



Evolutionary Approximation of Ternary Neurons for On-sensor Printed Neural Networks

Vojtech Mrazek
Brno University of Technology
Brno, Czech Republic
mrazek@fit.vutbr.cz

Argyris Kokkinis
Aristotle University of Thessaloniki
Thessaloniki, Greece
arkokkin@auth.gr

Panagiotis Papanikolaou
University of Michigan
Ann Arbor, USA
panagip@umich.edu

Zdenek Vasicek
Brno University of Technology
Brno, Czech Republic
vasicek@fit.vutbr.cz

Kostas Siozios
Aristotle University of Thessaloniki
Thessaloniki, Greece
ksiop@auth.gr

Georgios Tzimpragos
University of Michigan
Ann Arbor, USA
gtzimpra@umich.edu

Mehdi Tahoori
Karlsruhe Institute of Technology
Karlsruhe, Germany
mehdi.tahoori@kit.edu

Georgios Zervakis
University of Patras
Patras, Greece
zervakis@upatras.gr

ABSTRACT

Printed electronics offer ultra-low manufacturing costs and the potential for on-demand fabrication of flexible hardware. However, significant intrinsic constraints stemming from their large feature sizes and low integration density pose design challenges that hinder their practicality. In this work, we conduct a holistic exploration of printed neural network accelerators, starting from the analog-to-digital interface—a major area and power sink for sensor processing applications—and extending to networks of ternary neurons and their implementation. We propose bespoke ternary neural networks using approximate popcount and popcount-compare units, developed through a multi-phase evolutionary optimization approach and interfaced with sensors via customizable analog-to-binary converters. Our evaluation results show that the presented designs outperform the state of the art, achieving at least $6\times$ improvement in area and $19\times$ in power. To our knowledge, they represent the first open-source digital printed neural network classifiers capable of operating with existing printed energy harvesters.

CCS CONCEPTS

• **Hardware** → **Emerging technologies**; **Logic circuits**; • **Computing methodologies** → **Machine learning**.

KEYWORDS

Approximate Computing, Electrolyte-gated FET, Printed Electronics, Low-Power Classifiers, Ternary Neural Networks

ACM Reference Format:

Vojtech Mrazek, Argyris Kokkinis, Panagiotis Papanikolaou, Zdenek Vasicek, Kostas Siozios, Georgios Tzimpragos, Mehdi Tahoori, and Georgios



This work is licensed under a Creative Commons Attribution International 4.0 License. ICCAD '24, October 27–31, 2024, New York, NY, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1077-3/24/10
<https://doi.org/10.1145/3676536.3676728>

Zervakis. 2024. Evolutionary Approximation of Ternary Neurons for On-sensor Printed Neural Networks. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD '24)*, October 27–31, 2024, New York, NY, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3676536.3676728>

1 INTRODUCTION

Printed electronics hold significant promise for various applications, including smart packaging [30], in-situ monitoring [44], forensics [34], and accessible healthcare products and wearables [3, 19, 26, 28, 48]. A common denominator among these applications is the need for devices that are stretchable, porous, flexible, and conformal, all while adhering to strict area and power budgets. Another common feature is that processing typically involves classification, with recent research efforts directed towards the realization of printed bespoke machine learning (PBML) accelerators [22, 37].

The term bespoke denotes fully customized circuit implementations (e.g., hardwired model parameters [37, 38, 49]) as well as approximations [1, 2, 8–10, 24] per ML model and dataset. While this assumption may seem weak in other contexts, it is well-suited to printed electronics, where devices are considered disposable [11] due to their low fabrication cost and short turnaround times [14, 17, 37]. This specialization allows existing PBML designs, at least in theory, to operate with the power provided by available printed batteries [25]. However, the practicality of these solutions is often compromised by the overlooked cost of sensor-processor interfacing, a crucial aspect of the overall design since processing occurs directly on sensor data.

In this work, we address this problem holistically, starting our optimizations from the sensor boundary. Firstly, we replace costly analog-to-digital converters (ADCs) with efficient customized analog-to-binary converters (ABCs). Secondly, we explore ternary neural networks (TNNs) [5], where weights take only three values: $\{-1, 0, 1\}$, thus eliminating the need for multipliers. Thirdly, to further enhance hardware efficiency, we introduce approximate popcount and fused popcount-and-compare units. To facilitate their automated utilization, we implement a set of evolutionary multi-objective optimization techniques.

The result is approximate bespoke TNN designs that, on average, use 41% fewer resources and consume 42% less power than their exact counterparts, with no accuracy loss on common printed ML datasets [37, 43]. Compared to state-of-the-art digital approximate printed classifiers, our evaluation indicates an average reduction of 32× in area and 34× in power consumption, when factoring in the gains from replacing ADCs with the proposed ABC design.

In summary, the main **contributions** of this research are:

- A custom analog-to-binary converter design that uses only two resistors and an analog comparator.
- Approximate designs for popcount and fused popcount-and-comparison units, accompanied by a detailed Pareto analysis.
- The development of a multi-phase evolutionary optimization framework for the automated deployment of our approximation techniques and the search of the formed design space.
- The first, to the best of our knowledge, open-source complete digital printed classifier solution¹ that can be fully powered by either available printed batteries [25] or energy harvesters [4].

2 RELATED WORK

Printed electronics frequently targets applications centered on classification tasks. This focus, combined with the expanding fast-moving consumer goods market, has spurred research into printed machine learning accelerators. Various efforts have explored neural network models in this context, as summarized in Table 1.

The design choices behind the proposed solutions are shaped by the technology's unique attributes and the requirements imposed by its application domains. Printed electronics employ additive and mask-less manufacturing techniques like jet, screen, or gravure printing [17]. These methods, coupled with low capital equipment costs, enable rapid and inexpensive circuit production—sometimes for less than a cent per unit. However, these same characteristics lead to large feature sizes, low integration density (far below silicon VLSI), and increased device latencies [13, 27]. The latter is often not a concern due to the low sampling rates and relaxed performance requirements in typical applications [21]. As for the rest, to maximize functionality with minimum resources, bespoke designs are favored, especially since circuit programmability and reusability are less critical in this domain.

