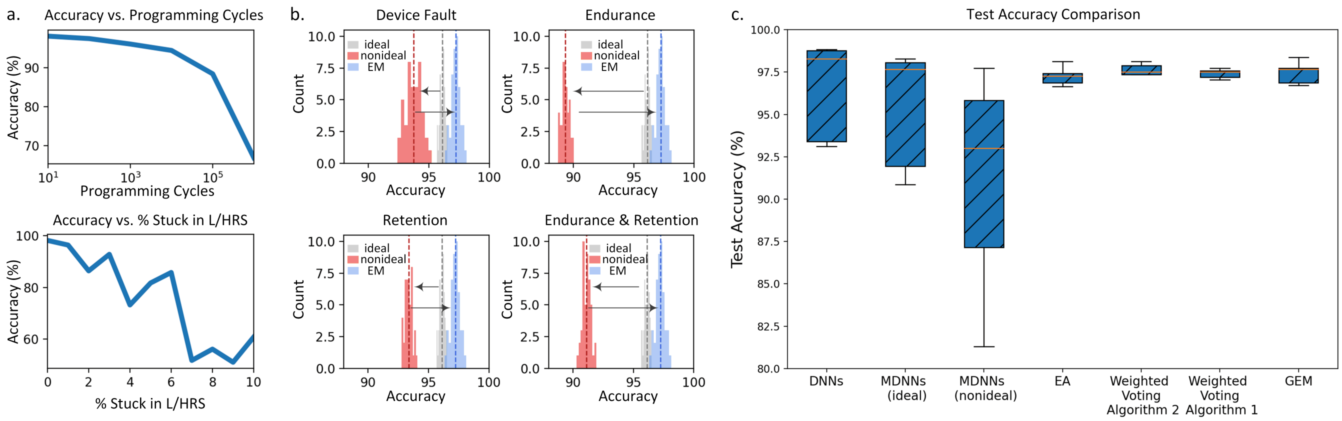

Fig. 3: (a) Variation of accuracy due to non-ideality parameters (fluctuation can be seen due to stochasticity of the non-idealities). (b) Distribution of accuracies over multiple trials, showing improvement due to ensemble methods. (c) Test accuracies of DNNs, ideal MDNNs, nonideal MDNNs, and four ensemble methods.

V. Results and Discussion

Firstly, we will investigate basic ensemble averaging (EA). As can be seen in Table 1, there is approximately a **6.3%** average improvement in performance when 5 MDNNs with non-idealities are combined using basic EA as opposed to a single MDNN with non-idealitites. These results are similar to the results obtained in [13].

For GEM with the ranks weighting scheme, there is an average improvement of approximately **6.5%** when 5 MDNNs with non-idealities are combined as opposed to a single MDNN with non-idealities. An interesting point to note is that this ensemble method performs better when the accuracies of the MDNNs are lower; this trend is somewhat evident in most of the other ensemble methods investigated as well. Moreover, this method performs best when the DNNs (and thus the MDNNs) are ”diverse”, meaning that they have been trained using different parameters (such as learning rate, momentum, etc.) and thus have different accuracies.

Normal, unweighted voting does not achieve good results, so we will only discuss weighted voting using the two weighting schemes discussed earlier. The first one, Algorithm 1, which uses the theoretically optimal weights overall achieves the third highest result, with an improvement in performance of **6.4%**. Although this is still better than basic EA, the reason that this is not much better is likely due to the fact that some of the assumptions made during the derivation of the optimal weights are not true in reality, such as the assumption of the outputs of each individual network being conditionally independent. The second weighting scheme, Algorithm 2, achieves better results overall, with an average improvement in performance of **6.6%**, thus making it the most effective ensemble method. Although the improvement in performance with GEM and Weighted Voting Algorithm 2 is only 0.2% and 0.3% greater than that of Basic EA, respectively, this is statistically significant because the results presented in Table 1 and Figure 3 are based on averages over 45 trials with low variance. Hypothesis testing was passed with $\alpha = 0.05$. The ensemble methods presented in this paper also have some benefits other than simply improving the accuracy. These include things such as reducing the computation time; although the results for such experiments are not presented in this paper, we found that by splitting up data and using multiple non-ideal MDNNs combined using an ensemble method rather than only using one single non-ideal MDNN, in addition to improving the accuracy, the total computation time was also reduced. In terms of the actual hardware implementation of memristors, there could be a trade-off between the improvement in performance and the additional hardware overhead when considering the use of ensemble methods. However, we suspect that this would be insignificant since similar results are obtained if multiple ”smaller” MDNNs (with less layers) are combined using ensemble methods instead of multiple MDNNs of the same size, and thus there will not be a large amount of hardware overhead since the total amount of memristor crossbars used would be the same as or slightly higher than one single ”larger” non-ideal MDNN. This is something that we plan on further investigating in the future, using actual hardware instead of software simulations.

Something worth noting is that although the improvements in performance achieved by the new ensemble methods tested in this paper compared to basic ensemble averaging do not, at face value, seem to have a significant increase just in terms of the % accuracy (see Table 1), this is simply due to the fact that the MNIST dataset was used where the accuracies are already very high. On other harder datasets, the difference would be more significant, but as we mentioned earlier, the results are statistically significant even on the MNIST dataset that we used for testing. The use of the methods presented in this paper to a broader range of applications (not just limited to other image classification datasets, but also completely different neural network architectures) is something that could be investigated further.

