Arbitrary Precision Printed Ternary Neural Networks with Holistic Evolutionary Approximation

Vojtech Mrazek, Konstantinos Balaskas, Paula Carolina Lozano Duarte, Zdenek Vasicek,
Mehdi B. Tahoori, Georgios Zervakis

Abstract—**Printed electronics offer a promising alternative for applications beyond silicon-based systems, requiring properties like flexibility, stretchability, conformality, and ultra-low fabrication costs. Despite the large feature sizes in printed electronics, printed neural networks have attracted attention for meeting target application requirements, though realizing complex circuits remains challenging. This work bridges the gap between classification accuracy and area efficiency in printed neural networks, covering the entire processing-near-sensor system design and co-optimization from the analog-to-digital interface–a major area and power bottleneck–to the digital classifier. We propose an automated framework for designing printed Ternary Neural Networks with arbitrary input precision, utilizing multi-objective optimization and holistic approximation. Our circuits outperform existing approximate printed neural networks by 17x in area and 59x in power on average, being the first to enable printed-battery-powered operation with under 5% accuracy loss while accounting for analog-to-digital interfacing costs.**

Index Terms—**Approximate computing, Electrolyte-gated FET, Printed Electronics, Ternary Neural Networks**

I. INTRODUCTION

Over the past several decades, the Very-Large-Scale Integration (VLSI) circuit industry has managed to fuel Moore's law, and feature scaling is expected to persist for the next ten years and beyond [1]. The common theme among these technological innovations is the never-ending pursuit of improving the power-performance-area (PPA) of transistors, and silicon-based systems overall. However, numerous application domains fall outside the PPA target, including forensics [2], disposables and smart packaging [3], [4], accessible healthcare products and wearables [5]–[8], and the ten trillion-dollar fast-moving consumer goods market [9]. Such applications require *ultra-low-cost fabrication, along with stretchability, porosity, flexibility, and conformality*—demands that rigid silicon systems cannot meet, positioning printed electronics at the forefront for integrating smart services in these domains that have yet to witness significant computing infiltration.

Printed electronics use mask-less, additive methods like jet, screen, or gravure printing, achieving ultra-low cost (even sub-cent) and fast-turnaround fabrication [10], [11] (see Section II-A for more information). Coupled with low capital equipment costs, printed electronics enable rapid and inexpensive manufacturing. On the other hand, the low-resolution printing results in large feature sizes, low integration density, and high device latencies [12], [13]. These limitations pose challenges for implementing printed Machine Learning (ML) classifiers (i.e., complex datapaths with increased gate count), but are essential for numerous applications in the aforementioned domains [14]. These applications frequently involve classifying analog sensor data to derive meaningful insights.

Extensive research has been dedicated to printed neural networks [14]–[20]. To address the inherent limitations of printed electronics—particularly in terms of area and power—two main design paradigms gain interest in printed ML: bespoke [21] and approximate [22]. Bespoke design involves fully customized circuits tailored to each ML model and dataset. This is enabled by the ultra-low fabrication and non-recurring engineering (NRE) costs of printed electronics, achieved by hardwiring model parameters into the circuit. Such a level of customization is unattainable in silicon-based systems.

Efficiency can be further enhanced by leveraging the error resilience of ML applications and employing approximate computing [23]. *Printed electronics and ML synergistically act as enablers for approximate computing, and vice versa.* Printed electronics are subject to strict physical constraints, such as tight area limits and restricted power availability, and can integrate only a very limited number of devices. The problem is exacerbated when ML circuits are considered, which inherently have significant hardware overheads. Approximation techniques that aggressively reduce the gate count are favored to enable practical implementations and deliver ultra-area-efficient solutions. On the other hand, extensive research in approximate computing emerged about fifteen years ago, but it saw limited adoption in real-world applications. This is largely due to the poor generalization of approximate circuits and the high fabrication costs of silicon systems, which discourage designers from creating systems prone to often significant or unforeseen errors or incurring the additional costs of verifying approximate designs. In contrast, as aforementioned, printed electronics offer negligible fabrication costs, fast on-field manufacturing, and a typically short circuit lifespan—often just up to a few days. These advantages empower designers to experiment with and adopt

unconventional computing paradigms, such as approximate computing. This is because, unlike traditional silicon systems where errors can lead to costly failures and re-fabrication, the low-risk context of printed circuits makes approximation-induced inaccuracies far more acceptable. Additionally, bespoke design in printed electronics enables ML circuit-specific approximations, addressing the generalization challenge of approximate circuits and maximizing hardware savings. Finally, the relatively low complexity of target printed ML applications makes the corresponding circuits even more resilient to errors [24], enabling more aggressive approximations and enhancing the applicability and benefits of approximate computing.

In this work, we focus on designing digital approximate printed neural networks using Electrolyte-Gated FET (EGFET) technology [9], which offers good mobility characteristics and enables operation at low supply voltages [25], making it ideal for printed battery-powered applications. Recognizing the prohibitive cost of printed multipliers [16], [17], the state of the art has shifted to multiplier-less networks, such as power-of-2 quantized Multilayer Perceptrons (MLPs) [18], [19] or Ternary Neural Networks (TNNs) [20]. This allows bespoke circuit implementations to use only adders/subtractors and re-wiring for shift operations. However, [16]–[19] require higher precision inputs to maintain acceptable accuracy, *overlooking the often dominant overheads of necessary analog-to-digital converters (ADCs) to process sensor data [26], [27]*. Using 1-bit inputs and simpler analog-to-binary converters (ABCs), in [20], we avoided interfacing costs. Existing approximate printed neural networks offer non-favorable accuracy-area trade-offs. The MLPs in [16], [17] offer low accuracy loss (e.g., 1-2%) but yield mainly modest hardware savings, with poor scalability for larger losses, while the multiplier-less 1-bit TNNs [20] and power-of-2 MLPs [18], [19] often suffer from high accuracy loss (e.g., $> 5\%$).

This paper extends our prior work in [20] and addresses the limitations mentioned above by *co-optimizing* the analog frontend (i.e., ADC) and digital classifier logic, proposing an automated framework for *designing approximate TNNs with arbitrary input precision*. Our approach refines the exploration of the accuracy-area trade-off in printed neural networks. To maximize area efficiency, we design, for the first time, approximate linear threshold gates (LTGs) to replace the TNN hidden neurons and use approximate popcount units in the output ones. Our multi-objective optimization identifies the optimal approximation for each neuron. Unlike the state of the art, our framework consistently delivers area-efficient solutions across all datasets and accuracy thresholds examined, achieving on average 17x lower area for similar accuracy.

To summarize, our novel contributions in this work are:
1) We are the first to investigate arbitrary input precision printed TNNs, co-designing and optimizing both the analog-to-digital interface and the digital classifier.
2) We propose a systematic methodology for approximating LTGs and a comprehensive TNN approximation[1] using our multi-objective optimization framework that incorporates

[1] https://github.com/ehw-fit/arbitrary-input-tnn.

our precise area model alongside our approximate LTG and popcount units. Additionally, we leverage formal verification and Binary-Decision Diagrams to enable the design and error analysis of large approximate LTG circuits.
3) For up to 5% accuracy loss, we enable, for the first time, printed-battery-powered operation of digital printed neural networks across all examined datasets, accounting also for the cost of analog-to-digital interfacing.

II. BACKGROUND & RELATED WORK

A. Printed Electronics

Printed electronics denote a class of fabrication techniques in which conductive materials are deposited onto flexible or rigid substrates using printing-based processes, such as inkjet, screen, and gravure printing [10]. Relying on resistor-nMOS logic, printed technology inherently differs from silicon-based CMOS by employing additive, maskless printing processes instead of complex, subtractive lithography-based fabrication, thus significantly reducing manufacturing costs. Therefore, it enables the infiltration of computing into low-cost applications (e.g., healthcare wearables, fast-moving consumer goods) where attributes such as form factor, porosity, conformality and non-toxicity are prioritized. This work employs the EGFET technology, which leverages electric double layer formation at the semiconductor–electrolyte interface to achieve high gate capacitance and sub-1V operation. Using inkjet-printed indium oxide channels and printed electrolytes, EGFETs combine high carrier mobility with mechanical flexibility, making them ideal for low-power, battery-operated printed systems [9].

