Fault Sensitivity Analysis of Printed Bespoke Multilayer Perceptron Classifiers

Priyanjana Pal*, Florentia Afentaki*†, Haibin Zhao*, Gurol Saglam*, Michael Hefenbrock‡,
Georgios Zervakis†, Michael Beigl*, Mehdi B. Tahoori*
*Karlsruhe Institute of Technology, DE, †University of Patras, GR, ‡RevoAI GmbH
*{priyanjana.pal, haibin.zhao, guerol.saglam, michael.beigl, mehdi.tahoori}@kit.edu,
†{afentaki, zervakis}@ceid.upatras.gr, ‡michael.hefenbrock@revoai.de,

***Abstract*—Printed Electronics (PE) is an emerging technology with flexible substrates and ultra-low-cost manufacturing, providing an appealing alternative to traditional wafer-scale silicon fabrication. With the increasing integration of various printed neural network (NN) architectures in diverse applications, the reliability of printed circuits has become a critical concern. This work provides a comprehensive analysis of the fault sensitivity on a variety of classification tasks for various digital and analog realizations of printed multilayer perceptrons (MLPs). We further evaluate different digital architectures, i.e., generic, bespoke, and approximate, to provide a comprehensive fault analysis on different benchmark datasets.**

I. Introduction

Printed Electronics (PE) has gained increasing attention in recent years, driven by its attractive benefits such as mechanical flexibility, cost-effectiveness, customizability, and on-demand fabrication. These advantages make PE suitable in various domains, including the Internet of Things (IoT), wearables, RFID tags, smart cards, smart labels, and smart sensors [1]. Functional PE circuits, incorporating both organic and inorganic printed transistors, have been developed through additive manufacturing processes that employ maskless and portable methods. The technology's key feature lies in its bespoke application-specific customization for both large and small quantities, owing to the cost-effective nature of printing compared to lithography processes.

To enable basic sensor processing tasks like classification in printed devices, incorporating foundational capacity requires various printed computing circuits [2, 3]. Printed analog neuromorphic circuits, utilizing resistor crossbars and inverter-based activation function circuitry, emulate artificial neural network (ANN) operations while directly operating on analog sensory inputs, eliminating costly analog to digital converters [3]. However, the digital realization of such a classifier can provide better noise immunity compared to analog counterparts, however faces major challenges due to large device counts and increased power consumption. To mitigate these limitations in digital approach, bespoke implementations can be exploited [4]–[6] in which the hardware is customized to the specific dataset and combines it with paradigms like Approximate Computing (AxC) [4, 5] or Stochastic Computing (SC) [6] in order to produce area and power adequate MLP circuits.

The main advantage of additive printing processes, the ultra-low-cost fabrication, comes from a simplified process. The downside, however, is the reduced control and resolution, and hence, increased process variability, which can lead to higher defectivity of the low-cost additive printing processes compared to high-precision lithography-based processes as used for silicon VLSI. Meanwhile, bespoke architectures used in the realization of printed MLPs, together with analog computing or aggressive digital approximation can have severe impacts on the fault sensitivity of printed MLPs. While the fault sensitivity of generic (model-agnostic) neural network (NN) hardware accelerators are extensively investigated [7], however, it is not explored for bespoke MLP hardware architectures, as commonly used in PE.

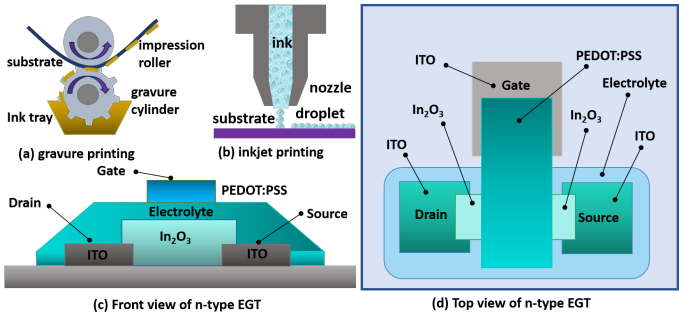

Fig. 1. Schematic of (a) gravure printing (b) inkjet printing; (c) front view and (d) top view of a printed N-type EGT.

Although significant efforts have been made to implement various printed bespoke machine learning (ML) classifiers, very few studies have yet been reported on fault sensitivity analysis with its printed design and implementation, ensuring the reliability and correct operation in diverse applications [8]. In this work, we provide a comprehensive analysis of fault injection in both printed analog multilayer perceptrons (p-AMLPs) and printed digital multilayer perceptrons (p-DMLPs) circuits. For the p-AMLPs, we propose a Monte Carlo injection of stuck-open and stuck-short transistor and resistor faults. For the p-DMLPs, we evaluate the impact of stuck-at-fault sensitivity on digital architecture with different customization levels; generic, exact, and bespoke designs while different levels of approximation were also considered. In short, the contributions of this work are as follows:

1) This is the first work that models the stuck-open and stuck-short faults on N-type printed electrolyte-gated transistor (EGT), and resistors considering Monte Carlo fault injection in p-AMLP.

