

LUT-NN: Empower Efficient Neural Network Inference with Centroid Learning and Table Lookup

Xiaohu Tang*
State Key Lab of CAD&CG,
Zhejiang University
Microsoft Research
tigertang@zju.edu.cn

Yang Wang
Microsoft Research
Yang.Wang92@microsoft.com

Ting Cao†
Microsoft Research
ting.cao@microsoft.com

Li Lyna Zhang
Microsoft Research
lzhani@microsoft.com

Qi Chen
Microsoft Research
cheqi@microsoft.com

Deng Cai
State Key Lab of CAD&CG,
Zhejiang University
dengcai@cad.zju.edu.cn

Yunxin Liu
Institute for AI Industry Research
(AIR), Tsinghua University
liuyunxin@air.tsinghua.edu.cn

Mao Yang
Microsoft Research
maoyang@microsoft.com

ABSTRACT

On-device Deep Neural Network (DNN) inference consumes significant computing resources and development efforts. To alleviate that, we propose LUT-NN, the first system to empower inference by table lookup, to reduce inference cost. LUT-NN learns the typical features for each operator, named centroid, and precompute the results for these centroids to save in lookup tables. During inference, the results of the closest centroids with the inputs can be read directly from the table, as the approximated outputs without computations.

LUT-NN integrates two major novel techniques: (1) differentiable centroid learning through backpropagation, which adapts three levels of approximation to minimize the accuracy impact by centroids; (2) table lookup inference execution, which comprehensively considers different levels of

parallelism, memory access reduction, and dedicated hardware units for optimal performance. LUT-NN is evaluated on multiple real tasks, covering image and speech recognition, and nature language processing. Compared to related work, LUT-NN improves accuracy by 66% to 92%, achieving similar level with the original models. LUT-NN reduces the cost at all dimensions, including FLOPs ($\leq 16\times$), model size ($\leq 7\times$), latency ($\leq 6.8\times$), memory ($\leq 6.5\times$), and power ($\leq 41.7\%$).

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence.**

KEYWORDS

Neural Network, Product Quantization, Table Lookup, Centroid Learning

ACM Reference Format:

Xiaohu Tang, Yang Wang, Ting Cao, Li Lyna Zhang, Qi Chen, Deng Cai, Yunxin Liu, and Mao Yang. 2023. LUT-NN: Empower Efficient Neural Network Inference with Centroid Learning and Table Lookup. In *The 29th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '23)*, October 2–6, 2023, Madrid, Spain. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3570361.3613285>

1 INTRODUCTION

DNNs are widely used in mobile applications, offering users unparalleled intelligent services. However, DNNs are computation hungry workloads, mainly composed of linear computation operators, heavily stressing the limited hardware resources on mobile devices.

*Contribution during internship at Microsoft Research

†Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. *ACM MobiCom '23*, October 2–6, 2023, Madrid, Spain
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9990-6/23/10...\$15.00

<https://doi.org/10.1145/3570361.3613285>

As a result, huge efforts have been taken for efficient and affordable DNN inference on mobile devices, such as model compression [4, 19], sophisticated computing operator optimization [18, 41, 46, 55], tensor compilers for operator generation [9, 39], and customized DNN accelerators [34, 35, 57]. These methods require redesigning or re-implementing model structures, computation operators, or accelerators repeatedly for diverse deployment scenarios.

Different from these directions, this paper explores a new possibility to replace computation operators of DNNs, for reduced inference cost, as well as the tedious efforts of operator development. To achieve this, we first rethink the essence of inference. Each layer of a DNN model is to output another level of features given the input features. For example, the front layers of an image recognition model output low-level features (e.g., edges and lines), and subsequent layers output high-level features (e.g., faces and objects). The fact is that the features of different images for each layer have semantic similarity. A vivid example is that though cats and horses are different input images for a model, their ear features result in similar output for the layer that extracts it. Same for language tasks, similar words of different sentences could have similar embeddings and output results.

Based on this similarity, the question is that whether the typical features can be learned for each computation operator of DNNs, so that the output of these typical features can represent the output of diverse features. If so, by precomputing and saving the output of typical features, the output of future inputs can be read directly without computation.

Towards answering this question, we propose LUT-NN, a novel system to empower DNN inference by table lookup. The system learns the typical features, named *centroid*, for each linear computation operator, and precompute the results for these centroids to save in lookup tables. During inference, the results of the closest centroids with the input features can be read directly from the table as the approximated output of this operator. The required computation cost for inference is significantly reduced. Fig. 1 shows an overview of LUT-NN using real data samples of a model. LUT-NN integrates two major techniques: (1) differentiable centroid learning through backpropagation, and (2) table-lookup inference execution.

MADDNESS [5] initiates the centroid learning for matrix multiplication (MM), leveraging the technique of Product Quantization (PQ). PQ is an effective vector quantization algorithm widely used for dataset compression [24, 33]. It compresses the dataset by clustering the vectors in the dataset first, and then learning the centroids to represent vectors in each cluster. By using PQ, MADDNESS learns the centroids from the training dataset for a single-operator image classifier, and precomputes the MM results for the centroids, as the approximation for future input matrices.

However, as we will show in the paper, directly applying MADDNESS to the computation operators of a DNN model results in poor accuracy (80% accuracy drops for CIFAR-10). A similar conclusion is also from McCarter *et al.* [44]. We first analyze the issue and expose the main reason for the poor result is *the different optimization goal of centroid learning for PQ and DNNs*. The goal of PQ is to learn the centroids that can minimize the total error/distance between centroids and vectors, while the goal of DNNs is to minimize the final loss function. The two goals have no direct relationship. Without considering the loss function, errors introduced by centroids in MADDNESS are accumulated layer by layer, resulting in poor accuracy, making the resulting DNN model unreliable and undeployable in real-world applications.

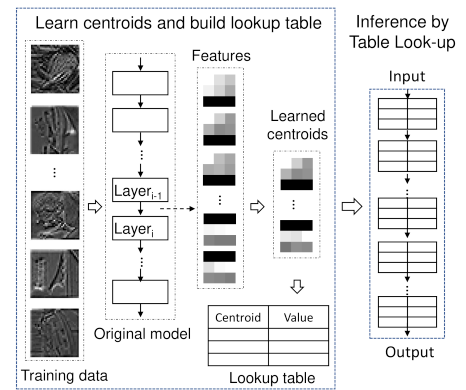


Figure 1: LUT-NN transforms model linear-computation layers to table lookup for inference.

Therefore, we identify the key factor for DNN centroid learning is to *pass the model loss to each operator through backpropagation*, and iteratively adjust the centroids by the gradients to minimize the model loss. However, the challenge is that currently it is not possible to learn centroids through backpropagation, since encoding a vector to the closest centroid e.g., argmin or hashing in MADDNESS, is not differentiable, and thus the gradients cannot be calculated.

To solve this challenge, we propose the novel **differentiable centroid learning** technique for DNN. Through backpropagation, it adapts three different levels of approximation introduced by centroids to model inference by three methods, and minimize the accuracy loss. (1) To enable gradient calculation and iteratively adjust centroids, LUT-NN uses *soft-PQ* method. It uses the continuous approximation of argmax , i.e., softmax , for backward pass. The forward pass still uses argmin for model loss calculation, since argmin will be used for inference. Soft-PQ can adapt the approximation introduced by centroids to reduce accuracy loss. (2) Since the use of softmax introduces approximation to argmin , LUT-NN uses *learned temperature* method, to learn the hyperparameter temperature of softmax for each operator, to tradeoff accuracy loss and learning convergence. (3) Since the lookup

tables for precomputed results are the main cost for LUT-NN, scalar quantization is used to reduce the table size. To adapt the approximation introduced by this scalar quantization, LUT-NN uses *quantization-aware training*, which uses quantized tables in the forward pass and real-value tables in the backward pass, to reduce accuracy loss. By the techniques above, the centroids can be learned, which empowers table lookup to replace the diverse DNN linear computations while maintain similar-level of model accuracy.

