Research

Cite this article: Mileiko S, Bunnam T, Xia F, Shafik R, Yakovlev A, Das S. 2019 Neural network design for energy-autonomous artificial intelligence applications using temporal encoding. *Phil. Trans. R. Soc. A* **378**: 20190166.
<http://dx.doi.org/10.1098/rsta.2019.0166>

Accepted: 16 October 2019

One contribution of 13 to a theme issue
'Harmonizing energy-autonomous computing
and intelligence'.

Subject Areas:

artificial intelligence, electrical engineering,
microsystems

Keywords:

neural networks, hardware design,
energy efficiency, energy autonomy

Author for correspondence:

Sergey Mileiko
e-mail: s.mileiko2@newcastle.ac.uk

Neural network design for energy-autonomous artificial intelligence applications using temporal encoding

Sergey Mileiko¹, Thanasin Bunnam¹, Fei Xia¹,
Rishad Shafik¹, Alex Yakovlev¹ and Shidhartha Das²

¹School of Engineering, Newcastle University,
Newcastle upon Tyne, UK

²Arm Ltd, Cambridge, UK

RAS, 0000-0001-5444-537X

Neural networks (NNs) are steering a new generation of artificial intelligence (AI) applications at the micro-edge. Examples include wireless sensors, wearables and cybernetic systems that collect data and process them to support real-world decisions and controls. For energy autonomy, these applications are typically powered by energy harvesters. As harvesters and other power sources which provide energy autonomy inevitably have power variations, the circuits need to robustly operate over a dynamic power envelope. In other words, the NN hardware needs to be able to function correctly under unpredictable and variable supply voltages. In this paper, we propose a novel NN design approach using the principle of pulse width modulation (PWM). PWM signals represent information with their duty cycle values which may be made independent of the voltages and frequencies of the carrier signals. We design a PWM-based perceptron which can serve as the fundamental building block for NNs, by using an entirely new method of realizing arithmetic in the PWM domain. We analyse the proposed approach building from a 3×3 perceptron circuit to a complex multi-layer NN. Using handwritten character recognition as an exemplar of AI applications, we demonstrate the power elasticity, resilience and efficiency of the proposed NN design in the presence of functional and parametric variations including large voltage variations in the power supply.

1. Introduction

Advances in sensing devices are causing a shift towards the fourth industrial revolution [1]. The large volumes of the data produced by these devices are enabling a new generation of artificial intelligence (AI) systems at the micro-edge that are designed to infer important decisions in the real world [2]. A promising direction of these AI Systems is the leap towards perpetual computability, allowing always available local AI service. To enable this, designers of pervasive AI systems are facing two grand challenges: *energy efficiency* and *energy autonomy* [3–6].

Energy efficiency refers to economizing the energy consumption of elementary compute operations. The aim is to prolong operating lifetime with a given energy budget, typically defined by the batteries. Reducing energy requires careful design considerations at device-, circuit- and system levels. Examples include reducing device geometry [7], scaling operating voltage [8] and designing circuits with reduced or approximate logic [9].

New generations of pervasive AI-based systems require maintenance-free long-life. As such traditional energy-efficient design principles applied in battery-operated systems are not feasible, as they need periodic re-charging and replacements. Portable energy harvesters, which produce electrical energy to supply to computation loads by scavenging energy from the environment, are gradually making inroads. Such a scheme of energy harvesting can remove the need of maintenance in favour of energy autonomy. However, mitigating their energy variations needs computational capability over a dynamic power envelope, otherwise known as power elasticity [10,11].

Despite advances in low-power design methodologies, the energy footprint of existing AI systems, such as Neural Networks (NNs), has generally remained high [12]. Our persistence in using arithmetic-heavy circuits with growing algorithmic complexities is a major contributor to this. For instance, object detection using deep NNs may require 100 to over 10 000 times the energy needed by the traditional histogram of oriented gradient techniques [13]. Due to such poor efficiency, the widespread adoption of energy-autonomous AI hardware at the micro-edge has proven challenging [14].

To appreciate the importance of efficient AI hardware design, we show the example of a perceptron, whose idea originates from Rosenblatt’s work of 1958 [15]. It is a basic building block of NNs used in AI applications [16–18]. It consists of an input vector, a set of weights and a bias to produce binary classification outcomes, as follows:

$$f(x) = \begin{cases} 1, & \text{if } \mathbf{w} \cdot \mathbf{x} + b > 0 \\ 0, & \text{otherwise,} \end{cases} \quad (1.1)$$

where w is a vector of real-valued weights, $\mathbf{w} \cdot \mathbf{x}$ is the dot product $\sum_{i=1}^m w_i x_i$ with m number of inputs, and b is the bias. The process of deciding the appropriate weights ($\mathbf{w}$), often also known as training, serves as the basic principle of supervised learning. When m becomes large, it approximates the behaviour of a biological neuron. Figure 1 shows the typical structure of a perceptron [19,20]. At its core is an adder that sums m weighted inputs. The result of the addition is compared with a reference during the training phase, during which the weights are updated to ensure the reference is matched. For hardware implementation, multiplication and addition are crucial arithmetic circuits in a perceptron [21]. Such arithmetic operations require significant area and power costs, which depend on the number of input-weight pairs, the precision of the multipliers/adders, their underlying technology nodes and algorithmic complexities.

Over the years, substantial research has been dedicated to improving the energy efficiency of AI hardware [22]. A vast body of this research has predominantly remained within the remits of Landauer’s logic boundaries for energy or power reduction [23]. Reducing threshold voltage

Figure 1. Structural organization of a perceptron, which is the basic building block of NNs.

that defines the logic boundaries and designing new low-complexity architectures are key to achieving this. Andri *et al.* [14] proposed an NN architecture that showed how high-performance NN operations can be achieved by parallel logic blocks. These blocks are designed using low-threshold technology nodes that are faster and ultra-low power. Prado *et al.* [24] showed a logic approximation method applied in parallel NNs. Due to low-complexity architecture the individual components are faster and more energy-efficient. Among others, Qiqieh *et al.* [9] proposed logic compression approaches for reducing power consumption, area and critical path delay of NNs. By combining the circuit-level approaches with online system-wide techniques, significant energy reduction was reported.

However, reducing power or energy alone using the above principles, does not solve the problem of energy-autonomous pervasive AI systems [10]. These systems will need to be able to not only work with limited power supplies but also survive extreme variations as power regulation and energy storage options are limited and expensive in low-end micro-edge devices [10,25]. Indeed, these systems will need to be built with natural power elasticity to operate over a large power domain [10,25,26].

Existing perceptron designs are predominantly digital, although a number of analogue implementations have been reported [27,28]. The digital designs can operate over a range of powers defined by paired voltages (V_{dd}) and frequencies (f). These designs are however vulnerable to dynamic power supply variations, for example, conditions where voltage of the power source changes in time and continuous V_{dd} and f pairing can prove expensive under limited energy budgets. As such, existing designs have poor power elasticity that prevents them from providing useful computation under unreliable or unstable power supply conditions.

In this paper, for the first time, we tackle the power-elastic AI hardware design. Our design underpins a radically new approach of duty-cycle based computing for pulse width modulated (PWM) signals using a number of parallel inverters. These building blocks are then integrated as part of a higher level analysis method to support system-wide investigations in the context of an exemplar application. Our key motivation to use the duty-cycle time-domain representation of data is due to its potential fundamental resilience to dynamic variations in the amplitude and frequency of the signal, which are inevitable for energy autonomous systems drawing energy from the environment. The other motivating factor is the natural ability of CMOS logic to perform multiplication and addition operation on the duty-cycled inputs. This is enabled by the inherent effects of proportionally ratioed current switching in CMOS networks between P and N subnets during the operational cycle of each gate. This gives way to implementing the PWM-based compute functions directly on CMOS logic gates. Thus, one of the important goals we pursue in this investigation is to verify our *hypothesis* that the combined use of PWM representation and CMOS logic, with minimal use of additional analogue (ideally, only passive components) electronics, will deliver the sought efficiency and robustness of AI hardware.

In more concrete terms, the aim of this paper is to design and demonstrate a voltage and frequency elastic perceptron, which performs its arithmetic computation in the PWM-coded

format for robustness to energy supply variations. For this purpose, a method of modelling and analysing such power and frequency elastic NN components also needs to be developed. The main *contributions* are:

- (i) a mixed-signal perceptron design using duty cycle-based temporal weight encoding and input switching via inverters,
- (ii) extensive validation experiments in Cadence Analog Design tool demonstrating the perceptron design's resilience in the presence of static or dynamic voltage and frequency variations,
- (iii) a mathematical model describing the input–output relations of the proposed perceptron for system-level design and analysis to support its use in NN design, and
- (iv) configurations, analyses and evaluations using an example PWM perceptron-based NN which solves the MNIST handwritten digits classification.

