

Computational and Storage Efficient Quadratic Neurons for Deep Neural Networks

Chuangtao Chen¹, Grace Li Zhang², Xunzhao Yin³, Cheng Zhuo³, Ulf Schlichtmann¹, Bing Li¹

¹Technical University of Munich, ²Technical University of Darmstadt, ³Zhejiang University
{chuangtao.chen, ulf.schlichtmann, b.li}@tum.de, grace.zhang@tu-darmstadt.de, {xzyin1, czhuo}@zju.edu.cn

Abstract—Deep neural networks (DNNs) have been widely deployed across diverse domains such as computer vision and natural language processing. However, the impressive accomplishments of DNNs have been realized alongside extensive computational demands, thereby impeding their applicability on resource-constrained devices. To address this challenge, many researchers have been focusing on basic neuron structures, the fundamental building blocks of neural networks, to alleviate the computational and storage cost. In this work, an efficient quadratic neuron architecture distinguished by its enhanced utilization of second-order computational information is introduced. By virtue of their better expressivity, DNNs employing the proposed quadratic neurons can attain similar accuracy with fewer neurons and computational cost. Experimental results have demonstrated that the proposed quadratic neuron structure exhibits superior computational and storage efficiency across various tasks when compared with both linear and non-linear neurons in prior work.

I. Introduction

Over the past decade, there has been an unprecedented success in the field of artificial neural networks. Neural networks, particularly those with deep layers, have been utilized in various tasks, including but not limited to, image classification [1], [2], natural language processing [3], [4] and generative tasks [5]. The excellent performance of deep neural networks (DNNs) are achieved at the cost of heavy computational and storage cost. These challenges have become a critical concern for devices with constrained computational and storage resources.

To mitigate the computational and storage burden of DNNs, there have been various research efforts including novel efficient network architectures [6], network pruning [7]–[9] and quantization [10]–[12]. Recently, some researchers have been focusing on structures of individual neurons, which map multiple inputs to a single output followed by an activation function. To increase the efficiency of neurons, several studies have explored the use of more sophisticated non-linear neurons [13], [14], particularly second-order neurons, which are also referred to as quadratic neurons [15], [16]. However, the parameter and computation costs associated with such non-linear neurons often make them unsuitable for deployment in practice. To reduce the parameter and computation overhead, some researchers have attempted to use quadratic neurons only in the initial or certain layers [13], [16], [17], or to optimize complex neurons into more manageable forms [17]–[19]. Nevertheless, there are still a large number of parameters and thus a high computation cost incurred by such quadratic neurons. In addition, the second-order information of neurons in these methods has not been exploited sufficiently, resulting in inefficient neuron implementations.

In this paper, we propose a novel efficient quadratic neuron

for DNNs to address the challenges mentioned above. The proposed quadratic neuron possesses the following features:

- By following a rigorous mathematical procedure rather than using a heuristic approach, we simplify quadratic neurons to reduce the incurred parameter and computation cost while preserving the expressivity.
- To maximally utilize the second-order information of quadratic neurons to enhance the expressivity of DNNs, the intermediate results in the second-order computation are used as the partial output of the quadratic neuron.
- The efficacy of the proposed neuron structure is evaluated in various tasks, such as convolutional neural networks (CNNs) and Transformers. Compared with state-of-the-art quadratic neurons [18], [20], neural networks with our neurons achieve a better or similar accuracy on CIFAR-10 dataset [21] but with at least 24% fewer parameters and 24% less computational cost. In Transformer architecture, networks equipped with our proposed neuron structure can reduce both the computational cost and storage cost by more than 20% with a better accuracy in the WMT14 German-English translation task.

This paper is structured as follows: Sec. II provides a more in-depth overview of non-linear neurons, encompassing their benefits as well as their limitations. In Sec. III, we present the construction of our proposed neuron, including a systematic simplifying procedure from a complex form. The effectiveness of our approach is evaluated by comparing the proposed neurons with the linear and other non-linear designs, and the experimental results as well as an in-depth analysis of the proposed quadratic neuron are reported in Sec. IV. Finally, Sec. V summarizes and concludes this paper.

II. Background and Motivation

This section presents the background and motivation of our work, highlighting various types of neurons explored by previous research, along with their respective advantages and drawbacks. We denote vectors with boldface lowercase letters and matrices with boldface uppercase letters. Scalars are represented by non-boldface lowercase letters. Unless otherwise stated, all vectors are assumed to be column vectors of length n and matrices have dimensions of $n \times n$. Subscripts are used for indexing elements in a vector or a matrix. For example, x_3 represents the third element (scalar) in vector \mathbf{x} .