LUT-NN achieves a high FLOPs reduction (up to 16×). To run LUT-NN on resource-limited devices, we design and implement the **table lookup inference execution** on commodity ARM and x86 CPUs. The challenge is that table lookup is a new inference paradigm, and no direct support from current systems and accelerators. Our design comprehensively considers different levels of hardware parallelism, memory access reduction, and dedicated hardware units, to fully utilize hardware resources. Specifically, we align the size of codebook and lookup table to the width of SIMD instructions for data-level parallelism, interleave the instructions of *argmin* reduction for instruction-level parallelism, reside the codebook in an inner loop for cache locality, and leverage the shuffle instruction for fast table access.

We evaluate LUT-NN on a range of tasks, including image recognition tasks: CIFAR [37], GTSRB [50], SVHN [47], and ImageNet [12]; speech recognition task: Google Speech Command [56]; nature language processing (NLP) tasks: GLUE [53] and a regression task: UTKFace [60] age prediction, and cover both CNN and BERT models. For all these tasks, compared to directly applying MADDNESS, LUT-NN can improve the accuracy by 66% to 92%. The accuracy difference between LUT-NN and the original DNN models ranges from +1.98% (Speech Command) to -2.42% (ImageNet).

For inference efficiency, LUT-NN can outperform baselines in all dimensions. The selected baselines include both the state-of-the-art hand-written inference system ONNX Runtime [46], and also the well-tuned TVM [9] generated kernels as the performance upper-bound of linear computations. The detailed results include: FLOPs reduces 5.7× ~ 16×; model (or disk) size reduces 3.4× ~ 7×; end-to-end inference latency reduces 1.3× ~ 6.8×; memory usage reduces 1.4× ~ 6.5×; and peak power reduces 15% ~ 41.7%. The real speedup is less than the FLOPs reduction due to the unfriendly support for table lookup in hardware. A customized table-lookup unit or accelerator could unleash the full potential of LUT-NN.

LUT-NN takes the first trial towards this new table-lookup inference paradigm, which potentially much simplifies the inference system and hardware design and improves efficiency. To sum up, the key contributions of this paper are as follows. The code is open sourced at <https://github.com/lutnn>.

- LUT-NN is the first to empower DNN inference by table lookup to reduce inference cost.
- It can achieve similar level of model accuracy by the novel differentiable centroid learning technique for DNNs through backpropagation.
- It designs the new table lookup operation to enable this new inference paradigm on commodity CPUs.
- We implement the learning and inference pipelines. Superior results are shown on all evaluated dimensions.

2 BACKGROUND AND MOTIVATION

Each data sample in the training set can be assumed as a vector. This section will first introduce the concept of vector quantization and its efficient solution, PQ, and then show the poor results of directly applying PQ to DNN inference.

2.1 Product Quantization: Background

Quantization has been well studied in information theory [20]. It reduces the cardinality of a dataset by using *centroids* to represent the data samples. The set of centroids is called a *codebook*. For vector quantization, the dataset is composed of D -dimension vectors. The vector quantizer can encode each vector to a centroid in the codebook.

To reduce the cost of vector quantization, PQ is proposed. The essence of it is to decompose the high-dimensional vector space into the Cartesian product of sub-vector spaces and then quantize these sub-vector spaces separately. As shown in Fig. 2(a), it splits the D -dimension vector into C distinct V -dimension sub-vectors ($D = C \cdot V$). The sub-vectors are quantized separately using C codebooks. The quantization result of the vector is the concatenation of the C centroids.

Centroid learning We now formulize the PQ process. To quantize a D -dimension vector $a \in \mathbb{R}^D$, PQ needs to learn the centroids from a training dataset $\hat{A} \in \mathbb{R}^{\tilde{N} \times D}$ composed of vectors with the same distribution as a . PQ first decomposes the vectors in the dataset into C distinct V -dimension sub-vectors, notated as $\hat{A}^c \in \mathbb{R}^{\tilde{N} \times V}$ (marked in different colors in Fig. 2). To make optimal quantization, the centroid learning process is to find the K centroids P^c (i.e., the c^{th} codebook) for \hat{A}^c by k -means [40], which can minimize the distance sum of each sub-vector \hat{A}_i^c and its nearest centroid P_k^c , as shown in Eq. 1.

$$\arg \min_P \sum_c \sum_i \|\hat{A}_i^c - P_k^c\|^2 \quad (1)$$

Sub-vector encoding With the learned centroids P , for an input vector a , PQ can encode it as the concatenation of the nearest centroids for each sub-vector. The encoding function for a sub-vector is shown in Eq. 2. By this vector decomposition method, the centroids can represent K^C different vectors by only $K \times C$ memory cost.

$$g^c(a^c) = \arg \min_k \|a^c - P_k^c\|^2 \quad (2)$$

Hashing for acceleration with bigger error Learning centroids is a NP-hard problem. Vanilla PQ uses k -means to learn centroids and encode sub-vectors. k -means satisfies Lloyd optimal conditions [33, 40] and can get a local optimal quantization error. However, the k -means encoding costs high to compute the Euclidean distance for each sub-vector with each centroid, as shown in Eq. 2.

To reduce the encoding cost, some works propose hashing methods to encode sub-vectors [5, 26], but *at the cost of higher quantization error*. Hashing hashes a sub-vector to one of the K buckets. For example, MADDNESS selects a 4-level balanced binary regression tree of the hashing function family, with each leaf as a hash bucket. A sub-vector is encoded by traversing the tree from the root, and moving to left or right child if the value of certain indices is above or below a threshold.

This paper will evaluate both distance-based encoding and hashing-based encoding methods.

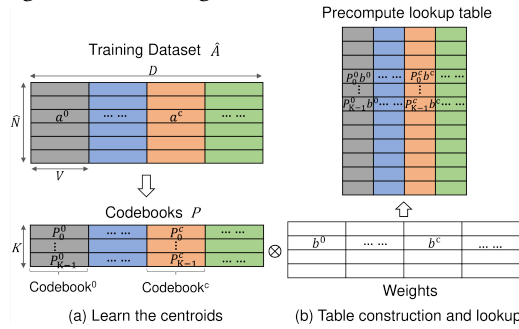


Figure 2: Product Quantization for AMM. Each color marks a sub-vector dataset.

2.2 PQ for AMM

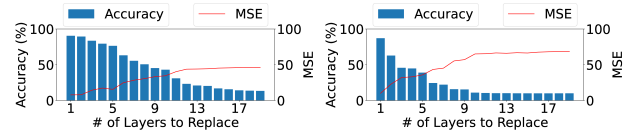
PQ can be used for approximated matrix multiplication (AMM) [5]. The essence is to approximate the matrix multiplication by the centroids' multiplication.

To formulize it, for matrix multiplication $A \times B^T$, a and b are the rows of A and B respectively. The centroid codebooks for A is P . For a layer of DNN, A can be the input feature maps, and B can be the weights (note that convolution can be computed as matrix multiplication too). Since B is constant, the multiplication of all the centroids and B can be precomputed to construct a lookup table, as shown in Fig. 2(b). The table construction function for b^c is shown in Eq. 3.

$$h^c(b^c) = [P_0^c \cdot b^c, P_1^c \cdot b^c, \dots, P_{K-1}^c \cdot b^c] \quad (3)$$

The matrix multiplication can then be approximated by looking up and aggregating the results of the nearest centroids in the precompute table, formulated in Eq. 4. Here considers the $g^c(a^c)$ function with *onehot* representation for *argmin*, i.e., the nearest centroid is marked as 1 and others as 0, $g^c(a^c) = \text{onehot}(\arg \min_k \|a^c - P_k^c\|^2) = (0, \dots, 0, 1, 0, \dots, 0)$.

$$a \cdot b = \sum_c \hat{a}^c b^c \approx \sum_c g^c(a^c) \cdot h^c(b^c) \quad (4)$$



(a) Vanilla PQ-based AMM

(b) MADDNESS AMM

Figure 3: Model accuracy keeps decreasing because MSE keeps increasing, while more MMs (from the last to the first) are replaced by PQ-based AMM, for ResNet20 on CIFAR-10. MADDNESS leads to more accuracy loss because using hashing for encoding.