Bespoke designs are customized for specific classifier models trained on particular datasets, eliminating generality-related overheads [15]. This specialization extends to developing and implementing tailored approximation techniques [20]. For instance, Armeniakos et al. [7] proposed post-training weight replacement with more hardware-friendly approximates to reduce the cost of the underlying multipliers. This effort was further enhanced by introducing voltage over-scaling techniques where necessary [9]. In subsequent works, Armeniakos et al. [8] incorporated hardware-friendly weight approximation into the training process and implemented post-training addition approximation using simple truncation. Afentaki et al. [1, 2] constrained weights to powers of 2—leveraging that, in bespoke circuits, a multiplication by a power of 2 is implemented simply by rewiring—approximated accumulations by pruning the adder trees, and employed bounded low-precision ReLU activation.

Table 1: State-of-the-art printed neural networks.

Ref.	Digital	Ax. Mult.	Ax. Add.	Low Prec. Act. Func.	Inp. Size/ ADC Cost
[37]	✓	✗	✗	✗	high
[7, 9]	✓	low	✗	✗	high
[8]	✓	moderate	low	✗	high
[1, 2]	✓	high	moderate	✓	high
[43]	✗	✗	✗	✓	n/a
Proposed	✓	high	high	✓	low

Lastly, an alternative approximation method, stochastic computing [29] neural networks, was explored by Weller et al. [43].

The work presented here distinguishes itself from existing literature in several ways: it avoids the potential for significant accuracy losses associated with stochastic computing approaches; it employs a multiplier-free design and approximate adders for maximum efficiency; it is purely digital, avoiding the noise and variability issues that can affect low-resolution printed analog designs [39, 45–47]; and it begins optimization at the sensor boundary, rather than post-ADC. To our knowledge, the only other printed classifier that considers sensor-processor interface optimizations is a unary decision tree accelerator [6]. However, this design relies on multiple analog comparators per analog-to-unary converter, indicating room for further improvement.

3 PROPOSED END-TO-END DESIGN

3.1 Sensor-Processor Interface

The sensor-processor interface typically consists of an array of ADCs. ADCs are often overlooked in existing optimization efforts despite their significant hardware costs. For example, a 4-bit printed flash ADC, as the one shown in Figure 1a, occupies 12 mm² and consumes 1 mW of power [6]. If incorporated into a recent printed multilayer perceptron design with 4-bit inputs [1], it increases its area by 1.11× to 44× and power consumption by 1.97× to 370×.

We propose to replace such ADCs with a custom analog-to-binary converter (ABC), presented in Figure 1b. Compared to the reference 4-bit ADC, our ABC design requires 8× fewer resistors, 15× less comparators, and does not require a priority encoder. This translates to an area of only 0.07 mm² and a power consumption of 0.03 mW—167× smaller and 34× more power efficient than the original ADC design. Additionally, we allow the use of resistors, with potentially different values R_1 and R_2 . Having more than one voltage rails is inefficient in printed electronics, so we assume a common reference voltage V_{ref} for all ABCs. Adjusting the R_1/R_2 ratio enables to define the voltage threshold at which a specific input feature transitions between 1 and 0.

3.2 Bespoke Ternary Neural Networks

This work focuses on ternary neural networks. TNNs provide accuracy levels that fall between binary and standard neural networks, are easy to implement, and eliminate the need for hardware multipliers. Previous printed multilayer perceptron designs also avoided multipliers by constraining synaptic weight values to powers of 2 [1, 2, 24]. However, this wider range of weight values can result in larger accumulation circuits compared to TNNs, where

¹Available at: <https://github.com/ehw-fit/approximate-popcount>




Figure 1: Panel a: A 4-bit flash ADC. Panel b: Proposed analog-to-binary converter. The ratio R_1/R_2 is used to regulate the voltage threshold at which an input feature becomes 1. V_{in} denotes the voltage output of the sensor, which feeds the two converters (a) or (b).

weights are restricted to only -1 , 0 , or 1 . For example, assuming 4-bit inputs I_j , the sum $8 \times I_1 - 2 \times I_2 - 4 \times I_3 - 2 \times I_4$ requires 0.67 cm^2 , while with ternary weights, the area of the respective sum is at worst 0.38 cm^2 , i.e., 43% smaller. These improvements are enhanced under the assumption of binary inputs ($I_j \in \{0, 1\}$), such as those provided by the proposed ABCs. In this case, the corresponding area reduces to 0.2 cm^2 , i.e., 3.4 \times smaller than the initial design with powers of 2 weights.

3.2.1 Inputs and Weights Quantization. We employ Qkeras [16] for quantization-aware training (QAT) of the TNN, allowing customization of quantizers for synaptic weights and inputs. For weights, we use the ternary quantizer with alpha set to 1, yielding coefficients in $\{-1, 0, 1\}$. Hidden-layer neuron inputs are quantized to $\{0, 1\}$ using the quantized_bits activation function. For first-layer neuron inputs, we extend the quantized_bits function to set specific quantization thresholds for each input feature. We normalize training set inputs to $[0, 1]$ and analyze their distribution to adjust the quantization threshold V_q based on whether the distribution is left-skewed, nearly normal, or right-skewed. This threshold determines whether an input is considered 1 or 0 in the binary representation.

In this work, we set V_q to the median of the normalized input distribution for each input, rather than treating it as a learnable parameter during training. However, process variations can affect the actual values of R_1 and R_2 , potentially altering the R_1/R_2 ratio in the ABC design. This makes V_q susceptible to process variations, which could negatively impact classification accuracy. Optimizing for this is beyond the scope of this paper, but we propose incorporating variation-aware training techniques into QAT [45, 47] to enhance the robustness of our TNNs.

3.2.2 Circuit Implementation. The trained TNN model is automatically translated into hardware by combining Icarus Verilog (iverilog) with a set of custom neuron templates designed for each layer. Adhering to bespoke design principles, the synaptic weights are hardcoded into the circuit description. For synaptic weights with a value of 0 between the input and hidden layer, the corresponding connection is removed from the hidden-layer neuron, effectively eliminating an addend from its accumulator. For non-zero weights,




Figure 2: Bespoke (3, 2, 2) TNN circuit with $[[[0, 1, -1], [-1, -1, 1]], [[1, -1], [1, 1]]]$ weights.

the sign determines the hardware operation—either addition or subtraction—over the respective input. Consequently, each hidden-layer neuron performs a multi-operand summation, with its output being the sign of the respective sum: 1 for a positive sum and 0 for a negative sum.