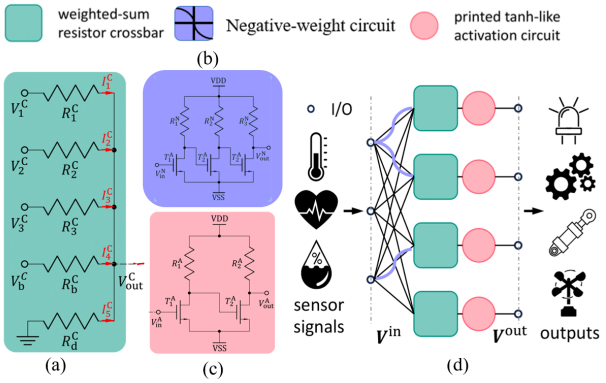

Fig. 2. Circuit primitives of the p-AMLP. (a) Example of a 3-input, 1-output printed resistor crossbar.(b) Inverter-based negative weight circuit. (c) Inverter-based printed tanh-like (ptanh) circuit.(d) Schematic of a p-AMLP [8] that receives sensor signals and yields outputs to subsequent devices.

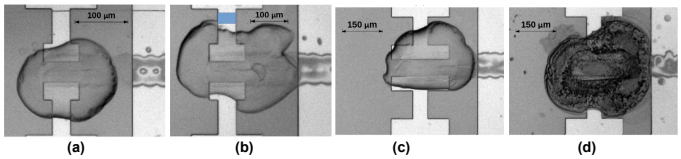

Fig. 3. Parametric Faults in: (a)-(b) Two transistors of width 100 µm with different appearances due to variation. (c) Functioning (nondefective) transistor compared to (d) transistor with exploded electrolyte (sourced from [9]).

2) This is the first time that a fault sensitivity analysis is conducted upon p-DMLP architectures with different levels of customization and approximation in PE.
3) We study the effectiveness of dropout in evaluating the robustness of p-AMLPs and p-DMLPs against faults.

The rest of this paper is structured as follows: Sec. II, provides the background of this work. Sec. III proposes the methodology of fault injection in MLP architectures. In Sec. IV, we evaluate the effectiveness of the fault sensitivity analysis on benchmark datasets and compare our results with different architectures. Finally, Sec. V concludes this paper.

II. Preliminaries

A. Printed Electronics

PE is an emerging technology which has achieved a $41.2 billion market in 2019, with a projected growth of $74 billion by 2030 [1], enabling diverse applications, including photovoltaics, RFID tags, sensor arrays, memory, displays, batteries, smart packaging, smart bandages, energy harvesting, and random number generators in the IoT applications [1].

PE employs an additive manufacturing approach, depositing materials layer by layer for active devices and interconnecting components. compared to traditional silicon electronics, it requires fewer steps and cost-effective processes. PE accommodates diverse materials for flexible and biocompatible electronics. It also involves various fabrication techniques, including high-volume replication (e.g., Fig. 1(a) gravure printing) and customized jet printing (e.g., Fig. 1(b) inkjet printing) [1]. Organic inkjet-printed FETs use structured semiconductors with higher voltages, while inorganic FETs operate at sub-1 V, promising low-power applications. A typical N-type EGT, as shown in Fig. 1(c) and Fig. 1(d), uses indium oxide (In_2O_3) as the semiconductor channel and a solid polymer electrolyte as the gate dielectric [10].

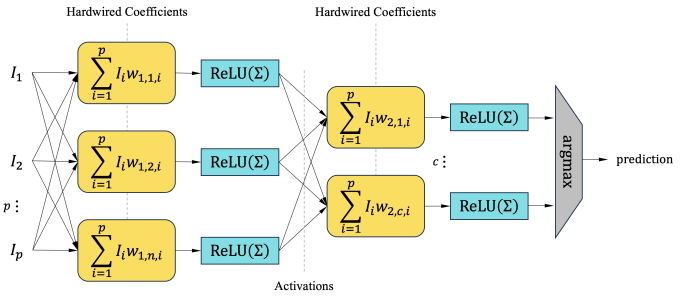

Fig. 4. Bespoke architecture of p-DMLP [13].