2.3 Motivation: Poor results of PQ for DNN

Since DNN models are composed of MM, a direct thought is that can we replace MM in a DNN model by the PQ-based AMM. However, results show that directly applying it to a DNN model leads to poor accuracy.

Fig. 3 shows the accuracy of using PQ-based AMM for ResNet-20 on CIFAR-10, as well as the Mean Square Error (MSE) of the replaced model and the original model. We replace the MM from the last to the first layer by PQ-based AMM. Fig. 3a uses vanilla PQ with k -means for encoding, while Fig. 3b uses MADDNESS with hashing for encoding.

Results show the accuracy keeps dropping while more layers are replaced by AMM, because the error of the AMMs is accumulated. As expected, vanilla PQ shows better results than MADDNESS, since k -means introduces smaller quantization error than hashing. MADDNESS can maintain accuracy when only the last layer is replaced. This is consistent with the MADDNESS paper, which only replaces the last layer, i.e., fully-connect, by AMM. However, if we replace the last two layers, the accuracy sharply drops by 30%, and finally ends in 10% accuracy. Vanilla PQ drops by 30% when the last six layers are replaced, and also ends in 10% accuracy.

Reason for poor accuracy We expose the key reason for the poor accuracy of PQ-based AMM is that *the optimization goal of PQ and DNN learning is different*. As shown in Eq. 1, the goal of PQ is to minimize quantization error, i.e., learn the centroids to minimize the distance of each sub-vector and its nearest centroid. On the other hand, the learning goal of DNN is to minimize the final loss function, through backpropagation to iteratively adjust the model parameters of each layer. Without considering the loss function, when more layers use AMM, the approximation error gets accumulated, as shown in Fig. 3. To get better accuracy, it is necessary to learn the centroids by the DNN training process.

However, the challenge is that PQ centroid learning and encoding functions shown in Eqs. 1 and 2 are not differentiable, and cannot use backpropagation to calculate gradient. This paper thus proposes the soft-PQ technique, which empowers the centroid learning for a DNN by backpropagation and gradient descent.

3 DIFFERENTIABLE CENTROID LEARNING FOR DNN

As exposed above, the key factor to empower table lookup for DNNs is to learn the centroids of each layer through backpropagation to minimize the model loss. However, the main challenge is that the argmin function in vanilla PQ, according to Eq. 2, is not differentiable. We therefore propose the differentiable centroid learning for DNN. The learning integrates three methods to adapt three levels of approximation.

Firstly, we propose soft-PQ method, which leverages the continuous and differentiable approximation to argmax function, i.e., softmax [32], for backpropagation. We approximate the *onehot* centroid result as the weighted sum of all centroid results (Sec. 3.1). Prior work has used softmax for differentiable PQ, but only for input embedding compression[8]. Our work is the first to enable soft-PQ for end-to-end DNNs.

Secondly, the use of softmax introduces another level of approximation to using argmin , which could also lead to reduced model accuracy. We propose a learnable temperature method to address this challenge to adjust the approximation error of softmax with argmin for each layer. The temperature of each layer's softmax can be learned through backpropagation (Sec. 3.2).

Thirdly, to reduce memory cost of the lookup tables, LUT-NN applies scalar quantization on the lookup tables. However, this also introduces a level of approximation. Similarly, we propose quantization-aware training to adapt this approximation during centroid learning (Sec. 3.3).

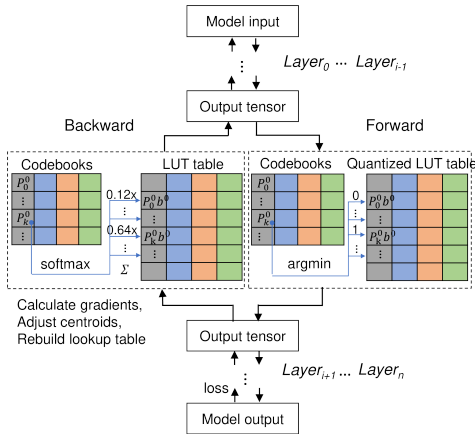


Figure 4: Soft PQ learns centroids. Forward pass uses argmin as the encoding function, and backward uses softmax . Each color marks a sub-vector dataset using one codebook. P_k^c is the k^{th} centroid for sub-vector a^c .

3.1 Backpropagation through soft-PQ

As discussed in Sec.2.1, vanilla PQ employs k -means to learn centroids from the dataset, and the encoding function $g^c(a^c)$ utilizes argmin to encode a sub-vector as the nearest centroid, represented by a onehot vector $(0, \dots, 1, \dots, 0)$ as shown in

Fig.4. The sub-vector AMM result can be read directly from the lookup table by $g^c(a^c) \cdot h^c(a^c)$.

However, to apply PQ to the entire DNN model and minimize the model loss, the centroids for each layer must be learned from backpropagation and gradient descent. We utilize the smooth argmax function, i.e., softmax , as the encoding function for backpropagation, as shown in Eq.5. Here, t represents the temperature hyperparameter, which will be further discussed in Sec. 3.2.

$$\tilde{g}^c(a^c) = \text{softmax}(-\|a^c - P_k^c\|^2 / t) \quad (5)$$

For a codebook with K centroids, the softmax function takes a vector of K distance results between the sub-vector a^c and each centroid P_k^c as input. It normalizes the input to a probability distribution that adds up to 1. According to the definition of softmax , each probability is proportional to the exponent of the distance, i.e., $\exp(-\|a^c - P_k^c\|^2 / t)$. The closer the centroid is to the sub-vector, the higher the probability will be. The encoding for a sub-vector is transformed from a deterministic onehot vector into a probability vector. For example, as illustrated in Fig. 4, the output of the softmax encoding function is $(0.12, \dots, 0.64, \dots)$, where 0.64 is the probability of the nearest centroid. The sub-vector AMM result is then obtained as the dot product of the probability vector and the lookup table entries.

Soft-PQ centroid learning: Using softmax , the centroid learning process of our soft-PQ for the entire model is illustrated in Fig.4. During the forward pass, the onehot argmin function is utilized to calculate the model output and loss, as model inference will also use argmin for simplicity. The backward pass utilizes softmax as the encoding function to calculate gradients, adjust centroids via gradient descent, and rebuild lookup tables with the updated centroids for the next training iteration. Based on Eq.4, the sub-vector AMM in soft-PQ is formulated as Eq. 6.

$$a^c b^c = \tilde{g}^c(a^c) \cdot h^c(b^c) - \text{sg}(\tilde{g}^c(a^c) \cdot h^c(b^c) - g^c(a^c) \cdot h^c(b^c)) \quad (6)$$

Here, sg represents the *stop gradient* operator. It serves as an identity function during the forward pass to enable the use of $g^c(a^c)$ encoding in argmin . During the backward pass, it drops gradients inside it to enable $\tilde{g}^c(a^c)$ to generate gradients via softmax .

The initial value is critical for learning convergence and accuracy. Therefore, we use the centroids learned by k -means from vanilla PQ to initialize the centroids and lookup tables.

3.2 Learned temperature

The temperature hyperparameter t of softmax [28] controls the approximation error of softmax to argmax . Shown in Fig. 5a, as $\text{softmax}(x)_i = \frac{\exp(x_i/t)}{\sum_{k=1}^K \exp(x_k/t)}$, when $t \rightarrow \infty$, $\text{softmax}(x)_i \rightarrow \frac{1}{K}$, i.e., the output probability distribution approaches uniform distribution. When $t \rightarrow 0$, $\text{softmax}(x) \rightarrow$

$\text{onehot}(\text{argmax}(x))$, i.e., the probability of the largest x approaching 1 and the others is 0.

Therefore, there is a tradeoff between small and large temperature. For small temperature, softmax is close to onehot argmax, but the training is difficult since the variance of gradients is large. For larger temperature, the approximation error is increased, but the variance of gradients is smaller.