The output layer implementation is customized to accommodate the different encoding of hidden-layer neuron outputs ($\{-1, 1\}$) compared to ABC outputs ($\{0, 1\}$). We employ XNOR-based neurons [40], which use XNOR gates instead of multiplications and a popcount operation to determine the number of high inputs.

Given the fixed, known weights, XNOR gates are simplified: becoming NOT gates for weights of -1 and simple wires for weights of 1 . The popcount operation is performed on $\{0, 1\}$ values, while our encoding assumes $\{-1, 1\}$. Thus, we apply a linear mapping that translates 0 to $\frac{1}{2}$. So, unlike the hidden layer connections, zero-valued weights are retained here as a $+0.5$ constant correction term. By ensuring, in training, the output neurons have the same number of zero-valued connections N , the correction term is consistent across all neurons and can be omitted without affecting the result of the classification, as $\text{argmax}(y + \frac{N}{2}, q + \frac{N}{2}) = \text{argmax}(y, q)$.

Lastly, to implement the argmax function, which identifies the output neuron with the highest score (i.e., highest popcount output), a set of comparators is employed. Figure 2 depicts a bespoke (3,2,2) TNN design. In this illustration, the hidden-layer neurons have weights $[0, 1, -1]$ and $[-1, -1, 1]$, while the output neurons have weights $[1, -1]$ and $[1, 1]$.

4 PROPOSED APPROXIMATION

The proposed approximation methodology, illustrated in Figure 3, consists of three main phases. In Phase 1, Cartesian genetic programming (CGP) [33] is used to evolve circuits that approximate the popcount function. In Phase 2, we identify Pareto-optimal combinations of these circuits to deliver the necessary approximate popcount-compare functionality for the hidden-layer neurons. Finally, in Phase 3, the NSGA-II evolutionary algorithm [18] is applied to integrate these approximate components into a bespoke TNN circuit, optimizing for maximum resource efficiency with minimal impact on accuracy.




Figure 3: Overview of the proposed three-phase TNN approximation framework, including Cartesian genetic programming for popcount circuit approximations in Phase 1, a Pareto analysis for identifying optimal combinations of popcount circuits for the construction of approximate popcount-compare circuits in Phase 2, and their integration into bespoke TNN circuits using the NSGA-II evolutionary algorithm in Phase 3. The underlying technology assumed is EGFET, operating at voltages as low as 0.6-1 V.

4.1 Approximate Popcount & Popcount-compare Circuits

For hidden-layer neurons, a binary step activation function is used. The computation per neuron is:

$$\sum_i w_i I_i \geq 0, \quad (1)$$

where w_i are synaptic weights and I_i are inputs. Given $w_i \in \{-1, 0, 1\}$, this transforms to:

$$\sum_{\forall i: w_i=1} I_i - \sum_{\forall i: w_i=-1} I_i \geq 0 \implies \sum_{\forall i: w_i=1} I_i \geq \sum_{\forall i: w_i=-1} I_i. \quad (2)$$

With $I_i \in \{0, 1\}$, each sum can be calculated by a popcount operation. Thus, each hidden-layer neuron requires two popcount circuits and a comparator. The comparator's complexity is minimal compared to this of the popcount operations. For example, with 8 inputs, the comparator uses only 16% of the hidden-layer neuron's total area; for 64 inputs, it's less than 3%. In the rest of this Section, we introduce approximate implementations of popcount-compare (for hidden-layer neurons) and popcount (for output-layer neurons) circuits to enhance area-efficiency.

4.1.1 Approximate Popcount Circuits. An n -bit popcount (PC) operation sums n binary inputs, typically implemented using a tree structure of adders. To approximate this exact circuit, we employ the evolutionary CGP algorithm [33]. In CGP, a simple search method is deployed over an integer address-based genetic representation of the circuit to create random mutations of an initial population

P that contains the accurate PC circuit. To evaluate the candidate circuits, we use a fitness function. Each member c in population P receives a fitness score $F(c)$, with the highest-scoring individual becoming the parent of the next population. This parent then generates λ new candidate solutions through mutation. The process terminates when either the number of iterations exceeds a set maximum, or a predefined time limit is reached.

The searching algorithm's optimization criterion is the estimated area for the target technology (e.g., EGFET PDK [12]), while maintaining the error ε below a specific threshold, consistent with previous error-oriented approximation research [35]. This error constraint τ is incorporated into the fitness function as follows:

$$F(c) = \begin{cases} \text{area}(c), & \text{if } \varepsilon(c) \leq \tau \\ \infty, & \text{otherwise.} \end{cases} \quad (3)$$

PC circuits produce standard arithmetic outputs, allowing the use of conventional error metrics in the approximation process: mean arithmetic error ε_{mae} and worst-case arithmetic error ε_{wcae} . These metrics calculate the average or maximum error across all input combinations. However, for a large number of inputs n in PC circuits, evaluating all 2^n input vectors becomes impractical. To address this, a formal verification approach using binary decision diagrams (BDD) [41] is used.

4.1.2 Approximate Popcount-compare Circuits. Hidden-layer neurons are implemented using popcount-compare (PCC) circuits to realize Equation (2). Each PCC consists of two PC circuits with n_{pos} and n_{neg} input bits, respectively, and a j -bit comparator, where j is equal to $\lceil \log_2 \max(n_{pos}, n_{neg}) \rceil$. To ensure a comprehensive analysis and evaluation, PCCs are created for all n_{pos} and n_{neg} configurations found in the target TNNs.

Distance metric: In approximating non-arithmetic circuits, Hamming distance is often considered the sole suitable error metric. However, this approach falls short in the context of approximate computing, as it may not accurately represent the error levels [23]. For example, in 4-bit output arithmetic circuits, an error in one bit position ($0 \rightarrow 8$) can be more severe than errors in all four bits ($7 \rightarrow 8$). This limitation extends to relational operators as well. Consider two approximate \geq_a and \geq_b circuits: if \geq_a erroneously determines that $0 \geq_a 1 = \text{TRUE}$ (a 1-position deviation) and \geq_b incorrectly asserts that $0 \geq_b 4 = \text{TRUE}$ (a 4-position deviation), both errors result in the same Hamming distance of 1 (FALSE \rightarrow TRUE), despite their clearly different magnitude of errors.