A. General Quadratic Neuron

A neuron can be perceived as a multivariate function of its inputs. We denote the conventional linear neuron, which is shown in Fig. 1a, as $f(\mathbf{x}) = \sigma(\mathbf{w}^T \mathbf{x} + b)$, where \mathbf{w} is the

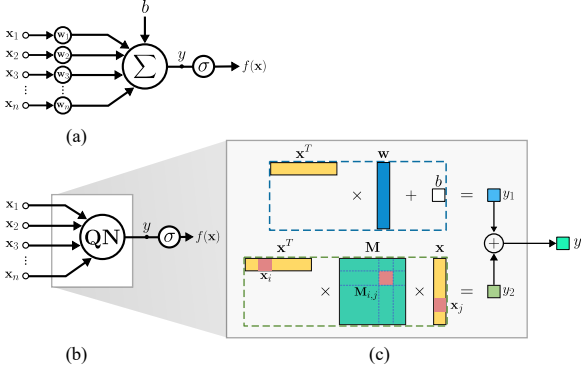


TABLE I
SUMMARY OF QUADRATIC NEURONS.

	Neuron Format	Parameter	Computation
[16]	$\mathbf{x}^T \mathbf{M} \mathbf{x} + \mathbf{w}^T \mathbf{x}$	$\mathcal{O}(n^2 + n)$	$\mathcal{O}(n^2 + 2n)$
[15]	$\mathbf{x}^T \mathbf{M} \mathbf{x}$	$\mathcal{O}(n^2)$	$\mathcal{O}(n^2 + n)$
[22]	$(\mathbf{w}_1^T \mathbf{x})(\mathbf{w}_2^T \mathbf{x}) + \mathbf{w}_3^T \mathbf{x}$	$\mathcal{O}(2n)$	$\mathcal{O}(2n)$
[17]	$\mathbf{x}^T \mathbf{Q}_1^k (\mathbf{Q}_2^k)^T \mathbf{x} + \mathbf{w}^T \mathbf{x}$	$\mathcal{O}(2kn + n)$	$\mathcal{O}(2kn + k)$
[18]	$(\mathbf{w}_1^T \mathbf{x})(\mathbf{w}_2^T \mathbf{x}) + \mathbf{w}_3^T (\mathbf{x}^{\odot 2})$	$\mathcal{O}(3n)$	$\mathcal{O}(4n)$
[20]	$(\mathbf{w}_1^T \mathbf{x})(\mathbf{w}_2^T \mathbf{x}) + \mathbf{w}_3^T \mathbf{x}$	$\mathcal{O}(3n)$	$\mathcal{O}(3n)$
Ours	$\{\mathbf{x}^T \mathbf{Q}^k \mathbf{\Lambda}^k (\mathbf{Q}^k)^T \mathbf{x} + \mathbf{w}^T \mathbf{x}, \mathbf{x}^T \mathbf{Q}^k \mathbf{x}\}$	$\mathcal{O}(n + \frac{k}{k+1})$	$\mathcal{O}(n + \frac{2k}{k+1})$

¹ Subscripts in this table are used to distinguish different items.

² $\mathbf{x}, \mathbf{w} \in \mathbb{R}^n, \mathbf{M} \in \mathbb{R}^{n \times n}$, where n is the number of neuron's inputs.

³ $\mathbf{Q}^k \in \mathbb{R}^{n \times k}, \mathbf{\Lambda}^k \in \mathbb{R}^{k \times k}$, where k is the rank of decomposition.

⁴ $(\cdot)^{\odot 2}$ is the element-wise square operation.

⁵ $\{\cdot\}$ is the concatenation operation.

Fig. 1. Structure of (a) the linear neuron and (b) the general quadratic neuron; (c) The computation process of the general quadratic neuron.

weight vector, \mathbf{x} is the input vector, b is the bias and $\sigma(\cdot)$ is the activation function. A linear neuron without the activation function can be viewed as a first-order polynomial function with respect to all the inputs with the bias serving as the zeroth-order term. In this paper, we only focus on the computation before the activation function in a neuron and the computational result is denoted as y for simplicity.

To enhance the expressivity of neural networks, a typical approach is to replace linear neurons with non-linear neurons, such as the quadratic neuron [15], [16], [18], [20]. The general quadratic neuron can be represented mathematically as:

$$y = \underbrace{\mathbf{x}^T \mathbf{M} \mathbf{x}}_{\text{quadratic part: } y_2} + \underbrace{\mathbf{w}^T \mathbf{x} + b}_{\text{linear part: } y_1}. \quad (1)$$