Works before normally set t as a fixed value (mostly 1), or anneal it from a large number to a small one during training [28, 32], but never analyze how to set it reasonably. This is because currently for a DNN model, softmax is only used by the output layer to produce class probability, or used by the input layer to produce symbol embedding. The approximation error barely impacts model accuracy. However, our soft-PQ employs softmax in every layer. The accumulated error can decrease accuracy without proper t settings.

We thus propose to learn the temperature for each layer, also during backpropagation while centroid learning. Fig. 5b shows the learned t for each layer for ResNet18. The value for each layer is different, and thus not practical to tune by hand. According to the CIFAR10 accuracy experiments, training with the learned temperature technique only spends $\frac{1}{10}$ iterations of training with temperature setting to 1 to achieve 85% accuracy. Detailed results will show in Sec. 6.

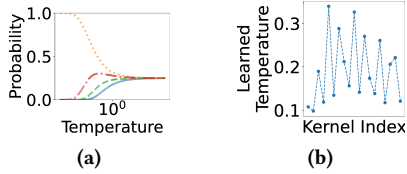


Figure 5: (a) The output probability distribution of $\text{softmax}(\frac{x}{t})$ at different temperature, for four centroids as an example. $t \rightarrow \infty$ approaches uniform distribution, and $t \rightarrow 0$ approaches argmax. (b) The learned t for each layer of ResNet18.

3.3 Scalar quantized lookup table

Lookup tables are the main disk and memory cost. We reduce the table size by scalar quantization (e.g., FP32 to INT8). We leverage the classic range-based linear quantization. The formula is $r = s(q - z)$ [29], in which r is the real value, s is the scaling factor, q is the quantized value, and z is the zero point. We use symmetric quantization, so z is forced to be 0, and the quantized range is $[-2^{n-1}, 2^{n-1} - 1]$. The scaling factor s is calculated as the max absolute value in the table divided by half of the range, i.e., $s = \frac{\max(|\text{value}|)}{2^{n-1} - 1}$.

Quantized lookup tables introduce another level of approximation. Similar to temperature, we thus quantize the tables during centroid learning, to minimize the loss function. Inspired by Jacob *et al.* [29], the backpropagation uses lookup tables in real values, so that they can be adjusted in small amounts. The forward pass uses quantized lookup tables as in the inference to calculate the loss. Results show

that by this learning method, the quantized lookup table has little impact on the model accuracy.

4 COST REDUCTION OF LUT-NN

Before introducing the inference design of LUT-NN, this section analyzes the theoretical FLOPs and model size reduction by LUT-NN. The two primary factors of LUT-NN are the number of centroids and the length of sub-vectors, representing a tradeoff between cost and accuracy. We conduct the analysis with these two factors.

According to the output formula Eq. 4 of a PQ-based AMM in LUT-NN, the main cost is from the encoding function $g^c(a^c)$, which calculates the Euclidean distance of the sub-vector with each centroid. After that, the cost is from table lookup with the encoding result (i.e., index of the closest centroid), and the result aggregation of sub-vectors. For file size, the major cost is from the lookup tables, which saves the dot product result of each centroid and the according sub-vectors in the weight matrix. The size of codebooks is relatively small, since the sub-vectors on the same column share one codebook.

Therefore, we analyze the FLOPs of encoding, table lookup and aggregation, as well as the size of lookup tables as the cost of a LUT-NN AMM, to compare it with normal MM in Table 1. Since convolution can be transformed to MM by im2col, its cost also follows these formulas. For a convolution, M is the number of output channels, D is the number of input channels \times filter size², and N is height \times width.

The number of centroids K and the sub-vector length V are two hyperparameters of LUT-NN. They are tradeoffs between accuracy and cost (refer to Sec. 6.3). The more centroids K and shorter sub-vector V may lead to higher accuracy, but will increase the cost of LUT-NN. Table 2 shows the calculated GFLOPs and model size by the formulas in Table 1 for different models, using typical (K, V) settings in LUT-NN. These typical settings can achieve comparable accuracy with the original model, and also align with the SIMD width for high performance. Similar to other hyperparameters in DNN training, K and V can be set by grid search, evolutionary search, or other popular searching methods.

It is clear that LUT-NN can achieve both computation and model size saving. The FLOPs saving is because K is normally smaller than M . For example, the number of output channels i.e., M , for ResNet50 is normally 128, 256, or 512, so the FLOPs can be reduced by 4 \times when $K = 8$. For ResNet20, the number of output channel is 16, 32, and 64, so the FLOPs is reduced by 2 \times when $K = 8$.

5 TABLE LOOKUP INFERENCE DESIGN

In this section, we present the design and optimization of the inference system that supports LUT-NN. Fig. 6 depicts the

$A \in \mathbb{R}^{N \times D}$: input matrix $B \in \mathbb{R}^{D \times M}$: weight matrix V : the length of a sub-vector a^c K : the number of centroids in a codebook for a^c		
	Ours	MM
FLOPs	$N \cdot D \cdot K + N \cdot M \cdot D/V$	$N \cdot D \cdot M$
Size	$4 \cdot D \cdot K + D \cdot M \cdot K/V$	$4 \cdot D \cdot M$

Table 1: FLOPs and disk size of a LUT-NN AMM compared to MM (FP32).

Model	GFLOPs		
	original	(8, 9)	(16, 9)
ResNet18 (CIFAR10)	0.555	0.098	0.132
SENet18 (CIFAR10)	0.556	0.098	0.131
VGG11 (CIFAR10)	0.606	0.085	0.102
ResNet18	1.814	0.412	0.515
SENet18	1.814	0.412	0.515
VGG11	7.609	1.132	1.334
	original	(16, 32)	(16, 16)
BERT	2.759	0.169	0.254

Model	Disk Size (MB)		
	original	(8, 9)	(16, 9)
ResNet18 (CIFAR10)	42.59	12.55	23.13
SENet18 (CIFAR10)	42.92	10.94	21.53
VGG11 (CIFAR10)	35.18	8.45	16.88
ResNet18	44.55	12.57	23.16
SENet18	44.88	12.91	23.49
VGG11	37.12	10.39	18.82
	original	(16, 32)	(16, 16)
BERT	162.26	23.05	43.30

Table 2: Theoretical GFLOPs and model size of typical (K, V), calculated by formulas in Table 1.

LUT-NN model inference architecture, comprising the Closest Centroid Search stage and the Table Read and Accumulation stage. In the Closest Centroid Search stage, LUT-NN first computes the distance between input tensors and centroids and searches the nearest centroids for each input tensor. In the Table Read and Accumulation stage, LUT-NN reads the pre-computed results from the lookup table and accumulate results. LUT-NN inference design fully takes advantage of the features of current hardware architectures to optimize the design and implementation of Table Lookup. These optimizations improve inference by utilizing memory hierarchy, utilizing shuffle instructions, and utilizing mixed-precision accumulation instructions.

5.1 Closest centroid search

The centroid search takes the most computation costs in LUT-NN, and efficient centroid search is critical for performance efficiency compared to conventional computational methods. The **Closest Centroid Search** stage first calculates the distance between input tensors and centroids. We can represent this calculation as matrix multiplications of

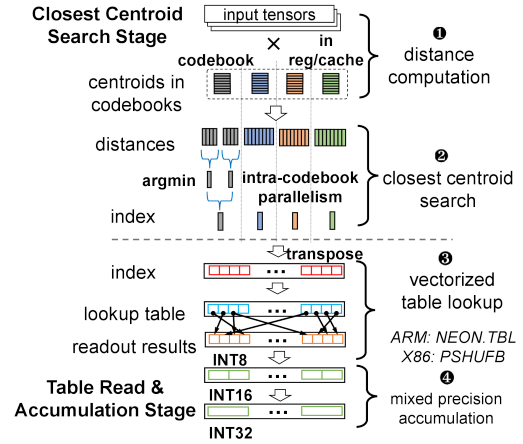


Figure 6: Table Lookup Inference Design

input tensors and codebook matrices. Then, it searches for the nearest centroids for each input sub-vector.

