

Article

A 28 nm FD-SOI Current-Mode Synaptic Weighting Cell for Low-Complexity Event-Based NILM MLP Inference

Zhiwei Ma , Yoann Charlon * , Erwin Franquet  and Gilles Jacquemod

Polytech'Lab, UPR UniCA 7498, Université Côte d'Azur, 06903 Sophia Antipolis, France;
zhiwei.ma@univ-cotedazur.fr (Z.M.); erwin.franquet@univ-cotedazur.fr (E.F.);
gilles.jacquemod@univ-cotedazur.fr (G.J.)

* Correspondence: yoann.charlon@univ-cotedazur.fr

Abstract

This paper presents a digitally controlled current-mode synaptic weighting cell in 28 nm FD-SOI CMOS for low-complexity Multi-Layer Perceptron (MLP) inference in event-based Non-Intrusive Load Monitoring (NILM). A compact bias-free [16,16] MLP is trained offline by backpropagation for fixed-weight feedforward inference. Using 616 ON/OFF events extracted from high-frequency REDD measurements, six appliance classes are characterized by four event-level features: active-power variation, reactive-power variation, current total harmonic distortion of the differential event signature, and event interval. With 16-level input quantization, the model achieves 92.9–93.4% test accuracy and 90.5–90.7% test Macro-F1, requiring 320 weighted-sum branches across two hidden layers. A four-input first-hidden-layer weighted-sum unit is selected as a representative circuit instance. Its computation is mapped to bounded current ranges using current-coded inputs, an 8-bit magnitude-controlled current-mode multiplier, sign-bit current steering, differential current accumulation, and signed-current scaling. The circuit contribution focuses on the weighted-sum datapath, particularly the repeated synaptic weighting cell; activation and complete classifier implementation are outside the scope of this work. Transistor-level PVT and supply-variation simulations validate the signed weighted-sum path, while post-layout extraction evaluates the repeated multiplier cell. The results demonstrate the block-level feasibility of digitally programmable current-mode synaptic weighting for compact event-based NILM inference.

Keywords: non-intrusive load monitoring; event-based NILM; current-mode synaptic weighting cell; signed weighted sum; 28 nm FD-SOI; analog neural network



Academic Editor: Costas Psychalinos

Received: 17 June 2026

Revised: 17 July 2026

Accepted: 18 July 2026

Published: 21 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Non-Intrusive Load Monitoring (NILM) estimates appliance-level activity from aggregate electrical measurements and is a key technique for energy feedback, demand-side management, and edge energy-monitoring systems [1–3]. Among NILM formulations, event-based NILM is particularly attractive for embedded implementation because it avoids continuous processing of long voltage or current waveforms. Instead, appliance activity is represented by switching events and by compact electrical descriptors extracted around each transition [4,5]. In this work, a six-class event-based NILM task is derived from high-frequency Reference Energy Disaggregation Dataset (REDD) measurements [6], where each retained ON/OFF event is represented by four features: active-power variation ΔP ,

reactive-power variation ΔQ , current total harmonic distortion of the differential event signature $\text{THD}_i(\Delta I)$, and event interval Δt [7].

Neural classifiers have been widely applied to NILM because their nonlinear decision boundaries can capture correlations among multidimensional electrical features [8–11]. For the present hardware-oriented study, a Multi-Layer Perceptron with Backpropagation (MLP-BP) is selected as a compact nonlinear model whose elementary inference operations are the weighted-sum computation and nonlinear activation [12]. The weighted-sum operation can be naturally mapped to current-domain analog computation, while the activation operation can be implemented by a compact nonlinear back end. Backpropagation is used in this paper only as an offline training algorithm to obtain fixed weights [13]. The targeted hardware operation is therefore feedforward analog inference, not on-chip learning. Accordingly, the circuit-level investigation is deliberately restricted to the repeated weighted-sum operation. The nonlinear activation and output decision stages are retained in the software model and are therefore beyond the transistor-level scope of this work.

Digital Field-Programmable Gate Array (FPGA) and System-on-Chip (SoC) implementations have demonstrated the feasibility of local NILM processing [7,14,15]. However, in very low-power edge nodes, the energy cost may still be affected by Analog-to-Digital Converter (ADC) activity, memory accesses, clocked multiply–accumulate operations, and data movement. This motivates current-mode analog inference approaches in which the weighted-sum operation is performed locally in the current domain through feature-to-current coding, magnitude-controlled weighting, sign handling, and current accumulation. In a current-mode neuron, feature values can be represented by input currents, weight magnitudes can be implemented by current scaling, and the weighted-sum operation can be obtained by accumulating signed-current contributions. Such current-domain principles are consistent with analog Very-Large-Scale Integration (VLSI) and neuromorphic computation, and with programmable current-mirror Vector–Matrix Multiplier (VMM) approaches [16,17]. Recent current-mode analog Artificial Neural Network (ANN) classifiers also confirm the relevance of current-domain neural computation for low-power inference [18].

Although neural-network-based NILM and embedded digital NILM inference have been widely investigated, and analog circuits have been reported for low-power classification [19], their combination at the NILM-oriented transistor-level design stage remains limited. In particular, prior works rarely provide an explicit mapping from compact event-based NILM features and fixed signed MLP weights. This work therefore focuses on the current-mode weighted-sum path required for a compact event-based NILM MLP-BP inference path. This path combines repeated weighted-sum branches, sign-bit current steering, and signed-current accumulation.

The circuit objective is to validate a compact MLP-compatible current-mode weighted-sum path, centered on a repeated synaptic weighting/multiplier cell and its associated signed-current processing principle. Based on previous analog MLP specification studies from our team [20], where non-ideal multiplier and activation blocks were injected at system level, the present Cadence-level design targets a relative error below $\pm 20\%$ for the current-mode weighted-sum path. This value is used as a first circuit-level specification for the weighted-sum path, not as a complete classifier-level robustness guarantee. The weight magnitude and weight sign are treated separately: a magnitude-controlled multiplier generates a non-negative weighted current, while a sign-bit-controlled steering stage and signed-current scaling translate the sign information into positive or negative current contributions around a midpoint current. This organization avoids a fully signed multiplier at branch level and keeps the activation input current within a prescribed range. The circuit is implemented in 28 nm Fully Depleted Silicon-On-Insulator (FD-SOI), selected for

compact repeated current-mode branches and low-voltage and low-power operation [21,22]. The body-bias capability of FD-SOI may provide a future calibration knob, although body-bias calibration is outside the scope of the present work [23]. The Sigmoid activation function is retained in the software model because it represents a suitable candidate for future hardware implementation: its characteristics are compatible with compact weak-inversion CMOS circuits, and it can provide a bounded nonlinear current response at the output of each hidden-layer stage [16,24–26]. However, the transistor-level design and validation of the Sigmoid activation circuit are beyond the scope of this work. Figure 1 summarizes the proposed modeling-to-circuit flow from event-based NILM preprocessing to the instantiated current-mode weighted-sum path and activation feasibility interface.