The Hamming distance's inability to capture error severity in the context of approximate computing is of particular importance to PCC approximation, which produces a one-bit output. To address this, we propose a new *distance metric* D , inspired by those used in approximate median filters and sorting networks [36]. This metric is defined as:

$$D(x, z) = \begin{cases} 0, & \text{if } \text{rel}(x, z) = \text{rel}'(x, z) \\ x - z, & \text{otherwise} \end{cases} \quad (4)$$

The distance D is a function of two inputs x and z , an accurate relational function rel with a Boolean output, and its approximate variant rel' . This provides a more nuanced measure of error magnitude, addressing the limitations of Hamming distance in this

context. Using D , we define the mean distance error (ϵ_{mde}) and worst-case distance error (ϵ_{wcde}) as follows:

$$\epsilon_{mde} = \frac{1}{|G|} \sum_{\forall (x,z) \in G} |D(x,z)|; \epsilon_{wcde} = \max_{\forall (x,z) \in G} |D(x,z)| \quad (5)$$

Here, G represents the input domain.

Pareto analysis: We create a library of approximate PCC circuits using the PC circuit approximation methodology described above. For each exact PCC circuit, we first generate all possible approximate PCC circuits by testing all combinations of n_{pos} and n_{neg} PC circuits from our approximate PC library, including exact PC circuits as zero-error designs. Subsequently, we identify Pareto-optimal PCC circuits using two criteria: (i) accuracy, evaluated through ϵ_{mde} for 10^6 random (x, z) pairs in Equations (4) and (5), and (ii) hardware cost, estimated by the combined area of the n_{pos} and n_{neg} PC circuits.

This search process is fully parallelizable and can be conducted at a high level (i.e., in Python) without the need for hardware evaluation. Consequently, the execution time remains well constrained, ensuring the scalability of our approximate PCC circuit search. Once the Pareto analysis is completed, the selected PCC circuits, including both comparators and approximate PC circuits, are synthesized to assess their actual hardware cost.

4.2 Approximate Circuit Integration

The last step is the integration of PC and PCC circuits from the developed libraries into a complete TNN circuit. The objective is to achieve the best area-accuracy trade-off. To this end, we deploy a multi-objective optimization based on the evolutionary NSGA-II algorithm [18]. As in prior optimizations, the two criteria are: accuracy and area. We encode the problem as an integer list, where each integer corresponds to one circuit from the PCC and PC library for the respective neuron. The sum of the areas of the selected PCC and PC circuits is used as a proxy for the overall area of the TNN. To evaluate the actual area and power results of the generated TNNs, we go again through a round of synthesis. The underlying technology assumed is EGFET, which allows operation at voltages as low as 0.6-1 V [32] and aligns well with low-power requirements of battery/self-powered IoT applications.

5 EVALUATION

The goal of the presented research is to develop bespoke TNN circuit designs that optimize hardware efficiency while minimizing any decrease in accuracy. Initially, we evaluate the proposed hardware implementations and the efficacy of our approximation methods. Subsequently, we conduct a comparative analysis against the current state-of-the-art printed neural networks.

Evaluation setup: For our evaluation, we utilize the five datasets from the UCI ML repository [31]. These datasets are chosen for two primary reasons: they enable direct comparisons with state-of-the-art designs, and they consist of sensor data suitable for printed electronics' applications, as discussed in studies [37, 43]. To perform digital circuit synthesis and analysis, we rely on Synopsys Design Compiler and PrimeTime tools, used in conjunction with the EGFET standard cell library [12]. For the ABC design, we utilize Cadence Virtuoso with the EGFET PDK, running SPICE simulations to assess

Table 2: TNN accuracy results.

Dataset	Exact MLP [37]			Our Exact TNN		
	A*	T†	I/W*	A*	T†	I/W*
Arrhythmia	0.62	(274,5,16)	4/8	0.60	(274,3,16)	1/2
Breast Cancer	0.98	(10,3,2)	4/8	0.98	(10,10,2)	1/2
Cardio	0.88	(21,3,3)	4/8	0.85	(21,3,3)	1/2
Redwine	0.56	(11,2,6)	4/8	0.56	(11,3,6)	1/2
Whitewine	0.54	(11,4,7)	4/8	0.50	(11,11,7)	1/2

*Inference accuracy. †Network Topology. *Bit-width Inputs/Weights.

its area and power consumption. The voltage supply for all circuits is set to 0.6 V, in line with EGFET capabilities and consistent with prior work in the field [1, 2]. In terms of operating speeds, all circuits are evaluated at a frequency of 5 Hz, although our designs can achieve frequencies up to 20 Hz. This decision is made to maintain consistency with existing solutions in the literature [8, 12] and facilitate a fair comparison.

TNN baseline: Exact TNN models are used as the baseline for accuracy comparisons. The selected datasets are split into 70% for training and 30% for testing. For training, the Adam optimizer is used with the number of epochs ranging from 10 to 20. The learning rate for each network is obtained through Bayesian optimization, aiming to maximize the model's inference accuracy with a maximum of 100 attempts. Learning rate parameters are selected from the range of 0.001 to 0.01. The final model is that with the highest accuracy among a set of 100 TNNs with various hyperparameters.

Regarding TNN hyperparameters, all models assume a single hidden layer. To determine the number of hidden-layer neurons, we perform an exhaustive search within the range of 1 to 40 neurons. For each iteration, we repeat the procedure described previously. Among the TNNs yielding the highest accuracy, we select the one with the fewest neurons. This exhaustive search takes less than one hour on an AMD EPYC 7552 server with 256 GB RAM.

Table 2 summarizes our findings across the examined datasets. For reference, we also include the accuracy numbers achieved by a baseline MLP network [37] utilizing 4-bit inputs and 8-bit weights. The comparison indicates 0-4% accuracy drop for our baseline TNNs. However, as will be discussed in the rest of this section, this compromise enables a significant boost in hardware efficiency, ensuring that the operating requirements of our printed TNNs fall within the integration and power budget of the target technology.

5.1 Approximate Circuits Evaluation

5.1.1 Popcount Circuits. We generate approximate PC circuits with the following constraints and configurations. The error limits τ_{mae} and τ_{wcae} are logarithmically distributed from 0.1 to $0.5 \cdot 2^m$ and from 1 to $0.5 \cdot 2^m$, with $m = \lceil \log_2 n \rceil$, resulting in 2,090 approximate circuits. We also set CGP search termination criteria to 30, 60, and 300 minutes for PC sizes $n < 16$, $n < 32$, and $n < 60$. Using BDD evaluation, we achieve average speeds of 7,759 and 1,362 evaluations per second for 16-bit and 32-bit PCs. For 60-bit PC circuits, this number falls to 25 evaluations per second. To account for this, we extend the time limit, thereby increasing the chances of identifying solutions with better area-accuracy trade-offs.