The general quadratic neuron in Eq. (1) is the most complex one and possesses the largest function space among all quadratic neurons. Fig. 1b depicts the structure of a general quadratic neuron, while its computation process is shown in Fig. 1c. As previously mentioned, the linear neuron on the left performs a summation of all weighted inputs and a bias term. In contrast, the computation of the general quadratic neuron is considerably more complex. Besides the parameters that also exist in a linear neuron, there is an additional parameter matrix \mathbf{M} with dimensions of $n \times n$ in the general quadratic neuron. The final output of the general quadratic neuron y is the sum of the linear term and the quadratic term. In Fig. 1c we divide the computation in the quadratic neuron into two sections enclosed by dotted-line rectangles: the upper section represents the linear part comprising of the same computation as in a linear neuron, where the result is denoted as $y_1 = \mathbf{w}^T \mathbf{x} + b$. The lower section represents the quadratic part, where the element at the i^{th} row and j^{th} column of the matrix corresponds to the coefficient of the cross term $x_i x_j$ (depicted in red in Fig. 1c). The quadratic term, represented by the equation $y_2 = \mathbf{x}^T \mathbf{M} \mathbf{x}$, is the summation of all cross terms, each of which is weighted by the corresponding element in the matrix \mathbf{M} . For example, if $n = 3$, the quadratic term is the summation of 9 cross terms:

$$\mathbf{x}^T \mathbf{M} \mathbf{x} = \sum_{i=1}^3 \sum_{j=1}^3 x_i M_{i,j} x_j. \quad (2)$$

Due to the inclusion of interactions between any two inputs, the general quadratic neuron exhibits strong non-linearity, yielding better expressivity than the linear neuron.

B. Complexity Reduction of the General Quadratic Neuron

Although the general quadratic neuron is theoretically superior, it introduces significant additional parameters and computation costs. In Table I we provide a comprehensive summary of various quadratic neurons, including the mathematical formulations as well as the associated parameter and computational complexities. For simplicity, we ignore the bias term in the formulation and complexity analysis. There has been attempt to integrate the general quadratic neuron shown in Eq. (1), which has a parameter complexity of $n^2 + n$, into modern DNNs [16]. To avoid excessive complexity, the general quadratic neurons are only deployed in the initial layer [16]. Another variant of the general quadratic neuron without the linear term has been utilized in [15]. Unfortunately, the quadratic complexity still makes it impractical to implement such quadratic neurons in modern neural networks directly.

Recently, some attention has been given to simplified quadratic neurons that can maintain the non-linearity of the general quadratic neuron without incurring excessive parameter overhead [17], [18], [20], [22]. One approach to address the complexity issues associated with quadratic neurons is to employ low-rank decomposition. In [17], the general quadratic form is simplified with low-rank decomposition, resulting in a reduction of complexity from n^2 to $2kn$, where k is the rank of decomposition ranging from 1 to n . [18], [20] has proposed a neuron that is similar to [17] but with a fixed rank of 1.

Our proposed quadratic neuron is presented in the last row of Table I. This neuron also incorporates a hyper-parameter k , which is the rank of decomposition. However, unlike prior work such as [17], the parameter and computational complexities of our design do not increase proportionally to k . This property allows for more flexible choices of decomposition ranks.

By investigating previous studies, we summarize the main challenges associated with quadratic neurons as follows:

- **Extra Parameter and Computational Cost:** Despite the efforts to simplify the general quadratic form to reduce its complexity to linear levels [17], [18], [20], there still exists considerable overhead compared to the linear neuron. Some work [16], [17] has attempted to mitigate these issues by only applying the complex neurons to the initial or some selected layers. However, this partial deployment impairs the deployment flexibility of these neurons.
- **Compromised Expressivity:** While some existing meth-

ods have leveraged simplified versions of general quadratic neurons, many of them have sacrificed the expressivity to reduce the neuron complexity [18], [20], [22]. For example, the quadratic neuron in [17] allows the adjustment of expressivity with a hyper-parameter k . However, as shown in Table I, the cost is proportional to the value of k , making it impractical to increase the neuron's expressivity with a large k . Notably, only neurons with low ranks such as $k = 1, 2$ and 3 have been explored in [17].

- **Underutilization of Internal Features:** Though quadratic neurons introduce complex computations to the inputs, the computational results are simply summed to a singular value, which forms the neuron's output. The direct summation of all intermediate terms may obscure the features extracted by different computational processes instead of exhaustively making use of their information.

III. Proposed Efficient Quadratic Neuron

In this section, we present our proposed novel quadratic neuron as a solution to the aforementioned issues in Sec. II. To maintain the neuron's expressivity, we follow a strict procedure to simplify the general quadratic neuron. Additionally, by incorporating the reuse of intermediate information within the neuron, the required number of neurons can be further reduced for the same size of feature maps, thereby lowering computational and storage complexities.

A. Quadratic Matrix Decomposition

As previously noted in Sec. II, the general quadratic neuron consists of two terms: the linear term and the quadratic term. In this section, we mainly discuss the quadratic term which comprises the majority of parameters.