B. Printed Analog Multilayer Perceptrons (p-AMLP)

p-AMLPs receive substantial attention, as they can directly operate sensory data in the analog domain and thus require significantly reduced hardware footprints compared to digital counterparts. Meanwhile, the p-AMLPs targeted in the state-of-the-art [8, 11] is composed only of simple operations like crossbar-arrays for weighted-sum, analog inverters for negative weights and printed tanh for nonlinear activations as shown in Fig. 2. In this work, we adopt multiple fault injection training methodologies i.e., nominal training, variation-aware [8], dropout [12], and both (variation-aware + dropout) training to provide a comprehensive fault sensitivity analysis and evaluate the classification accuracy on various benchmark datasets.

C. Printed Digital Multilayer Perceptrons (p-DMLP)

Low-cost embedded ML systems are typically task-specific, prompting the need for model-specific ML circuits to overcome limitations. The work in [3] demonstrated the customization potential of low-cost printed circuits, by designing bespoke ML circuits, i.e. designs highly tailored to the specific dataset and model. Although the achieved gains were immense, the resulting circuit overheads were still inadequate for printed applications [14].

The AxC paradigm reduces hardware overheads at the expense of accuracy. In [14] AxC was introduced in printed MLP classifiers, through post-training weight approximation and gate-pruning. However, [14] acknowledged the prohibitive hardware overheads of MLPs in printed circuits, emphasizing the ongoing need for additional optimizations. In [4], the combination of the bespoke architecture with the AxC was fully exploited by approximating all the core components of the MLP; multipliers, accumulators, and activation function. These holistic approximate bespoke MLP circuits achieved remarkable area and power savings that proved sufficient for the realization of complex MLP classifiers w.r.t. both area and power overheads. In this work, we adopt various digital architectures (p-DMLP) to provide a comprehensive fault sensitivity analysis w.r.t. different levels of approximations and also different levels of customization. To this end, this work adopts baseline MLP with the identical topology as described in [3] and injects stuck-at-faults into the designs. As shown in Fig. 4, for the design of the corresponding p-DMLP circuit, either approximate or accurate, we employ the efficiency of bespoke design paradigm [3, 13]. In addition, we follow the different approximation architectures of [4].

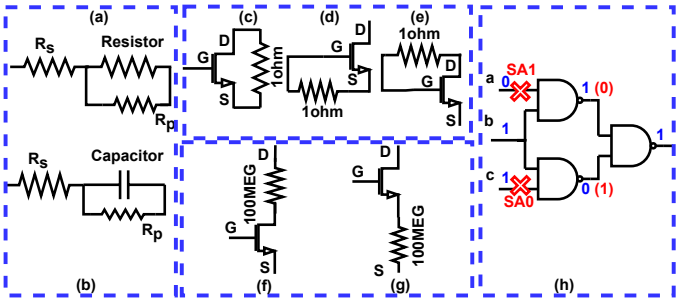

Fig. 5. Catastrophic Fault Models in: (a) Resistor and (b) Capacitor; (c, d, e) Stuck-short and (f, g) Stuck-open in N-type EGT; (h) Stuck-at-zero/Stuck-at-one in logic cells

D. Fault Models in Analog and Digital Printed Circuits

Circuit faults, arising from defective components, signal line breaks, short circuits, and delays, impair their correct functionality. Fig. 3 shows the evidence of defects that can occur typically in an n-type EGT and the inkjet printing process [9].

Typically, these faults are represented using a model that captures the alterations induced in circuit signals. Faults that happen in any analog circuits can be categorized into permanent faults, which are observable during testing due to their prolonged existence, and temporary faults, appearing and disappearing rapidly. Additionally, delay faults also impact the operating speed of the circuit.

1) Catastrophic Faults: Catastrophic faults are categorized into stuck-open and stuck-short faults. Stuck-open faults, also called hard faults, occur when component terminals lose contact with the circuit, creating high resistance. Simulating stuck-open faults involves adding a high series resistance, e.g., $(R_s = 100\,\mathrm{M}\Omega)$, to the faulty component. Conversely, stuck-short faults involve a short between component terminals, effectively bypassing the component. Simulating stuck-short faults includes connecting a small parallel resistor, e.g., $(R_p = 1\,\Omega)$ to the component. Both faults can be simulated in resistors, capacitors, or transistors as shown in Fig. 5. While catastrophic faults often alter circuit functionality, some only affect circuit specifications without impacting overall operation [15].

2) Parametric Faults: Parametric faults, impacting only parameter values like resistors, capacitors, and transistor values, arise from local or global defects. Global parametric faults are caused due to imperfections in manufacturing processes, affecting all components during production. Conversely, local parametric faults result from specific defect mechanisms, such as particles that affect a transistor's channel length.

Similarly, digital circuits also experience various types of faults, each posing unique challenges to the system's functionality. In this work, only stuck-at-faults are used as fault models for digital MLP classifiers, as it is arguably the most critical reason for circuit failure.