Figure 1. Conceptual overview of the proposed low-complexity event-based NILM MLP and the corresponding instantiated 28 nm FD-SOI current-mode synaptic weighting path.

The main contributions of this work are as follows. First, a compact no-bias [16, 16] event-based NILM MLP is selected as a hardware-oriented inference target by jointly considering classification performance, input quantization, bias removal, and hidden-layer branch count. Second, the selected MLP is mapped to bounded current targets for current-mode implementation, including 0–75 nA input currents and a 0–6 μ A activation-input current range. Third, a current-mode weighted-sum path is proposed for fixed-weight MLP-BP inference, combining repeated weighted-sum branches, sign-bit current steering, and signed-current accumulation. Fourth, the proposed first-hidden-layer weighted-sum path is evaluated in 28 nm FD-SOI through transistor-level Process, Voltage, and Temperature (PVT) variations simulations and post-layout verification of the multiplier.

The remainder of the paper is organized as follows. Section 2 presents the event-based NILM dataset, preprocessing, quantization, and compact MLP selection. Section 3 describes the instantiated current-mode weighted-sum path architecture and its bounded current mapping. Section 4 reports transistor-level PVT validation of the weighted-sum path and multiplier-core post-layout validation results. Section 5 discusses the implementation scope and remaining integration steps, and Section 6 concludes the paper.

2. Low-Complexity Event-Based NILM Model

The purpose of the modeling stage is to define a compact and hardware-compatible inference target. The event detection and feature extraction follow the referenced event-based NILM methodology [7], while the present work focuses on the subsequent selection

of a low-complexity MLP and on translating its inputs and fixed signed weights into bounded current-domain quantities. Consequently, the software model is evaluated only to verify that the retained architecture remains a reasonable classification target under input quantization and bias removal.

2.1. Event Dataset and Feature Definition

The event set is obtained from the high-frequency part of the REDD dataset after event detection and feature extraction [6]. Switching events are detected using a Cumulative-Sum (CUSUM) change detector, following the event-based NILM methodology used in previous work [7]. This preprocessing is used to obtain a compact event-level dataset for hardware-oriented MLP selection.

After event localization, only events for which stable pre-event and post-event windows can be identified are retained. Each retained event is represented by four quantities: active-power variation ΔP , reactive-power variation ΔQ , current total harmonic distortion of the differential event signature $\text{THD}_i(\Delta I)$, and event interval Δt . The interval feature is expressed as a number of samples. ON and OFF events are kept, and different sub-power operating modes of the same appliance are merged into a single appliance class. This process gives 616 ON/OFF events distributed over six classes: Bath-GFI, Electronics, Furnace, Kitchen outlet, Microwave, and Washer–Dryer [7]. The distribution of these events among the six appliance classes is summarized in Table 1.

Table 1. Class distribution of the event-based NILM dataset.

Appliance Class	Samples	Percentage
Bath-GFI	31	5.0%
Electronics	100	16.2%
Furnace	192	31.2%
Kitchen outlet	91	14.8%
Microwave	64	10.4%
Washer–Dryer	138	22.4%
Total	616	100.0%

The dataset is imbalanced: Furnace and Washer–Dryer dominate the samples, whereas Bath-GFI and Microwave are less represented. This motivates two choices in the modeling flow. First, balancing is applied only during training. Second, Macro-F1 is evaluated together with accuracy when selecting the classifier, because Macro-F1 gives equal importance to minority and majority classes [27].

Because the retained event set contains only 616 events, the classification results should be interpreted as an event-level hardware-oriented model selection result rather than as a definitive NILM generalization benchmark. The objective of this stage is to identify a compact MLP architecture that remains compatible with the subsequent current-mode weighted-sum path.

2.2. Leakage-Free Preprocessing and Current-Compatible Quantization

The preprocessing stage is designed to avoid data leakage and to produce bounded feature values that can later be represented by input currents. Polarity normalization is applied first. For ON and OFF transitions of the same appliance, the signs of ΔP , ΔQ , and $\text{THD}_i(\Delta I)$ are opposite; the three differential features are therefore multiplied by -1 for OFF events and left unchanged for ON events. The interval feature Δt does not carry a sign and is kept as measured.

The dataset is split before oversampling using three train/validation/test ratios: 0.6/0.2/0.2, 0.5/0.3/0.2, and 0.5/0.2/0.3. The corresponding split sizes and class counts before training-only oversampling are reported in Table 2.

Table 2. Train/validation/test split sizes and class counts before training-only oversampling.

Split Ratio	Subset	Total	Bath-GFI	Electronics	Furnace	Kitchen-Outlet	Microwave	Washer-Dryer
0.6/0.2/0.2	Train	369	19	60	115	55	38	82
	Validation	124	6	20	39	18	13	28
	Test	123	6	20	38	18	13	28
0.5/0.3/0.2	Train	309	16	50	96	46	32	69
	Validation	184	9	30	58	27	19	41
	Test	123	6	20	38	18	13	28
0.5/0.2/0.3	Train	309	16	50	96	46	32	69
	Validation	123	6	20	38	18	13	28
	Test	184	9	30	58	27	19	41

For each split ratio, the partition is performed at the event level before any oversampling or preprocessing step. Controlled bootstrap oversampling is applied only to the training subset, where each class is expanded to 500 samples, while the validation and test subsets are kept unchanged. Therefore, oversampled duplicates cannot appear in validation or test data.

The present evaluation is based on an event-level split derived from a single REDD household event set. Therefore, it does not constitute a household-independent NILM evaluation. Since episode-independent splitting was not enforced, temporally related events from the same appliance operation may appear in different subsets. The reported accuracy and Macro-F1 values are therefore used to select a compact hardware-oriented MLP target under the specified leakage-free preprocessing protocol, rather than to claim general NILM robustness.