Building on these developments, printed neural networks–especially fully connected multilayer perceptron (MLP) circuits realized in printed electronics–have emerged as a way of embedding ML into ultra-low-cost systems. These networks form the core of printed ML applications, which target resource-constrained domains such as smart packaging, diagnostic strips, and wearable healthcare devices. In this context, terms like printed multipliers and digital approximate printed neural networks refer to specific circuit components or optimized implementations within these systems, all built using the same printed technology. ML essentially enhances such printed systems with decision-making capabilities, increasingly critical in targeted far-edge applications.

B. 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, mostly favoring bespoke designs.

Bespoke circuits are customized for specific classifier models trained on particular datasets, eliminating generality-related overheads [21]. This specialization extends to developing and implementing tailored approximation techniques [22]. For instance, Armeniakos et al. [16] proposed post-training weight replacement with more hardware-friendly approximates to reduce the cost of the underlying multipliers. In subsequent

Fig. 1. Overview of a printed classification-based system.

Fig. 2. a) Comparator in EGFET, b) N-bit Flash ADC and c) N-bit SAR ADC. EGFET is a resistor-nMOS only technology.

works, Armeniakos et al. [17] incorporated hardware-friendly weight approximation into the training process and implemented post-training addition approximation using simple truncation. Afentaki et al. [18], [19] 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. The authors in [26], [27] optimize the sensor-processor interface by reducing input representations and pruning the ADC circuit. However, they do not explore classifier design optimizations, and high input precision is primarily needed to enable ADC pruning. Lastly, an alternative approximation method, stochastic computing neural networks, was explored by Weller et al. [15] that frequently leads, however, to an unacceptable drop in accuracy.

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 approximates LTGs and popcounts towards a holistic approximation; it is purely digital, avoiding the noise and variability issues that can affect low-resolution printed analog designs [28]; it considers arbitrary input precision and it begins optimization at the sensor boundary, rather than post-ADC.

III. ARBITRARY PRECISION PRINTED TNNs

This section details the design of our bespoke exact TNN implementation, which serves as the baseline for our approximation in Section IV, and briefly discusses the cost of the printed ADCs needed to support arbitrary precision.

Fig. 1 presents an overview of a classification-based system in printed electronics, as would be deployed in a real-world printed ML application. In our work, we address the sensor interfacing and classifier blocks, by designing arbitrary-precision ADCs and approximate TNNs, respectively.

TABLE I
EVALUATION OF ADCs IN EGFET

ADC		$A^\dagger$	P^*		$A^\dagger$	P^*		$A^\dagger$	P^*
Flash	2-bit	5.3	0.04	3-bit	9.9	0.13	4-bit	24.2	0.32
SAR		19.0	0.43		30.1	0.76		35.8	1.03

†Area (mm^2). *Power (mW) at 0.6V and 200ms latency.

A. Printed ADCs

Research on ADC design in printed electronics, particularly in EGFET technology, remains limited. Therefore, we design and evaluate two popular ADC architectures: Flash [29] and Successive Approximation Register (SAR) [30]. We also evaluated Binary [31] and Sigma-Delta [32] ADCs; however, they are not functional in EGFET technology due to factors like high parasitic effects and the absence of pMOS transistors.

Flash ADCs (Fig. 2b) are ideal for printed applications due to their regularity and simplicity. They operate by comparing the input voltage to multiple reference voltages using an array of analog comparators. The comparator, i.e., the basic building block of the Flash ADC, is presented in Fig. 2a. This parallel, regular structure allows easy implementation. The comparators' output represents the digital value in thermometer code, which is then converted to a binary encoding by counting the number of '1's with a digital parallel counter. An N-bit ADC requires $2^N - 1$ analog comparators.

SAR ADCs (Fig. 2c) sequentially compare the input signal to reference levels, adjusting bit-by-bit, from most to least significant bits. They require only one analog comparator, but also need a Digital-to-Analog Converter (DAC) and more complex digital logic (SAR logic) to store intermediate results and set reference values. Moreover, SAR ADCs require registers that are highly area-inefficient in EGFET. Hence, in SAR logic we use only an N-bit register for the output and a $\log_2 N$ counter, with all other signals generated through counter multiplexing.

Table I reports the hardware overheads of the examined ADCs while the cost of an ABC [20] is 0.005mm^2 and 0.001mW. Measurements are obtained using the EGFET Process Design Kit (PDK) [9] and Cadence Virtuoso. For the examined precisions, Flash outperforms SAR in EGFET and will be used next.

B. Arbitrary Precision Bespoke TNN Circuit

Focusing on area efficiency, a primary constraint in printed circuits [14], we use a single hidden layer with input precision ranging from 1 to 4 bits. Low-bit inputs allow for significantly smaller and less power-intensive ADCs, while also reducing the input precision requirements of the classifier (i.e., TNN). Throughout recent research efforts in ML circuits for printed electronics [16], [17], [27], [33], low input precision (i.e., below 4 bits) has been commonly used since target domains have low performance and precision needs, enabling printed circuits with acceptable area and power overheads [23].

The `sign` activation function is used for the hidden layer, and the output layer uses `argmax`. The `sign` function computes the sign of the weighted sum of each hidden neuron, with `sign(0)=1` [34], [35], while the `argmax` function identifies

Fig. 3. Bespoke exact TNN circuit overview.

the output neuron with the highest value. Ternary weights $\{-1, 0, 1\}$ are used in both the hidden and output layers.

An abstract overview of our exact bespoke TNN circuits is illustrated in Fig. 3. We adopt a fully parallel bespoke architecture, as it has proven more efficient than sequential and/or conventional ones for printed ML circuits [14]. Fully parallel refers to unfolded architectures, with all computations being performed simultaneously (i.e., purely combinational architecture). Bespoke refers to highly customized implementations, specific to the model and dataset, where coefficients are hardwired in the circuit implementation. Such a degree of customization would be infeasible in silicon-based conventional designs (i.e., non-bespoke), due to high fabrication and NRE costs. In brief, each neuron has its own dedicated hardware, incorporating hardwired weights and eliminating the need for registers or memory that are highly inefficient in EGFET. Connections with zero weights are removed from the corresponding neuron, while the sign of non-zero weights determines at design time whether the respective input is added or subtracted.

We implement the hidden neurons using LTGs and the output y of each hidden neuron is given by:

$$y = \begin{cases} 1 & \text{if } \sum_{i=1}^{n} w_i x_i \geq 0 \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where w_i are the neuron weights, x_i are the input activations, and n is the number of inputs. This reduces the precision required to represent y, with -1 being encoded as 0 while 1 remains 1. Note that (1) applies to any input precision, including 1-bit inputs.