III. Evaluation Framework

A. Fault Injection in p-AMLP

To assess the fault sensitivity of p-AMLP circuits, a Monte Carlo-based fault injection methodology is employed. The first step in this work involves defining potential fault types, such as resistor open-short or transistor open or gate-drain (G-D) / gate-source (G-S) / drain-source (D-S) shorts-open, specifying their characteristic behaviors. For simulating the faults in the resistor crossbars, we utilized an analytical expression, i.e.,

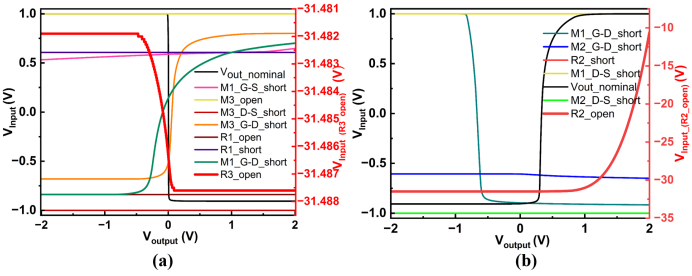

Fig. 6. SPICE simulated stuck-open and stuck-short faults in (a) negative-weight and (b) p-tanh activation circuit of p-AMLP hardware primitives.

$$\tilde{\boldsymbol{g}} = M_g \odot \boldsymbol{g}, \tag{1}$$

where $\boldsymbol{g}$ collects the conductance values and M_g functions as a mask to emulate by different faults by multiplying the conductances with 1 for fault-free, 0 for open, and ∞ for short. Moreover, $\odot$ denotes an element-wise multiplication, and therefore, $\tilde{\boldsymbol{g}}$ indicates the conductances with fault injection.

However, the impact of the fault in the nonlinear circuits, i.e., negative weight circuit and activation circuit, is a much more sophisticated mechanism. Even though there has been works on estimating the transfer characteristics of those nonlinear circuits through ML-based *surrogate nonlinear circuit* models [16], they can not be used for fault sensitivity analysis, they can only provide confident estimations when the component values are within a reasonable range. Therefore, we consider the nonlinear subcircuits as sub-systems and inject faults into different components. Thereafter, we conduct a SPICE simulation based on the pPDK to obtain the characteristic curves for the corresponding faults. Fig. 6 (a) and Fig. 6 (b) show the possible single fault that can occur in the negative weight circuit and ptanh circuit. In addition to being stuck at VSS/GND and stuck at supply voltage VDD, we also notice various unexpected faulty behaviors in the circuit which results in an abnormal rise in the negative voltage level when the resistor R_3 is kept open in Fig. 6 (a) and R_2 in Fig. 6 (b) respectively due to a discontinuity in the closed circuit connection, further disrupting the expected outputs. All these fault scenarios lead to the reduced classification accuracy of various datasets in the p-AMLP architecture and are therefore a critical concern.

Subsequently, we use Monte Carlo sampling to draw the fault components within a print p-AMLP. As the training framework of the p-AMLP is an ML-based model, we employ a top-down sampling strategy to keep a consistent fault rate for each component, i.e., we first draw N_l faults for each layer l with the probability

$$p(l) = \frac{N_l}{\sum_l N_l}, \forall l, \tag{2}$$

with

$$N_l = 6N_l^N + 4N_l^A + N_l^R, \tag{3}$$

where N_l^N, N_l^A, and N_l^N refer to the number of negative weight circuits, ptanh activation circuits, and the number of resistors of the crossbar in the l-th layer. Afterward, each layer is requested to sample N_l faults for three circuit primitives

proportional to their device counts. Note that, for each non-linear circuit, we consider maximally one fault, justifying the ignorance of multiple faults occurring simultaneously in the same subcircuit.

B. Different training approaches in p-AMLP

1) Nominal Training: We adopt nominal training as a design approach which is gradient-based training of p-AMLP [8]. It can efficiently train the parameters in a p-AMLP, fulfilling the constraints on the printed devices, e.g., limited printable resistances.

2) Variation-Aware Training: In addition to the nominal training, variation-aware training takes the parametric faults (manufacturing errors) of printed components into account during the training [8] by modeling the fabrication through a stochastic variable. As it aims to improve the robustness against parametric faults, we tested the p-AMLPs from variation-aware training to analyze whether it can bring advantages for catastrophic faults.

3) Dropout: Dropout is a common training method in deep learning [12]. As it randomly turns off some neurons/inputs, which is similar to the mechanism of a stuck-open, it is hypothesized that it can also improve the robustness of the circuits against catastrophic faults.