VI. Conclusion

In this paper, we tested numerous methods and presented two ensemble methods that outperform basic ensemble averaging as ways of improving the performance of MDNNs with non-idealities. Moreover, these methods work universally as they are not dependent on hardware or specific types of non-idealities; this work demonstrates ensemble methods as a universal way to deal with MDNN non-idealities.

References

[1] D. Strukov, G. Snider, D. Stewart, and S. Williams, "The missing memristor found," *Nature*, vol. 453, pp. 80–3, 06 2008.

[2] G. Adam, A. Khiat, and T. Prodromakis, "Challenges hindering memristive neuromorphic hardware from going mainstream," *Nature Communications*, vol. 9, 12 2018.

[3] K. Van Pham and K.-S. Min, "Non-ideal effects of memristor-cmos hybrid circuits for realizing multiple-layer neural networks," in *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2019, pp. 1–5.

[4] S. M. Laube and N. Taherinejad, "Device variability analysis for memristive material implication," *CoRR*, vol. abs/2101.07231, 2021. [Online]. Available: https://arxiv.org/abs/2101.07231

[5] S. Stavrinides, M. Hanias, M. Gonzalez, F. Campabadal, Y. Contoyiannis, S. POTIRAKIS, M. M. Al Chawa, C. Benito, R. Tetzlaff, R. Picos, and L. Chua, "On the chaotic nature of random telegraph noise in unipolar rram memristor devices," *Chaos, Solitons Fractals*, vol. 160, p. 112224, 07 2022.

[6] A. Amirsoleimani, T. Liu, F. Alibart, S. Eccofey, Y.-F. Chang, D. Drouin, and R. Genov, "Mitigating state-drift in memristor crossbar arrays for vector matrix multiplication," in *Memristor*, Y.-F. Chang, Ed. Rijeka: IntechOpen, 2021, ch. 4. [Online]. Available: https://doi.org/10.5772/intechopen.100246

[7] O. Krestinskaya, A. Irmanova, and A. James, "Memristive non-idealities: Is there any practical implications for designing neural network chips?" 05 2019.

[8] J. Woo, K. Moon, J. Song, S. Lee, M. Kwak, J. Park, and H. Hwang, "Improved synaptic behavior under identical pulses using alox/hfo2 bilayer rram array for neuromorphic systems," *IEEE Electron Device Letters*, vol. 37, no. 8, pp. 994–997, 2016.

[9] C. Li, M. Hu, Y. Li, H. Jiang, N. Ge, E. Montgomery, J. Zhang, W. Song, N. Dávila, C. E. Graves, Z. Li, J. P. Strachan, P. Lin, Z. Wang, M. D. Barnell, Q. Wu, R. S. Williams, J. J. Yang, and Q. Xia, "Analogue signal and image processing with large memristor crossbars," *Nature Electronics*, vol. 1, pp. 52–59, 2018.

[10] Y. Fang, Z. Yu, Z. Wang, T. Zhang, Y. Yang, Y. Cai, and R. Huang, "Improvement of hfox-based rram device variation by inserting ald tin buffer layer," *IEEE Electron Device Letters*, vol. 39, no. 6, pp. 819–822, 2018.

[11] A. Kenyon, W. Ng, M. Munde, M. Buckwell, D. Joksas, and A. Mehonic, "The interplay between structure and function in redox-based resistance switching," *Faraday Discussions*, vol. 213, 07 2018.

[12] W. Wu, H. Wu, B. Gao, P. Yao, X. Zhang, X. Peng, S. Yu, and H. Qian, "A methodology to improve linearity of analog rram for neuromorphic computing," in *2018 IEEE Symposium on VLSI Technology*, 2018, pp. 103–104.

[13] D. Joksas, P. Freitas, Z. Chai, W. Ng, M. Buckwell, C. Li, W. Zhang, Q. Xia, A. Kenyon, and A. Mehonic, "Committee machines—a universal method to deal with non-idealities in memristor-based neural networks," *Nature Communications*, vol. 11, 08 2020.

[14] V. Tresp, *Committee machines*, 01 2001, pp. 5–1.

[15] Z.-H. Zhou, *Ensemble Methods: Foundations and Algorithms*, 2012.

[16] C. Lammie, W. Xiang, B. Linares-Barranco, and M. Rahimi Azghadi, "Memtorch: An open-source simulation framework for memristive deep learning systems," *Neurocomputing*, vol. 485, pp. 124–133, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0925231222002053

[17] S. Kvatinsky, M. Ramadan, E. G. Friedman, and A. Kolodny, "Vteam: A general model for voltage-controlled memristors," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 62, no. 8, pp. 786–790, 2015.

[18] J. I. Canales-Verdial, W. Woods, C. Teuscher, M. Osinski, and P. Zarkesh-Ha, "Impact of memristor defects in a neuromorphic radionuclide identification system," in *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2020, pp. 1–5.

[19] A. Fadeev and K. Rudenko, "To the issue of the memristor's hrs and lrs states degradation and data retention time," *Russian Microelectronics*, vol. 50, pp. 311–325, 09 2021.

[20] Z. Guo, Y. Zhang, S. Xu, Z. Wu, and W. Zhao, "A multi-conductance states memristor-based cnn circuit using quantization method for digital recognition," in *2021 IEEE 14th International Conference on ASIC (ASICON)*, 2021, pp. 1–4.