Compared with previous methods which have fixed and compromised expressivity [18], [20], [22] or do not fully exploit the characteristics of the quadratic neuron, such as symmetry [17], we adopt a systematic approach that utilizes the symmetry of the quadratic matrix. Starting from the general quadratic neuron, we manage to reduce the number of parameters in the quadratic neuron while maintaining its expressive ability to the fullest extent possible.

To facilitate the subsequent construction, we present the following lemma to illustrate the property of quadratic neurons:

Lemma 1. *For any real-valued $n \times n$ matrix \mathbf{M} , there exists a real-valued **symmetric** matrix \mathbf{M}' such that:*

$$\mathbf{x}^T \mathbf{M} \mathbf{x} = \mathbf{x}^T \mathbf{M}' \mathbf{x}, \forall \mathbf{x} \in \mathbb{R}^n. \quad (3)$$

Proof. To prove this, we can split the quadratic term into two halves and transpose one of them:

$$\begin{aligned} \mathbf{x}^T \mathbf{M} \mathbf{x} &= \frac{1}{2} \left(\mathbf{x}^T \mathbf{M} \mathbf{x} + (\mathbf{x}^T \mathbf{M} \mathbf{x})^T \right) \\ &= \mathbf{x}^T \left(\frac{\mathbf{M} + \mathbf{M}^T}{2} \right) \mathbf{x}. \end{aligned} \quad (4)$$

By defining $\mathbf{M}' \triangleq (\mathbf{M} + \mathbf{M}^T)/2$, which is a real-valued symmetric matrix, we prove the lemma. \square

With Lemma 1, the matrix \mathbf{M} can always be replaced by a symmetric matrix without affecting the output. Therefore, for

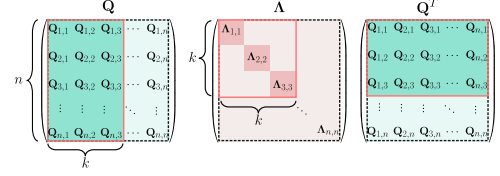


Fig. 2. Top- k selection of the decomposed matrices \mathbf{Q} and $\mathbf{\Lambda}$.

simplicity in notation, we assume that the matrix \mathbf{M} is already symmetric in the following discussions.

To simplify the quadratic neuron, we use a low-rank decomposition method which is similar to [17]. However, there are two key differences in our methodology. Firstly, we utilize eigendecomposition to factorize \mathbf{M} into three matrices. Secondly, we leverage the matrix symmetry from Lemma 1 to further reduce the complexity by half compared with [17].

According to the spectral theorem that a real symmetric matrix can be diagonalized using eigendecomposition with an orthonormal matrix, we simplify the quadratic matrix as:

$$\mathbf{M} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^T, \quad (5)$$

where \mathbf{Q} is an orthonormal matrix and $\mathbf{\Lambda}$ is a diagonal matrix.

The eigendecomposition possesses a crucial property that the order of eigenvalues in matrix \mathbf{M} may be interchanged, if the corresponding eigenvectors are also swapped in the same order. This property is leveraged in principal component analysis (PCA) by sorting the eigenvalues in a descending order based on their magnitudes [23]. For simplicity in notation, we make the assumption that the elements of matrix $\mathbf{\Lambda}$ have already been sorted in descending order.

In this study, we adopt a method that retains the top- k eigenvalues in terms of their magnitude. This method achieves the optimal rank- k approximation for a matrix in terms of the Frobenius norm according to the Eckart-Young-Mirsky theorem [24]. Fig. 2 illustrates the top- k selection from the original matrices, resulting in an approximation of \mathbf{M} denoted as \mathbf{M}^k . The approximation can be formulated as:

$$\mathbf{M} \approx \mathbf{M}^k = \mathbf{Q}^k \mathbf{\Lambda}^k (\mathbf{Q}^k)^T, \quad (6)$$

where \mathbf{Q}^k is a $n \times k$ matrix pruned from the first k columns of \mathbf{Q} , and $\mathbf{\Lambda}^k$ is a $k \times k$ matrix with top- k eigenvalues placed at its diagonal.

We denote the quadratic term with rank- k approximation as:

$$y_2 \approx y_2^k = \mathbf{x}^T \mathbf{Q}^k \mathbf{\Lambda}^k (\mathbf{Q}^k)^T \mathbf{x}. \quad (7)$$

With the use of inherent symmetry in the quadratic matrix, our approximation can reduce the complexity in terms of the number of parameters from n^2 to $kn + k$, where k is an adjustable hyper-parameter, allowing us to control the neuron's expressivity. Compared with previous works using low-rank approximation [17], our method reduces the complexity significantly while fully preserving the expressive ability.