4) Variation-aware+Dropout: Finally, a combined approach of variation-aware training and dropout is considered to comprehensively analyze their effectiveness in enhancing the performance of p-AMLP.

C. Fault Injection in p-DMLP

Similar as explained in Sec. III-A, to assess a fault sensitivity simulation on the digital designs a Monte Carlo-based simulation is utilized. In this process, we obtain all the wires described in the post-synthesis netlist design, and randomly wires are selected. These obtained wires are the faulty wires of the circuit and are randomly stuck-at 0 or 1. Afterward, every defective wire is intentionally introduced into the gate-level netlist for each respective architecture. Then the gate-level simulation is performed to evaluate the impact of these faults on the classification accuracy of the MLP. Also, we simulate the entire test dataset on the fault-injected netlist to obtain the impact on the classification accuracy.

D. Architectures in p-DMLP

1) Different customization architecture: We adopt two architectures of different customization levels; the conventional (i.e., generic) and the bespoke architecture. Both of these architectures use full fixed-point precision for their computations. The conventional architecture utilizes general-purpose multipliers with two operands; one for the input and one for the weights. On the other hand, the bespoke design paradigm utilizes multipliers specific for each weight, in the MLP circuit. Subsequently, each multiplier produces a specific product based on the input for each weight within each neuron of the MLP.

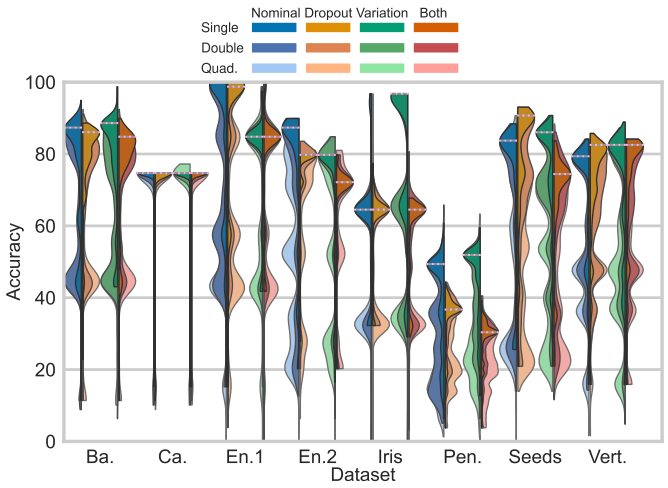

Fig. 7. Evaluation of various analog design approaches after single, double, and quadruple fault injection in p-AMLP. The dotted line represents the fault-free accuracy.

2) Different approximation tensity: We consider three different bespoke approximate MLP architectures with different approximation tensity. The first approximate MLP architecture uses only power-of-2 (pow2) quantization of weights and thus, performs a multiplier-less inference. The second approximate MLP architecture uses pow2 approximation and a fine-grain approximate accumulation (pow2+axAcc) [4]. The third MLP architecture (axAll) [4], approximates all the components of the MLP by using an argmax approximation with pow2+axAcc.

3) Dropout: As mentioned before, dropout is a common training method in ML. This approach randomly nullifies some inputs during training, which in turn helps prevent overfitting. After training, certain inputs have larger weights that dramatically change the outcome. Without dropout, the circuit is expected to be more prone to failure, if a stuck-at-fault occurs which could affect the outcome more.

IV. Fault Sensitivity Analysis

To analyze the fault sensitivity of both p-AMLP and p-DMLP, we conduct the following experiment. The code is available at GitHub repository[1].

A. Experiment Setup

1) Simulation Setup: All analog hardware primitives were designed based on the n-EGT pPDK. We used Cadence Virtuoso[2] tool to simulate fault injection (Fig. 6) in SPICE and trained the p-AMLPs afterward.

Digital circuits are synthesized using Synopsys Design Compiler S-2021.06 with the printed n-EGT library [17]. For simulation and power analysis, VCS T-2022.06 and PrimeTime T-2022.03 are employed. Accuracy is reported on the test dataset, with synthesis adhering to directives in [4] to align delay values with typical PE performance [18]. The p-DMLPs architecture mirrors [3, 11], and bespoke exact p-DMLPs circuits, as outlined in [3], utilizes 8-bit fixed point weights and 4-bit inputs.