However, the design of centroid search in LUT-NN presents challenges that involve leveraging the features of existing hardware architectures. Firstly, LUT-NN distance computation is irregular-shaped (tall-and-skinny) and difficult to optimize using BLAS libraries. For example, the XNNPACK library achieves only 23.0 GFLOP/s on Pixel 6 for the second layer of LUT-NN based ResNet18, which accounts for only 25.7% of peak performance. The height of the input tensor (N) is usually much larger than the length of the sub-vector ($N \gg V$) and the number of centroids ($N \gg K$) on each codebook. Therefore, the operation intensity of the distance calculation can be approximated by $\frac{2NVK}{NV+KV+NK} \approx \frac{2}{1/K+1/V}$ FLOP/Byte [55]. Since the length of the sub-vector and the number of centroids are small in LUT-NN, the operation intensity $\frac{2}{1/K+1/V}$ is also small. The distance computation becomes a memory-intensive matrix multiplication.

Therefore, we focus primarily on optimizing memory access for centroid distance computations (1 in Fig. 6) to leverage memory hierarchy in current hardware. To optimize memory access for centroid distance computations (1 in Fig. 6), we reduce memory access overhead by keeping frequently accessed data in registers and caches as much as possible. So we design a *centroid-stationary computation* scheme to reside centroid matrices in registers and reorder centroid matrix loads in the inner loop to keep them in the cache as long as possible. The centroid-stationary computation keeps K centroids in the cache for each codebook, which only requires reading these centroids once from DRAM. Consequently, it only requires reading an $N \cdot V$ input tensor from DRAM once, reducing memory bandwidth costs and improving performance.

Second, after computing centroid distances, the Closest Centroid Search stage must identify the nearest centroid for each input sub-vector and generate the centroid index (2 in Fig. 6). It can be represented by an `argmin` function,

which finds an index with the shortest distance. However, searching for the nearest centroid for each input sub-vector is a data-dependent operation. To find the closest one, we must compare each distance sequentially, which is RAW (Read After Write) dependent and hard to be parallelized on CPUs.

We propose *intra-codebook parallelism* to optimize the Closest Centroid Search stage. Intra-codebook parallelism searches the nearest centroid for the input sub-vector on a codebook in parallel. We slice a codebook into multiple sub-codebooks and compare each distance between sub-vector and centroids in each sub-codebook. Then, we interleave each sub-codebook's execution, and avoid data dependency in closest centroid search. We merge the compared distances by reduction and find the index corresponding to the closest centroid. It leverages instruction-level parallelism to improve hardware utilization and performance.

5.2 Table read and accumulation

LUT-NN obtains the indices of the nearest centroids for the input sub-vectors after closest centroid search, and leverages **Table Read and Accumulation** stage to compute the final results. This stage first reads the pre-computed results from the corresponding lookup table through indices (ⓐ in Fig. 6) and completes the computation by accumulation operation (ⓑ in Fig. 6). For example, convolution operators directly read out filter's outputs from lookup tables and accumulate each input channel's result for output channels.

However, Table Read and Accumulation stage introduces additional overhead in inference. We enhance inference efficiency by skillfully utilizing widely available and supported instructions in commercial hardware. First, table read is difficult to parallelize and introduces additional indirect memory accesses, which exaggerate the memory overhead on lookup tables. Since we quantize lookup table into INT8 in Sec. 3.3, we leverage widely supported SIMD *shuffle* instructions (x86: pshfb and ARM: tbnl) to achieve parallel and efficient table read. We demonstrate the implementation of table read using shuffle instructions in Fig. 6. The shuffle instruction [6] permutes each byte of a vector based on an index vector and stores the shuffled bytes in the result vector register in each clock cycle. On 128-bit wide SIMD, a vectorized table read instruction handles 16 sub-vector lookups ($128/8 = 16$) on 16 results ($128/8 = 16$) simultaneously, greatly simplifying table read and reducing overheads.

Second, the accumulation still has comparable computation costs to the entire process. For example, when a codebook handles N index lookups on a $K \cdot M$ lookup table, it costs $N \cdot M$ table reads ($K = 16$) and $N \cdot M$ accumulation adds. Therefore, the vectorized lookup table leaves accumulation operations as the performance bottleneck of table reads. The number of parallel processing units within a SIMD instruction is called SIMD lanes. We typically set the number of

centroids to 16 ($K=16$) to maximize the utilization of all SIMD instruction lanes. Since a higher number of SIMD lanes leads to higher throughput on the same width of SIMD instruction (e.g., the vectorized INT16 add instruction has the half throughput of INT8 on a 128-bit SIMD), we maximize accumulation throughput by mixed precision accumulation. It first accumulates results in INT16 to utilize more SIMD lanes and then gathers INT16 to INT32 to avoid overflow.

5.3 Utilization of memory hierarchy

As LUT-NN is composed of many memory accesses, we explain more details for its memory hierarchy usage.

The frequent-accessed codebook (for centroids) is kept in the inner loop during distance calculation, so that the hardware cache can leverage the locality to keep the codebook in cache. The size of a codebook for one layer is normally small, in KB level (calculated by sub-vector length \times number of centroids \times number of codebooks \times number of bytes). For example, the largest codebook in our tested models is 288 KB. It can be held in L2 cache on many mobile devices.

The lookup table (for precomputed results) is larger, in MB level. It can be held in L3 cache. Even if it is larger than L3 cache, the memory access overhead could still be hidden in our implementation. LUT-NN uses Shuffle instructions to parallelize table lookup, turning the pattern of reading index numbers and lookup tables from random to sequential. The memory access is thus predictable and easy to prefetch. Furthermore, we interleave computation and memory access instructions within the kernel, which can overlap computation and memory access time.

6 EVALUATION

6.1 Experiment methodologies and settings

Dataset and Models Our experiment tests LUT-NN's effectiveness on image recognition, speech recognition, and NLP tasks. For image recognition and speech recognition, we evaluate VGG, ResNet, and SENet family models on different datasets, including CIFAR-10 [37], GTSRB (German Traffic Sign Recognition Benchmark) [50], Google Speech Command [56], SVHN (Street View House Numbers) [47], UTKFace [60], and the large-scale complex ImageNet [12] dataset. The CIFAR-10 dataset consists of 60K images in 10 classes, with 50K for training and 10K for validation. The GTSRB dataset has 43 classes of traffic signs, with 39K training images and 12K test images. Google Speech Command is an audio dataset containing 65K recordings of 30 words, designed for keyword spotting on edge devices. SVHN is a dataset including 600K real-world images of house numbers for 10 classes. UTKFace dataset consists of 20K face images labeled by age, gender, and ethnicity. We use the age prediction task to test the regression ability of LUT-NN. We will report Mean Average Error (MAE) for this task. ImageNet dataset has 1.28M images for training and 50K for validation.

Dataset	Model	Centroid Learning Rate	Temperature Learning Rate	Weight Decay	Batch Size	# of Epochs	Optimizer	LR Scheduler
CIFAR10	ResNet18 SENet18 VGG11	1×10^{-3}	1×10^{-1}	0	256	200	Adam	Cosine Annealing
GRSRB								
Speech Commands		200						
UTKFace		150						
SVHN		1×10^{-3}						
ImageNet								
GLUE	BERT	$\{5, 4, 3, 2\} \times 10^{-5}$		1×10^{-2}	32	3	AdamW	Constant

Table 3: Soft-PQ training settings for centroid and temperature learning. We choose the best fine-tuning learning rate among $\{5 \times 10^{-5}, 4 \times 10^{-5}, 3 \times 10^{-5}, 2 \times 10^{-5}\}$ based on the accuracy of the validation set for BERT, following [15].