This rest of the paper is organized as follows: §2 establishes the method of designing the PWM-based perceptron: the idea of the PWM to voltage conversion, performing the arithmetic operations and the implementation of the voltage to PWM converter. Section 3 validates the approach using a number of parametric sweeps to demonstrate the frequency elasticity and power resilience. Section 4 discusses the strengths and weaknesses of the proposed approach, and the ideas of the future improvements. Finally, §5 concludes the paper.

2. Method

This section focuses on the design of the PWM-based perceptron, including the fundamental theories, the circuits of its constituent parts, methods of PWM-based arithmetic, leading to the construction of NNs. The design methods form the basis of extensive analysis supporting the validation of the perceptrons integrated in a NN.

A perceptron capable of voltage and frequency elasticity may be constructed by exploiting the fact that relative temporal properties, such as a PWM's duty cycle, are resilient to voltage and frequency variations. As the supply voltage reduces, any oscillatory activity, such as a clock signal, may show a reduced amplitude and reduced frequency. However, the ratio between the time within a period when the clock signal is high and the time within a period when the clock signal is low stays the same as both would increase at the same rate.

Our method, therefore, is dedicated to finding ways of exploiting this fact by transferring computation from the digital domain, which is affected by voltage and frequency variations, to the relative temporal domain, which is not. This means making use of PWM-based techniques.

(a) Principles of duty cycle to voltage conversion

Figure 2 shows an inverter-based PWM to voltage converter, which produces an output voltage whose value represents the value carried by the input PWM signal, i.e. its duty cycle. Here we exploit the principle that if the input of an inverter is a periodic signal, such as a clock, the average voltage on its output is inversely proportional to the duty cycle of the input signal. In other words, the analogue average value of the inverter's output voltage encodes the value of the duty cycle of the input signal. Since an inverter is a digital component, whose output equals to logic '0' or '1' at any moment in time, it needs to be 'analogized' (i.e. transcoded) in order to convert the input duty cycle into the output voltage that is a corresponding proportion of the supply voltage. This may be achieved by the following ways:

- increasing the input switching frequency,
- increasing the output capacitance,
- limiting the output current.

Figure 2. A CMOS-based inverter circuit.

Figure 3. Inverter output with PWM-coded input. (Online version in colour.)

For the inverter-based PWM to voltage converter shown in [figure 2](#), with the input clock duty cycle at 50%, the average output voltage is around $V_{dd}/2$ ([figure 3](#)). This is due to the fact that during the interval of time when the input is low the output capacitance is charged with current from the power source via the PMOS transistor, and during the interval of input being high the capacitance is discharged via the NMOS transistor. With a 50% duty cycle these two periods of time are the same length and, assuming the transistors are balanced, their voltages average out to half the supply voltage. When the duty cycle deviates from 50% the average value of the output voltage deviates from $V_{dd}/2$ proportionally in the same direction.

If the frequency is high enough that the output capacitor is never fully charged or discharged, the inverter may be equivalently represented as a resistive voltage divider ([figure 4](#)). The output voltage of such a divider can be calculated using the following equation:

$$V_{out} = (V_{dd} - \text{GND}) \cdot \frac{R_n^* + R_{out}^*}{(R_n^* + R_{out}^*) + (R_p^* + R_{out}^*)}, \quad (2.1)$$

where R_n and R_p are parasitic resistances of NMOS and PMOS transistors, respectively. During the charging phase (t_{low}) the input of the inverter is low and current passes through the PMOS and the output resistor. During the discharging phase (t_{high}), the input of the inverter is high and current goes through the output resistor and NMOS. As a result, the resistance values can be calculated from the lengths of time of each phase:

$$R_n^* + R_{out}^* = (R_n + R_{out}) \cdot \frac{t_{low} + t_{high}}{t_{high}} \quad (2.2)$$

and

$$R_p^* + R_{out}^* = (R_p + R_{out}) \cdot \frac{t_{low} + t_{high}}{t_{low}}. \quad (2.3)$$

Figure 4. PWM inverter equivalent circuit, approximated as a voltage divider.

Figure 5. Output voltage of the PWM inverter versus input duty cycle. (Online version in colour.)

Assuming that $R_n \approx R_p$ (this transistor balancing can be achieved by the appropriate relative sizing of the PMOS and NMOS transistors, for instance, by setting the PMOS width to 2.7 times the NMOS width for the UMC65nm technology) and $GND = 0$, equation (2.1) is simplified to:

$$V_{out} = V_{dd} \cdot \frac{t_{low}}{t_{low} + t_{high}} = V_{dd} \cdot (1 - DC), \quad (2.4)$$

where DC is the input duty cycle—the ratio between the length of time when the input clock is high during a clock period and the length of the clock period.

Figure 5 shows the relationship between the input duty cycle and the output voltage of the PWM inverter. In the case when there is no output resistor, the dependency of the output voltage on the input is not linear. The reason of this nonlinearity is that the PMOS and NMOS resistances change with the change of their drain voltages. Thus, $R_p \neq R_n$ when the value of V_{out} is different from $V_{dd}/2$. This nonlinearity, given the arithmetic functional requirements of a perceptron, is undesirable and needs to be either removed or compensated for. Compensation means very high per-inverter overheads which need to be precise in the analogue domain. However, by adding an output resistor $R_{out} \gg (R_p, R_n)$, the difference between PMOS and NMOS resistances no longer affects the output, and the input duty cycle to output voltage relationship becomes completely linear. This requires no high-precision tuning in the analogue domain.

Figure 6. PWM adder circuit performed by parallel inverters, with outputs connected via a capacitor.

(b) Pulse width modulation arithmetic

A perceptron needs to perform arithmetic operations. Converting from PWM to voltage is not the only function of the PWM inverters. They can also be used to construct arithmetic units, such as adders and weighted accumulators. Below we discuss these two operations and their circuits relevant to NNs.

(i) PWM Adder

The circuit of a PWM adder is shown in figure 6. To add n PWM-coded numbers, we use n inverters connected in parallel. Each inverter has an output resistor. The result is stored in the output capacitor in the form of its average voltage.

This kind of adder works on the principle of current summation and charge (i.e. voltage) accumulation. In other words, the values encoded in the input PWM signals are accumulated in the voltage on the output capacitor, and such circuits can be called voltage accumulators (VACs). To calculate the VAC output voltage, we use the principle of current summation and rewrite (2.1) using conductances instead of resistances. The following equation is for a single PWM inverter:

$$V_{out} = V_{dd} \cdot \frac{G_p^*}{G_p^* + G_n^*}, \quad (2.5)$$

where $G_p^* = 1/R_p^* + R_{out}^*$ and $G_n^* = 1/R_n^* + R_{out}^*$.

Likewise equation (2.4) can be expressed as follows:

$$G_p^* = G \cdot \frac{t_{low}}{t_{low} + t_{high}} = G \cdot (1 - DC) \quad (2.6)$$

and

$$G_n^* = G \cdot \frac{t_{high}}{t_{low} + t_{high}} = G \cdot DC, \quad (2.7)$$

where $G = 1/R_p + R_{out} = 1/R_{out} + R_n$.

Since the inverters in figure 6 are connected in parallel, the output voltage of a multi-inverter VAC can be given by

$$V_{out} = V_{dd} \cdot \frac{\sum_{i=1}^n G_{pi}^*}{\sum_{i=1}^n (G_{pi}^* + G_{ni}^*)}. \quad (2.8)$$

Using equations (2.6) and (2.7), equation (2.8) can be simplified as

$$V_{out} = V_{dd} \cdot \left(1 - \frac{\sum_{i=1}^n DC_i}{n}\right). \quad (2.9)$$

Figure 7. A single cell of the PWM weighted adder, based on a NAND gate.

Figure 8. PWM weighted addition VAC with three inputs and 3-bit weights.

In simple terms, the output voltage of a multi-inverter VAC is inversely proportional to the average value of the duty cycles of its inputs, which is exactly what is required.

(ii) Weighted PWM Accumulation

In order to design a perceptron, the ability to integrate weighted additions is another crucial design requirement. The VACs must be capable of programming the input weights, when required. This is performed by replacing the inverters by two-input NAND gates (figure 7). One input of this gate is the PWM-coded signal, and the other is a digital switch signal for enabling or disabling this cell. The output of a disabled cell is always connected to V_{dd} having the same effect as an enabled cell with zero input duty cycle. In this way, the perceptron can be programmed to determine which NAND gates participate in the accumulation. This programming may be carried out in the digital domain without affecting the voltage and frequency elasticity of the computation.