We implement the output neurons using popcount units. The output o of each neuron is given by:

$$o = \sum_{j=1}^{m} w_j y_j, \tag{2}$$

where w_j are the neurons' weights and y_j are the outputs of the hidden layer. However, considering the aforementioned encoding for y_j, negation is achieved with a single NOT gate ($0 \rightarrow 1$ and $1 \rightarrow 0$). In addition, due to the bespoke implementation, if a weight is -1, the corresponding activation is inverted before added to the respective popcount, whereas if the weight is 0, the corresponding connection is removed. The product of a zero weight by the respective activation is always 0 (numerical value). The (numerical) 0 is in the middle of -1 and 1 and, thus, in our encoding 0 should be mapped to $\frac{1}{2}$ (i.e., $\frac{-1+1}{2} \rightarrow \frac{0+1}{2}$). Therefore, a $\frac{1}{2}$ is missing from
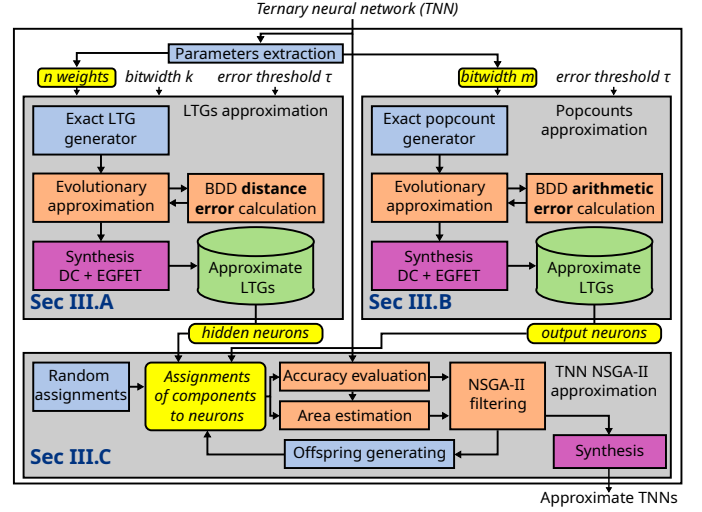

Fig. 4. Overview of the proposed two-phase TNN approximation framework, including Cartesian genetic programming for popcount and LTG approximations in phase one, and their integration into bespoke TNN circuits using the NSGA-II evolutionary algorithm in phase two. The underlying technology assumed is EGFET.

the result for each removed connection due to zero weight. Hence, the output neuron can be implemented by a simple popcount unit (the sum in (2) accumulates only zeros or ones) and to obtain the correct output, a constant correction term $\frac{Z}{2}$ must be added to the popcount result, where Z equals the number of zero weights in that neuron. Instead, we right-shift the popcount result P and add Z: $2P + Z$. Note that Z is a constant known post-training. Since the activation function of the output neuron is argmax, $2P + Z$ gives the same result as $P + \frac{Z}{2}$ without any rounding errors from integer division. As a result, the output neuron computes the following:

$$o = 2 \sum_{\substack{j=1 \\ w_j \neq 0}}^{m} p_j + Z, \text{ where } p_j = \begin{cases} y_j, & \text{if } w_j = 1 \\ \overline{y_j}, & \text{if } w_j = -1 \end{cases}. \tag{3}$$

Similar to w_j, Z values are also known at design time and are thus hardwired in the bespoke circuit implementation.

IV. HOLISTIC TNN APPROXIMATION

Given a dataset, we train a TNN for each considered input precision (1-4 bits), as different precisions can result in varying TNN sizes (i.e., number of hidden neurons) as well as different sensitivity to hardware approximation. For each trained TNN, we generate its exact bespoke circuit, and to maximize area efficiency, *we approximate all its major components*.

Our proposed approximation methodology, illustrated in Fig. 4, consists of two main phases. We first use Cartesian Genetic Programming (CGP) [36] to evolve circuits that approximate the respective LTG and popcount functions, and form our approximation library of area-error Pareto-optimal approximate units for the hidden and output neurons. In the second phase, the Non-dominated Sorting Genetic Algorithm (NSGA-II) [37] is applied to integrate these approximate components into a bespoke TNN circuit, optimizing for maximum resource efficiency with minimal impact on accuracy. Factoring in the analog interfacing cost for each precision,

we conduct a Pareto analysis to identify optimal designs that combine input precision and appropriate approximation. This co-design of the analog front-end and digital classifier allows us to achieve maximum area efficiency for each target accuracy constraint.

Specifically, we replace the LTGs in the hidden layer with our approximate ones and also substitute the popcount units in the output layer with their approximate counterparts. In bespoke circuits, where constant parameters are hardcoded, addition with a fixed term (i.e., Z in (3)) of just a few bits is fairly cheap, making popcount the most resource-intensive part of the output neurons. Approximate popcount units can be used across different neurons or TNNs as long as they require the same number of inputs. Though, some output neurons may require popcount units of the same size but different Z values. Since the cost of adding Z is negligible, generating approximate popcount units for every Z and input size pair would unnecessarily increase our framework's complexity. Lastly, we do not approximate `argmax`, as errors this close to the output are difficult to recover from, while the respective area gains are limited [18].

Despite significant research on arithmetic circuits [22], research on approximate LTG units is limited, as the state of the art mainly approximates the multipliers or adders of ML circuits [24]. This is attributed to the fact that silicon-based implementations do not allow for per-neuron circuit optimizations. Additionally, the binary output of LTGs makes generating and evaluating approximate LTGs more challenging compared to adders or multipliers with multi-bit outputs. Similarly, despite extensive research on approximate adders, dedicated approximations for popcount circuits remain scarce.

A. Approximate LTG Circuit Design

1) LTG Netlist Approximation: Approximating LTGs as a single function is highly efficient, as it enables extensive optimization. As derived from (1), the key requirement is determining whether the multi-operand sum generates a borrow. Thus, approximating the entire LTG in one-shot achieves superior area-error trade-offs compared to approximating and combining its components independently, as usually done in the state of the art for approximating dataflows. An LTG configuration is defined by the number of inputs to add, subtract, and the input precision.

We employ the evolutionary Cartesian Genetic Programming (CGP) algorithm, proven effective for generating approximate circuits [36]. Starting with a gate-level implementation of the exact LTG circuit, our algorithm modifies its structure to produce approximate variants, ensuring the error stays acceptable, i.e., below a predefined threshold τ, fixed to empirically set values (see Section V-A1 for more information on the explored range of τ). The circuit is encoded as an integer-based netlist, with mutation as the sole genetic operator. Each mutation randomly changes a chosen integer, which can affect the connection of a gate's inputs, the gate function, or its output connection.

Operating on a population of $\lambda + 1$ candidates, the initial population includes the exact circuit and λ mutated versions.

A fitness value $F(c)$, reflecting how well each candidate approximates the target circuit while adhering to the error threshold, guides the search. Through generations, the algorithm evolves the population, by selecting the fittest candidates and exploring/refining the solution space via mutation. $F(c)$ represents the area of the approximate LTG circuit when implemented as a printed circuit, and is set to ∞ if the accuracy threshold (τ) is exceeded, discarding that solution:

$$F(c) = \begin{cases} area(c), & \text{if } \varepsilon(c) \leq \tau \\ \infty, & otherwise. \end{cases} \quad (4)$$

where $\varepsilon(c)$ denotes the error of the candidate approximate LTG c. To avoid costly circuit synthesis, $F(c)$ is set as the sum of the area of the netlist's gates, obtained from the EGFET PDK. Each member c in the population 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.

2) LTG Error Analysis: LTGs are basically relational operators with Boolean output, where the output is determined by comparing a weighted sum with zero; hence, Hamming-based error metrics are unsuitable [38]. Given a set of weights $\boldsymbol{w}$, the exact LTG ($\text{LTG}_{\boldsymbol{w}}$) and its approximate variant ($\hat{\text{LTG}}_{\boldsymbol{w}}$), we adopt a more appropriate distance error metric for such functions:

$$D(\boldsymbol{x}) = \begin{cases} 0, & \text{if } \text{LTG}_{\boldsymbol{w}}(\boldsymbol{x}) = \hat{\text{LTG}}_{\boldsymbol{w}}(\boldsymbol{x}) \\ |\sum_{i=0}^{n} w_i x_i| + 1, & \text{otherwise.} \end{cases} \quad (5)$$

If $\text{LTG}_{\boldsymbol{w}}$ and $\hat{\text{LTG}}_{\boldsymbol{w}}$ differ, it means that the two functions produced outputs of opposite signs–that is, one computed a positive weighted sum and the other a negative one. However, since $\hat{\text{LTG}}_{\boldsymbol{w}}$ is implemented as a black-box function via CGP, we only observe its binary output (0 or 1), without any direct insight into how closely it approximated the actual weighted sum. To capture the severity of such mismatches, we define an error metric $D(\boldsymbol{x})$ that quantifies the difference/distance between the weighted sum of $\text{LTG}_{\boldsymbol{w}}$ and the decision threshold at 0, using the exact LTG as the reference value. This provides a meaningful measure of how far the approximate function was from producing the correct output, assuming that $\hat{\text{LTG}}_{\boldsymbol{w}}$ flipped the sign of the sum. An offset of 1 is added to handle cases where the sum is zero, but $\hat{\text{LTG}}_{\boldsymbol{w}}$ wrongly produced a negative sign.