Standardization, clipping, normalization, and quantization parameters are determined from the training subset and then kept fixed for validation and test evaluation. This leakage-free preprocessing prevents information from the validation or test subsets from being used during model preparation, while allowing different input quantization resolutions to be compared under the same protocol.

2.3. Compact MLP Selection

A Multilayer Perceptron (MLP) is retained as a compact nonlinear classifier. The model is trained offline in software using backpropagation and is subsequently used for fixed-weight feedforward inference. The hardware study reported here implements only the repeated weighted-sum operation; it does not implement the complete MLP inference chain.

The MLP maps the four event-level features onto six appliance classes. Sigmoid functions are used in the two hidden layers, and a Softmax function produces the six output probabilities in the software model. These nonlinear and decision functions are required for complete model inference, but their circuit implementation is outside the scope of this work.

The explored model space includes 6, 8, 10, 12, 16, 20, and 32 input quantization levels, one to three hidden layers, 4, 8, 16, 32, and 64 neurons per hidden layer, and both biased and bias-free variants. For each tested configuration and split ratio, seven independent training runs were performed with different random initializations. The reported accuracy and Macro-F1 values correspond to the mean test performance over these repeated runs.

For all results in Table 3, the trained weights are quantized using the circuit-compatible representation: an 8-bit magnitude code plus one sign bit. The final selection is not based only on the highest software accuracy. It also accounts for analog complexity, represented here by the number of hidden-layer weighted-sum branches, the number of additional bias current sources, and the resolution required for input-current generation.

Table 3. Quantization-level and compact MLP selection results across different dataset splits and seven independent training runs.

Levels	Configuration	Bias	Branches	Bias Src.	Mean Test Acc.	Mean Test Macro-F1
6	[16, 16]	No	320	0	79.9%	76.0%
8	[16, 16]	No	320	0	81.0%	76.7%
10	[16, 16]	No	320	0	80.2%	75.8%
12	[16, 16]	No	320	0	86.7%	83.6%
20	[16, 16]	No	320	0	89.7%	88.5%
16	[16, 16]	No	320	0	93.2%	90.6%
16	[16, 16]	Yes	320	32	93.3%	90.0%
16	[32, 32]	No	1152	0	93.5%	90.9%
16	[32, 64, 32]	No	4224	0	92.9%	89.9%
16	[64, 64, 32]	No	6400	0	93.0%	90.0%

Table 3 highlights the accuracy–complexity trade-off used to select the hardware-oriented MLP. Among the tested configurations, the 16-level no-bias [16, 16] network provides the most favorable compromise, reaching 93.2% test accuracy and 90.6% test Macro-F1 with only 320 hidden-layer weighted-sum branches. Increasing the network size to [32, 32] or adding bias terms provides only marginal performance gains, while significantly increasing hardware complexity through more weighted-sum branches or additional bias current sources. The no-bias [16, 16] architecture is therefore retained as the compact hardware-oriented design point.

The retained no-bias [16, 16] network was also evaluated under the three proposed train/validation/test partitions to verify that the selected architecture was not tied to a single data split. Across these partitions, the test accuracy remained between 92.9% and 93.4%, and the test Macro-F1 stayed in the 90.5–90.7% range.

Figure 2 reports one test-set confusion matrix obtained for the retained no-bias [16, 16] architecture under the 0.5/0.2/0.3 train/validation/test split. Most predictions remain on the diagonal, giving a test accuracy of 94.6%. The remaining errors are mainly associated with the Kitchen outlet class, while Electronics and Washer–Dryer are correctly classified in this run.

Consequently, the model passed to the circuit-design stage contains four input features, 16 input quantization levels, two hidden layers with 16 Sigmoid neurons per layer, no neural bias terms, and a six-class Softmax output kept at the system level. This configuration is used as a compact hardware-oriented design point, fixing the main current-mode neuron requirements: four first-layer inputs, 320 hidden-layer weighted-sum branches, finite input-current coding, and a reusable nonlinear activation range.

Confusion (NO-BIAS) | L=16 | hidden=[16 16] | testAcc=94.6%

Output Class	Bath-gfi	8 4.3%	0 0.0%	0 0.0%	2 1.1%	0 0.0%	0 0.0%	80.0% 20.0%
	Electronics	0 0.0%	30 16.3%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	100% 0.0%
	Furnace	1 0.5%	0 0.0%	57 31.0%	2 1.1%	1 0.5%	0 0.0%	93.4% 6.6%
	Kitchen outlet	0 0.0%	0 0.0%	0 0.0%	20 10.9%	0 0.0%	0 0.0%	100% 0.0%
	Microwave	0 0.0%	0 0.0%	1 0.5%	3 1.6%	18 9.8%	0 0.0%	81.8% 18.2%
	Washer-Dryer	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	41 22.3%	100% 0.0%
		88.9% 11.1%	100% 0.0%	98.3% 1.7%	74.1% 25.9%	94.7% 5.3%	100% 0.0%	94.6% 5.4%
		Bath-gfi	Electronics	Furnace	Kitchen outlet	Microwave	Washer-Dryer	
		Target Class						

Figure 2. Test confusion matrix of the final architecture.

3. Current-Mode Low-Complexity Weighted-Sum Path Architecture

This section presents the current-mode weighted-sum path derived from the compact no-bias [16,16] MLP selected in Section 2. The objective is to validate the repeated hardware branch used to implement the MLP weighted-sum operation during fixed-weight feedforward analog inference. The circuit prototype corresponds to a four-input weighted-sum unit of the first hidden layer. The same principle can be extended to the second hidden layer by increasing the number of weighted-sum branches and adapting the signed-current scaling ratios.

3.1. Circuit-Level Accuracy Target

The current-mode weighted-sum path is designed with a first circuit-level relative-error target of $\pm 20\%$. This target is derived from previous analog MLP specification work from our team [20], where non-ideal multiplier and activation models were injected into the neural-network model to evaluate their impact on classification performance. In that study, the multiplier and weighting-related errors were shown to be less critical than activation-function errors, which provides a useful first specification for the present weighted-sum path. In this work, the $\pm 20\%$ criterion is therefore used as an initial Cadence-level design target for the current-mode weighted-sum path only. It does not replace a future classifier-level reinjection of extracted circuit nonidealities into the NILM MLP-BP model. This target is used to interpret the PVT and post-layout results of the proposed weighted-sum path and to identify which support blocks require further stabilization or calibration in a complete implementation.