Figure 8 shows a perceptron arithmetic VAC architecture for 3×3 weighted addition based on these types of gates. As can be seen, the circuit adds 3 PWM-coded inputs multiplied by 3-bit weights. Every weight bit is implemented on a separate cell. The least significant bit goes to the cells with the smallest transistor sizes and the largest output resistors (cells ' $\times 1$ '). The second bit is computed at the cells with doubled transistor widths and halved output resistances (cells ' $\times 2$ '). And the most significant bit is coded with four times the transistor widths, and one-fourth times the output resistances (cells ' $\times 4$ ').

The output voltage of the 3×3 weighted addition VAC can be calculated using (2.9), considering the $\times 2$ and $\times 4$ cells as 2 and 4 single cells, respectively.

$$V_{out} = V_{dd} \cdot \left(1 - \frac{\sum_{i=1}^n DC_i \cdot W_i}{n \cdot (2^k - 1)} \right). \quad (2.10)$$

Figure 9. The ring oscillator-based voltage to PWM converter.

where n is the number of inputs, k is the number of bits of the weight, DC_i is the duty cycle of the input i , and W_i is the weight of the input i .

In the case of the 3×3 weighted addition VAC, where $n = 3$ and $k = 3$, the output voltage is

$$V_{out} = V_{dd} \cdot \left(1 - \frac{\sum_{i=1}^3 DC_i \cdot W_i}{21}\right). \quad (2.11)$$

The arithmetic part of this equation is the weighted sum of duty cycles DC_{sum} :

$$DC_{sum} = \frac{\sum_{i=1}^3 DC_i \cdot W_i}{21}. \quad (2.12)$$

Thus, the definition of the 3×3 weighted addition VAC is that its output voltage is proportional to the weighted sum of its input duty cycles, which is exactly as required

$$V_{out} = V_{dd} \cdot (1 - DC_{sum}). \quad (2.13)$$

(c) Voltage to pulse width modulation conversion

In order to design a perceptron based on the type of VAC described in §2b, we need to provide an output interface for it. The output of the perceptron must be used as an input for the perceptrons of any subsequent layer in an NN. Therefore, the output voltage of the PWM arithmetic unit (its VAC) should be converted back to the PWM format.

The schematics of the voltage to PWM converter is shown in figure 9. The converter circuit was proposed originally by Vaisband *et al.* [29]. The converter is a ring oscillator with different power supplies: the odd-numbered inverters are supplied with a voltage of $V_{dd}/2$, and the even-numbered inverters are supplied with the input voltage, which is the output voltage of the VAC. The difference between the supply voltages of the odd- and even-numbered inverters determines the output duty cycle. If the input voltage equals $V_{dd}/2$, the inverters have equal delay and the output duty cycle is 50%. If the input voltage increases, the period of switching from 0 to 1 increases, and the output duty cycle goes down. If the input voltage is lower than $V_{dd}/2$, the switching from 1 to 0 takes more time and the output duty cycle goes up.

Given that the VAC theoretically achieves a linear relationship between its input duty cycle and its output analogue average voltage, the voltage to PWM converter should also ideally achieve a linear conversion relationship. In that case, the overall relationship between the input duty cycle signal and the output duty cycle signal would also be linear, for the simple case where the perceptron is programmed to do no arithmetic processing. In theory, the inverter chain-based voltage to PWM converter should be able to achieve this if the inverters are set to work in the linear regions of their transistors.

(d) Pulse width modulation-coded perceptron design

The PWM-based perceptron consists of two main parts. The first part is the PWM arithmetic unit in the form of a VAC. This converts the PWM-coded inputs to a voltage which encodes the result

Figure 10. Structure of the perceptron: PWM adder, voltage to PWM converter and compensation transistor.

Figure 11. Neural network for MNIST. The DC_{sum} signals are voltages. *in* and *out* signals are duty cycles.

of the computation as programmed by the enable signals. The second part then converts this voltage result to PWM format for use as inputs by subsequent perceptrons as their inputs.

This structure is shown in figure 10, with the 3×3 weighted addition VAC as an example PWM-based arithmetic unit. Any desired VAC arithmetic unit can be put in this place to satisfy specific perceptron functionality requirements. The simple glue logic consisting of a PMOS transistor between the two blocks will be discussed in detail in §3a.

The size of such a perceptron is such that its design may be entirely analysed and validated through simulations within the VLSI CAD environment in which it is implemented. At least some of this analysis must be conducted in the analogue signal domain as the voltage signal between the two parts of the perceptron holds the computation results in its analogue value. As a result, simulations in a VLSI CAD tool environment that support mixed-signal or analogue studies are the best way of analysing and validating such designs. In this work, we implement our perceptron and analyse it using the Cadence Analogue Design Environment. Detailed results will be shown in §3a.

(e) Pulse width modulation-coded neural network design

The proposed PWM perceptrons can be used in constructing traditional NNs such as the example shown in figure 11. In this NN, the input vector (*in*) is fed to the input layer, and the activity propagates through a number of hidden layers to reach the output layer, where the output vector (*out*) is generated. Then, the output vector is compared to the target vector and the error is back propagated to update the weights of each layer using gradient descent. This procedure is iterated with respect to the specified epoch.

In this work, the *in* and *out* signals are of the PWM-type. The value of such a signal, which is between 0 and 1, is represented by its duty cycle value between 0% and 100%. The VAC arithmetic units then compute on such *in* values. This is illustrated by the example described by equation (2.12), where each *in* is multiplied by its weight and all results are accumulated

Figure 12. Integer weight training.

by the VAC in the DC_{sum} voltages. In other words, the weight and sum blocks in figure 11 are implemented by the proposed perceptron's VAC. Then, every DC_{sum} , which is an analogue voltage across a capacitor, is fed to the activation function (AF) whose output is in PWM format to be used as the input of the next layer. To include the AF, equation (2.12) can be modified as expressed in equation (2.14). This requires that the voltage to PWM conversion also implements the AF. Potential modifications from the basic ring oscillator may be necessary, although the basic ring oscillator already approximates a popular AF. This will be discussed in detail in §3a.

Finally, *out* is obtained, the error is calculated and every weight is adjusted by the back-propagation (BP) algorithm. The comparison to target vector is not necessarily implemented with a perceptron-like device and may be implemented by some external controller, which is outside the scope of this paper.

For the PWM-coded NNs, a number of design choices must be made: weight encoding, maximum weight, AF and number of layers, among others. This section establishes a method of making the best use of the proposed PWM-based perceptron to construct NNs to perform specific computational tasks. We will explore aspects of NN design, including weight types, AF, maximum weight and number of layers. We will use the well-known handwriting digit (MNIST) classification problem [30], which is widely used for machine learning algorithm testing [31], as the benchmark application and case study for this investigation. The goal is to suitably determine the best NN configurations for the proposed PWM-coded NN.

$$out = f(DC_{sum}) = f\left(\frac{\sum_{i=1}^n in_i \cdot W_i}{n \cdot (2^k - 1)}\right). \quad (2.14)$$

(i) Integer weight and training

Regarding the circuit design, the weight is discretized to an integer value. This is different from most related work where floating-point (FP) numbers are used for weights. As the circuit size depends on the bit-width of the weight, it is crucial to find the smallest bit-width that still provides the specified error rate tolerance.

The integer weight training can be designed as illustrated in figure 12. The MNIST input vector (*in*) is multiplied by the integer weight (*W*) and the results are accumulated as *t*. Then, *out* is divided by $n \cdot (2^k - 1)$ (i.e. normalizing), which yields the final value of *out* between 0 and 1. Consequently, *out* is scaled to the same range and comparable to the target vector. Then, *out* passes the AF, and the output vector is obtained and compared to the target vector. Next, the FP update is computed from the gradient descent, the learning rate and the error. To adjust the integer weight, the update is scaled back to the integer number by multiplying by $n \cdot (2^k - 1)$ and rounding.

Figure 13. ReLU function.

Next, the integer weight is updated and capped if it exceeds the specified bit-width (e.g. the example 3×3 weighted addition VAC in figure 8 has 3-bit weights). Finally, the training process iterates until the number of specified epochs is reached. Note that the weight capping can be disabled to allow unlimited weight adjustments to mimic FP training.

(ii) Activation function

The AF is necessary in an NN-based learning process because it provides nonlinearity to the computation so that the learning is not limited to linear problems. It also helps map the resulting values in a certain range, depending on the function.