Using $D(\boldsymbol{x})$ we define three statistical error measures, i.e., error probability ε_{ep}, mean distance error ε_{mde}, worst-case distance error ε_{wcde}, and normalized distance error ε_{epmde} involving the erroneous outputs only:

$$\varepsilon_{ep} = \frac{1}{K} \sum_{\forall \boldsymbol{x}} [\![D(\boldsymbol{x}) \neq 0]\!], \quad \varepsilon_{mde} = \frac{1}{K} \sum_{\forall \boldsymbol{x}} D(\boldsymbol{x}),$$
$$\varepsilon_{wcde} = \max_{\forall \boldsymbol{x}} D(\boldsymbol{x}), \quad \varepsilon_{epmde} = \frac{\varepsilon_{mde}}{\varepsilon_{ep}} \quad (6)$$

where K represents the size of the input domain for $\boldsymbol{x}$ and $[\![\cdot]\!]$ denotes the Iverson bracket.

Fig. 5. Miter circuit for calculating distance error $D(\boldsymbol{x})$.

For k-bit inputs (x_i), evaluating the full input space involves $K = 2^{nk}$ stimuli. To expedite error evaluation, Binary Decision Diagrams (BDDs) [39] with a miter circuit (Fig. 5) are used to efficiently compute ε_{ep}, ε_{mde}, and ε_{wcde}.

B. Approximate Popcount Circuits

A popcount configuration is defined by the number of inputs to count. An m-bit popcount operation sums m binary inputs, typically implemented using a tree structure of adders. To approximate this exact circuit, we also employ the evolutionary CGP algorithm. Similarly, the searching algorithm's optimization criterion is the estimated area for the target technology (e.g., EGFET PDK [9]), while maintaining the error ε below a specific threshold, consistent with previous error-oriented approximation research [36]. This error constraint τ is again incorporated into the fitness function as in (4).

Popcount 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. Assuming $P(\boldsymbol{p})$ is the output of the exact popcount for a set of m binary inputs $\boldsymbol{p}$ and $\hat{P}(\boldsymbol{p})$ is its approximate counterpart, ε_{mae} and ε_{wcae} are calculated as follows:

$$\varepsilon_{mae} = \frac{1}{2^m} \sum_{\forall \boldsymbol{p}} |P(\boldsymbol{p}) - \hat{P}(\boldsymbol{p})|,$$
$$\varepsilon_{wcae} = \max_{\forall \boldsymbol{p}} |P(\boldsymbol{p}) - \hat{P}(\boldsymbol{p})| \quad (7)$$

However, for a large number of inputs m in popcount circuits, evaluating all 2^m input vectors becomes impractical. Hence, BDDs [39] are also used in this case for error assessment.

C. Approximate TNN Design

1) Multi-objective optimization: Our framework aims to select the optimal approximate unit for each neuron in the exact TNN, i.e., to identify the most suitable replacement for each exact LTG and popcount from the approximate units in our library. If a required LTG or popcount configuration is not already available in our library, we generate the respective approximate units offline as a one-time effort.

We transform the approximation selection into a multi-objective optimization and solve it using the NSGA-II algorithm [37], with objectives to minimize accuracy loss and area estimation. In printed electronics, leakage power dominates power consumption. Hence, minimizing area minimizes power as well. We encode our optimization problem as an integer list, where each integer selects a component from the library for

the respective neuron. All components, including the TNN, (approximate) LTGs, and (approximate) popcounts, are described in both Verilog and Python, enabling efficient accuracy estimation in Python. Similarly, area estimation is conducted in Python using our surrogate area model. The Pareto-optimal circuits identified by NSGA-II are synthesized and evaluated using Electronic Design Automation (EDA) tools to determine the actual area and power values of our approximate TNNs.

2) Area Estimator: Since area is additive, we use as a surrogate area model the sum of the areas of the selected approximate units. During library generation, all approximate units are synthesized to determine their area. Implementing a holistic approximation that covers all major TNN components enables our estimator to achieve high accuracy. The only non-approximated components are `argmax` and constant addition by Z, which are insignificant in size.

We validate our estimator across 500 approximate TNNs with various datasets and approximation levels. It demonstrates a *Pearson correlation of* 0.995 *and an R²* *value of* 0.969, indicating both near-perfect linear correlation, enabling us to effectively guide our optimization toward more area-efficient solutions, but also precise area estimation.

V. RESULTS AND ANALYSIS

The goal of the presented research is to develop arbitrary input precision bespoke TNN circuit designs that maximize 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.

Experimental Setup: For our evaluation, we use eight datasets from the UCI ML repository [40], selected for two reasons: they allow direct comparison with the state of the art [14], [17]–[20], and are well-suited for sensor-based printed applications as detailed in [14], [15]. Such neural networks are orders of magnitude simpler than contemporary Convolutional Neural Networks and Large Language Models targeted by silicon systems. *However, printed circuits can integrate only a few hundred to thousand gates per chip, making them both representative and highly complex examples for printed electronics to address.* For instance, Pendigits is the most complex dataset explored by the current state of the art [14]–[20], yet no implementation to date has achieved an adequate trade-off in terms of area, power, and accuracy. We use Synopsys Design Compiler, PrimeTime, and the EGFET standard cell library [9] for synthesis and hardware evaluation. The voltage supply for all our circuits is set to 0.6 V, in line with EGFET capabilities and consistent with prior work in the field [18], [19]. The Verilog description of our TNNs is generated directly from the Python model using the description of the LTGs and popcounts in our library.

TNN Training: Exact TNN models are used as the baseline for accuracy comparisons. Input features are normalized within $[0, 1]$, datasets follow a $70\%/30\%$ split for training/testing and TNNs are trained using QKeras [41] and its `quantized_bits` functionality. We employ the Adam
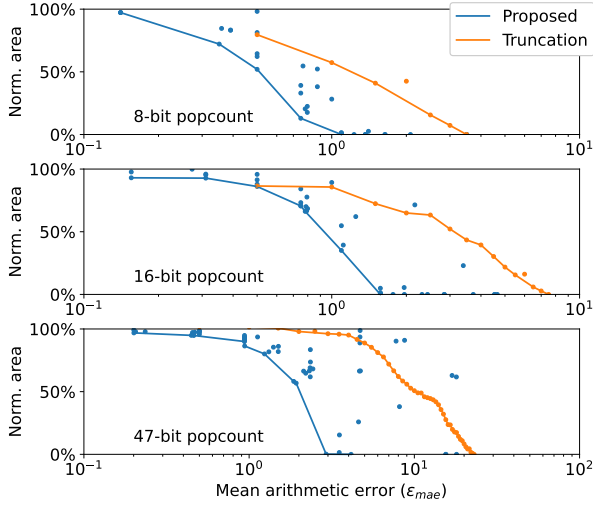
Fig. 6. Comparison between the proposed approximation approach and the previously used truncation technique for popcount circuits of various sizes. The results displayed are based on post-synthesis area measurements.
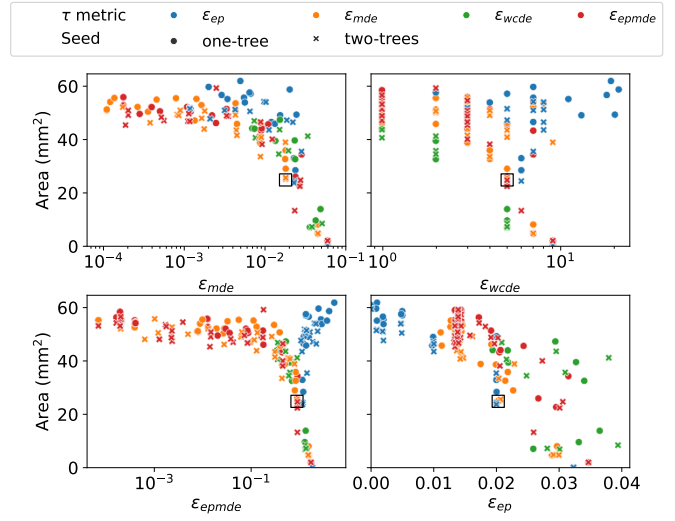

Fig. 7. Approximate LTGs realizing Pendigits' first hidden neuron with ten 2-bit inputs. Exact LTG area: 59.1mm^2. The results displayed are based on post-synthesis area measurements. Highlighted design points are in rectangles.

optimizer for up to 30 epochs of training (with early stopping). 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. Hidden neurons are selected via grid search within the $[1, 50]$ range. For each dataset, we train a TNN for input precisions from 1 to 4 bits and, for each precision, select the TNN with the smallest hidden-layer size that achieves accuracy closest to the exact MLP [14].