[1]https://github.com/PrintedElectronics/Fault_Sensitivity_Analysis
[2]https://www.cadence.com/en_US/home.html

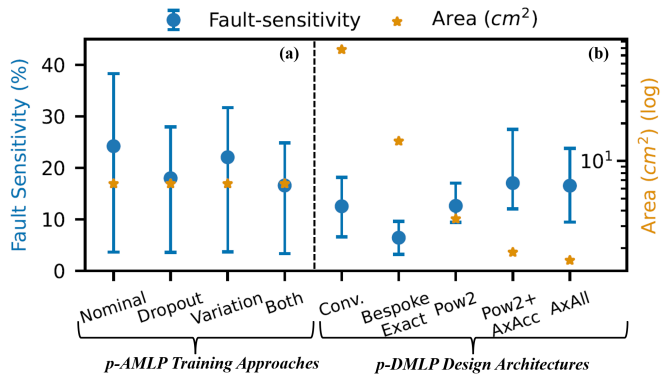

Fig. 8. Evaluation of fault sensitivity w.r.t (a) design approaches in p-AMLP (The area for p-AMLP is same in all training process) and (b) design architectures in p-DMLP. (Area is in logarithmic scale).

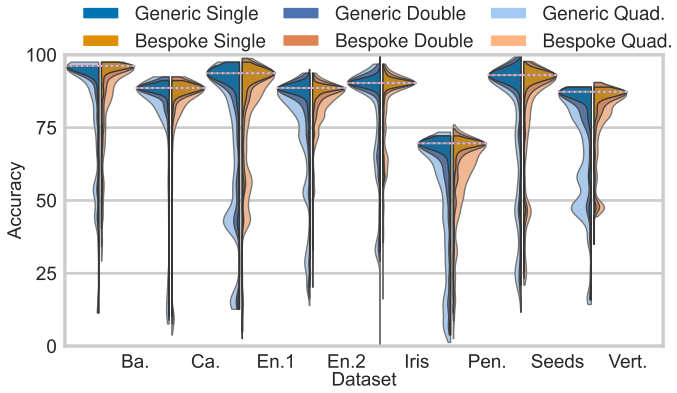

Fig. 9. Evaluation of conventional and bespoke architecture after single, double, and quadruple fault injection. The bespoke architecture is as in [3]. The dotted line represents the fault-free accuracy.

2) Training Setup: We employed eight preprocessed datasets, namely Balance Scale (Ba.), Cardiotocography (Ca.), Energy1 (En1.), Energy2 (En2.), Iris, Pendigits (Pen.), Seeds, and Vertebral 3 Columns (Vert.), from [19] whose task complexity match the target applications of PE, and normalized their inputs to $[0, 1]$ to simulate the electrical signals from sensors with limited range. To this, we randomly split datasets into training (60%), validation (20%), and test (20%) sets.

B. Evaluation Results of p-AMLPs

We utilized various non-linear components in the p-AMLP circuit primitives. However, their reliable functioning in analog circuits is highly susceptible to faults due to several factors like input variations, printing process variations, device-geometry, and variations in ink compositions and substrates. The accuracy distribution with single, double and quadruple fault injection of p-AMLPs on 8 benchmark datasets using various training approaches are shown in Fig. 7. It is evident that none of the methods can effectively mitigate the impact of catastrophic faults, resulting in a significant reduction in classification accuracy when faults are introduced. Additionally, with the increasing number of faults injected, the classification accuracy is reduced continuously. Both variation-aware training and dropout training can marginally contribute to robustness against fault and their effectiveness is dataset-dependent. However, variation-aware training stands out by achieving higher accuracy in fault-free testing scenarios, while dropout training yields lower accuracy under fault-free conditions. We speculate that, as variation-aware training introduces stochastic variables with continuous probabilistic distribution, the objective function was smoothed during training. As a result, the gradient can guide the training process more informatively. In contrast, as dropout introduces non-differentiability (turning-off vs. turning-on) into the training process, it creates difficulties for the gradient-based training process resulting in lower accuracies in fault-free cases.

Unexpectedly, the combination of both dropout and variation-aware training leads to even worse results than nominal training. In this regard, we argue that the introduction of both variation and dropout forces excessive perturbations to the parameters, so that the circuits are unable to produce promising accuracy while overcoming the large perturbations in the small optimization search space (solely 1 hidden layer with 3 neurons). Thus, p-AMLPs inherently operates with limited precision due to continuous voltage levels resulting in reduced classification accuracy for different training approaches.

Fig. 8 depicts the summary of techniques used in this work, for both analog and digital MLPs w.r.t. fault sensitivity and area. For each technique, Fig. 8 presents the min-max range of fault sensitivity, the blue dot represents the average of this range, while the yellow star is the area of the technique in cm^2. In this work, as fault sensitivity, we consider the average accuracy drop (classification miss) for each dataset and each fault injection type i.e., single, double, quadruple, normalized with the respective fault-free accuracy.

C. Evaluation Results of p-DMLPs

Fig. 9 depicts the accuracy distribution with single, double, and quadruple fault injection for both conventional and bespoke designs. The same distribution is shown for bespoke exact and approximate architectures in Fig. 10. As discussed in Sec. IV-B, increasing the number of faults in the circuit results in a reduction of classification accuracy. In some cases, the approximate MLP circuits achieve higher accuracy than their exact counterparts due to the benefits of approximation in terms of area, power gains, and improved generalization [20].