In this work, the input and output ranges of the AF are a main concern because they need to match the output format of the problem and the circuit behaviour. In other words, depending on the purpose of the NN, it may expect its input and output variables to take values within certain ranges. These ranges then need to be mapped onto the working signal range of our perceptron, which is restricted by the duty-cycle representation between 0 and 100%. Here, we take popular MNIST benchmark [32] as an exemplar to explore this aspect of NN design using our perceptron as the basic building block.

In the context of MNIST, the AFs are needed to provide a fully positive output to comply with the target vector [30]. Also, our perceptron design stores the VAC result as the voltage across C_{out} between its two blocks, which means that the DC_{sum} signals are entirely positive voltages. And such a voltage gets converted to a PWM duty cycle, which is also entirely positive. For these reasons, the well-known AF ReLU [16,33], which has an entirely positive output range, is best suited.

Certain other popular AFs are less suitable for this initial investigation. For instance, the sigmoid function is clearly nonlinear across an input range between -5 and 5 [16], which requires representation of negative values. The nonlinearity also means that major modifications to the voltage to PWM part need to be investigated for implementing such AFs. Hence, we decided to concentrate on trying to mimic the ReLU AF using our perceptron's voltage to PWM converter.

The ReLU function in equation (2.15) [33] is depicted in figure 13. One of its attractions is that it is easily differentiable, facilitating gradient descent. To mimic the output of the VAC, it is better than the sigmoid function because the charge in the output capacitor (C_{out}) is emptied when the VAC result is negative. Otherwise, the capacitor is charged and the positive result is obtained. However, the output of this function must be capped at 1 to represent the limit of the PWM range as shown in equation (2.16) and figure 14. This work will attempt to construct an AF that approximates the capped ReLU function.

The size of an entire NN designed for the MNIST problem is such that it is not possible to analyse it entirely within a VLSI CAD environment. For instance, to analyse an image of 784 pixels (cf. the example in figure 11) there need to be 784 perceptrons in the first layer of the NN

Figure 14. Capped ReLU function.

alone and this is clearly beyond analogue simulations at the VLSI level. Effort must be expended in building models in a higher-level language to investigate the design properly.

$$f(x) = \begin{cases} 0 & , x < 0 \\ x & , x > 0 \end{cases} \quad (2.15)$$

and

$$f(x) = \begin{cases} 0 & , x < 0 \\ x & , 0 < x < 1 \\ 1 & , x > 1. \end{cases} \quad (2.16)$$

3. Results

This section reports experimental results of the PWM-based perceptron, leading to an NN architecture. These results validate the design methods at both circuit- and architecture-level.

(a) Analysis and validation of pulse width modulation-coded perceptron

A prototype circuit of the PWM perceptron is designed using UMC65 nm technology and simulated in the Cadence Analog Design Environment tool.¹ We used the high voltage transistors (with 2.5 V nominal voltage) in purpose of better observation. Below we analyse the behaviour of the perceptron circuit under different parametric variations, generated by the design tool.

(i) Voltage accumulators validation

The first constituent part of the perceptron is the VAC. Figure 15 shows the charging of the capacitor in the VAC based on three inverters connected in parallel as shown in figure 6. The frequencies and duty cycles of the inputs are: $f_1 = 140$ MHz, $DC_1 = 70\%$, $f_1 = 120$ MHz, $DC_1 = 30\%$, $f_1 = 100$ MHz, $DC_1 = 50\%$. The capacitor have been charged to the voltage value, proportional to the average duty cycle of the inputs. The charging time of the capacitor depends on the RC value, and the input frequency does not affect it. However, if the frequency is too low, it may result in too high ripple of the output voltage, and, thereafter, reduction of accuracy.

To support our VAC design based on inverters/NANDs and voltage summation on a capacitor, we implemented the 3×3 weighted addition VAC shown in figure 8 in Cadence and ran simulation experiments on it. The results of these simulations are compared to theoretical results obtained from 2.11 and compared in table 1. The differences between the theoretical and simulation results do not exceed 10%. These results validate the correctness of the PWM-based weighted addition VAC design.

¹URL: <https://tinyurl.com/y6k73k4t>

Figure 15. Capacitor charging in the three inverters VAC. (Online version in colour.)

Table 1. Experimental and theoretical results of the 3×3 weighted adder.

DC1 (%)	W_1	DC2 (%)	W_2	DC3 (%)	W_3	V_{out} theoretical (V)	V_{out} simulation (V)
70	7	80	7	90	7	0.50	0.51
50	1	50	2	50	4	2.08	2.11
20	5	60	6	80	7	1.29	1.33
95	7	90	6	80	6	0.50	0.45
30	1	40	4	50	2	2.16	2.21
80	7	20	3	50	4	1.54	1.61

(ii) Validation of voltage to pulse width modulation conversion

The second constituent part of the proposed perceptron is the voltage to PWM converter. This converts the result of VAC arithmetic computation stored as an analogue voltage (a DC_{sum} signal) back to the PWM format for output to subsequent perceptrons, as presented in §2c. The Cadence Analog Environment simulation results of the voltage to PWM converter is shown in figure 16.

Figure 16. Output duty cycle of the voltage to PWM converter. (Online version in colour.)

Figure 17. Output versus input duty cycle of the perceptron. (Online version in colour.)

Ideally, the voltage to PWM conversion should be linear. The real relationship between the output and the input is almost linear for input voltages between 0.7 V and 2.3 V. However, outside this range the ring oscillator stops oscillating. The reason for this will be discussed below.

Another interesting effect is that the linearity of the output increases with increasing the number of inverters, but the difference between 9 and 13 inverters is small. Thus, a chain of nine inverters should be considered a reasonable voltage to PWM converter and we use this design in our subsequent studies.

(iii) Perceptron experimental results and design adjustments

Figure 17 shows the combined operation of both parts of the perceptron: the 3×3 weighted PWM addition VAC (figure 8) connected to the voltage to PWM converter (figure 9). The three inputs of the perceptron are connected together, and all the weights are 7 (all the cells are enabled). The line labelled ‘ideal’ is obtained through the equations in §2. Analysing these results, we can say that:

- In this simulation for the ideal case, we expect the output duty cycle to be equal to the input duty cycle. However, the real output is slightly different from the ideal; and this difference increases with the input duty cycle above 50%.
- The output duty cycles for different supply voltages are similar. The difference does not exceed 10%. This indicates voltage variation resilience in the perceptron design.
- The input duty cycle has limited range—from 20 to 70%. Beyond this range the output stops oscillating and becomes a constant signal.

Figure 18. Output versus input duty cycle of the perceptron with and without compensation. (Online version in colour.)

The observed reduction of operational range and loss of linearity in the voltage to PWM converter are caused by the fact that the voltage DC_{sum} powers the voltage to PWM converter. When the input duty cycle is above 70%, DC_{sum} is below 30% of V_{dd} . For $V_{dd} = 2.5V$ this is below the threshold voltage. And in this case the NMOS transistors of the ring oscillator are always off, and the output stops oscillating. In other words, there is a mismatch between the voltage ranges of the two parts of the perceptron. The output voltage range of the PWM weighted addition VAC is from 0 V to 2.5 V (figure 5); and the input voltage range of the voltage to PWM converter is from 0.7 V to 2.3V (figure 16).

We may limit the range of the output voltage of the PWM weighted addition VAC. This can be done by adding a small glue logic between the two blocks of the perceptron. This may consist of no more than a compensation PMOS transistor, whose gate and drain are connected to the capacitor as shown in figure 10. In this case, when the voltage on the capacitor goes below the threshold, the PMOS starts charging this capacitor, and when the voltage is above the threshold, the PMOS is off.

The input and output duty cycles of the perceptron with compensation are depicted in figure 18. The output is closer to the ideal, and its range is much wider: from 10 to 90%.

(iv) Power elasticity and resilience

To demonstrate the perceptron's resilience to power variations, we simulated the 3×3 PWM-based weighted addition VAC circuit (figure 8) with different values of supply voltage and input signal amplitude. The results are shown in figure 19. As can be seen, the output voltage grows almost linearly with increased V_{dd} . As expected, higher duty cycle show lower output voltages, and vice versa. In the case of the unstable supply voltage, the absolute value of the output voltage does not bear any reliable information. In this case, we should consider the relative relationship between the output voltage and the supply voltage. This relationship should be proportional to the input duty cycle independently from V_{dd} . This is demonstrated by figure 20 where the y -axis represents not the absolute value of V_{out} , but the ratio between V_{out} and V_{dd} that is more relevant for unstable power conditions.