3.2. Weighted-Sum Path Architecture

Figure 3 shows the proposed four-input current-mode implementation of the MLP hidden-layer operation. As highlighted in the figure, the left part implements the weighted-sum operation in current mode, while the I2V-Sigmoid stage is retained only as a conceptual functional block in the system-level signal-flow diagram to indicate the activation

operation required for complete MLP inference. Each quantized NILM feature is represented by an input current and weighted by an 8-bit magnitude-controlled current-mode multiplier. The multiplier output is routed by a sign-bit-controlled current-steering block toward the positive or negative accumulation path. The signed contributions are then accumulated around a midpoint current and processed by the I2V–Sigmoid interface. In this work, the current-mode weighted-sum path includes the repeated input-current generator, the magnitude-controlled multiplier, the sign-bit current steering stage, and the signed-current accumulation block. Table 4 summarizes the current targets for the first-hidden-layer weighted-sum path implementation.




Figure 3. Architecture of the current-mode hidden-layer neuron.

Table 4. Current targets of the proposed hidden-layer weighted-sum path.

Block/Quantity	Input Range	Output Range	Role in the Architecture
Input-current generator	Digital code	0–75 nA	Four-bit current coding of the quantized NILM feature.
Magnitude multiplier	0–75 nA	0–1912.5 nA	8-bit magnitude-controlled weighted-current generation.
Sign-bit controlled current steering	0–1912.5 nA	Positive or negative current path	Branch-level handling of the weight sign.
Signed-current scaling stage	Steered branch current	± 750 nA max. contribution	Limits each first-layer branch contribution around I_{mid} .
Current summation/I2V input	3 μ A plus four signed contributions	0–6 μ A	Midpoint-biased signed accumulation before activation.

3.3. Current-Coded Inputs and Magnitude-Controlled Multiplier

The 16-level input quantization retained at the MLP level is implemented with binary-weighted current branches. With a least significant current step of $I_0 = 5$ nA, the input-current generator produces

$$I_{in}[j] = q_j I_0, \quad q_j \in \{0, 1, \dots, 15\}, \quad I_0 = 5 \text{ nA}, \quad (1)$$

where $j = 1, \dots, 4$ for the first-hidden-layer prototype. This gives the input-current range $I_{in}[j] \in [0, 75]$ nA. In circuit form, the four control bits switch binary-weighted current branches of 40 nA, 20 nA, 10 nA, and 5 nA.

The magnitude-controlled multiplier is implemented as a digitally controlled switched-current-mirror block. The MLP weights are trained offline and then represented by digital control words during analog inference. At the branch level, this block forms the weighting core of the repeated weighted-sum branch. Each enabled branch copies and scales the input current according to its binary weight, and the enabled branch currents are added in the

current domain. In this way, the multiplier generates a non-negative current proportional to the input current and to the stored weight magnitude:


$$I_{\text{mul}}[j] = \frac{|W_j| I_{\text{in}}[j]}{10}, \quad 0 \leq |W_j| \leq 255. \quad (2)$$

Here, $|W_j|$ is the 8-bit digital magnitude of the synaptic weight, while the weight sign is handled by the following sign-bit controlled current steering block. This magnitude-plus-sign representation is the same as that used in the MATLAB R2025a MLP evaluation. The factor 10 is introduced by the current-mirror sizing and normalization used to set the multiplier output-current range before the signed-current steering and scaling stages. The values $I_0 = 5$ nA, the normalization factor of 10, and $I_{\text{mid}} = 3$ μ A are hardware design parameters chosen from a power-area trade-off. The 5 nA step keeps the 16-level input range limited to 0–75 nA, while the factor 10 restricts the multiplier output before signed scaling. After scaling, the four-input accumulation is bounded within ± 3 μ A, and the 3 μ A midpoint current shifts it to the 0–6 μ A input range required by the I2V-Sigmoid back end. These values can be retuned by changing the unit current source and current-mirror ratios.


With the maximum input current and the maximum weight magnitude, the multiplier output reaches

$$I_{\text{mul,max}} = \frac{255 \times 75 \text{ nA}}{10} = 1912.5 \text{ nA}. \quad (3)$$

Thus, the multiplier core provides the programmable magnitude of the weighted current, whereas the subsequent steering and scaling stages determine the signed contribution delivered to the I2V-Sigmoid back end. The corresponding input-current generator and magnitude multiplier are shown in Figure 4.



(a) Binary-weighted input-current generator.



(b) An 8-bit magnitude-controlled current-mode multiplier.

Figure 4. Current-mode front-end for feature and weight-magnitude currents.

3.4. Sign-Bit Controlled Current Steering and Signed-Current Scaling

The multiplier provides only the non-negative weight magnitude current. The weight sign is handled by a sign-bit controlled current-steering block placed after the multiplier, as shown in Figure 5a. For each branch, the steered current is routed either to a positive path or to a negative path. The positive-path currents are first summed together, and the negative-path currents are summed separately:

$$I_{+, \Sigma} = \sum_{j=1}^4 I_{+,j}, \quad I_{-, \Sigma} = \sum_{j=1}^4 I_{-,j}. \quad (4)$$

The two accumulated currents are then scaled by matched current mirrors, as shown in Figure 5b, and subtracted around a fixed midpoint current:

$$I_{I2V} = I_{\text{mid}} + \frac{I_{+, \Sigma} - I_{-, \Sigma}}{K_{\text{signed}}^{(1)}}, \quad I_{\text{mid}} = 3 \mu\text{A}. \quad (5)$$

Thus, a positive weight increases the I2V input current above the midpoint, a negative weight decreases it below the midpoint, and equal positive and negative contributions cancel around I_{mid} .

The scaling factor is chosen so that the four-input neuron covers the full 0–6 μA I2V range without exceeding it. Since one multiplier branch can deliver up to

$$I_{\text{mul}, \text{max}} = \frac{255 \times 75 \text{ nA}}{10} = 1912.5 \text{ nA}, \quad (6)$$

the four maximum positive branches produce $4I_{\text{mul}, \text{max}} = 7.65 \mu\text{A}$ before scaling. This current must correspond to a +3 μA displacement from the midpoint. Therefore,

$$K_{\text{signed}}^{(1)} = \frac{4I_{\text{mul}, \text{max}}}{3 \mu\text{A}} = 2.55, \quad (7)$$

or equivalently a mirror ratio of about $1/2.55 \simeq 0.392$.

With this choice, all four branches at maximum positive weight give $I_{I2V} = 6 \mu\text{A}$, all four branches at maximum negative weight give $I_{I2V} = 0$, and balanced positive and negative contributions keep the I2V input close to 3 μA .