B. Information Utilization in Quadratic Neurons

Complex neurons in prior research try to improve the expressive ability of a neural network by introducing additional parameters. Consequently, these modifications lead to the generation of more intermediate results during the computation, including the cross terms. However, most of them simply

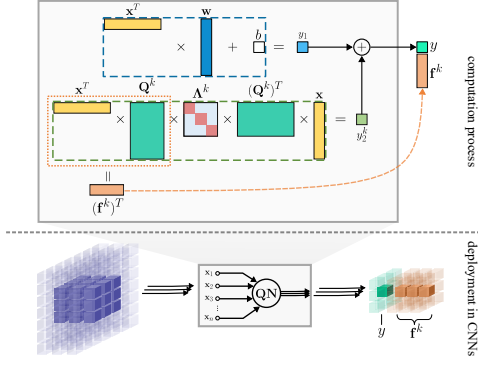


Fig. 3. Computation process of the proposed quadratic neuron and its deployment in a convolutional layer.

collapse multiple terms into a singular output using additions. Such a simplistic method of reduction can result in a significant loss of information, particularly when the number of intermediate results is substantial. For example, the general quadratic neuron generates n^2 cross terms, each of which may represent some kind of feature. The simple summation may inhibit the expression of these features, causing the underutilization of internal information in a neuron.

To address this issue, we propose a scheme that outputs the intermediate features of a quadratic neuron to the next layer. It can be observed that the computation process of $(Q^k)^T x$ in Eq. (7) is mathematically equivalent to k linear neurons with each column of Q^k serving as the linear weights. By defining the intermediate product as:

$$f^k \triangleq (Q^k)^T x, \quad (8)$$

where f^k is a vector with the length of k , each element of f^k can be viewed as the output of a linear neuron. These intermediate features may also contain useful information which can be used by the following layers. Therefore, to achieve exhaustive utilization of this internal information, we take f^k as additional outputs besides the original output y , which is the sum of y_1 and y_2^k , to the next layer, resulting in a neuron with a vectorized output. Fig. 3 illustrates the structure of the proposed quadratic neuron, where y and f^k are concatenated as the final output, and its deployment in a CNN. When deployed in a convolutional layer of a CNN, the additional outputs are placed along the channel dimension, which means that each convolutional filter can output multiple channels.

C. Parameter and Computational Complexity Analysis

In this section, we discuss the complexity of our proposed neuron in terms of **storage cost** (number of *parameters*) and **computational cost** (number of *multiply-accumulate operations* (MACs)). For simplicity, we ignore the cost of the bias term, which is negligible compared to other operations.

Parameter: Each quadratic neuron contains the following trainable parameters: $Q^k \in \mathbb{R}^{n \times k}$, $\Lambda^k \in \mathbb{R}^{k \times k}$, and $w \in \mathbb{R}^{n \times 1}$. Note that Λ^k is a diagonal matrix thus it has only k parameters. The total number of parameters is the sum of all mentioned terms above:

$$\#Parameter = (k+1)n + k. \quad (9)$$

Computation: The computation process of our proposed neuron contains two parts as shown in Fig. 3: the lin-

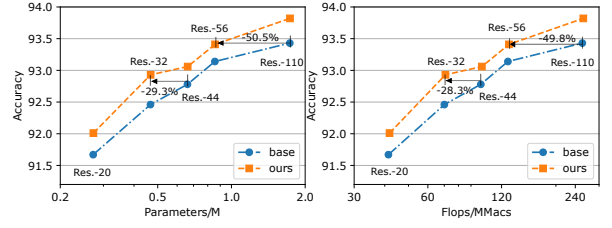


Fig. 4. Accuracy on CIFAR-10, number of parameters and FLOPs of ResNets (20, 32, 44, 56 and 110) with linear and quadratic neurons.

ear part $w^T x$ with the cost of n , and the quadratic part $x^T Q^k \Lambda^k (Q^k)^T x$ in Eq. (7) which consists of four matrix multiplications.

The overall computational cost of the quadratic part depends on the order of matrix multiplications. We first compute the product of $(Q^k)^T$ and x . Then the next step is to compute $y_2^k = (f^k)^T \Lambda^k f^k$. The computation cost of $f^k = (Q^k)^T x$ is kn , and the computation of $(f^k)^T \Lambda^k$ which contains k element-wise multiplications has a cost of k . $y_2^k = (f^k)^T \Lambda^k f^k$, which requires twice such multiplications, has a cost of $2k$. Accordingly, the overall cost of our proposed neuron structure is the sum of n from the linear part, kn from $f^k = (Q^k)^T x$, and $2k$ from $y_2^k = (f^k)^T \Lambda^k f^k$:

$$\#MAC = (k+1)n + 2k. \quad (10)$$

Averaged Complexity: As we mentioned in Sec. III-B, each of our neurons generates $k+1$ outputs. Therefore, when deploying them to a DNN, fewer neurons are required to obtain the original sizes of feature maps. To reflect the cost of our neurons in practice more accurately, we analyze the complexity of a neuron averaged by the size of its output. With Eqs. (9) and (10), the averaged parameter and computation cost of our proposed neuron per output are $n + \frac{k}{k+1}$ and $n + \frac{2k}{k+1}$, respectively. Considering that n is large in DNNs, the additional overhead is negligible compared with the linear neurons.