A visualization of the obtained results for three PC circuit sizes is provided in Figure 4. The featured circuits are synthesized and

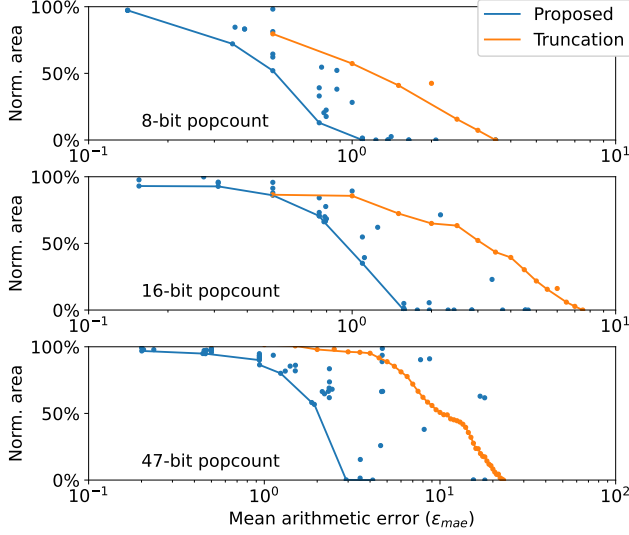


Figure 4: Comparison between the proposed approximation approach and the previously used truncation technique for PCC circuits of various sizes. The results displayed are based on post-synthesis area measurements.

evaluated based the error metrics and methodology described above. For easier comparison, area results are normalized relative to exact PC circuit of the same size. Our results indicate that as the error limit increases, CGP optimization identifies increasingly area-efficient solutions. Additionally, we assess the efficacy of our approximation approach through a comparison of our PC circuits with variants derived using the truncation approximation approach, which has previously been applied in a wide range of approximate circuits, including multilayer perceptrons [42]. In this case, we achieve about $2\times$ area reduction with ϵ_{mae} of only 0.5 for an 8-bit PC circuit, 1.1 for a 16-bit PC circuit, and 1.9 for a 47-bit PC circuit. Similar results to those shown in Figure 4 are observed for all PC circuit sizes.

5.1.2 Popcount-compare Circuits. Using the methods outlined in Section 4, we generate 5,841 approximate PCC circuits for our five target datasets. Figure 5a presents an area-versus-error analysis for three distinct PCC circuits customized for the Arrhythmia dataset. The results reveal significant overlap in designs, indicating that the approximate PCC circuit design space can be substantially pruned. This pruning reduces the size of the design space for TNN approximation. Consequently, we limit the approximate PCC library to 134 circuits across 22 different input sizes.

Figure 5b presents a detailed error analysis of the design points on the black Pareto frontier from Figure 5a. The logarithmic-scale histograms depict the distribution of distance error D for $|G| = 10^6$ random input samples. Hence, the x-axis represents the distance metric D introduced in Section 4, while the y-axis shows normalized area results, with area savings increasing as you move down the axis. The analyzed PCC circuit features $n_{pos} = 45$, $n_{neg} = 39$, and a 6-bit comparator, corresponding to the top hidden-layer neuron of our bespoke Arrhythmia TNN model.

Each row in the plot represents a Pareto-optimal approximation of the PCC circuit. At the top, the fully accurate design’s histogram




Figure 5: Panel a: Trade-off illustration for three PCC circuits applied to the Arrhythmia dataset. The circuits are characterized by their sizes (n_{pos}, n_{neg}): black (45, 39), purple (47, 30), and pink (60, 29). Panel b: Distribution of distance error for Pareto optimal PCC circuits of size $n_{pos} = 45, n_{neg} = 39$ (corresponding to the black line in Panel a). Area results are post-synthesis and normalized relative to the exact circuit.




Figure 6: Comparison between area estimates and post-synthesis results for various PCC circuits. Design points with similar colors indicate the use of the same exact design as a starting point.

is entirely concentrated in the bin for $D = 0$, indicating no errors. As we move down the rows, accuracy is traded for reduced area, causing the probability mass to gradually spread to bins with distances further from 0. If our exploration allows differences between positive and negative PC circuits up to $\epsilon_{wcde} = 4$, it results in: 12.6% area reduction (6th row from the top in Figure 5b), 95.57% error-free PCC operations, and a mean distance ϵ_{mde} of only 0.06. Moving down the y-axis, for $\epsilon_{wcde} \leq 8$, area savings increase to approximately 30%. For approximations at the bottom of the plot, area can be reduced to as little as 25% of the accurate design while still maintaining over 80% correct results. These results are consistent across all PCC circuit configurations examined.

Given that the results presented in the first leg of Pareto analysis are based on area estimates, the final step is to evaluate how closely these estimates match the actual synthesis results. Figure 6 provides this comparison, illustrating the relationship between estimated and synthesized area values for the PCC circuits. While our estimation tends to underestimate smaller PCCs due to not accounting for comparator cost, occasional overestimations occur due to synthesis optimizations. However, overall, there is a good enough correlation between the actual and estimated area.




Figure 7: Post-synthesis area-versus-accuracy evaluation for Pareto-optimal approximate TNNs.

5.2 Approximate TNN Evaluation

The final stage of our evolutionary approximation approach utilizes the NSGA-II algorithm to construct a TNN design. This design incorporates components from our PCC (hidden neurons) and PC (output neurons) circuit libraries. We implement NSGA-II using the pymoo library, employing an integer-based representation. Each integer in this representation corresponds to the index of an approximate component for a given neuron. Depending on the dataset under examination, chromosome lengths range from 6 to 19 integers.

To navigate the design search space, which ranges from $2.4 \cdot 10^3$ approximate candidates for Cardio to $6.8 \cdot 10^{11}$ for Arrhythmia, we employ binary crossover with polynomial mutation. If we were to consider PC approximation for every popcount operation in the hidden-layer neurons instead of combining them in Pareto-optimal PCC circuits, these values would increase to $5.2 \cdot 10^5$ and $1.7 \cdot 10^{14}$, respectively, complicating the required exploration. The results of the design process for the Arrhythmia TNN (our largest TNN) are shown in Figure 8. The algorithm was run for 200 generations, taking less than 5 minutes to complete, with substantial progress achieved within the first 50 generations.