The circuit shows high resilience to static supply voltage variations. Starting from 1 to 1.5 V the ratio V_{out} and V_{dd} remains the same for each duty cycle value of the input signal.

Further simulation experiments are carried out to investigate the VAC's resilience to static frequency variations. Two sizes of the 3×3 VAC are investigated: the small—with the output capacitor $C_{out} = 10$ pF and the output resistors of each cell $R_{out} = 100$ K Ω ; and the large—with $C_{out} = 100$ pF and $R_{out} = 1$ M Ω . The duty cycle of all the inputs is 50%, and all the weights equal to 7 (all the cells are enabled). Figure 21 shows that both VACs produce the output 1.25 V, that equals to $V_{dd}/2$. The average output voltage remains the same on the simulated range of frequencies: from 1 kHz to 1 GHz.

Figure 19. Output voltage (absolute values) versus static variation of power supply. (Online version in colour.)

Figure 20. Output voltage (relative values) versus static variation of power supply. (Online version in colour.)

Figure 21. Output voltage versus static variation of input frequency. (Online version in colour.)

On the other hand, the value of C_{out} does affect other aspects of perceptron performance. C_{out} contributes to the RC time constant of the VAC circuit, providing a low-pass filter effect on the voltage DC_{sum} . As a result, a larger C_{out} is less suitable than a smaller C_{out} for fast response, but would provide better robustness in the presence of frequency variations. In addition, as the voltage to PWM converter depends on the charge on C_{out} for energy, a smaller C_{out} may encounter difficulties in keeping DC_{sum} constant enough to complete the conversion.

Figure 22 shows the DC_{sum} voltage swing in the presence of static frequency variations. As can be seen, with reduced input frequency the voltage swing increases, and at some point the VAC operates as a simple inverter with the output voltage DC_{sum} oscillating between V_{dd} and GND.

Figure 22. Output voltage swing versus frequency. (Online version in colour.)

Figure 23. Power versus frequency of the 3×3 VAC. (Online version in colour.)

Ideally, we would like the voltage swing to be not larger than 0.2 V. It means the frequency of the input PWM signals should not be lower than 1 MHz for the large VAC and 100 MHz for the small VAC.

In addition, figure 23 shows that VAC size and frequency also affect power consumption. The small VAC has higher power consumption. This is due to the output resistor limiting the charging current. The resistor is $10\times$ larger in the large VAC, and the current and the power consumption are smaller.

In the large VAC, we increase the size of C_{out} and reduce the charging current. This increases the charging time of the capacitor. To investigate this, we simulated the time when the voltage on C_{out} reaches the average output value (which is $V_{dd}/2 = 1.25$ V for the 50% input duty cycle). The capacitor is initially charged to $V_{dd} = 2.5$ V. The charging time of the capacitor is around $0.14\ \mu\text{s}$ for the small VAC and $14.5\ \mu\text{s}$ for the large VAC, which is true for the entire range of frequencies. This approximately $100\times$ ratio is because the RC product is $100\times$ as large for the large VAC as for the small VAC.

Figure 24 shows the operation of the perceptron with dynamic supply voltage variations to investigate the dynamic power elasticity. The simulations have the following parameters: the cell output resistance $R_{out} = 100\ \text{K}\Omega$; the VAC output capacitor $C_{out} = 100\ \text{pF}$; the supply voltage varies from 1.8 V to 3.2 V with a period of $10\ \mu\text{s}$; input duty cycle $DC_1 = DC_2 = DC_3 = 50\%$; the weights are $W_1 = W_2 = W_3 = 7$ in Region A, and then change to $W_1 = W_2 = W_3 = 2$ in Region B.

This simulation illustrates the behaviour of the perceptron under very rapid voltage variations at 100 kHz frequency (for illustration purposes only). The time period of this change is the same value as the system time constant $RC = 10\ \mu\text{s}$. The voltage value swing is also very large—a variation amplitude of 1.4 V for a voltage whose nominal value is 2.5 V (also used for

Figure 24. The operation of the perceptron with dynamic power supply voltage variations using an AC supply with 100 kHz frequency (for illustration purposes only). (Online version in colour.)

demonstrating extreme variations). Even under these extreme conditions $V_{\text{cap}}/V_{\text{dd}}$ still maintains a high degree of resilience. After putting the VAC together with the voltage to PWM converter, however, the combination fares less well, with the output duty cycle changing up to 47% in Region B (Observation 2). This may be mitigated by improving the compensation mechanism in the voltage to PWM conversion circuit (§2) or by *in situ* voltage regulation (not discussed in this paper for brevity).

(b) Validation and analysis of pulse width modulation-coded neural network

This section explores an NN system built using the proposed PWM-based perceptron. This NN is designed for solving the MNIST problem and has the structure shown in figure 11. Firstly, a high-level model of the perceptron is constructed so that analysis can be carried out in MATLAB, at a higher level than analogue VLSI simulations, which is impractical for systems of this size. Then this model is used in MATLAB investigations on system properties to validate our NN-design approach.

(i) Pulse width modulation model and voltage to pulse width modulation converter serving as ReLU AF

In this section, the duty cycle output of the perceptron is modelled in the form of a mathematical equation with parameters. Then, the model and the voltage to PWM converter itself are studied to verify that the device approximately incorporates the capped ReLU AF.

The equations in §2e pertain to ideal cases. These can be used for comparing with how the implemented perceptron actually delivers. In order to make this comparison at the whole system level, we need to generate a high-level mathematical model based on observations made while experimenting with the perceptron circuit at low level.

We experimented in the Cadence Analog Design Environment with a single perceptron, two perceptrons connected in series, mimicking the simplest two-layer NN and three perceptrons connected in series, emulating the simplest NN with a depth of 3. This is as far as analogue VLSI

Figure 25. PWM output. (Online version in colour.)

Figure 26. Output versus model stage 1. (Online version in colour.)

simulations could practically go, as the three-layer study took many hours on a competitively specified server machine.

The outputs of these perceptron connection topologies are shown in figure 25. In the ideal case, the input and output duty cycles should be equal when every weight is at the maximum value (dashed line). However, there is a nonlinear relationship between the input and output of the single perceptron (red line) and the degree of nonlinearity increases when the depth of the NN is increased (blue and green lines). In addition, the output begins to saturate in the last (third) stage when the input (DC_{sum}) reaches 0.85.

To model this relationship, a third-order polynomial equation, which is easy to differentiate, is curve-fitted to the response of the single perceptron using basic regression in MATLAB. The result is shown in equation (3.1). Note that the saturation point of the model is set at the maximum output duty cycle, which is 98%. Then, the model is connected in the same two- and three-stage series topologies as in the Cadence experiments and their outputs are plotted in figure 26–28, together with the relevant Cadence results for comparison. All figures show that this model accurately estimates the input–output relationship of the perceptron. The accuracy can be obtained as the R^2 -values of stages one, two and three, which are 99.88%, 99.33% and 97.66%, respectively.

$$DC_{out} = 107.27V_{Cout}^3 - 53.25V_{Cout}^2 + 52.92V_{Cout} + 13.44 \quad (3.1)$$

Figure 27. Output versus model stage 2. (Online version in colour.)

Figure 28. Output versus model stage 3. (Online version in colour.)

and

$$f(x) = \begin{cases} 0 & , x < 0 \\ DC_{out} & , 0 < x < 1 \\ 0.98 & , x > 0.98. \end{cases} \quad (3.2)$$

(ii) Pulse width modulation-based neural network simulations

We use the perceptron model in equation (3.1) to assemble models of large-size MNIST NNs then simulate these systems in full in MATLAB. The model plot is shown in figure 29. We also create a capped ReLU function with offset (Ofc.ReLU), expressed in equation (3.3). As can be seen in the equation, this offset ReLU function takes the constant 13.44 from the perceptron model in equation (3.1) to have the same offset nonlinearity as the perceptron model in figure 29, but otherwise has a similar straight line behaviour to the non-offset capped ReLU in figure 14. The plot of this offset ReLU can be found in figure 30.

This function is used to investigate whether the step nonlinearity or the curvature nonlinearity higher in the curve of equation (3.1) is more important when it comes to NN performance, through

Figure 29. PWM perceptron output model function.