IV. Experimental Results

This section presents a series of experiments that evaluate the efficacy of our proposed neuron when deployed in existing DNNs. In Sec. IV-A and IV-B, we assess the impact of our neuron on the accuracy as well as computational and parameter costs of DNNs in image classification tasks and WMT14 English-German translation task. In Sec. IV-C, we analyze the proposed quadratic neuron, including parameter distributions, and visualization of the neuron's response in CNNs.

A. Image Classification

Within this subsection, we assess the effectiveness of our proposed neuron by implementing it across all convolutional layers in a variety of CNNs. Our evaluation consists of analyzing the accuracy of the resulting CNNs in image classification tasks. In addition, we also conduct an analysis of the associated costs of parameters and computations.

1) Evaluation on Efficiency

The CIFAR-10 dataset [21] consists of $60k$ 32×32 images categorized into 10 classes. Prior to training, we applied a random crop to each image to 32×32 from 40×40 images with paddings, as well as a random horizontal flip as data augmentation. All networks were trained with SGD optimizer

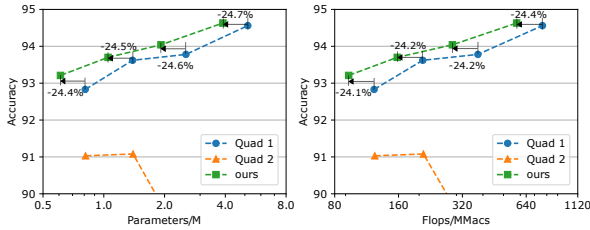


Fig. 5. Accuracy on CIFAR-10, number of parameters and FLOPs of ResNets (20, 32, 56 and 110) with quadratic neurons from [18] (Quad 1), [19] (Quad 2), and this work.

for 180 epochs with a training batch size of 128. The learning rate was initially set to 0.1 and was decayed by a factor of 0.1 at epochs 90 and 135. The learning rate of Λ^k in our proposed neuron was initially set to 10^{-4} and the rank of decomposition k was fixed to 9.

As shown in Fig 4, ResNets with our proposed neurons can achieve significant improvement in terms of accuracy and the computational and storage efficiency compared to the baseline networks. For example, ResNet-32 with the proposed quadratic neurons can achieve better accuracy than ResNet-44 with conventional linear neurons with 29.3% fewer parameters and 28.3% less computational cost. The improvement is even more significant for deeper networks: ResNet-56 with quadratic neurons can reduce the computational cost and storage cost approximately by 50% while having a similar accuracy with baseline ResNet-110.

We also compared our proposed neuron with previous quadratic neurons. To ensure a fair comparison, we expanded the networks equipped with our proposed neurons by adding channels to provide a slight accuracy advantage. This enabled us to compare different neurons based on their parameter and computation costs. We reproduced the results of [20] using the same configuration as ours, while the results of [18] were obtained from the original paper.

Fig 5 presents a comparison of the accuracy, parameter, and computational costs of networks with quadratic neurons from previous research [18], [20] and this work. Among these designs, our proposed neuron structure demonstrates the best efficiency in terms of storage and computations. Notably, as the network depth increases, the design from [20] performs worse on the test set, with the accuracy falling below 90%. Compared to [18], our proposed design is able to achieve significant better accuracy with at least 24.4% less storage cost and 24.1% less computational cost.

2) Training Stability

Kervolutional neuron (KNN) [13] is a non-linear design without any additional parameters. In this section, we demonstrate that, despite the advantages of KNN in the number of parameters, the proposed quadratic neuron can guarantee a more stable training process and results in better accuracy. We trained ResNet-18 with different configurations on the ILSVRC-2012 dataset [25]. The networks were trained for 100 epochs with a batch size of 256. The learning rate was set 0.1, and was decayed by 0.1 at epochs 30, 60, and 90. Regarding our proposed neuron, the learning rate of Λ^k was set to 10^{-5} , and the rank of decomposition k was fixed to 9. We equipped the first n convolutional layers with neurons from [13] (referred

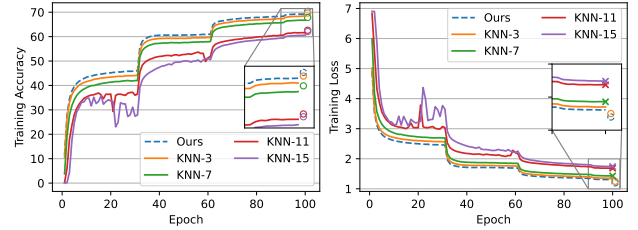


Fig. 6. Training stability analysis of ResNet-18 with KNN [13] and the proposed neurons on ImageNet. “KNN- n ” refers to the network equipped with neurons from [13] in first n layers.