The second hidden layer follows the same signed-current principle but has a different fan-in. Since it contains 16 input branches per neuron, the maximum signed displacement per branch must be reduced to

$$\frac{3 \mu\text{A}}{16} = 187.5 \text{ nA}$$

if the same 0–6 μA I2V–Sigmoid input range is to be preserved. The exact second-layer scaling factor therefore depends on the selected second-layer multiplier current range, while the midpoint-biased I2V–Sigmoid back end can remain unchanged.

Before the I2V–Sigmoid interface, the current-mode weighted-sum path remains fully current-domain: quantized inputs are converted into 0–75 nA feature currents, weighted by the magnitude-controlled multiplier, routed according to the sign bit, and accumulated through the signed-current scaling stage. The resulting current I_{I2V} is then applied to the I2V converter.

This four-input implementation therefore validates the first-hidden-layer current-mode weighted-sum path up to the midpoint-biased output I_{I2V} . Extension to the second hidden layer would require 16 input branches and adapted signed-current scaling, while the downstream I2V–Sigmoid implementation remains outside the scope of this work.



(a) Sign-bit controlled current steering module.

(b) Signed-current scaling module.

Figure 5. Sign-bit controlled current steering and signed-current scaling structure.

4. Circuit Implementation and Simulation Results

This section evaluates the proposed four-input current-mode weighted-sum path at the transistor level. The simulations focus on the first-hidden-layer prototype defined in Section 3. The results are organized from the current-domain front-end blocks to the signed-current path, and the post-layout verification of the multiplier core. The objective is to assess the block-level circuit behavior of the proposed weighted-sum implementation.

4.1. Simulation Setup

The circuits are designed and simulated in a 28 nm FD-SOI technology with a nominal supply voltage of 1.0 V. Table 5 summarizes the PVT conditions used in the transistor-level validation. The TT corner is used as the nominal operating condition. The FF and SS corners are combined with temperature and supply-voltage shifts in order to stress the current-mode behavior under faster and slower device operation. The supply voltage is swept over a $\pm 10\%$ range around the nominal 1.0 V supply in order to evaluate the robustness of the proposed current-mode weighted-sum path under wider voltage conditions. These conditions are used as a block-level PVT validation and do not constitute a complete classifier-level robustness assessment.

The signed-current accumulation path is simulated under the same PVT sweep for several representative weight magnitudes and for both positive and negative steering configurations. Post-layout verification is carried out for the multiplier without the signed-current steering and scaling stages, because the multiplier is the repeated magnitude-generation core of the weighted-sum branch.

In each figure, the TT response is shown as the nominal reference together with the two PVT responses that produce the largest deviation for the corresponding block.

4.2. Block-Level Validation

The first validation step evaluates the main current-processing blocks: the input-current generator, the magnitude-controlled multiplier, and the sign-bit controlled signed-current path.

The input-current generator provides 16 current levels by controlling four binary-weighted current branches, covering 0–75 nA with a 5 nA step. As shown in Figure 6a, the simulated currents in the 5, 10, 20, and 40 nA elementary branches remain close to their

target values under the displayed PVT conditions. The maximum observed error over all simulated branches/codes is below 17%, and the generated currents therefore preserve the intended 16-level feature-coding range used by the following multiplier stage.

Table 5. PVT conditions used for transistor-level circuit validation.

Corner	Process Condition	Temperature	Supply Voltage
TT	Typical–Typical	27 °C	1.00 V
FF	Fast–Fast	85 °C	1.1 V
FF	Fast–Fast	85 °C	0.9 V
FF	Fast–Fast	−40 °C	1.1 V
FF	Fast–Fast	−40 °C	0.9 V
SS	Slow–Slow	85 °C	1.1 V
SS	Slow–Slow	85 °C	0.9 V
SS	Slow–Slow	−40 °C	1.1 V
SS	Slow–Slow	−40 °C	0.9 V

The magnitude-controlled multiplier is evaluated before the signed-current path. Its input-current range is 0–75 nA. When the weight magnitude is set to 255, the target output-current range is 0–1912.5 nA. Figure 6b shows that the multiplier preserves the expected current-scaling behavior under PVT variations. The largest deviation with respect to the TT response remains below 2% over the simulated range. This result supports the use of the multiplier as a non-negative current-magnitude generator before sign steering under the considered PVT conditions.

The signed-current path between the multiplier and the I2V input is verified for representative nonzero weight magnitudes. The verification focuses on the PVT-induced deviation of the final signed-current difference,

$$I_{\Delta,s} = I_{+,s} - I_{-,s}, \quad (8)$$

where $I_{+,s}$ and $I_{-,s}$ are the scaled positive-path and negative-path currents. In this test, the input current is kept constant at its maximum value of 75 nA. For each tested weight magnitude, the worst-case PVT deviation is computed with respect to the nominal TT value of $I_{\Delta,s}$. The resulting worst-case deviations are summarized in Table 6.

Table 6. Worst-case PVT deviation of the signed-current difference.

Weight Magnitude	$\Delta I_{\Delta,\max}$ (nA)
$ W = 16$	3.1
$ W = 32$	7.8
$ W = 64$	10.3
$ W = 128$	18.1
$ W = 192$	25.1
$ W = 255$	31.5

The maximum signed-current deviation is 31.5 nA at $|W| = 255$, which corresponds to approximately 0.53% of the 0–6 uA signed-current path output-current window. Therefore, although the positive and negative contributions use different steering paths and mirrors, the signed-current path remains within a small deviation range with respect to the output-current window.

4.3. Cascaded Current-Mode Weighted-Sum Path Validation

In addition to the individual validation of the input-current generator, the magnitude-controlled multiplier, and the signed-current accumulation path, the complete cascaded

current-mode weighted-sum path up to the I2V input was also verified. This path includes the input-current generation block, the 8-bit magnitude-controlled current-mode multiplier, the sign-bit-controlled current-steering block, and the signed-current scaling/accumulation stage. This validation checks whether the cascaded module-level errors remain within the $\pm 20\%$ block-level target [20], and whether the final current delivered to the I2V converter stays in the expected operating range.




Figure 6. PVT simulation results of (a) input-current generator, (b) magnitude-controlled multiplier.

Figure 7 shows the complete PVT simulation results of the cascaded weighted-sum path for $I_{in} = 75$ nA, with the weight swept from -255 to 255 . Table 7 summarizes the current values and relative deviations corresponding to the worst-case PVT deviations.