A. Approximate Components Evaluation

1) Popcount Circuits: We generate approximate popcount circuits with the following constraints and configurations. The error limits τ_{mae} and τ_{wcae} are logarithmically distributed from 0.1 to $0.5 \cdot 2^g$ and from 1 to $0.5 \cdot 2^m$, with $g = \lceil \log_2 m \rceil$, resulting in 2,090 approximate circuits. We also set CGP search termination criteria to 30, 60, and 300 minutes for popcount sizes $m < 16$, $m < 32$, and $m < 60$, respectively. Using BDD evaluation, we achieve average speeds of 7,759 and 1,362 evaluations per second for 16-bit and 32-bit popcounts. For 60-bit popcount 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 popcount circuit sizes is provided in Fig. 6. The featured circuits are synthesized and evaluated based on the error metrics and methodology described above. For easier comparison, area results are normalized relative to the exact popcount 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 popcount circuits with variants derived using the truncation approximation approach, which has previously been applied

in a wide range of approximate circuits, including MLPs [42]. As shown in Fig. 6, our approximation achieves significantly superior trade-offs compared to truncation, and delivers about 2x area reduction for ε_{mae} of only 0.5, 1.1, and 1.9 for 8-bit, 16-bit, and 47-bit popcounts, respectively. Similar results to those shown in Fig. 6 are observed for all popcount circuit sizes.

2) LTG Circuits: The exact LTG can be implemented either as one adder tree that sums or subtracts inputs and checks for a borrow, or as two separate adder trees—one to sum the inputs to be added and another to sum the inputs to be subtracted—whose results are then subtracted. We generate our approximate LTGs (Section IV-A) for these exact implementations and optimize them against the error metrics of (6). Fig. 7 depicts the error-area trade-off for all our approximate LTGs, using a ten 2-bit inputs (n=10, k=2) neuron as an example. As shown, the implementation with two adder trees provides more efficient solutions (✖ are mainly below ● for similar error). This is also confirmed by the inverted hypervolume (i.e., area under the curve (AUC), where lower values indicate better performance) for the same neurons, presented in Fig. 8. On average, the two adder tree implementation achieves an $8.43 \pm 4.74\%$ improvement over the single case, and is therefore used next. Moreover, Fig. 7 provides a descriptive enough illustration of the gains of our LTG approximation. For example, for $\varepsilon_{mde} \leq 0.018$ in Fig. 7, 58% area reduction is achieved compared to the exact LTG (indicated by black rectangles in all subplots). Similar results are obtained for all LTG configurations required by other datasets and neurons.

For the TNNs examined, the LTGs required are relatively large, with $n \cdot k$ averaging 22. When $n \cdot k = 32$ and 40, the BDD-based LTG error evaluation takes at most 0.09s and 0.21s, respectively, which is 450x and 1850x faster than vectorized simulation [43], requiring 42s and 392s. This rapid error assessment enables efficient traversal of the approxima-

Fig. 8. Inverted hypervolume (i.e., AUC from [0, 0] point) for 2-bit LTGs realizing the hidden neurons of our Pendigits TNN. Both implementations from Fig. 7 (i.e., one adder tree or two adder tree architecture) are considered.


Fig. 9. Approximate 2-bit Pendigits TNNs using different LTG libraries, each optimized against specific error metric.

B. 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 LTG (hidden neurons) and popcount (output neurons) circuit libraries. We apply our proposed approximation to all examined datasets. For each TNN, if the required LTG and popcount configurations are not already in our library, they are generated as described above.

Fig. 9 shows the results of our multi-objective optimization for the 2-bit Pendigits dataset using various approximate LTG libraries, each evolved with a distinct error metric. Libraries containing only Pareto-optimal LTG units ("pareto") yield superior outcomes. For high accuracy, LTG designs optimized for the ε_{wcde} metric are better, while those evolved for the ε_{mde} metric are preferable for maximizing area gains. Similar findings to Fig. 9 apply to other datasets. Hence, for each TNN, the required LTG circuits are approximated against both ε_{mde} and ε_{wced} to populate our library. In Fig. 10, we further investigate the role of our D error metric in approximating LTG units for achieving area-efficient and high-accuracy approximate TNNs. Fig. 10 presents a correlation analysis between ε_{mde} and classification accuracy, when all TNN hidden neurons are replaced with approximate LTGs of identical ε_{mde}, for a wide range of error values. Two TNNs are used in this analysis: our most complex one (Pendigits) and a relatively small one (RedWine). A strong non-linear correlation between our ε_{mde} metric and the classification accuracy can be observed, especially as the network's knowledge capacity increases–i.e., when moving from smaller (RedWine) to larger (Pendigits)

Fig. 10. Correlation analysis between our distance error metric (ε_{mde}) and classification accuracy for (a) Pendigits and (b) RedWine TNNs, when approximate LTGs of the same ε_{mde} error are used for their hidden neurons.

networks, or from lower (1-bit) to higher (4-bit) input precision. Fig. 10 validates that we obtain the expected correlation between neural network inference accuracy and approximate arithmetic units, and confirms that our ε_{mde} metric is suitable for effectively guiding approximation exploration toward high-accuracy regions. Specifically, as ε_{mde} (and thus area savings) increases, there exists a broad region where minimal accuracy degradation is observed at the TNN level, maximizing, thus, the area efficiency of our approximate TNNs.

Fig. 11 shows the accuracy-area trade-off for all explored solutions during our multi-objective optimization, with the Pareto fronts (per input precision) annotated by lines. As Fig. 11 includes 15,000 designs, synthesizing all is impractical, so we report the estimated area using our precise surrogate model. Post-synthesis results for key points are presented in Table II. Note that the area estimation in Fig. 11 reflects only the TNN classifier area, excluding any interfacing costs. In Fig. 11, we include only the input precisions that result in at least one Pareto-optimal point for each dataset, since precisions that lead to only dominated solutions (i.e., with higher area and lower accuracy) are suboptimal in our evolutionary-based optimization. As shown, our framework consistently delivers a smooth accuracy-area trade-off. Supporting arbitrary input precision allows us to effectively populate the accuracy-area Pareto front, as it is primarily composed of designs with varying precisions, with only one precision, especially the low cost 1 bit, rarely dominating.

Two cases require further discussion. For Pendigits with 1-bit input precision, our framework maintains a smooth accuracy-area trade-off despite a significant initial accuracy loss w.r.t the exact 1-bit TNN. Pendigits, being the most complex dataset examined, becomes highly sensitive to approximation with only 1-bit inputs. This is evident from the 1-bit exact Pendigits TNN achieving just 76% accuracy, while the target MLP reaches 94% [14]. This further highlights that it is mandatory to support arbitrary input precision, as 1-bit precision only (e.g., [20]) may be insufficient in many cases.