TABLE II
PERFORMANCE ON ENGLISH-TO-GERMAN NEWSTEST2014 DATASET
AND PARAMETER COST OF TRANSFORMERS.

	Eval. Setting		Baseline	Quadratic		
	Tokenization	Cased		1E-4	1E-5	1E-6
BLEU ¹	13a	Yes	25.3	25.2	25.2	25.4
	13a	No	25.7	25.7	25.7	25.9
	International	Yes	26.0	25.9	25.9	26.1
	International	No	26.4	26.3	26.4	26.6
#Params/M			15.7	12.6 (-20.3%)		

¹ Higher is better.

to as “KNN- n ”), as well as deploying our proposed neurons to all convolutional layers for comparison.

The training accuracy and loss over time is shown in Fig. 6, where the circles represent the validation accuracy and loss after training, and a cross mark indicates that the loss does not converge on the validation set. As can be found in Fig. 6, when only deployed in the first 3 layers, [13] can achieve stable training and similar accuracy to our proposed design. However, as with more deployed layers such as 11 and 15, the training process becomes unstable as obvious fluctuation can be found, resulting in bad performance in the end. Also, extreme values can be found during the testing process, indicating the inference on unseen test data is also non-stable. In summary, compared to [13], our proposed neuron can achieve stable training with the highest accuracy while being applied to all layers. This brings superior flexibility for the deployment of our neuron.

B. Machine Translation with Transformers

In this section, we deploy the proposed quadratic neuron into the Transformer architecture, including all linear projection operators in the multi-head attention blocks. The experiments followed the same settings as in [3] unless otherwise stated. The network performance was evaluated on the English-to-German newstest2014 dataset and we used BLEU scores to evaluate the translation quality (higher BLEU score is better). Table II shows the BLEU scores in the translation task with different evaluation settings. Since the number of FLOPs averaged by the input length of a Transformer is approximately twice the number of parameters [26], we only report the number of network parameters in Table II. The last three columns present the results of networks with the proposed quadratic neurons with different learning rate configurations (10^{-4} , 10^{-5} and 10^{-6}) to quadratic parameters Λ^k . Compared with the baseline, the quadratic transformer can achieve better performance with significantly fewer (20.3%) parameters and FLOPs.

C. Analysis of the Quadratic Neuron

In this section, to comprehensively analyze the proposed quadratic neuron, experiments of the ResNet-20 network on

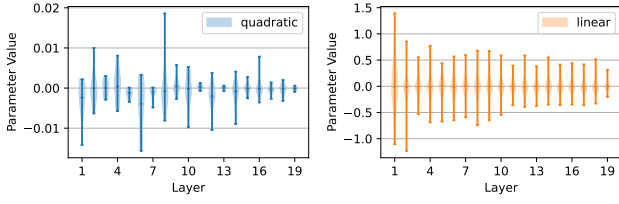


Fig. 7. Distribution of linear and quadratic parameters in ResNet-20 after training on CIFAR-100.

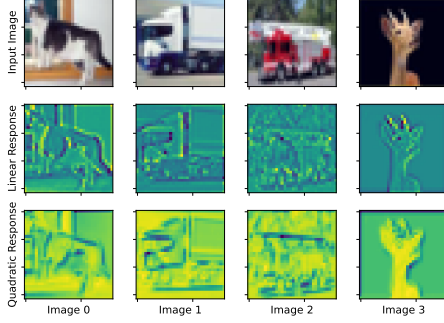


Fig. 8. Response visualization of linear and quadratic neurons.

the CIFAR-100 dataset [21] are conducted. The networks are trained with a batch size of 64 for 250 epochs.

1) Parameter Distribution

Fig. 7 presents the distribution of linear and quadratic parameters in different layers of ResNet-20 after the training process. It can be observed that, in contrast to linear parameters, the distribution range of quadratic parameters has a larger variance with respect to depth. Specifically, the values of quadratic parameters are more significant in layers 1, 6, and 8, while in some layers such as 11, 13, and 19, the quadratic parameters tend to be more centered around zero. This observation suggests that, unlike linear neurons, quadratic neurons may not be essential for every layer. However, our findings also suggest that the application of high-order neurons only in the first layer [13], [16] may not be the optimal choice.

2) Neuron Response Analysis

To gain further insights into the functionality of quadratic neurons in CNNs, we conducted a visualization experiment to examine their outputs. Fig. 8 illustrates the responses of both linear and quadratic parts of our neurons, where the linear and quadratic responses correspond to $\mathbf{w}^T \mathbf{x} + b$ and y_2^k , respectively. Our analysis indicates that, in contrast to linear parts that tend to extract edges, quadratic neurons focus more on the entire object. This property is particularly evident in the deer image, where the whole shape is clearly distinguished from the background. We also discover that quadratic neurons tend to extract low-frequency information such as the shape of an object, while high-frequency details such as texture or noise are less significant in the quadratic response.