Model	ResNet18			SENet18			VGG11		
Dataset	LUT-NN	MADNESS	baseline	LUT-NN	MADNESS	baseline	LUT-NN	MADNESS	baseline
CIFAR10	94.40	10.01	95.26	94.48	10.65	95.47	93.89	22.87	95.04
GTSRB	98.73	4.53	98.80	98.36	5.68	98.84	98.55	5.70	99.22
Speech Commands	93.70	1.49	91.72	93.04	1.49	94.36	93.38	1.49	93.11
SVHN	96.00	20.68	96.67	96.22	20.12	96.60	96.23	29.97	96.62
UTKFace	4.91	10.51	5.57	4.74	11.02	5.46	5.69	24.57	5.85
ImageNet	67.38	0.10	69.76	68.21	0.17	70.63	68.04	0.16	68.33

Table 4: LUT-NN achieves comparable accuracy with the original models, significantly outperforming MADNESS. UTKFace is the lower the better. Others are the higher the better.

Dataset Task	Single Sentence	Similarity and Paraphrase	Natural Language Inference		Average
	SST-2	QQP	QNLI	RTE	
Training Dataset Size	67k	364k	105k	2.5k	
Test Dataset Size	1.8k	391k	5.4k	3k	
BERT base (%)	93.5	71.2	90.5	66.4	80.4
LUT-NN (%)	92.4	69.6	87.4	64.7	78.5

Table 5: LUT-NN achieves comparable accuracy to BERT-base on NLP tasks selected from GLUE.

On the large ImageNet dataset, ResNet18 and SENet18 models use 7×7 convolution for the first layer and 3×3 max-pooling, while other datasets use 3×3 convolution for the first layer and no max-pooling. The final three dense layers of the original VGG models are replaced by an average pooling layer and a dense layer to reduce model size for mobile deployment, following the general practice for VGG deployment [55]. Except for ImageNet, the first max pooling layer is removed for VGG. In evaluation results, the models for ImageNet use the model names, and the models for other datasets are marked as "Model (CIFAR10)" to differentiate.

To evaluate LUT-NN for NLP tasks, we used BERT models and the tasks from GLUE benchmark [53].

Soft-PQ training setting Table 3 lists the detailed training hyperparameter settings. Except for the initial learning rate, we followed standard practices for all other training procedures [15, 27]. As the learned temperature requires a larger learning rate to converge quickly to a sub-optimal solution, we use different learning rates for the centroid and temperature learning. Prior to soft-PQ training, we initialize centroids using k -means clustering. Specifically, we apply

the original model inference to a randomly sampled sub-dataset (consisting of 1024 training samples) and collect each operator’s inputs. We then utilize the k -means algorithm to cluster each operator’s inputs and obtain the initial centroids. **LUT-NN system settings** In order to balance accuracy and cost, we can specify various settings for LUT-NN, including the number of centroids in a codebook (K), the length of the sub-vector (V) (as shown in Table 1), and the operators to be replaced by table lookup. The default setting for K and V is to align with both the feature size and the length of SIMD instructions. Specifically, we set $(K, V) = (16, 9)$ for 3×3 convolution, $(K, V) = (16, 4)$ for 1×1 convolution. We will also evaluate other settings in Sec. 6.3 ablation study.

For operators to replace, we replace all convolution operators for CNN models by table lookup except the first one. This is because as explained by Zhou *et al.* [61], the layer interacted with the model input is very sensitive to accuracy. We also observe this phenomenon. Replacing the first layer by table-lookup can lead to obvious accuracy drop. This issue is even more serious for BERT. Replacing the first two layers results in 80% accuracy loss (refer to Fig. 13). Therefore, we

do not replace these layers by table lookup. For BERT, other than explicitly stated, we replace the fully connected operators of the last 6 layers. The results with different replaced layers for BERT will be shown in Section 6.3 ablation study.

Evaluation platforms LUT-NN is implemented on ARM Neon and Intel SIMD ISA (Instruction Set Architecture), and can run in a single or multi threads. We evaluate LUT-NN on two mobile devices Google Pixel 4 and 6, which are equipped with Cortex-A76 (2.42 GHz) and Cortex-X1 (2.80 GHz), respectively. In addition, we evaluate LUT-NN on a desktop CPU (Intel Core i7-4790) running at 3.60 GHz and a server CPU (Xeon Silver 4210) running at 2.20 GHz.

Comparison baselines The accuracy baselines are (1) original models and (2) MADDNESS [5], respectively. The inference performance baselines are ONNX Runtime (ORT) v1.12.1 [46] and TVM v0.9.0 [9], as the SOTA hand-written performance, and the auto-tuned performance upper bound for linear computation operators. We tune the TVM baselines using AutoScheduler and AutoTVM in 1500 iterations.

6.2 Accuracy, latency, memory and power

Accuracy Table 4 and 5 summarize the accuracy comparison. Remarkably, to compare with directly applying MADDNESS, LUT-NN can improve accuracy by 66% to 92%, achieving similar-level of accuracy with the original models. LUT-NN behaves particularly well on GTSRB ($\leq 0.67\%$ difference), SVHN ($\leq 0.67\%$ difference), CIFAR-10 ($\leq 1.15\%$ difference), Speech Command ($\leq 1.32\%$ difference), and UTKFace (LUT-NN is even better). For the Speech Command task, LUT-NN can even improve accuracy by 1.98% for ResNet18. Also for the UTKFace dataset, LUT-NN can even reduce MAE by at most 13.2%. Expectedly, the accuracy difference for the more complex tasks, ImageNet and NLP, is a bit larger compared to the tasks above, with 0.9%~2.42% difference for ImageNet, and an averaged 1.9% difference for NLP tasks.

Operator speedup Fig. 7 shows the operator speedup selected from three models. The operator speedup for VGG11 ranges from 4.28 \times to 5.43 \times on ARM CPU, and from 3.82 \times to 3.86 \times on x86 CPU. The best speedup on BERT kernels are 12.5 \times (ARM CPU) and 10.3 \times (x86 CPU).

BERT operators achieve higher speedups due to larger $M = 768$ or 3072 , and longer sub-vector $V = 32$. According to Table 1, the FLOPs of our proposed LUT-NN is $ND(K + M/V)$, and the FLOPs of a normal MM is NDM . The FLOPs reduction by LUT-NN can be calculated by $M/(K + M/V)$. Therefore, the reduction is particularly large when $M \gg K$ and $V \gg 1$. The kernel speedup gradually improves as the layer index increases for CNN models. The reason is that the number of output channels increases from 64 to 128, 256, and 512 as the layer index increases. The increased number of output channels M contributes to the speedup in these layers.

Model speedup Fig. 8 shows the end-to-end model latency. LUT-NN achieves 1.30~4.23 \times speedup on ResNet18, SENet18

and VGG11 compared to TVM and ORT. The best speedup is 4.23 \times in VGG11 (for CIFAR10) on Pixel6 ARM CPU over ONNX Runtime. The BERT model has higher throughput, and the speedups are 5.6 \times and 6.8 \times on the ARM CPU and the x86 CPU, respectively. Compared to CNN models ($M \leq 512$), the BERT model has a larger input tensor ($M = 768, 3072$) to acquire better performance gains. It also has a longer sub-vector length ($V = 32$) compared to CNN models ($V \leq 9$).

Multi-thread speedup To evaluate the scaling capability of LUT-NN, Fig. 9 shows the speedup with the thread increases, i.e., one to four on x86 CPUs and one to two on the Pixel 6 ARM CPU. Note that we evaluate the scalability on the big cores (cores with the highest frequency) of ARM CPUs, to avoid the interference of asymmetric big and little cores. Pixel 6 has two big cores. Since Pixel 4 only has one big core, we skip it in this evaluation. The results show that LUT-NN achieves better scalability compared to the baselines. It can reach 2.24~2.49 \times speedup as the number of threads increases to 4. Meanwhile, LUT-NN achieves a 2.17~2.82 \times speedup over other baselines under the same number of threads.

Model memory footprint Fig. 10 compares the memory footprint of LUT-NN with baselines. The results show that LUT-NN can save 1.43~2.83 \times memory for CNN models and 4.76~6.49 \times for BERT. For these CNN models, the sub-vector length is 9 or 4 for different operators, while the length is set to 16 for BERT. Because longer sub-vectors result in a higher compression ratio on feature maps, BERT achieves lower memory costs than other models.