Figure 30. Capped ReLU function with PWM-like offset.

comparisons with both the perceptron model in equation (3.1) and the capped ReLU function without offset in figure 30.

$$f(x) = \begin{cases} 0 & , x < 0 \\ \text{DC}_{out} + 13.44 & , 0 < x < 1 \\ 1 & , x > 0.86.56. \end{cases} \quad (3.3)$$

There are two groups of simulations using MATLAB: without/with limiting the maximum weight. All implement the training procedure described in figure 12 for the MNIST problem, which is selected as our benchmark. Without defining the maximum weight, the weight is adjusted freely like the basic FP training while the proposed NN is demonstrated by the limited weight simulation. Four AFs: ReLU, capped ReLU (Cap.ReLU), capped ReLU with offset (OfReLU) and PWM perceptron (PWM percept.) are applied in three network configurations: two (784/10), three (784/300/10) and four (784/300/100/10) layers. The PWM perceptron AF is implemented by the PWM perceptron on its own unmodified—the justification being that it may be considered as an approximation of the capped ReLU (cf. figure 29 and figure 14).

As can be seen from these data, these systems being simulated include hundreds of perceptrons and are well beyond analysing in the VLSI design domain.

Table 2. Simulation result of the floating-point weight neural network.

no.	no. perceptron	activation function	learning rate	error
1	784/10	ReLU	0.010	1.40
2	784/10	Cap.ReLU	0.008	1.75
3	784/10	Oft.ReLU	0.009	5.09
4	784/10	PWM percept.	0.004	8.54
5	784/300/10	ReLU	0.040	1.63
6	784/300/10	Cap.ReLU	0.009	1.91
7	784/300/10	Oft.ReLU	0.002	79.54
8	784/300/10	PWM percept.	0.004	27.01
9	784/300/100/10	ReLU	0.040	2.07
10	784/300/100/10	Cap.ReLU	0.010	3.60
11	784/300/100/10	Oft.ReLU	0.010	90.20
12	784/300/100/10	PWM percept.	0.090	79.07

Table 3. Simulation result of the integer weight neural network.

no.	no. perceptron	activation function	learning rate	initial weight	max weight	error
1	784/10	ReLU	0.030	± 3	± 31	9.28
2	784/10	Cap.ReLU	0.040	± 3	± 63	6.12
3	784/10	Oft.ReLU	0.004	± 7	± 255	7.10
4	784/10	PWM percept.	0.030	± 1	± 255	9.98
5	784/300/10	ReLU	0.020	± 255	± 255	79.49
6	784/300/10	Cap.ReLU	0.020	± 255	± 255	79.49
7	784/300/10	Oft.ReLU	0.020	± 3	± 255	18.35
8	784/300/10	PWM percept.	0.010	± 15	± 255	25.17
9	784/300/100/10	ReLU	0.010	± 31	± 255	88.50
10	784/300/100/10	Cap.ReLU	0.010	± 31	± 255	88.50
11	784/300/100/10	Oft.ReLU	0.020	± 127	± 255	64.09
12	784/300/100/10	PWM percept.	0.010	± 63	± 255	53.25

For the unlimited weight simulation, the learning rate is swept from 0.001 to 0.1 for every AF. The configurations with the smallest error are listed in table 2. The limited weight simulation is carried out in the same way except that the initial weight is swept from ± 1 to ± 255 . This is because that a small weight causes a small update which can keep the final weight within the specified range. Then, the configurations with higher than 90% accuracy are selected to sweep their maximum weights down until the accuracy is nearly equal to 90%. This is to save the circuit area by using the smallest bit-width. The simulation results are listed in table 3.

4. Discussions

This section discusses the results, and highlights the challenges and opportunities in the proposed approach. The section is organized hierarchically from perceptron circuit design to NN experiments validating through the MNIST benchmark application. Towards the end of the

section, we relate to our original hypothesis and summarize the key features of the proposed design approach.

(a) Pulse width modulation-coded perceptron

The design of the PWM-based VAC is shown in figure 20 to have satisfied its main design aim, which is to provide resilience in the presence of power supply uncertainty. Adopting a PWM-based approach in order to transfer computation from the digital domain to the relative temporal domain resulted in a device which is essentially independent of the value of the supply V_{dd} .

The use of an analogue voltage across a capacitor to represent the result of perceptron arithmetic computations allowed the use of the simplest digital gate, the inverter, to be the fundamental building block for both parts of the perceptron. Programmability, required for the FP functionality of NNs, for instance, can then be realized by using the next simplest digital gate, a two-input NAND gate. This results in an approach which implements digital computations using the smallest digital components working in the relative temporal domain on analogue values. Trading potential loss of precision for power resilience in this way is acceptable for NN applications, as they tend to be accuracy resilient at the point of any particular perceptron. The reduction of circuit complexity and avoidance of conventional multipliers and adders should also contribute to savings in both circuit size and energy consumption.

Both parts of the perceptron have been shown to have acceptable quality, including being reasonably linear within their operating ranges. However, after putting both together, the perceptron as a whole has certain range problems because the VAC output voltage may be outside the linear region of transistors in the voltage to PWM converter, which is a design assumption for that part. A single-transistor glue logic is then shown to help mitigate this problem.

The entire perceptron, however, shows weaker power supply variation resilience than the VAC on its own, under extreme dynamic V_{dd} variation conditions. This is primarily because of the following two factors:

- as can be seen from figure 9, the oscillator used to convert the voltage on the capacitor between the two parts of the perceptron (DC_{sum}) to PWM (DC_{out}) draws different amounts of power from the capacitor under different values of V_{dd} , and
- the simple glue logic in the form of a PMOS transistor does not fully compensate for this because its main function is to keep the current going once the voltage across the capacitor DC_{sum} gets down to threshold.

Computationally, if the correct value of DC_{sum} should be below threshold, which is entirely possible coming from the VAC, this glue logic will cause an error in that value by keeping DC_{sum} at threshold, leading to DC_{out} becoming inaccurate. This is preferable to having the voltage to PWM part dying but computation correctness is still lost.

The above discussion is supported by observing the behaviour shown in figure 24. The most likely places for errors to appear are when V_{dd} becomes low. In only some of these cases (Observation 2 but not Observation 1), the computation in the VAC would lead to a low DC_{sum} value which might dip below threshold. This will cause the output duty cycle DC_{out} to become incorrect. During Observation 1, although V_{dd} dips low, the computation because of the weights, etc., produces a high relative value result that manages to keep DC_{sum} above threshold. This is not the case during Observation 2. On the other hand, if V_{dd} itself reduces below threshold, no matter what value the VAC produces the voltage to PWM part will produce relatively large errors as the perceptron's glue logic would kick in anyway.

Also of concern is the accumulation of nonlinearity after both parts of the perceptron have been put together, and this nonlinearity continues to increase once the perceptron is connected in series across multiple stages of an NN of non-trivial depth.

Another issue is the fact that there may be difficulties for the perceptron to implement AFs other than flavours of ReLU. Of particular concern is that the perceptron, because of the use of

the voltage across a grounded capacitor to represent a crucial value, only works in the positive value domain.

Future research topics include the better matching between the constituent parts of the perceptron to overcome the threshold and nonlinearity problems and the extension of the perceptron to cover a larger set of AFs.

(b) Pulse width modulation-coded neural network

In the single perceptron, two-perceptron and three-perceptron experiments, both the resultant model and the Cadence simulations indicate that the voltage to PWM converter, without modifications, may serve as an approximate capped ReLU AF, qualitatively. In addition, the three-perceptron, three-stage full analogue simulation analysis shows that the single-perceptron MATLAB model can be used in multi-stage system analysis without worrying about the fidelity of high-level MATLAB models when multiple layers of perceptrons exist in a system.

Quantitatively, however, the use of a nonlinear perceptron to approximate linear behaviour becomes increasingly problematic when the depth of the network increases, as the nonlinearity accumulates. This is shown to be true by the subsequent whole-NN experiments.

The unlimited weight simulation result in table 2 gives us traditional NN examples which contain FP weights. It shows that both ReLU and capped ReLU functions give less than 4% errors at every depth. The results for the PWM perceptron AF are similar to those from the capped ReLU with offset. They obtain small error rates at the shallowest NN depth, while the capped ReLU without offset outperforms all others at every depth. This confirms that it is mainly the step transition at $DC_{sum} = 0$ that causes the convergence problem in our NN, more than the curvature nonlinearity higher in the curve of figure 29. Therefore, compensating the circuit to shift the output duty cycle back to 0% appears to be a promising route of investigation. This will be a subject in our future work.

Table 3 shows the results with weight limitations. All results at two-layer NN are worse than the ones in table 2 due to the weight capping and rounding, except for the PWM perceptron, which does better. The PWM perceptron continues to perform better at higher layer depths than in the unlimited weight case, confirming an advantage for it when weight is limited and represented by an integer. However, it again fails to approximate the ReLU function quantitatively at higher layer depths, by returning obviously better performances than the latter.