The PVT corners mainly introduce a gain shift in the current-weight transfer, while preserving its monotonic and approximately linear behavior. Under the TT condition, when $W = 32, 64, 128$, the positive-path currents are 378.6, 753.4, and 1502.6 nA, while the negative-path currents are -375.2 , -751.1 , and -1504.2 nA, both close to a 1:2:4 ratio. Similar proportional scaling is also observed under the other two PVT corners, indicating that PVT variations mainly affect the transfer gain rather than the linear weighting behavior.

The maximum PVT error of the complete cascaded weighted-sum path occurs when the positive current path is fully enabled. The corresponding error is 19.4%, which remains below the $\pm 20\%$ block-level error target defined in this work. Although the errors from the current mirrors and switching paths accumulate after cascading, the final output current still remains within the expected tolerance range.

Table 7. Cascaded signed-current path validation at $I_{in} = 75$ nA for the maximum positive and negative weights.

Weight	TT Reference	SS, $T = 85$ °C, $V_{DD} = 0.9$ V	Rel. Dev. vs. TT	FF, $T = -40$ °C, $V_{DD} = 1.1$ V	Rel. Dev. vs. TT
$W = +255$	+2872.9 nA	+2461.5 nA	−14.3%	+3428.6 nA	+19.4%
$W = -255$	−2987.4 nA	−2615.0 nA	+12.5%	−3455.7 nA	−15.7%

4.4. Post-Layout Multiplier Verification

The post-layout multiplier deviations are interpreted using the $\pm 20\%$ block-level target defined in Section 3. The magnitude-controlled multiplier is further verified after layout

extraction. This block is selected for post-layout validation because it is the repeated current-magnitude generation element of the hidden-layer array. Figure 8a shows the extracted multiplier layout. The layout area is

$$30.4 \mu\text{m} \times 30.4 \mu\text{m} = 924 \mu\text{m}^2. \quad (9)$$




Figure 7. PVT simulation results of the complete cascaded current-mode weighted-sum path before the I2V input.

The reported layout corresponds to a first conservative implementation intended for post-layout PVT validation of the repeated weighting cell. The device dimensions were selected with design margin: $L = 300$ nm was used to reduce short-channel effects (SCE) and drain-induced barrier lowering (DIBL), and a 200 nm unit width was chosen for the smallest branch, although $W_{\min} = 104$ nm, in order to preserve current accuracy and matching margin.

The large devices were not implemented as single ultra-wide transistors. Instead, they were realized by combining multi-finger structures with device multiplicity, i.e., parallel repetitions of identical unit MOS devices. In addition, five dummy MOS devices identical to the basic multiplicity MOS unit were placed at the array boundary to reduce edge effects and to preserve a more uniform layout environment. The main analog current paths are kept as comparable as possible across branches, and the digital control lines are routed separately from the weighted-current nodes to limit unwanted coupling. The reported area should therefore be interpreted as a conservative prototype-layout area, rather than as a fully optimized array cell. A future compact version could investigate reducing the unit width, provided that mismatch, current accuracy, and PVT robustness remain acceptable.

Figure 8b compares, at the maximum weight value of 255, the pre-layout TT response with the post-layout TT response and the two post-layout PVT responses that produce the largest output-current deviations.

Table 8 summarizes the post-layout deviations at the maximum input current, $I_{\text{in}} = 75$ nA. The post-layout deviations remain below 1% with respect to the TT post-layout response, and the TT post-layout/pre-layout difference is about 1.03%. The multiplier remains well within the $\pm 20\%$ block-level target under the considered PVT conditions. These results indicate limited parasitic impact on the extracted multiplier core under the reported maximum-input test condition.

In addition to the post-layout PVT verification, a post-layout Monte Carlo analysis was performed at the maximum input current, $I_{\text{in}} = 75$ nA, to evaluate the statistical dispersion of the extracted multiplier output. Figure 9 shows the resulting output-current distribution. The mean output current is $\mu = 1920.79$ nA, with a standard deviation of

$\sigma = 93.57$ nA, corresponding to a coefficient of variation of 4.87%. The MC mean remains very close to the nominal post-layout TT simulation, with a difference of only -3.85 nA, corresponding to approximately -0.20% and the 3σ spread corresponds to approximately 14.6% of the mean value, remaining below the $\pm 20\%$ block-level target. This indicates that the extracted multiplier keeps acceptable statistical robustness under the considered maximum-input condition.




Figure 8. Post-layout verification of the magnitude-controlled multiplier: (a) extracted layout and (b) pre-layout TT response, post-layout TT response, and the two largest-deviation post-layout PVT responses.

Table 8. Post-layout multiplier deviations at $I_{in} = 75$ nA.

Comparison	Current Deviation	Relative Deviation
Upper PVT post-layout vs. TT post-layout	+4.29 nA	+0.22%
Lower PVT post-layout vs. TT post-layout	−8.25 nA	−0.43%
TT post-layout vs. TT pre-layout	+19.55 nA	+1.03%




Figure 9. Post-layout Monte Carlo distribution of the multiplier output current at $I_{in} = 75$ nA.

Overall, the post-layout PVT and Monte Carlo simulations provide circuit-level evidence for the feasibility of the proposed current-mode synaptic weighting cell and identify layout area and statistical robustness as important aspects toward a complete MLP-BP inference array.

5. Discussion

The results demonstrate the feasibility of the proposed current-mode weighted-sum path at the building-block level. Starting from the selected no-bias [16,16] MLP-BP model, the first-hidden-layer operation is mapped into explicit circuit quantities: 0–75 nA current-coded inputs, 8-bit magnitude-controlled current weighting, sign-bit-controlled positive/negative current steering, midpoint-biased signed-current accumulation over a 0–6 μ A activation-input range. Although a nonlinear activation operation is required for complete MLP inference, its circuit implementation is outside the scope of the present work.

The main contribution of this work is the current-mode weighted-sum path, centered on the repeated synaptic weighting/multiplier cell and its associated signed-current accumulation principle. Compared with earlier analog MLP implementations based on analog weight storage or general-purpose four-quadrant multipliers [20], the proposed approach keeps the trained weights in digital form and performs the weighting operation through switched current-domain branches. This avoids analog weight storage, removes the need for a per-weight DAC, and simplifies the repeated weighted-sum branch.