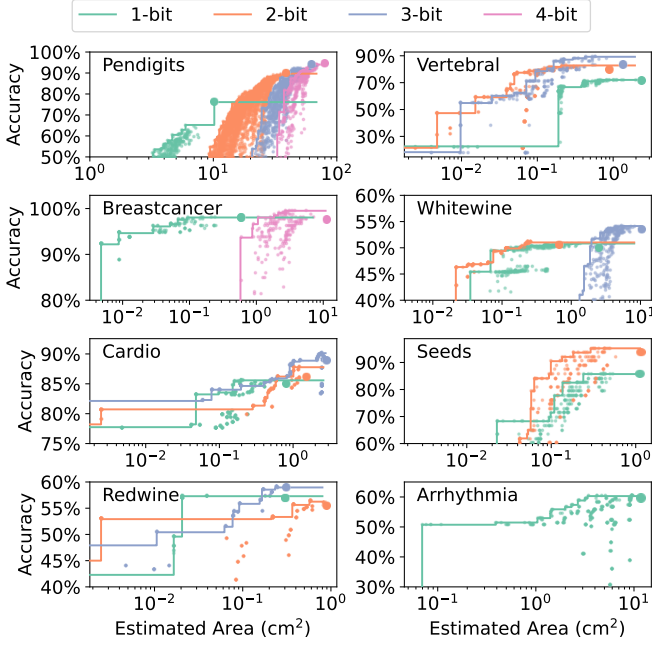For Arrhythmia, we present results only for 1-bit precision,

improved readability, for each dataset and accuracy threshold, we report only the most area-efficient approximate printed MLP [17]–[19]. The table also includes the input precision used by each neural network, emphasizing the significance of our framework's support for arbitrary input precision. As shown in Table II, different datasets and accuracy thresholds mandate varying precision for our TNNs. The reported area and power values include the analog-to-digital interface cost, which our approach minimizes by selecting the most area-efficient solution with respect to both digital classifier and interfacing costs. This often involves choosing the minimum necessary precision due to high ADC overheads. For consistency, the interfacing costs reported in Section III-A are used in all designs in Table II. Hereafter, references to ADC or interfacing costs refer to both ADC and/or ABC.

As shown in Table II, with a maximum 2% accuracy loss compared to the respective exact TNN, our approximate ones achieve on average 2.9x and 4.7x lower area and power, including ADC costs. These values increase to 10.2x and 11.1x when excluding ADC costs, i.e., with respect to only the TNN classifier, demonstrating the high efficiency of our approximation. Compared to the baseline exact MLPs [14], our TNNs feature 88x and 783x, on average, lower area and power (including ADC costs) for up to 2% accuracy loss.

Compared to the most efficient approximate MLP [17]–[19] in each case, our approximate TNNs achieve, on average, 21x lower area and 67x lower power for the 2% accuracy loss threshold. For the 5% threshold, these gains increase to 36x and 139x, highlighting the limited scalability of existing printed MLP approximation frameworks. Excluding ADC costs, the latter gains become 24x and 53x. This demonstrates two key points: i) our framework can apply high degree of approximation and achieve significant area efficiency, even with limited input precision; and ii) minimizing the cost of ADCs is crucial, as for half of the printed neural networks in Table II, including ADCs increases area by over than 47%. Finally, it is noteworthy that [18], [19] fail to meet the 2% accuracy loss constraint for Pendigits and the 5% loss for Seeds and Vertebral.

The comparison with our prior work on approximate TNNs [20] further highlights the effectiveness of this solution. For 1-bit inputs (as in [20]) and identical accuracy, our TNNs achieve, on average, 1.8x and 2.9x lower area and power (including ADC costs), and 5.5x and 5.2x savings when considering only the TNN classifier. These gains (especially the latter) underline the superiority of our approximation strategy. [20] relies on approximate popcount units in both hidden and output layers, requiring multi-bit exact comparators in the hidden layer. In contrast, *we fully approximate each hidden neuron as a single function, improving area-accuracy trade-offs for both the hidden neurons and the TNN overall*. Moreover, [20] is restricted to 1-bit inputs due to its popcount-only approximation and uses only binary weights in the output layer for implementation simplicity. Therefore, [20] fails to meet the 2% constraint for Cardio and WhiteWine, and does not even achieve a 5% accuracy loss for Pendigits, Seeds, and Vertebral. Supporting arbitrary input precision and ternary weights in both layers, *our framework delivers hardware-*
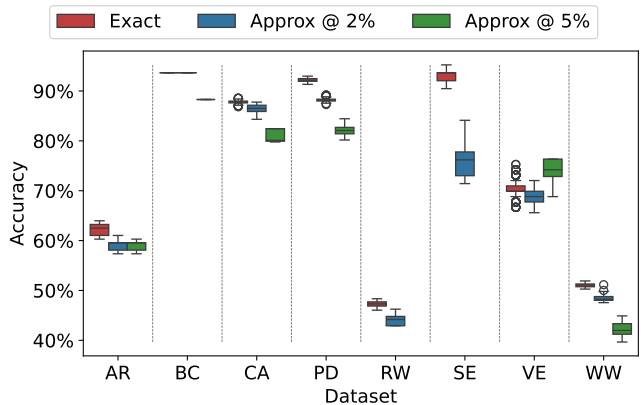

Fig. 12. Classification accuracy of our TNNs in Table II under 10% process variation in the analog interfacing.

efficient solutions across all datasets and accuracy thresholds.

Importantly, *only our TNNs* meet the 30mW power constraint across all datasets, enabling battery operation with an existing printed battery, e.g., Molex 30mW [14]. In printed applications, design feasibility is of paramount concern, taking precedence over maximizing classification accuracy, making the 5% accuracy loss a mostly acceptable constraint [18].

Table II presents a comparative study against the printed neural networks state-of-the-art. For completeness, we also measure the accuracy that could be achieved by a silicon system, e.g., a deep MLP (more than 5 hidden layers and more than 50 neurons per layer) that could run on a TinyML microcontroller. We observed that our TNNs remain in most cases within 5% accuracy loss of the respective FP32 silicon-oriented MLP. The maximum accuracy drop is observed for the Arrhythmia dataset (i.e., 16%), which, due to its increased number of input features, poses a priori a significant challenge for printed implementations [14]. From this analysis, we conclude that, given the immense computational limitations of printed technologies, our TNNs deliver reasonable accuracy compared to the silicon-oriented MLPs, while maintaining realistic area requirements and power consumption compatible with existing printed batteries—thus enabling feasible implementations. It is also noteworthy that these silicon-oriented MLPs require over 24,000 MAC operations, which would translate to unrealistic area requirements in printed electronics.

Finally, we investigate the impact of our approximations on classification boundaries by analyzing the confidence margins of our approximate TNNs relative to the exact ones. By calculating the confidence margins for our approximate circuits of Table II, we observe that they are always maintained above 1 and often remain comparable to—or even better than—the exact TNN, demonstrating that classification boundaries are preserved under the applied approximations. Note, classification boundaries can be seamlessly incorporated into our optimization, by modifying (4) accordingly.

D. Discussion on Process Variation

Process variation is an inherent and significant phenomenon in printed electronics, stemming from the low-resolution print-

ing process. Nevertheless, our TNN classifiers are purely digital and their functionality and output may only be affected by improper timing caused by process variations. In addition, our TNN classifiers implement a fully parallel, unfolded architecture without needing any sequential elements. Therefore, the absence of sequential logic, combined with the large clock slack we employ, makes the functionality our digital classifiers robust to timing variations caused by variability.

Nevertheless, variability may affect the required ADCs/ABCs, which, as analog components, are more susceptible to process variations and can therefore be subject to degradation under imperfections during the manufacturing process. Variations in the analog interfacing may eventually lead to misclassifications in the digital part. For example, variations in the resistor ladder of the Flash ADC/ABC can alter the reference voltage levels of the quantization segments, potentially causing quantization errors at the ADC/ABC output that manifest as noise (or input perturbations) at the input layer of our TNNs. To evaluate the impact of process variation in the analog part on classification accuracy, we perform a Monte Carlo analysis assuming high variation levels (i.e., 10% in the analog components). Specifically, for our exact and Pareto-optimal approximate TNNs of Table II, we run a 200-point Monte Carlo simulation, capturing the inference accuracy on the test dataset under process variation in the ADCs/ABCs.