These results confirm the discussions in the previous section. Our proposed PWM-based perceptron's nonlinearity as well as not being able to properly extend to the low-percentage range of the PWM duty cycle (it starts from approximately 13% rather than 0% as shown in figures 26 and 29) makes it less suitable for deeper NNs. Its lack of support for negative values also limits its wider usability as the fundamental element of NNs, without further modification to better incorporate established AFs. The approximation of ReLU, although qualitatively promising, proves to be quantitatively unsatisfactory at higher NN depths, although in some cases this results in the perceptron's AF being better than the ReLU AF. The high error rate also comes from the resolution loss in basic weight rounding which may be solved by implementing a rounding technique and PF inference quantization presented in [34] and [35], respectively.

In other words, even if the negative value representation problem is solved, computing AFs in the analogue and relative temporal domains remains a challenge that must be solved. As a result, a future work direction is the incorporation of more general arithmetic operations in these domains, which is needed to improve the accuracy of AF implementations.

A related and interesting unsolved problem is the 'comparing with target vector' function in figure 11, which is currently relegated to external controllers. It is a duty cycle in and digital out block and can potentially be designed by extending the methods in this work.

Table 4 summarizes our design compared to related work. The work in [36] quantizes the entire NN and yields the lowest error. However, it is designed for a digital-based processing unit which contains the CPU-memory bottleneck issue implying extra power budget and latency. Furthermore, real power measurement is missing as it estimates the power consumption from

Table 4. Performance comparison.

work	weight type	MNIST	NN conf.	error	power (μ W)	power elastic	hardware
[36]	integer	quantized	WAGE	0.4	n.a.	N	MCU
[34]	integer	binarized	MLP	<3%	n.a.	N	n.a.
[37]	n.a.	reduced	MLP	13.5%	53 000	N	memristor
		quantized			(NN)		crossbar
[31]	fixed-point	original	MLP	~5%	14 800	N	Spec. CMOS
					(NN)		
this study	integer	original	MLP	<10%	14-1080	Y	Std. CMOS
					(VAC)		

the literature. Weight rounding methods are proposed in [34]. Although they achieve the second lowest error, they require binarized input data which is a challenge for analogue applications. Moreover, it does not include an investigation of hardware implementation. The memristor crossbar NN in [37] acquires the lowest accuracy with the highest power consumption even its input image size is reduced. Charge trap transistor-based NN which performs MNIST classification from the original data are presented in [31]. It mainly aims to save power and requires a specific CMOS technology to fabricate the special transistors.

Even though the error of our design is higher, it is still within the same order of magnitude, and we have yet to investigate more sophisticated techniques for compensating the voltage to PWM part to improve the duty cycle coverage and eliminate the step nonlinearity at $DC_{sum} = 0$, which promises to reduce error. We also do not investigate beyond standard CMOS technology as that is out of scope for this investigation. Our focus is on tolerating unstable and unpredictable power supply voltages, a necessity for energy autonomous AI devices. In this regard, this solution is unique as existing research in the literature invariably requires stable and known voltages and operating frequencies. Note that the power figures in table 4 are not directly comparable. The figure for [34] is obtained from measurements carried out on a fabricated chip at the 28 nm technology node with special non-CMOS transistor techniques aiming to showcase the advantage of that technology in low-power operations, while that for this work is obtained from simulating one VAC at the 65 nm technology node with a deliberately high (2.5 V) nominal supply voltage to facilitate studying voltage instability scenarios. This is similar to [37] where the circuit is implemented with non-CMOS technology. A fair power consumption comparison with [34] and [37] is not yet possible without fabricating and testing real chips, preferably at the same VLSI technology node, as whole-NN simulations at the VLSI level is not practical. In addition, the power figures for [34] and [37] are themselves not comparable with each other as [37] covers the entire system including peripherals and [34] covers the NN engine only.

(c) Overall summary

Our design methods, supported by extensive analysis and validations, have proven the original hypothesis and demonstrated the following features:

- *power elasticity and resilience* across a dynamic range of V_{dd} (statistically varying by $5\times$) and f (statistically varying by up to 6 orders of magnitude). Such elasticity is achieved for the VAC without requiring any voltage regulator circuit and clock pairing between V_{dd} and f . Dynamic power supply variations also show good resilience. However, further compensation will be required at lower voltages to avoid large errors at the whole-perceptron level;

- minimal use of additional analogue (ideally, only passive components) electronics, coupled with low-complexity PWM-coded arithmetic using primarily digital components, making our approach highly power efficient and suitable for low-cost fabrication;
- extensive validation and analysis using multi-layer PWM-coded NNs show good scalability of the proposed approach; however, deeper NNs may need circuit-level compensation after each layer or high-precision representation techniques to improve the overall accuracy and efficiency.

5. Conclusion

We propose the first mixed-signal (analogue/digital/relative temporal) perceptron design using the principles of PWM. Central to our design are a number of parallel inverters that suitably transcode the input-weight pairs from the spatial domain to the relative temporal domain. This approach aims to deliver high resilience to amplitude and frequency variations in the supply voltage, exploiting the fact that PWM-based solutions are typically agnostic to such variations.

Another advantage of the proposed design is its simplicity. While conventional implementations of the perceptron require complex logic to perform multiplication and addition, the proposed approach uses only one gate (either an inverter or a two-input NAND) per bit for every input. Thus, for the 3×3 weighted addition VAC we used only 54 transistors. This significantly reduces the logic requirement and, therefore, the power consumption of the entire device.

Extensive experimentation on the perceptron and its use in neural networks of relatively significant sizes helps to explore the perceptron design's advantages, usability and limitations. Also through experimental studies, design improvements are found which further strengthen the perceptron's case. These experimental explorations also lead to further insights into the design and provide guidance on potential future work.

The perceptron's arithmetic unit design is shown to fully accomplish its design aim of power and frequency resilience. It is also shown to be working within reasonable boundaries in pragmatic applications, especially in NNs of limited depth which are nevertheless of significant size. Future improvements should be concentrated on improving its linearity, its threshold voltage independence and its representation of negative values to increase its usability of NNs of greater depth and more sophisticated AFs.

This work points to potentially exciting research to extend the computation capabilities of devices working in the relative temporal and analogue domains.

Machine learning is finding more applications at the micro-edge, where power variation from energy harvesters is becoming commonplace. We believe the proposed perceptron will find practical implementations in these applications as it is highly robust to these variations.

Data accessibility. All supplementary data have been submitted as electronic supplementary material.

Authors' contributions. S.M. was responsible for the low-level perceptron and NN circuit design, experiments and analysis, and also led the writing of the paper. T.B. was responsible for the scale-up NN circuit modelling, analysis and experiments using MNIST benchmark application, and co-led the writing of the paper. F.X. contributed to the writing of the paper, led its editing and participated in technical discussions. R.S. contributed in the scale-up models, tied up with the low-level circuits, co-supervised the circuit- and system-level works of S.M. and T.B. and contributed to writing/editing. A.Y. contributed by proposing the idea of using CMOS logic for duty-cycle based computing for power elasticity and robustness, supervising the work of S. Mileiko and T. Bunnam, and all technical discussions. S.D. has contributed in the circuit design aspects, particularly reflecting on how averaging could affect the overall accuracy.

Competing interests. The authors declare that they have no competing interests.

Funding. This work was supported by the EPSRC STRATA project (grant no. EP/N023641/1).

Acknowledgements. The authors thank Mr Jonathan Edwards, CTO of Temporal Computing Ltd., for his useful suggestions in the early stages.