Figure 7 presents the accuracy-versus-area results of our explorations, featuring designs found on the Pareto optimal curve. Our findings reveal two key points: first, the generated approximate TNN designs can match the accuracy of their exact TNN counterparts while achieving an average area reduction of 41%; second, if we allow for up to a 5% decrease in accuracy compared to the exact TNN, the area savings increase to an average of 67%.

Comparison with State-of-the-art:

Table 3 presents a comparison of our results against the current state-of-the-art in printed neural networks. Our approximate TNN

Table 3: Comparison against the State of the Art. All circuits target the EGFET technology.

Dataset	Ref.	Accuracy (%)	w/o ADC cost		w/ ADC cost	
			Area (cm²)	Power (mW)	Area (cm²)	Power (mW)
Arrhythmia	Exact MLP [37]	62	266.00	998.00	298.06	1273.92
Arrhythmia	Ax. MLP [2]	60	13.51	12.80	45.57	288.72
Arrhythmia	Our Exact TNN	60	8.87	8.09	9.06	16.31
Arrhythmia	Our Ax. TNN	60	7.73	7.12	7.92	15.34
Arrhythmia	Our Ax. TNN	57	5.20	4.95	5.39	13.17
BreastCancer	Exact MLP [37]	98	12.00	40.00	13.17	50.07
BreastCancer	Ax. MLP [8]	97	5.10	18.00	6.27	28.07
BreastCancer	Ax. MLP [8]	93	1.50	5.69	2.67	15.76
BreastCancer	Ax. MLP [2]	94	0.08	0.08	1.25	10.15
BreastCancer	Ax. MLP [1]	94	0.03	0.03	1.20	10.10
BreastCancer	Our Exact TNN	98	0.29	0.31	0.30	0.61
BreastCancer	Our Ax. TNN	98	0.05	0.04	0.06	0.34
BreastCancer	Our Ax. TNN	93	0.04	0.04	0.05	0.34
Cardio	Exact MLP [37]	88	33.40	124.20	35.86	145.35
Cardio	Ax. MLP [8]	87	6.08	20.80	8.54	41.95
Cardio	Ax. MLP [8]	83	4.92	16.80	7.38	37.95
Cardio	Ax. MLP [2]	85	1.35	1.57	3.81	22.72
Cardio	Ax. MLP [1]	87	1.46	1.70	3.92	22.85
Cardio	Our Exact TNN	85	0.75	0.91	0.76	1.54
Cardio	Our Ax. TNN	85	0.36	0.42	0.37	1.05
Cardio	Our Ax. TNN	84	0.27	0.30	0.28	0.93
RedWine	Exact MLP [37]	56	17.60	73.50	18.89	84.58
RedWine	Ax. MLP [8]	55	1.06	3.95	2.35	15.02
RedWine	Ax. MLP [8]	51	0.87	3.25	2.16	14.33
RedWine	Ax. MLP [2]	55	0.03	0.02	1.32	11.10
RedWine	Ax. MLP [1]	52	0.03	0.03	1.32	11.11
RedWine	Our Exact TNN	56	0.08	0.09	0.09	0.42
RedWine	Our Ax. TNN	56	0.03	0.03	0.04	0.36
WhiteWine	Exact MLP [37]	54	31.20	126.40	32.49	137.48
WhiteWine	Ax. MLP [8]	53	6.47	21.30	7.76	32.38
WhiteWine	Ax. MLP [8]	49	0.79	2.77	2.07	13.84
WhiteWine	Ax. MLP [2]	50	0.25	0.25	1.54	11.33
WhiteWine	Ax. MLP [1]	51	0.23	0.25	1.52	11.33
WhiteWine	Our Exact TNN	50	0.16	0.18	0.17	0.51
WhiteWine	Our Ax. TNN	50	0.11	0.12	0.12	0.45
WhiteWine	Our Ax. TNN	45	0.06	0.07	0.07	0.40




Figure 8: Evolutionary optimization of the Arrhythmia TNN over 200 generations.

designs demonstrate a superior accuracy-area trade-off compared to approximate multilayer perceptrons, even without considering the sensor-processor interface cost. Within a 1% accuracy variation, our TNNs achieve, on average, 42 \times , 2.17 \times , and 1.99 \times lower area compared to previous works on approximate multilayer perceptrons [1, 2, 8]. Power consumption is similarly reduced by 162 \times , 2.05 \times , and 1.97 \times , on average, respectively. When sensor-processor interface costs (ADCs and ABCs) are factored in, these gains grow further, with area reductions exceeding 6 \times and power savings of over 19 \times compared to the most efficient reported multilayer perceptron design [1]. When compared to an exact multilayer perceptron implementation [37], our approximate TNNs achieve 55 \times lower area and 97 \times less power at the expense of a 5% accuracy drop.

Notice that all our TNN circuits, except those targeting the Arrhythmia dataset, consume well below 2 mW, even when considering the power consumed by the ABCs and the power overhead of the sensors (approximately 5 μ W based on previous studies [12]). Thus, they can be powered by a printed energy harvester [4]. To the best of our knowledge, this is the first digital printed classifier to achieve this. As for our Arrhythmia TNN designs, they can be powered by either a Zinergy (15 mW) or Molex (30 mW) printed battery [37]. This represents a significant advancement over previous solutions [2, 37] targeting this application that exceeding the power budget of printed batteries.

6 CONCLUSION

Printed electronics, particularly classifier circuits, offer transformative potential in applications requiring flexible substrates and ultra-low costs. Despite recent advancements, practical implementation faces challenges due to low integration density and limited power from printed batteries and harvesters. Our research approaches these issues holistically, focusing on optimizations from the sensor-processor interface to the processor itself. To our knowledge, we present the first open-source end-to-end digital printed classifier that meets both resource and power constraints. Key contributions include: replacing ADCs with more efficient ABCs, customizable to specific input features; implementing bespoke ternary neurons, without costly multipliers and with approximate popcount and popcount-compare units; developing and deploying a multi-stage evolutionary optimization method that assists both in the approximation of the popcount and popcount-compare circuits and in their integration into a complete TNN accelerators. When sensor-processor interface costs are factored in, our results show a 32 \times improvement in area and 34 \times reduction in power compared to current state-of-the-art approximate digital printed neural network implementations, while maintaining comparable accuracy across diverse datasets.