Fig. 12 illustrates the obtained Monte Carlo results. As expected, classification accuracy is influenced by process variation, since small variances in the resistor ladder can slightly shift the voltage reference levels, potentially leading to significant errors due to the quantization effect of low-precision ADCs. Nevertheless, the accuracy variation observed in Fig. 12 remains well constrained in most cases, demonstrating the inherent robustness of TNNs to input perturbations. For example, the average standard deviation across all TNNs in Fig. 12 is only 0.8%. The largest dispersion is observed in the approximate Seeds TNN, where the range of obtained accuracies (i.e., maximum minus minimum accuracy measured) in the Monte Carlo analysis is only 13%, while the average accuracy range across all datasets is just 4%. The worst-case accuracy drop compared to the respective ideal (variation-free) TNN is, on average, 7%. Fig. 12 demonstrates that even with the very high process variation inherent in printed electronics, the impact on classification accuracy is not catastrophic and remains well bounded in most cases. Moreover, since in our work the classifier is purely digital, the impact of process variation is significantly smaller compared to the current state-of-the-art analog printed classifiers [46].

Although variability-induced effects are shown to be well bounded in our TNNs, approaches such as Monte Carlo simulations or variation-aware training may be used to enhance our system's tolerance to process. Numerous approaches exist that include various robustness-aware training methodologies that account for manufacturing variations [46], aging [47], reliability [28], and sensing uncertainty [48] during the design process [49]. While addressing process variation mitigation is beyond the scope of our work, such methods are orthogonal to our approach and can be seamlessly integrated.

E. Execution Time Discussion

Training our arbitrary input precision TNNs took a maximum of 25 minutes. After training, we identify the required LTG and popcount configurations; if any are missing from our library, we generate the corresponding approximate components. Fairly large LTGs are required with $n \cdot k$ reaching up to 89. For each LTG approximation (i.e., threshold τ), we perform three independent CGP runs, lasting between 30min to 2h, based on the exact LTG's complexity ($k \cdot n$). Similar timings are used to generate the approximate popcount units. However, all LTG and popcount approximations run concurrently, not impeding the scalability of our library generation. Finally, our multi-objective optimization, which approximates the TNN using our library of approximate components and NSGA-II, is very fast as it performs all evaluations in Python without requiring time-consuming hardware synthesis or simulations. In the worst case, it required only 6 minutes for Pendigits with 4-bit inputs. Experiments ran on a AMD Ryzen 5 3600 with 32GB RAM.

VI. CONCLUSIONS

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. Specifically, we propose an automated framework for designing hardware-efficient approximate printed TNNs by co-optimizing the analog front-end and digital classifier. Supporting arbitrary input precision and holistic approximation, we show the importance of approximating neurons as a single function and the entire network overall. Our approach enables fine-tuned area-accuracy trade-offs while minimizing analog-to-digital interfacing overhead, crucial in sensor-based printed applications. Our TNNs achieve 88x lower area than the state-of-the-art exact baseline, making them ideal for resource-limited environments, being the only solution enabling printed-battery-powered operation with up to 5% accuracy loss.

REFERENCES

[1] P. A. Gargini, "Overcoming semiconductor and electronics crises with irds: Planning for the future," *IEEE Electron Devices Magazine*, vol. 1, no. 3, pp. 32–47, 2023.
[2] R. K. Mishra *et al.*, "Simultaneous detection of salivary δ9-tetrahydrocannabinol and alcohol using a wearable electrochemical ring sensor," *Talanta*, vol. 211, p. 120757, 2020.
[3] X. Luo, "Application of inkjet-printing technology in developing indicators/sensors for intelligent packaging systems," *Current Opinion in Food Science*, vol. 46, p. 100868, 2022.
[4] A. Beniwal, P. Ganguly, A. K. Aliyana, G. Khandelwal, and R. Dahiya, "Screen-printed graphene-carbon ink based disposable humidity sensor with wireless communication," *Sensors and Actuators B: Chemical*, vol. 374, p. 132731, 2023.
[5] J.-W. Lee *et al.*, "High sensitivity flexible paper temperature sensor and body-attachable patch for thermometers," *Sensors and Actuators A: Physical*, 2020.

[6] Y. Li *et al.*, "The soft-strain effect enabled high-performance flexible pressure sensor and its application in monitoring pulse waves," *Research*, 2022.

[7] J. Gao *et al.*, "Ultra-robust and extensible fibrous mechanical sensors for wearable smart healthcare," *Advanced Materials*, vol. 34, no. 20, p. 2107511, 2022.

[8] K. Zhou, R. Ding, X. Ma, and Y. Lin, "Printable and flexible integrated sensing systems for wireless healthcare," *Nanoscale*, vol. 16, pp. 7264–7286, 2024.

[9] N. Bleier *et al.*, "Printed microprocessors," in *Annu. Int. Symp. Computer Architecture (ISCA)*, jun 2020, pp. 213–226.

[10] Z. Cui, *Printed electronics: materials, technologies and applications*. John Wiley & Sons, 2016.

[11] J. S. Chang, A. F. Facchetti, and R. Reuss, "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.

[12] T. Lei *et al.*, "Low-voltage high-performance flexible digital and analog circuits based on ultrahigh-purity semiconducting carbon nanotubes," *Nature communications*, vol. 10, no. 1, p. 2161, 2019.

[13] G. Cadilha 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.

[14] M. H. Mubarik *et al.*, "Printed machine learning classifiers," in *Annu. Int. Symp. Microarchitecture (MICRO)*, 2020, pp. 73–87.

[15] D. D. Weller *et al.*, "Printed stochastic computing neural networks," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, 2021, pp. 914–919.

[16] G. Armeniakos, G. Zervakis, D. Soudris, M. B. Tahoori, and J. Henkel, "Cross-layer approximation for printed machine learning circuits," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2022, pp. 190–195.

[17] G. Armeniakos, G. Zervakis, D. Soudris, M. B. Tahoori, and J. Henkel, "Co-design of approximate multilayer perceptron for ultra-resource constrained printed circuits," *IEEE Trans. Comput.*, pp. 1–8, 2023.

[18] F. Afentaki *et al.*, "Bespoke approximation of multiplication-accumulation and activation targeting printed multilayer perceptrons," in *Int. Conf. on Computer Aided Design (ICCAD)*, 11 2023, pp. 1–9.

[19] F. Afentaki, M. Hefenbrock, G. Zervakis, and M. B. Tahoori, "Embedding hardware approximations in discrete genetic-based training for printed mlps," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 3 2024.

[20] V. Mrazek *et al.*, "Evolutionary approximation of ternary neurons for on-sensor printed neural networks," in *International Conference On Computer Aided Design (ICCAD)*, 2024.

[21] H. Cherupalli, H. Duwe, W. Ye, R. Kumar, and J. Sartori, "Bespoke processors for applications with ultra-low area and power constraints," in *Annu. Int. Symp. Computer Architecture (ISCA)*, 2017, pp. 41–54.

[22] H. Jiang, F. J. H. Santiago, H. Mo, L. Liu, and J. Han, "Approximate arithmetic circuits: A survey, characterization, and recent applications," *Proceedings of the IEEE*, vol. 108, no. 12, pp. 2108–2135, 2020.

[23] J. Henkel *et al.*, "Approximate computing and the efficient machine learning expedition," in *International Conference On Computer Aided Design (ICCAD)*, 2022, pp. 1–9.

[24] G. Armeniakos, G. Zervakis, D. Soudris, and J. Henkel, "Hardware approximate techniques for deep neural network accelerators: A survey," *ACM Comput. Surv.*, vol. 55, no. 4, nov 2022.

[25] C. Marques *et al.*, "Progress Report on "From Printed Electrolyte-Gated Metal-Oxide Devices to Circuits"," *Advanced Materials*, vol. 31, 2019.