V. Conclusion

In this paper, we introduce an efficient quadratic neuron for DNNs using a systematic approximation approach and the vectorized-output technique. Experimental results demonstrate its superior computational and storage efficiency, training

stability, and deploy-flexibility compared to prior non-linear neuron designs.

Acknowledgment

This work is supported by TUM International Graduate School of Science and Engineering (IGSSE), NSFC (Grant No. 62034007, 62141404), the Major Program of Zhejiang Provincial NSF D24F040002, and the SGC Cooperation Project (Grant No. M-0612).

References

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [2] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [3] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez *et al.*, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [4] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [5] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," *arXiv preprint arXiv:2204.06125*, 2022.
- [6] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [7] G. Yuan, P. Behnam, Y. Cai, A. Shafiee, J. Fu, Z. Liao *et al.*, "Tinyadc: Peripheral circuit-aware weight pruning framework for mixed-signal dnn accelerators," in *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2021, pp. 926–931.
- [8] R. Petri, G. L. Zhang, Y. Chen, U. Schlichtmann, and B. Li, "Powerpruning: Selecting weights and activations for power-efficient neural network acceleration," in *Design Automation Conference (DAC)*, 2023.
- [9] M. Jiang, J. Wang, A. Eldebiky, X. Yin, C. Zhuo, I.-C. Lin *et al.*, "Class-aware pruning for efficient neural networks," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2024.
- [10] M. Li and X. S. Hu, "A quantization framework for neural network adaption at the edge," in *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2021, pp. 402–407.
- [11] S. Lee, H. Sim, J. Choi, and J. Lee, "Successive log quantization for cost-efficient neural networks using stochastic computing," in *2019 56th ACM/IEEE Design Automation Conference (DAC)*, 2019, pp. 1–6.
- [12] W. Sun, G. L. Zhang, H. Gu, B. Lil, and U. Schlichtmann, "Class-based quantization for neural networks," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023.
- [13] C. Wang, J. Yang, L. Xie, and J. Yuan, "Kervolutional neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 31–40.
- [14] M. Mahmoudi, A. Chetouani, F. Boufera, and H. Tabia, "Taylor series kernelized layer for fine-grained recognition," in *2021 IEEE International Conference on Image Processing (ICIP)*, 2021, pp. 1914–1918.
- [15] P. Mantini and S. K. Shah, "Cqnn: Convolutional quadratic neural networks," in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 9819–9826.
- [16] G. Zoumpourlis, A. Doumanoglou, N. Vretos, and P. Daras, "Non-linear convolution filters for cnn-based learning," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 4761–4769.
- [17] Y. Jiang, F. Yang, H. Zhu, D. Zhou, and X. Zeng, "Nonlinear cnn: improving cnns with quadratic convolutions," *Neural Computing and Applications*, vol. 32, pp. 8507–8516, 2020.
- [18] F.-L. Fan, M. Li, F. Wang, R. Lai, and G. Wang, "Expressivity and trainability of quadratic networks," *arXiv preprint arXiv:2110.06081*, 2021.
- [19] Z. Xu, J. Xiong, F. Yu, and X. Chen, "Efficient neural network implementation with quadratic neuron," *arXiv preprint arXiv:2011.10813*, 2020.
- [20] Z. Xu, F. Yu, J. Xiong, and X. Chen, "Quadrallib: A performant quadratic neural network library for architecture optimization and design exploration," in *Proceedings of Machine Learning and Systems*, D. Marculescu, Y. Chi, and C. Wu, Eds., vol. 4, 2022, pp. 503–514.
- [21] A. Krizhevsky, "Learning multiple layers of features from tiny images," *Tech Report*, 2009. [Online]. Available: <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- [22] J. Bu and A. Karpatne, "Quadratic residual networks: A new class of neural networks for solving forward and inverse problems in physics involving pdes," in *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*. SIAM, 2021, pp. 675–683.
- [23] I. Jolliffe, *Principal Component Analysis*. John Wiley & Sons, Ltd, 2005, ch. 1, pp. 1–9. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/0470013192.bsa501>
- [24] C. Eckart and G. Young, "The approximation of one matrix by another of lower rank," *Psychometrika*, vol. 1, no. 3, pp. 211–218, 1936.
- [25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma *et al.*, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [26] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child *et al.*, "Scaling laws for neural language models," *CoRR*, vol. abs/2001.08361, 2020. [Online]. Available: <https://arxiv.org/abs/2001.08361>