For different customization architectures, Fig. 8 showcases that the bespoke architecture is more fault-tolerant than the conventional model-agnostic architecture. As previously mentioned, the bespoke architecture uses fully customized multipliers that generate a specific product, while its conventional counterpart utilizes general-purpose multipliers with weight as one of its inputs. That means that if a fault occurs in the netlist of the conventional MLP architecture, the multiplier will result in a different input's product. On the other hand, the bespoke netlist is more fault-tolerant since the multiplier's netlist is dedicated to the specific value of the weight.

Furthermore for different approximation tensity, Fig. 8 showcases that pow2 is more fault-tolerant than the other more fine-grain approximations. More precisely, while approximate circuits are generally assumed to be more fault tolerant than their exact counterpart [20], this is not the case for bespoke circuits as depicted in Fig. 8. The approximate bespoke designs are more fault-sensitive than their bespoke exact counterparts. The aggressive fine-grain approximations in these designs result in significant area reduction, minimizing the gate count, meaning that any remaining redundancy is removed from the circuit

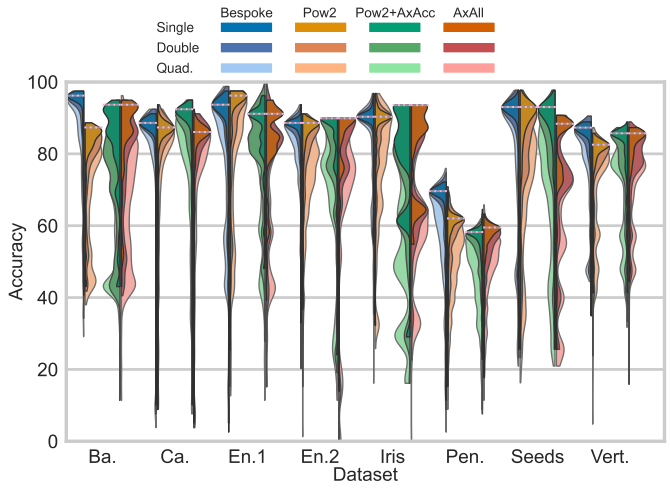

Fig. 10. Evaluation of approximate architecture with single, double, and quadruple fault injection. All the architectures follow [4]. The dotted line represents the fault-free accuracy.

and every gate is crucial for final classification. The decreasing area in Fig. 8 demonstrates a trade-off between optimization intensity and fault sensitivity, emphasizing the criticality of even a simple permanent fault in highly optimized circuits.

Analyzing gate-level netlists post single fault injection, also reveals that errors near the output (argmax circuit) significantly degraded accuracy. A stuck-at-fault within the argmax circuit propagates to the output, causing one MLP output bit to be stuck-at 0/1. Consequently, the classification output becomes confined to a specific subset of available MLP classes, emphasizing the criticality of faults in proximity to the output for accuracy degradation.

D. Evaluation Results of Dropout

Fig.11 shows the fault sensitivity for p-AMLP nominal, dropout, p-DMLP AxAll, and AxAll+dropout approaches. The p-AMLP operates on analog sensory inputs, and for a fair analog-digital area comparison, p-DMLP areas are reported with ADC areas included. The digital area is nearly half of its analog equivalent due to a more approximate design with all p-DMLP components approximated. Dropout during p-DMLP training alters model coefficients and circuits, affecting min-max fault-tolerance ranges on AxAll dropout. However, on average, AxAll p-DMLP dropout is more fault-tolerant than without dropout during training. Overall, dropout increases the fault tolerance in both p-AMLPs and p-DMLPs on average.

V. Conclusion

In this work, we provided the fault sensitivity analysis of printed MLP classifiers. Specifically, we evaluated analog approaches and customized digital printed MLP architectures and evaluated both the classification accuracy drop and fault sensitivity on 8 benchmark datasets for single and multiple faults. We conclude that the bespoke digital architectures had higher fault tolerance compared to other different architectures i.e. conventional, bespoke, and approximate MLP circuits, while dropout improves fault tolerance on average in both analog and digital MLPs. Also, the thorough analysis in p-AMLPs reveals that other methodologies for efficient fault-tolerant MLP circuit design and fault-aware training strategies are required to be explored in future research.