1. Schwab K. 2017 *The fourth industrial revolution*. London, UK: Portfolio Penguin.
2. Javaid N, Sher A, Nasir H, Guizani N. 2018 Intelligence in IOT-based 5g networks: opportunities and challenges. *IEEE Commun. Mag.* **56**, 94–100. (doi:10.1109/MCOM.2018.1800036)
3. Biswas A, Chandrakasan AP. 2018 Conv-ram. An energy-efficient SRAM with embedded convolution computation for low-power CNN-based machine learning applications. In *2018 IEEE Int. Solid - State Circuits Conf. - (ISSCC), 11–15 Feb, 2018, San Francisco, CA*, pp. 488–490. New York, NY: IEEE. (doi:10.1109/isscc.2018.8310397)
4. Chen M, Miao Y, Jian X, Wang X, Humar I. 2018 Cognitive-LPWAN. Towards intelligent wireless services in hybrid low power wide area networks. (<http://arxiv.org/abs/1810.00300>).
5. Sharbati MT, Du Y, Torres J, Ardolino ND, Yun M, Xiong F. 2018 Low-power, electrochemically tunable graphene synapses for neuromorphic computing. *Adv. Mater.* **30**, 1802353. (doi:10.1002/adma.201802353)
6. Neftci EO. 2018 Data and power efficient intelligence with neuromorphic learning machines. *iScience* **5**, 52–68. (doi:10.1016/j.isci.2018.06.010)
7. Shafik R, Yakovlev A. 2019 Chapter: from power-efficient to power-driven computing. In *Many-core computing: hardware and software* (eds G Merrett, BM Al-Hashimi), pp. 271–293. London, UK: IET. (doi:10.1049/pbpc022e_ch11)
8. Shafik RA, Yang S, Das A, Maeda-Nunez LA, Merrett GV, Al-Hashimi BM. 2016 Learning transfer-based adaptive energy minimization in embedded systems. *IEEE Trans. Computer-Aided Design Integr. Circuits Syst.* **35**, 877–890. (doi:10.1109/TCAD.2015.2481867)
9. Qiqieh I, Shafik R, Tarawneh G, Sokolov D, Yakovlev A. 2017 Energy-efficient approximate multiplier design using bit significance-driven logic compression. In *Proc. of the Conf. on Design, Automation & Test in Europe, 27–31 March, Lausanne, Switzerland*, ser. DATE '17. European Design and Automation Association, pp. 7–12. See <http://dl.acm.org/citation.cfm?id=3130379.3130382>.
10. Shafik R, Yakovlev A, Das S. 2018 Real-power computing. *IEEE Trans. Comput.* **67**, 1445–1461. (doi:10.1109/TC.12)
11. Beeby S, White NM. 2010 *Energy harvesting for autonomous systems*. Boston, MA; London, UK: Artech House.
12. Yang T-J, Chen Y-H, Sze V. 2017 Designing energy-efficient convolutional neural networks using energy-aware pruning. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition, Honolulu, HI, 21–26 Jul 2017*, pp. 5687–5695. New York NY: IEEE. (doi:10.1109/cvpr.2017.643)
13. Suleiman A, Chen Y-H, Emer J, Sze V. 2017 Towards closing the energy gap between HOG and CNN features for embedded vision. In *2017 IEEE Int. Symp. on Circuits and Systems (ISCAS), Baltimore, MD, 28–31 2017*, IEEE, pp. 1–4. New York, NY: IEEE. (doi:10.1109/iscas.2017.8050341)
14. Andri R, Cavigelli L, Rossi D, Benini L. 2018 Yodann: an architecture for ultralow power binary-weight CNN acceleration. *IEEE Trans. Computer-Aided Design Integr. Circuits Syst.* **37**, 48–60. (doi:10.1109/TCAD.2017.2682138)
15. Rosenblatt F. 1958 The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* **65**, 386–408. (doi:10.1037/h0042519)
16. Hagan MT, Demuth HB, Beale M. 1996 *Neural network design*. Boston, MA: PWS Publishing Co.
17. Wilson E, Tufts DW. 1994 Multilayer perceptron design algorithm. In *Proc. of IEEE Workshop on Neural Networks for Signal Processing, Ermioni, Greece, 6–8 Sep 1994*, pp. 61–68. New York, NY: IEEE. (doi:10.1109/nnspp.1994.366063)
18. Adeli H, Yeh C. 1989 Perceptron learning in engineering design. *Computer-Aided Civil Infrastruct. Eng.* **4**, 247–256. (doi:10.1111/mice.1989.4.issue-4)
19. Hung S, Adeli H. 1991 A model of perceptron learning with a hidden layer for engineering design. *Neurocomputing* **3**, 3–14. (doi:10.1016/0925-2312(91)90016-5)
20. Jeong B, Lee YH. 1994 Design of weighted order statistic filters using the perceptron algorithm. *IEEE Trans. Signal Process.* **42**, 3264–3269. (doi:10.1109/78.330393)
21. Qinruo W, Bo Y, Yun X, Bingru Liu. 2003 The hardware structure design of perceptron with FPGA implementation. In *SMC'03 Conf. Proc. 2003 IEEE Int. Conf. on Systems, Man and*

- Cybernetics. Conference Theme - System Security and Assurance (Cat. No.03CH37483)*, vol. 1, Washington, DC, 5–8 Oct 2003, pp. 762–767. New York, NY: IEEE. (doi:10.1109/icsmc.2003.1243906)
22. hsin Chen Y, Yang T-J, Emer JS. 2018 Understanding the limitations of existing energy-efficient design approaches for deep neural networks. *Energy* **2**, L3.
  23. Keyes RW, Landauer R. 1970 Minimal energy dissipation in logic. *IBM J. Res. Dev.* **14**, 152–157. (doi:10.1147/rd.142.0152)
  24. de Prado M, Denna M, Benini L, Pazos N. 2018 Quenn: Quantization engine for low-power neural networks. In *Proc. of the 15th ACM Int. Conf. on Computing Frontiers, Ischia, Italy, 8–10 May 2018*, pp. 36–44. New York, NY: ACM. (doi:10.1145/3203217.3203282)
  25. Yakovlev A. 2011 Energy-modulated computing. In *2011 Design, Automation Test in Europe, Grenoble, France, 14–18 Mar 2011*, pp. 1–6. New York, NY: IEEE. (doi:10.1109/date.2011.5763216)
  26. Shang D, Zhang X, Xia F, Yakovlev A. 2014 Asynchronous design for new on-chip wide dynamic range power electronics. In *2014 Design, Automation Test in Europe Conf. Exhibition, Grenoble, France, 24–28 Mar 2014*, March 2014, pp. 1–6. New York, NY: IEEE. (doi:10.7873/date.2014.151)
  27. LiKamWa R, Hou Y, Gao Y, Polansky M, Zhong L. 2016 Redeye: Analog convnet image sensor architecture for continuous mobile vision. In *2016 ACM/IEEE 43rd Annual Int. Symp. on Computer Architecture (ISCA), Seoul, Korea, 18–22 Jun 2016*, pp. 255–266. New York, NY: IEEE/ACS. (doi:10.1109/isca.2016.31)
  28. Chen YH, Krishna T, Emer JS, Sze V. 2016 Eyeriss: an energy-efficient reconfigurable accelerator for deep convolutional neural networks. In *IEEE Int. Solid-State Circuits Conf., ISSCC 2016, Digest of Technical Papers, San Francisco, CA, 31 Jan–4 Feb 2016*, pp. 262–263. New York, NY: IEEE. (doi:10.1109/isscc.2016.7418007)
  29. Vaisband I, Azhar M, Friedman EG, Köse S. 2014 Digitally controlled pulse width modulator for on-chip power management. *IEEE Trans. Very Large Scale Integr. VLSI Syst.* **22**, 2527–2534. (doi:10.1109/TVLSI.2013.2294402)
  30. Lecun Y, Bottou L, Bengio Y, Haffner P. 1998 Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324. (doi:10.1109/5.726791)
  31. Du Y, Du L, Gu X, Du J, Wang XS, Hu B, Jiang M, Chen X, Iyer SS, Chang MF. 2018 An analog neural network computing engine using CMOS-compatible charge-trap-transistor (CTT). *IEEE Trans. Computer-Aided Design Integr. Circuits Systems* **38**, 1–1. (doi:10.1109/tcad.2018.2859237)
  32. Deng L. 2012 The MNIST database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Process Mag.* **29**, 141–142. (doi:10.1109/MSP.2012.2211477)
  33. Zeiler MD, et al. 2013 On rectified linear units for speech processing. In *2013 IEEE Int. Conf. on Acoustics, Speech and Signal Processing, Vancouver, Canada, 26–31 May 2013*, pp. 3517–3521. New York, NY: IEEE.
  34. Muller LK, Indiveri G. 2015 Rounding methods for neural networks with low resolution synaptic weights. (<http://arxiv.org/abs/1504.05767>).
  35. Jacob B, Kligys S, Chen B, Zhu M, Tang M, Howard A, Adam H, Kalenichenko D. 2017 Quantization and training of neural networks for efficient integer-arithmetic-only inference. See <https://ui.adsabs.harvard.edu/#abs/2017arXiv171205877J>.
  36. Wu S, Li G, Chen F, Shi L. 2018 Training and inference with integers in deep neural networks. See <https://ui.adsabs.harvard.edu/#abs/2018arXiv180204680W>.
  37. Jiang H, et al. 2018 Pulse-width modulation based dot-product engine for neuromorphic computing system using memristor crossbar array. In *2018 IEEE Int. Symp. on Circuits and Systems (ISCAS), Conf. Proc., Florence, Italy, 27–30 May 2018*, pp. 1–4. New York, NY: IEEE. (doi:10.1109/iscas.2018.8351276)