[26] F. Afentaki, P. C. L. Duarte, G. Zervakis, and M. B. Tahoori, "Reducing adc front-end costs during training of on-sensor printed multilayer perceptrons," *IEEE Embedded Systems Letters*, vol. 16, no. 4, pp. 353–356, 2024.

[27] P. C. Lozano Duarte, F. Afentaki, G. Zervakis, and M. B. Tahoori, "Design and In-training Optimization of Binary Search ADC for Flexible Classifiers," in *Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2025.

[28] H. Zhao *et al.*, "Highly-bespoke robust printed neuromorphic circuits," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023.

[29] A. Kumary and S. Rao, "Design techniques of flash adc: Review," in *Advances in Communication, Signal Processing, VLSI, and Embedded Systems*, S. Kalya, M. Kulkarni, and K. Shivaprakasha, Eds. Singapore: Springer Singapore, 2020, pp. 89–95.

[30] X. Tang *et al.*, "Low-power sar adc design: Overview and survey of state-of-the-art techniques," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 69, no. 6, pp. 2249–2262, 2022.

[31] S. P. M. Bhai and D. N. Gaonkar, "Design of binary search adc using n comparators," in *2016 IEEE First International Conference on Control, Measurement and Instrumentation (CMI)*, 2016, pp. 499–502.

[32] O. Bajdechi, G. Gielen, and J. Huijsing, "Systematic design exploration of delta-sigma adcs," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 51, no. 1, pp. 86–95, 2004.

[33] G. Armeniakos *et al.*, "On-sensor printed machine learning classification via bespoke adc and decision tree co-design," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 3 2024.

[34] M. Courbariaux, Y. Bengio, and J.-P. David, "Binaryconnect: training deep neural networks with binary weights during propagations," in *International Conference on Neural Information Processing Systems (NeurIPS)*, 2015, p. 3123–3131.

[35] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or -1," *arXiv:1602.02830*, 2016.

[36] V. Mrazek, S. S. Sarwar, L. Sekanina, Z. Vasicek, and K. Roy, "Design of power-efficient approximate multipliers for approximate artificial neural networks," in *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2016, pp. 1–7.

[37] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, April 2002.

[38] V. Mrazek and Z. Vasicek, "Axmed: Formal analysis and automated design of approximate median filters using bdds," in *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2025, p. 5.

[39] V. Mrazek, "Optimization of bdd-based approximation error metrics calculations," in *2022 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2022, pp. 86–91.

[40] D. Dua and C. Graff, "UCI machine learning repository," 2017.

[41] C. N. Coelho *et al.*, "Ultra low-latency, low-area inference accelerators using heterogeneous deep quantization with qkeras and hls4ml," *arXiv preprint arXiv:2006.10159*, p. 108, 2020.

[42] S. Venkataramani, A. Ranjan, K. Roy, and A. Raghunathan, "Axnn: Energy-efficient neuromorphic systems using approximate computing," in *Int. Symp. on Low Power Electronics and Design (ISLPED)*, 2014, pp. 27–32.

[43] L. Sekanina, Z. Vasicek, and V. Mrazek, *Automated Search-Based Functional Approximation for Digital Circuits*. Cham: Springer International Publishing, 2019, pp. 175–203.

[44] V. Mrazek, M. A. Hanif, Z. Vasicek, L. Sekanina, and M. Shafique, "autoAx," in *Proceedings of the 56th Annual Design Automation Conference 2019*. New York, NY, USA: ACM, 6 2019, pp. 1–6.

[45] G. Zervakis, S. Xydis, D. Soudris, and K. Pekmestzi, "Multi-level approximate accelerator synthesis under voltage island constraints," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 66, no. 4, pp. 607–611, 2019.

[46] P. Pal *et al.*, "Neural architecture search for highly bespoke robust printed neuromorphic circuits," in *Proceedings of the 42nd IEEE/ACM International Conference on Computer-Aided Design*, 2024.

[47] H. Zhao, M. Hefenbrock, M. Beigl, and M. B. Tahoori, "Aging-aware training for printed neuromorphic circuits," in *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2022, pp. 1–9.

[48] H. Zhao, M. Hefenbrock, M. Beigl, and M. B. Tahoori, "Highly-dependable printed neuromorphic circuits based on additive manufacturing," *Flexible and Printed Electronics*, vol. 8, no. 2, p. 025018, 2023.

[49] M. B. Tahoori, E. Ozer, G. Zervakis, K. Balaskas, and P. Pal, "Computing with printed and flexible electronics," in *European Test Symposium*, 2025.

Vojtech Mrazek is an Assistant Professor at the Brno University of Technology. He received M.Sc. and Ph.D. degrees in information technology from the Faculty of Information Technology, Brno University of Technology, Czech Republic, in 2014 and 2018, respectively. Before he was also a visiting post-doc researcher at Institute of Computer Engineering, Technische Universität Wien (TU Wien), Vienna, Austria (2018-2019). His research interests are approximate computing, genetic programming, formal verification and machine learning. He has authored or co-authored over 60 conference/journal papers focused on approximate computing and evolvable hardware.

Konstantinos Balaskas is a post-doctoral researcher at the University of Patras. He received his Ph.D. degree from the Aristotle University of Thessaloniki in 2024, and was a research associate at the Chair for Embedded Systems (CES), at the Karlsruhe Institute of Technology (KIT), between 2021 and 2024. He received his Bachelor Degree in Physics and Master Degree in Electronic Physics from the Aristotle University of Thessaloniki in 2018 and 2020, respectively. His main research interests include embedded machine learning, electronic design automation and physical-driven approximate computing.

Georgios Zervakis is an Assistant Professor at the University of Patras. Before that he was a Research Group Leader at the Chair for Embedded Systems (CES), at the Karlsruhe Institute of Technology (KIT) from 2019 to 2022. He received the Diploma and Ph.D. degrees from the School of Electrical and Computer Engineering (ECE), National Technical University of Athens (NTUA), Greece, in 2012 and 2018, respectively. Dr. Zervakis serves as a reviewer in many IEEE and ACM journals and is also a member of the technical program committee of several major design conferences. He has received one best paper nomination at DATE 2022. His main research interests include low-power design, accelerator microarchitectures, approximate computing, and machine learning.

Paula Carolina Lozano Duarte is a Ph.D. student at the Chair of Dependable Nano-Computing at the Karlsruhe Institute of Technology, Germany. She received her Bachelor's degree (2021) and Master's degree (2023) in Telecommunications Engineering from the Public University of Navarra, Spain. Her main research interests include printed and flexible electronics, machine learning classifiers, and neuromorphic computing.

Zdenek Vasicek received all his degrees from Brno University of Technology, Czech Republic, where he is currently an Associate professor. He holds a Ph.D. (2012) and an M.S. equivalent (2006) in Computer Science and Engineering. His research interests include formal verification techniques and application of evolutionary approaches in areas related to the design and optimization of complex digital circuits and systems. He is an active PC member of several evolutionary conferences such as EuroGP, GECCO, and ICES. Dr. Vasicek received the Silver and Gold medals at HUMIES, in 2011 and 2015, respectively.

Mehdi B. Tahoori (M'03, SM'08, F'21) is a Professor at the Chair of Dependable Nano-Computing at Karlsruhe Institute of Technology, Germany. He received the B.S. degree in computer engineering from Sharif University of Technology, Iran, in 2000, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2002 and 2003, respectively. He is currently the deputy editor-in-chief of IEEE Design and Test Magazine. He was the editor-in-chief of Microelectronic Reliability journal. He was the program chair of VLSI Test Symposium in (VTS) 2021 and 2018, and General Chair of European Test Symposium (ETS) in 2019. He is the chair of the IEEE European Test Technology Technical Council (eTTTC). Prof. Tahoori was a recipient of the US National Science Foundation Early Faculty Development (CAREER) Award in 2008. He has received a number of best paper nominations and awards at various conferences and journals. He is a recipient of European Research Council (ERC) Advanced Grant.