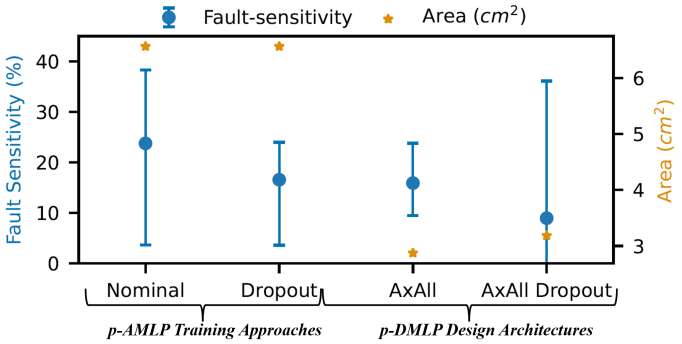

Fig. 11. Evaluation of fault sensitivity using dropout for both p-AMLP and p-DMLP. (The area for p-AMLP is same in all training process)

Acknowledgment

This work is partially supported by the European Research Council (ERC) and co-funded by the H.F.R.I call "Basic Research Financing (Horizontal support of all Sciences)" under the National Recovery and Resilience Plan "Greece 2.0" (H.F.R.I. Project Number: 17048) and Carl-Zeiss-Foundation as part of "stay young with robots" (JuBot).

References

[1] M. A. Leenen *et al.*, "Printable electronics: Flexibility for the future," *physica status solidi (a)*, vol. 206, no. 4, pp. 588–597, 2009.
[2] P. Pal *et al.*, "Analog printed spiking neuromorphic circuit," in *IEEE DATE*, 2024, p. 6 S.
[3] M. H. Mubarik *et al.*, "Printed machine learning classifiers," in *2020 53rd Annual IEEE/ACM MICRO*, 2020, pp. 73–87.
[4] F. Afentaki *et al.*, "Bespoke approximation of multiplication-accumulation and activation targeting printed multilayer perceptrons," in *2023 IEEE/ACM (ICCAD)*, 2023, pp. 1–9.
[5] G. Armeniakos *et al.*, "Co-design of approximate multilayer perceptron for ultra-resource constrained printed circuits," *IEEE Trans. on Computers*, vol. 72, no. 9, pp. 2717–2725, 2023.
[6] D. D. Weller *et al.*, "Printed stochastic computing neural networks," in *2021 DATE*, 2021, pp. 914–919.
[7] F. Su, C. Liu, and H.-G. Stratigopoulos, "Testability and dependability of ai hardware: Survey, trends, challenges, and perspectives," *IEEE Design & Test*, 2023.
[8] H. Zhao *et al.*, "Highly-dependable printed neuromorphic circuits based on additive manufacturing," *Flexible and Printed Electronics*, vol. 8, no. 2, p. 025018, jun 2023.
[9] A. T. Erozan *et al.*, "Defect detection in transparent printed electronics using learning-based optical inspection," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 29, no. 8, pp. 1505–1517, 2021.
[10] Chang *et al.*, "A circuits and systems perspective of organic/printed electronics: Review, challenges, and contemporary and emerging design approaches," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 7, no. 1, pp. 7–26, 2017.
[11] H. Zhao *et al.*, "Power-aware training for energy-efficient printed neuromorphic circuits," in *2023 IEEE/ACM ICCAD*, 2023, pp. 1–9.
[12] Tingting *et al.*, "Improved convolutional neural network fault diagnosis method based on dropout," in *2020 7th IFEEA*, 2020, pp. 753–758.
[13] G. Armeniakos *et al.*, "Model-to-circuit cross-approximation for printed machine learning classifiers," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 11, pp. 3532–3544, 2023.
[14] G. Armeniakos *et al.*, "Cross-layer approximation for printed machine learning circuits," in *2022 DATE*, 03 2022, pp. 190–195.
[15] L. Milor *et al.*, "Detection of catastrophic faults in analog integrated circuits," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 8, no. 2, pp. 114–130, 1989.
[16] H. Zhao *et al.*, "Highly-bespoke robust printed neuromorphic circuits," in *2023 DATE*. IEEE, 2023, pp. 1–6.
[17] N. Bleier *et al.*, "Printed microprocessors," in *Annu. Int. Symp. Computer Architecture (ISCA)*, jun 2020, pp. 213–226.
[18] C. Marques *et al.*, "Digital power and performance analysis of inkjet printed ring oscillators based on electrolyte-gated oxide electronics," *Applied Physics Letters*, vol. 111, no. 10, p. 102103, 2017.
[19] M. Fernández-Delgado *et al.*, "Do We Need Hundreds of Classifiers To Solve Real World Classification Problems?" *The journal of machine learning research*, vol. 15, no. 1, pp. 3133–3181, 2014.
[20] M. Ahmadilivani *et al.*, "Special session: Approximation and fault resiliency of dnn accelerators," in *2023 IEEE 41st (VTS)*, 2023, pp. 1–10.