Power We evaluate the power consumption of LUT-NN on Pixel 4. The experiment results in Table 6 demonstrate that LUT-NN reduces power consumption by 15%~41.7% compared to TVM, which reveal that LUT-NN is an energy-efficient algorithm and system design.

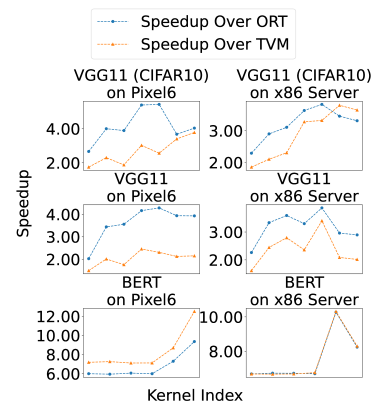


Figure 7: LUT-NN operator speedup over ORT and TVM on Pixel 6 and the x86 server CPU.

6.3 Ablation Study

We evaluate the effectiveness of the learned temperature and hyperparameters in the ablation study. The hyperparameters

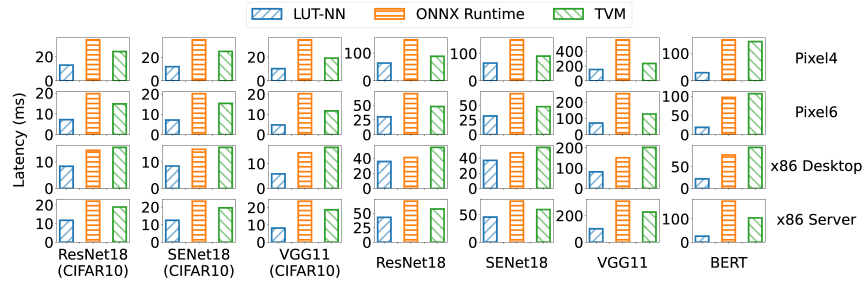


Figure 8: End-to-end latency comparison between LUT-NN, ORT, and TVM (lower is better).

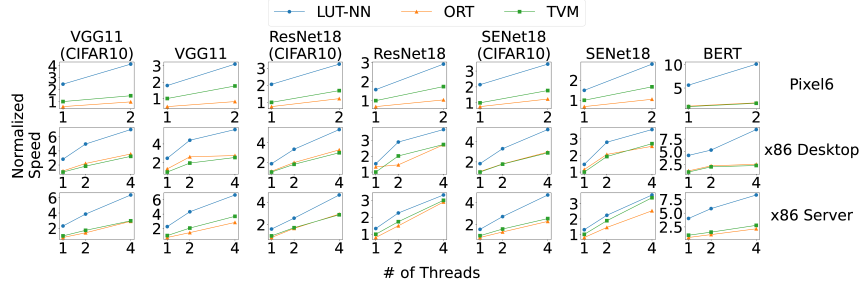


Figure 9: Normalized multithread speedup (over TVM one thread) of LUT-NN, ORT, and TVM (higher is better).

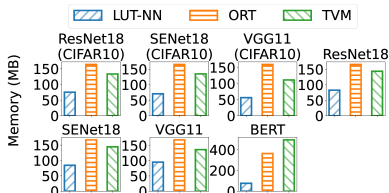


Figure 10: Model memory consumption of LUT-NN, ORT, and TVM on the x86 server CPU.

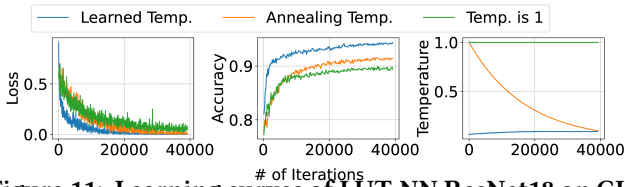


Figure 11: Learning curves of LUT-NN ResNet18 on CIFAR10 with temperature tuning methods. “Annealing Temperature” (Orange) refers to manually annealing temperature from 1 to 1e-1. Our learned temperature technique reaches 94.4% accuracy, higher than the annealing temperature (91.55%) and setting temperature as 1 (89.85%).

include the number of centroids, the length of sub-vector, and the number of replaced linear computation operators. **Learned temperature** In Sec. 3.2, we use gradient descent to learn the temperature. To evaluate the effectiveness of our method, we compare three temperature tuning strategies together: learned temperature, statically setting the temperature as 1, and annealing temperature from 1 to 0.1. We compare the training curves in Fig. 11. The figure suggests that our proposed learned temperature reaches the highest accuracy (94.4%) and outperforms the statically setting

(89.85%) and the annealing temperature (91.55%). In addition, the learned temperature has a faster convergence speed.

Impact of centroid number and sub-vector length The centroid number and sub-vector length do affect not only the inference throughput but also the model accuracy. We present the ablation study for the impact of these two hyperparameters. We follow the same experiment setting as Sec. 6.1 on ResNet18.

Figs. 12 collect accuracy and FLOPs for the variant of centroid numbers and sub-vector lengths in ResNet18. For this model, the sub-vector length significantly affects the model accuracy and worsens as the sub-vector length increases. Each codebook has to handle higher dimensions as the sub-vector length grows, which cannot be accurately classified and harms the model accuracy.

The centroid number also affects accuracy and performance. As the centroid number increases, each codebook can classify sub-vector in more fine-grained granularity, and it improves the accuracy of ResNet18. To balance the model accuracy with performance, we set $K = 16$ and $V = 9$ for better accuracy and fewer GFLOPs than the vanilla ResNet18.

The optimal hyperparameter setting for DNNs is a long-standing research topic. Some neural network search techniques, such as weight sharing [7] and search space search [48], could apply to LUT-NN to reduce the searching cost.

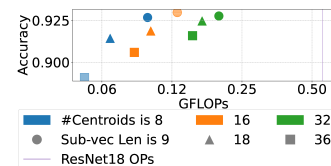


Figure 12: The scaling of centroid number and sub-vector length on ResNet18. x-axis is logarithmic.

The number of replaced operators As stated, to replace the front layers of BERT by table lookup can greatly impact accuracy, and we replace the last six layers in our default settings. Fig. 13 shows the accuracy impact by replacing more layers. We gradually replace the operators of more layers, from the last to the front layer, for BERT on the Semantic Textual Similarity Benchmark (STS-B) task. We can see the accuracy drops obviously for the front three layers.

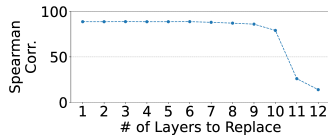


Figure 13: The accuracy of LUT-NN based BERT with respect to the number of layers to replace.

Speedup breakdown In Sec. 5, we propose four inference optimization techniques, which include: ❶ memory optimization on centroid distance computations, ❷ intra-codebook parallelism, ❸ shuffle instructions table lookup, and ❹ mix-precision accumulation in Fig. 6. To better understand performance gains, we evaluate the execution of a convolution operator ($C_{in} = C_{out} = 64, k = 3, s = 1, H = W = 56$, the second layer of ResNet18) on Pixel 6 when different techniques are enabled. The results show that ❸ (44.6%) saves the most execution time, followed by methods ❶ (18.5%) and ❷ (16.4%). Finally, method ❹ delivers a minor improvement (4.1%). The results show that our inference design and optimization can take advantage of ARM/x86 ISA’s special instructions, which outperform conventional operators.

Scalar quantization level for lookup table As explained in Sec. 3.3, LUT-NN leverages scalar quantization to reduce the table size. In this subsection, we evaluate the accuracy impact of different quantization levels by using ResNet18 for the CIFAR10 dataset as an example. The accuracy of LUT-NN is 94.44% with FP32 lookup table, 94.40% with INT8 lookup table, and 94.27% with INT4 lookup table. The results show that our quantization-aware training can achieve similar accuracy by using lower bits compared to FP32. We select INT8 in our current implementation as its the tradeoff of accuracy and speedup. INT4, on the other hand, is not supported by the hardware SIMD units directly. The exploration of lower-bit lookup table could be a potential future work.