ACKNOWLEDGMENTS

This work was supported by the Czech Science Foundation project 22-02067S, by the European Research Council (ERC), and 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). The authors thank Ampere for server support.

REFERENCES

- [1] Florentia Afentaki, Michael Hefenbrock, Georgios Zervakis, and Mehdi B. Tahoori. 2024. Embedding Hardware Approximations in Discrete Genetic-based Training for Printed MLPs. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*.
- [2] Florentia Afentaki, Gurol Saglam, Argyris Kokkinis, Kostas Siozios, Georgios Zervakis, and Mehdi B. Tahoori. 2023. Bespoke Approximation of Multiplication-Accumulation and Activation Targeting Printed Multilayer Perceptrons. In *IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. 1–9.
- [3] Shweta Agarwala, Guo Liang Goh, Truong-Son Dinh Le, Jianing An, Zhen Kai Peh, Wai Yee Yeong, and Young-Jin Kim. 2019. Wearable Bandage-Based Strain Sensor for Home Healthcare: Combining 3D Aerosol Jet Printing and Laser Sintering. *ACS Sensors* 4, 1 (2019), 218–226.
- [4] Ammar Ahmed and others. 2021. Additively manufactured nano-mechanical energy harvesting systems: advancements, potential applications, challenges and future perspectives. *Nano Convergence* 8, 37 (2021).
- [5] Hande Alemdar, Vincent Leroy, and Frederic Petrot. 2017. Ternary neural networks for resource-efficient AI applications. In *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2547–2554.
- [6] Giorgos Armeniakos and others. 2024. On-sensor Printed Machine Learning Classification via Bespoke ADC and Decision Tree Co-Design. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*.
- [7] Giorgos Armeniakos, Georgios Zervakis, Dimitrios Soudris, Mehdi B. Tahoori, and Jörg Henkel. 2022. Cross-Layer Approximation For Printed Machine Learning Circuits. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 190–195.
- [8] Giorgos Armeniakos, Georgios Zervakis, Dimitrios Soudris, Mehdi B. Tahoori, and Jörg Henkel. 2023. Co-Design of Approximate Multilayer Perceptron for Ultra-Resource Constrained Printed Circuits. *IEEE Trans. Comput.* (2023), 1–8.
- [9] Giorgos Armeniakos, Georgios Zervakis, Dimitrios Soudris, Mehdi B. Tahoori, and Jörg Henkel. 2023. Model-to-Circuit Cross-Approximation For Printed Machine Learning Classifiers. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2023), 1–1.
- [10] Konstantinos Balaskas, Georgios Zervakis, Kostas Siozios, Mehdi B. Tahoori, and Jörg Henkel. 2022. Approximate Decision Trees For Machine Learning Classification on Tiny Printed Circuits. In *Int. Symp. Quality Electronic Design*.
- [11] Ajay Beniwal, Priyanka Ganguly, Akshaya Kumar Aliyana, Gaurav Khandelwal, and Ravinder Dahiya. 2023. Screen-printed graphene-carbon ink based disposable humidity sensor with wireless communication. *Sensors and Actuators B: Chemical* 374 (2023), 132731.
- [12] N. Bleier, M. Mubarik, F. Rasheed, J. Aghassi-Hagmann, M. B. Tahoori, and R. Kumar. 2020. Printed Microprocessors. In *Annu. Int. Symp. Computer Architecture (ISCA)*. 213–226.
- [13] Gabriel Cadilha Marques, Suresh Kumar Garlapati, Simone Dehm, Subho Dasgupta, Horst Hahn, Mehdi Tahoori, and Jasmin Aghassi-Hagmann. 2017. Digital power and performance analysis of inkjet printed ring oscillators based on electrolyte-gated oxide electronics. *Applied Physics Letters* 111, 10 (2017), 102103.
- [14] Joseph S Chang, Antonio F Facchetti, and Robert Reuss. 2017. 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* 7, 1 (2017), 7–26.
- [15] Hari Cherupalli, Henry Duwe, Weidong Ye, Rakesh Kumar, and John Sartori. 2017. Bespoke processors for applications with ultra-low area and power constraints. In *Annu. Int. Symp. Computer Architecture (ISCA)*. 41–54.
- [16] Claudionor Coelho, Aki Kuusela, Hao Zhuang, Thea Aarrestad, Vladimir Loncar, Jennifer Ngadiuba, Maurizio Pierini, and Sioni Summers. 2021. Ultra Low-latency, Low-area Inference Accelerators using Heterogeneous Deep Quantization with QKeras and hls4ml. *arXiv:2006.10159* (2021).
- [17] Zheng Cui. 2016. *Printed electronics: materials, technologies and applications*. John Wiley & Sons.
- [18] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- [19] Jiuwei Gao, et al. 2022. Ultra-Robust and Extensible Fibrous Mechanical Sensors for Wearable Smart Healthcare. *Advanced Materials* 34, 20 (2022), 2107511.
- [20] Jie Han and Michael Orshansky. 2013. Approximate computing: An emerging paradigm for energy-efficient design. In *IEEE European Test Symposium (ETS)*.
- [21] Jörg Henkel, Hai Li, Anand Raghunathan, Mehdi B. Tahoori, Swagath Venkataramani, Xiaoxuan Yang, and Georgios Zervakis. 2022. Approximate Computing and the Efficient Machine Learning Expedition. In *International Conference On Computer Aided Design (ICCAD)*. 1–9.
- [22] Konstantinos Iordanou, Timothy Atkinson, Emre Ozer, Jędrzej Kufel, Grace Ali-gada, John Biggs, Gavin Brown, and Mikel Luján. 2024. Low-cost and efficient prediction hardware for tabular data using tiny classifier circuits. *Nature Electronics* (2024), 1–9.
- [23] Honglan Jiang, Cong Liu, Leibo Liu, Fabrizio Lombardi, and Jie Han. 2017. A Review, Classification, and Comparative Evaluation of Approximate Arithmetic