The separation between weight magnitude and sign is central to the proposed cell. The multiplier generates only a non-negative magnitude-controlled current, while the sign-bit steering stage routes it toward the positive or negative accumulation path. This avoids a fully signed four-quadrant multiplier at each branch and enables simple local signed current.

At the circuit level, the current-mode weighted-sum path preserves the intended operating behavior under the considered PVT conditions. The multiplier is the only block verified after layout extraction and is identified as the most mature and most repeated element of the future MLP-BP array. The signed-current path was validated at transistor level before layout extraction.

Existing analog neural-network hardware has explored several complementary implementation directions. Programmable current-mirror VMMs have demonstrated current-domain vector-matrix computation using floating-gate devices for programmable analog weight storage [17]. At the synaptic-cell level, compact floating-gate circuits have been developed to combine local analog weight storage with weighted-current generation [28]. More recently, multi-bit analog Multiply-Accumulate (MAC) cells implemented in FD-SOI have combined dynamic capacitive weight storage, back-gate weight control, multiplication, and accumulation within compact computing cells [29]. For the multiplier mismatch analysis, Nägele et al. [29] report a relative standard deviation of approximately 11–12% at the maximum positive and negative weights. By comparison, the post-layout Monte Carlo analysis of the proposed multiplier gives a coefficient of variation of 4.87% at the maximum weight magnitude under the reported maximum-input condition. These studies demonstrate the feasibility of current-domain neural computation across VMM-array, synaptic-cell, and integrated-MAC levels, and provide the broader hardware context for the development of compact analog weighting structures. However, they differ in target function, technology node, signal domain, weight implementation, and validation level. In this context, the present work focuses on directly mapping a trained compact application-specific NILM MLP-BP inference model into a transistor-level, digitally controlled, current-mode signed weighted-sum path. Building on this model-to-circuit mapping, the main contributions include a current-mode architecture with separate magnitude and sign processing and its

circuit-level robustness verification, including post-layout PVT and Monte Carlo analyses of the repeated multiplier.

This specialization is both an advantage and a limitation. It reduces the complexity of the repeated branch compared with a fully analog four-quadrant multiplier and avoids analog weight storage, but it does not provide the flexibility of a general-purpose analog multiplier or a fully programmable VMM. The proposed cell is therefore best suited to inference engines where the topology and weights are fixed after offline training. This matches the targeted NILM MLP–BP use case, where training is performed offline and the embedded hardware only executes feedforward inference.

The extracted multiplier area is $924\ \mu\text{m}^2$, and the reported block-level power is $2\ \mu\text{W}$. These values provide an initial quantitative basis for positioning the proposed weighting cell, but they should not be interpreted as complete neuron or classifier metrics. Further integration will require array-level layout, optimized support blocks, and classifier-level reinjection of circuit nonidealities.

Overall, the proposed block should be viewed as a step from system-level analog MLP specifications toward a transistor-level current-mode weighted-sum implementation in advanced FD-SOI technology. The next integration step is to extend this block toward compact full-neuron and array-level implementations.

6. Conclusions

This work presented a digitally controlled current-mode weighted-sum path in 28 nm FD-SOI as a building block for low-complexity event-based NILM MLP inference. A no-bias [16, 16] MLP trained offline by backpropagation was selected as a compact system-level target. This model provides the software-level inference target and the circuit constraints for the proposed current-mode implementation.

The selected model was translated into bounded current-domain quantities using a representative first-hidden-layer four-branch weighted-sum path. The proposed chain combines current-coded inputs, an 8-bit magnitude-controlled current-mode multiplier, sign-bit-controlled positive/negative current steering, and midpoint-biased signed accumulation.

Circuit simulations validate the main elements of the proposed current-mode weighted-sum path. The multiplier shows limited deviation under the considered PVT conditions and remains the most mature block of the implementation. Post-layout extraction confirms the limited parasitic impact on the multiplier core, which occupies $924\ \mu\text{m}^2$ and consumes $2\ \mu\text{W}$ at block level.

The I2V–Sigmoid stage is shown only as a conceptual functional block in the system-level signal-flow diagram to indicate the connection between midpoint-biased current accumulation and the subsequent voltage-domain activation operation. Its circuit implementation and validation are outside the scope of the present work. Further integration will focus on optimized support circuitry, compact full-neuron and array-level implementations, complete post-layout verification of the signed-current path, and classifier-level reinjection of the extracted circuit nonidealities.

Overall, the proposed block represents a step toward compact current-mode MLP–BP inference hardware by combining digital weight storage, local current-domain weighting, signed-current accumulation, and post-layout validation of the repeated multiplier.

Author Contributions: Conceptualization, Z.M., Y.C., E.F. and G.J.; methodology, Z.M., Y.C., E.F. and G.J.; software, Z.M.; validation, Z.M. and Y.C.; formal analysis, Z.M. and Y.C.; investigation, Z.M. and Y.C.; resources, G.J.; data curation, Z.M.; writing—original draft preparation, Z.M.; writing—review and editing, Z.M., Y.C., E.F. and G.J.; visualization, Z.M.; supervision, Y.C., E.F. and G.J.; project administration, Y.C. and G.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study uses a previously released electrical-load dataset and does not involve human participants or animal subjects.

Informed Consent Statement: Not applicable.