- Circuits. *J. Emerg. Technol. Comput. Syst.* 13, 4, Article 60 (aug 2017), 34 pages.
- [24] Argyris Kokkinis, Georgios Zervakis, Kostas Siozios, Mehdi B. Tahoori, and Jörg Henkel. 2023. Hardware-Aware Automated Neural Minimization for Printed Multilayer Perceptrons. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 1–2.
- [25] Senentxu Lanceros-Méndez and Carlos Miguel Costa. 2018. *Printed Batteries: Materials, Technologies and Applications*. Wiley.
- [26] Jin-Woo Lee, Younguk Choi, Jaehoon Jang, Se-Hyuk Yeom, Wanghoon Lee, and Byeong-Kwon Ju. 2020. High sensitivity flexible paper temperature sensor and body-attachable patch for thermometers. *Sensors and Actuators A: Physical* (2020).
- [27] Ting Lei, et al. 2019. Low-voltage high-performance flexible digital and analog circuits based on ultrahigh-purity semiconducting carbon nanotubes. *Nature communications* 10, 1 (2019), 2161.
- [28] Yue Li, et al. 2022. The Soft-Strain Effect Enabled High-Performance Flexible Pressure Sensor and Its Application in Monitoring Pulse Waves. *Research* (2022).
- [29] Yidong Liu, Siting Liu, Yanzhi Wang, Fabrizio Lombardi, and Jie Han. 2021. A Survey of Stochastic Computing Neural Networks for Machine Learning Applications. *IEEE Trans. Neural Netw. Learn. Syst.* 32, 7 (2021), 2809–2824.
- [30] Xiaoyu Luo. 2022. Application of inkjet-printing technology in developing indicators/sensors for intelligent packaging systems. *Current Opinion in Food Science* 46 (2022), 100868.
- [31] Kolby Nottingham Markelle Kelly, Rachel Longjohn. [n. d.]. The UCI Machine Learning Repository. ([n. d.]). <https://archive.ics.uci.edu>
- [32] Cadilha Marques and others. 2019. Progress Report on “From Printed Electrolyte-Gated Metal-Oxide Devices to Circuits”. *Advanced Materials* 31 (2019).
- [33] Julian F. Miller. 2011. *Cartesian Genetic Programming*. Springer-Verlag.
- [34] Rupesh K. Mishra, et al. 2020. Simultaneous detection of salivary Δ^9 -tetrahydrocannabinol and alcohol using a Wearable Electrochemical Ring Sensor. *Talanta* 211 (2020), 120757.
- [35] Vojtech Mrazek, Syed Shakib Sarwar, Lukas Sekanina, Zdenek Vasicek, and Kaushik Roy. 2016. Design of power-efficient approximate multipliers for approximate artificial neural networks. In *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 1–7.
- [36] Vojtech Mrazek and Zdenek Vasicek. 2016. Automatic design of arbitrary-size approximate sorting networks with error guarantee. In *Int. Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS)*. 221–228.
- [37] Muhammad Husnain Mubarak, Dennis D. Weller, Nathaniel Bleier, Matthew Tomei, Jasmin Aghassi-Hagmann, Mehdi B. Tahoori, and Rakesh Kumar. 2020. Printed Machine Learning Classifiers. In *Annu. Int. Symp. Microarchitecture (MICRO)*. 73–87.
- [38] Emre Ozer, Jędrzej Kufel, John Biggs, Gavin Brown, James Myers, Anjit Rana, Antony Sou, and Catherine Ramsdale. 2019. Bespoke Machine Learning Processor Development Framework on Flexible Substrates. In *Int. Conf. Flexible and Printable Sensors and Systems (FLEPS)*. 1–3.
- [39] Priyanjana Pal, Haibin Zhao, Maha Shatta, Michael Hefenbrock, Sina B. Mamaghani, Sani Nassif, Michael Beigl, and Mehdi B. Tahoori. 2024. Analog Printed Spiking Neuromorphic Circuit. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 6 S.
- [40] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. 2016. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks. In *Computer Vision – ECCV 2016*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 525–542.
- [41] Mathias Soeken, Daniel Große, Arun Chandrasekharan, and Rolf Drechsler. 2016. BDD minimization for approximate computing. In *2016 21st Asia and South Pacific Design Automation Conference (ASP-DAC)* (Macao, Macao). IEEE Press, 474–479.
- [42] Swagath Venkataramani, Ashish Ranjan, Kaushik Roy, and Anand Raghunathan. 2014. AxNN: Energy-efficient neuromorphic systems using approximate computing. In *Int. Symp. on Low Power Electronics and Design (ISLPED)*. 27–32.
- [43] Dennis D. Weller, Nathaniel Bleier, Michael Hefenbrock, Jasmin Aghassi-Hagmann, Michael Beigl, Rakesh Kumar, and Mehdi B. Tahoori. 2021. Printed Stochastic Computing Neural Networks. In *Design, Automation Test in Europe Conference Exhibition (DATE)*. 914–919.
- [44] Xi Zhang, Xuechuan Shan, and Jun Wei. 2017. Hybrid flexible smart temperature tag with NFC technology for smart packaging. In *IEEE Electronics Packaging Technology Conference (EPTC)*. 1–5.
- [45] Haibin Zhao, Michael Hefenbrock, Michael Beigl, and Mehdi B. Tahoori. 2022. Aging-Aware Training for Printed Neuromorphic Circuits. In *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. 1–9.
- [46] Haibin Zhao, Priyanjana Pal, Michael Hefenbrock, Michael Beigl, and Mehdi B. Tahoori. 2023. Power-Aware Training for Energy-Efficient Printed Neuromorphic Circuits. In *International Conference on Computer Aided Design (ICCAD)*. 1–9.
- [47] Haibin Zhao, Brojogopal Sapui, Michael Hefenbrock, Zhidong Yang, Michael Beigl, and Mehdi B. Tahoori. 2023. Highly-Bespoke Robust Printed Neuromorphic Circuits. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*.
- [48] Kemeng Zhou, Ruochen Ding, Xiaohao Ma, and Yuanjing Lin. 2024. Printable and flexible integrated sensing systems for wireless healthcare. *Nanoscale* 16 (2024), 7264–7286. Issue 15.
- [49] Emre Özer, Jędrzej Kufel, James Myers, John Biggs, Gavin Brown, Anjit Rana, Antony Sou, Catherine Ramsdale, and Scott White. 2020. A hardwired machine learning processing engine fabricated with submicron metal-oxide thin-film transistors on a flexible substrate. *Nature Electronics* 3 (07 2020), 1–7.