Data Availability Statement: All data are available from the authors upon request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Hart, G.W. Nonintrusive appliance load monitoring. *Proc. IEEE* **1992**, *80*, 1870–1891. [\[CrossRef\]](#)
- Zoha, A.; Gluhak, A.; Imran, M.A.; Rajasegarar, S. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors* **2012**, *12*, 16838–16866. [\[CrossRef\]](#) [\[PubMed\]](#)
- Kaseli, M.; Doulamis, N.; Voulodimos, A.; Protopapadakis, E.; Doulamis, A. Towards trustworthy energy disaggregation: A review of challenges, methods, and perspectives for non-intrusive load monitoring. *Sensors* **2022**, *22*, 5872. [\[CrossRef\]](#) [\[PubMed\]](#)
- Laughman, C.; Lee, K.; Cox, R.; Shaw, S.; Leeb, S.; Norford, L.; Armstrong, P. Power signature analysis. *IEEE Power Energy Mag.* **2003**, *1*, 56–63. [\[CrossRef\]](#)
- Liang, J.; Ng, S.K.K.; Kendall, G.; Cheng, J.W.M. Load signature study—Part I: Basic concept, structure, and methodology. *IEEE Trans. Power Deliv.* **2010**, *25*, 551–560. [\[CrossRef\]](#)
- Kolter, J.Z.; Johnson, M.J. REDD: A public data set for energy disaggregation research. In *Proceedings of the Workshop on Data Mining Applications in Sustainability (SustKDD)*; ACM Press: New York, NY, USA, 2011.
- Nguyen, T.K.; Dekneuve, E.; Jacquemod, G.; Nicolle, B.; Zammit, O.; Nguyen, V.C. Development of a real-time non-intrusive appliance load monitoring system: An application level model. *Int. J. Electr. Power Energy Syst.* **2017**, *90*, 168–180. [\[CrossRef\]](#)
- Kelly, J.; Knottenbelt, W. Neural NILM: Deep neural networks applied to energy disaggregation. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments, BuildSys '15*; ACM Press: New York, NY, USA, 2015; pp. 55–64. [\[CrossRef\]](#)
- Mauch, L.; Yang, B. A new approach for supervised power disaggregation by using a deep recurrent LSTM network. In *Proceedings of the 2015 IEEE Global Conference on Signal and Information Processing, GlobalSIP*; IEEE: New York, NY, USA, 2015; pp. 63–67. [\[CrossRef\]](#)
- Massidda, L.; Marrocu, M.; Manca, S. Non-intrusive load disaggregation by convolutional neural network and multilabel classification. *Appl. Sci.* **2020**, *10*, 1454. [\[CrossRef\]](#)
- Huber, P.; Calatroni, A.; Rumsch, A.; Paice, A. Review on deep neural networks applied to low-frequency NILM. *Energies* **2021**, *14*, 2390. [\[CrossRef\]](#)
- Heidari, M.; Shamsi, H. Analog programmable neuron and case study on VLSI implementation of Multi-Layer Perceptron (MLP). *Microelectron. J.* **2019**, *84*, 36–47. [\[CrossRef\]](#)
- Rumelhart, D.E.; Hinton, G.E.; Williams, R.J. Learning representations by back-propagating errors. *Nature* **1986**, *323*, 533–536. [\[CrossRef\]](#)
- Hernandez, A.; Nieto, R.; Fuentes, D.; Ureña, J. Design of a SoC architecture for the edge computing of NILM techniques. In *Proceedings of the 2020 XXXV Conference on Design of Circuits and Integrated Systems (DCIS)*; IEEE: New York, NY, USA, 2020; pp. 1–6. [\[CrossRef\]](#)
- Tapiador, M.; de Diego-Otón, L.; Hernández, A.; Nieto, R. Implementing a CNN in FPGA programmable logic for NILM application. In *Proceedings of the 2023 38th Conference on Design of Circuits and Integrated Systems (DCIS)*; IEEE: New York, NY, USA, 2023; pp. 1–6. [\[CrossRef\]](#)
- Andreou, A.G.; Boahen, K.A. Translinear circuits in subthreshold MOS. *Analog Integr. Circuits Signal Process.* **1996**, *9*, 141–166. [\[CrossRef\]](#)
- Paliy, M.; Strangio, S.; Ruiu, P.; Rizzo, T.; Iannaccone, G. Analog vector-matrix multiplier based on programmable current mirrors for neural network integrated circuits. *IEEE Access* **2020**, *8*, 203525–203537. [\[CrossRef\]](#)
- Alimisis, V.; Papathanasiou, A.; Georgakilas, E.; Eleftheriou, N.P.; Sotiriadis, P.P. An ultra-low power adjustable current-mode analog integrated general purpose artificial neural network classifier. *AEU Int. J. Electron. Commun.* **2024**, *186*, 155467. [\[CrossRef\]](#)
- Bellasi, D.; Crescentini, M.; Cristaudo, D.; Romani, A.; Tartagni, M.; Benini, L. A Broadband Multi-Mode Compressive Sensing Current Sensor SoC in 0.16 μm CMOS. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2019**, *66*, 105–118. [\[CrossRef\]](#)
- Jouni, H.; Harb, A.; Jacquemod, G.; Leduc, Y. Design and specification of analog artificial neural network. *SN Appl. Sci.* **2019**, *1*, 1508. [\[CrossRef\]](#)

21. Planes, N.; Weber, O.; Barral, V.; Haendler, S.; Noblet, D.; Croain, D.; Bocat, M.; Sassoulas, P.-O.; Federspiel, X.; Cros, A.; et al. 28 nm FDSOI technology platform for high-speed low-voltage digital applications. In *Proceedings of the 2012 Symposium on VLSI Technology (VLSIT)*; IEEE: New York, NY, USA, 2012; pp. 133–134. [[CrossRef](#)]
22. Magarshack, P.; Flatresse, P.; Cesana, G. UTBB FD-SOI: A process/design symbiosis for breakthrough energy-efficiency. In *Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE)*; IEEE: New York, NY, USA, 2013; pp. 952–957. [[CrossRef](#)]
23. Charlon, Y.; Jacquemod, G.; Jouni, H.; Harb, A. New implementation of analog artificial neural network for breast cancer classification. In *Proceedings of the 2024 International Conference on Artificial Intelligence, Computer, Data Sciences and Applications (ACDSA)*; IEEE: New York, NY, USA, 2024; pp. 1–6. [[CrossRef](#)]
24. Enz, C.C.; Vittoz, E.A. *Charge-Based MOS Transistor Modeling: The EKV Model for Low-Power and RF IC Design*; John Wiley & Sons: Chichester, UK, 2006. [[CrossRef](#)]
25. Khodabandehloo, G.; Mirhassani, M.; Ahmadi, M. Analog implementation of a novel resistive-type sigmoidal neuron. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **2012**, *20*, 750–754. [[CrossRef](#)]
26. Maheshwari, S. Current-mode design of sigmoid activation function circuit. *Circuits Syst. Signal Process.* **2023**, *42*, 6347–6354. [[CrossRef](#)]
27. He, H.; Garcia, E.A. Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.* **2009**, *21*, 1263–1284. [[CrossRef](#)]
28. Chen, M.; Sonnadara, C.; Shah, S. Open-source floating-gate cell for analogue synapses. *Electron. Lett.* **2024**, *60*, e70036. [[CrossRef](#)]
29. Nägele, R.; Finkbeiner, J.; Stadlander, V.; Grözing, M.; Berroth, M. Analog Multiply-Accumulate Cell With Multi-Bit Resolution for All-Analog AI Inference Accelerators. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2023**, *70*, 3509–3521. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.