Model	LUT-NN v.s. TVM Avg. power (W)
BERT	2.6/3.7
ResNet18	2.6/3.0
ResNet18 (CIFAR)	2.6/3.3
SENET18	2.6/2.9
SENET18 (CIFAR)	2.8/3.2
VGG11	2.3/2.9
VGG11 (CIFAR)	2.7/3.3

Table 6: Power (deducted SoC idle power) on Pixel 4.

7 RELATED WORK

Approximated matrix multiplication Traditional approximated matrix multiplication techniques focus on minimizing the difference between the ground truth and the approximated output. For example, Magen et al. [42] use random projection to map two matrices to a lower-dimensional subspace, and then perform accurate multiplication on the mapped matrices. MADDNESS [5] uses product quantization to learn the most typical features (also known as centroids) in the input distribution. MADDNESS then stores the products of centroids and the weight matrix in advance in a lookup table, so that the product of any input matrix and weight matrix can be approximated by the result stored in the table. To conclude, traditional methods will suffer from severe accuracy loss when used in neural network inference if the approximation is too harsh. Our method can avoid this defect by integrating itself into the end-to-end training process.

Also, previously, some attempts have been made to approximate the weights of NN with low-rank tensor expansion [13, 14, 30]. Specifically, LCNN [2] is also a kind of low-rank approximation. With LCNN, a dictionary is trained and the NN weights can be represented by linear combinations of the dictionary items. Afterward, LCNN transforms a vanilla convolution into lookup operations in the dictionary. Our work is different from LCNN in that our work learns centroids for each layer’s input feature map but not the weight.

Product quantization for DNN Product Quantization [43] is a popular and successful method for large-scale approximated nearest neighbor search [33]. In recent years, researchers have integrated product quantization into neural networks for various purposes. For example, [31, 36, 58, 59] incorporated product quantization as a layer in a convolutional neural network to obtain a compact and discriminative image representation. [8] utilized Product Quantization to compress the embedding layer in NLP models. [10, 17, 51] compressed the weight matrix for neural networks with Product Quantization. However, our approach differs from the above methods that our purpose is to achieve end-to-end neural network inference acceleration.

Scalar quantization Scalar quantization methods aim to reduce scalar bits in neural networks, e.g., from float representation to INT8 representation. They can both compress the neural network bits and accelerate the inference. Jacob et al. [29] proposed an INT8 quantization method with both a training scheme and an efficient kernel implementation. Other researchers [1, 3, 11, 21, 45, 62] further proved that 1/2/4-bit quantization is sufficient for image classification. More recently, Wang et al. [54] proposes a new approach to quantization using trainable lookup tables which can fit the distributions in different layers and have small additional computational costs. Although we leverage scalar quantization to compress the size of the lookup table (Sec. 3.3), the

main idea of our work is different from these quantization methods. Our work focuses on learning the most typical features of each NN layer and computing the results of these features in advance for future input tensors, in order to both compress NN weights and accelerate NN inference.

Inference result reuse There are a range of works that exploit the duplicated inputs, to reuse their inference results and skip further model inference. Cachier [16] reuses results among nearby users. Potluck [23] reuses results across applications. Euphrate [63] reuses results of temporal pixel motion for continuous vision. FoggyCache [22] reuses results across devices. DeepCache reuses results of similar regions in recent frames of consecutive videos. MCDNN [25] reuses results across applications. Semantic Memory [38] reuses results from some layers for continuous vision. These works rely on specific application inputs. Our work, by comparison, optimizes the model itself to reduce size and FLOPS.

Lookup table for DNN Different from our work that learns and stores the centroid results in the table, there are works that store other results in the tables. LogicNets [52] maps artificial neurons to truth tables in FPGA. It also compresses the model by available pruning and scalar quantization techniques. Ramanathan et al. [49] store the results of bitline computing in Processing in Cache hardware in the table, to accelerate DNN. These works do not target neural network optimization or training techniques like LUT-NN.

8 FUTURE OPPORTUNITY DISCUSSION

Hardware implication Potentially, LUT-NN can greatly reduce FLOPs, e.g., 16× for BERT. However, current hardware restrains LUT-NN from showing its full potential. The reason is that current hardware is optimized for linear computation, but not table lookup. Compared to direct Multiply-Add support for MM in hardware, LUT-NN has to run `argmin` partially sequentially to return the index of the nearest centroid, and then lookup in the table followed by aggregation. Besides, the SIMD width limits the options of the number of centroids. These are obviously inefficient.

Even on this unfriendly hardware, LUT-NN has shown the benefits of table lookup on all dimensions, particularly for latency, power, and memory reduction. An accelerator or function unit supporting the first-class parallel table-lookup pipeline could even more significantly reduce the inference cost. LUT-NN opens new opportunity for hardware design.

Replace more operator types Currently LUT-NN replaces linear computation operators by table lookup, including MM, convolution, and fully connected. The scaled dot-product attention (< 2% of total latency) in attention layers have no weights, and cannot precompute the results with centroids to save in tables, so we do not replace it by lookup table in this work. A possible way to replace it by lookup table is to precompute and save the production results of centroids. In

addition, how to replace the non-linear operators, such as the activation functions, by LUT-NN can be future work.

Learning for hashing As introduced in Sec. 2.1, hashing can be used to avoid the Euclidean distance calculation of k -means encoding, but at the cost of higher quantization error (refer to Fig. 3b). We evaluate the potential of hashing for DNN inference. Since hashing is not differentiable, LUT-NN uses hashing after the centroids are learned. Our current results show that to achieve similar model accuracy as Euclidean distance for encoding, we have to use a 12-level decision tree for hashing. Compared to distance calculation, it can further reduce FLOPs by 30% to 3×. It could be possible to also integrate hashing into the backpropagation to learn the hashing functions. This could reduce the depth of the tree and further reduce the encoding cost.

On current commodity hardware without direct hashing support, hashing costs even higher than Euclidean distance calculation, since the tree traversing is sequential without SIMD support. Therefore, we still use Euclidean distance for encoding in LUT-NN now. Future hardware can also integrate hashing units for further speedup.

LUT-NN on GPUs and FPGAs NVIDIA GPUs provide the efficient shuffle instruction, i.e., “`prmt`”, achieving Tera-OPs throughput. Similar as the LUT-NN implementation on the CPUs, which heavily depends on the shuffle instructions, this GPU shuffle instruction can be leveraged for parallelly looking up the approximated computation results in the tables. FPGAs contain a vast array of lookup tables implemented through reconfigurable memory, which LUT-NN can utilize for the storage and retrieval of codebooks. Additionally, FPGAs are equipped with Block RAM, a primary memory resource that enables efficient data storage. It can be employed for LUT-NN’s lookup tables.

9 CONCLUSION

This paper takes the first step towards DNN inference by table lookup. A new paradigm potentially brings significant benefits to the DNN inference ecosystem, simplifying the inference software and hardware design and decoupling with the DNN algorithm updates. By the centroid learning technique for DNN, LUT-NN achieves comparable accuracy for complex tasks with much less resource cost. However, LUT-NN still has space to be improved, such as the learning technique to improve accuracy and hardware support for table lookup, which calls for future efforts from the community.

10 ACKNOWLEDGEMENT

The authors thank the anonymous shepherd and reviewers for their valuable comments. The work of Yunxin Liu was supported by the National Key R&D Program of China (No. 2022YFF0604501) and Xiaomi Foundation. The work of Deng Cai was supported by the National Nature Science Foundation of China (Grant Nos: 62273302, 62036009, 61936006).

REFERENCES

- [1] Hande Alemdar, Vincent Leroy, Adrien Prost-Boucle, and Frédéric Pétrot. 2017. Ternary neural networks for resource-efficient AI applications. In *2017 international joint conference on neural networks (IJCNN)*. IEEE, 2547–2554.
- [2] Hessam Bagherinezhad, Mohammad Rastegari, and Ali Farhadi. 2017. Lcnn: Lookup-based convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7120–7129.
- [3] Ron Banner, Yury Nahshan, and Daniel Soudry. 2019. Post training 4-bit quantization of convolutional networks for rapid-deployment. *Advances in Neural Information Processing Systems* 32 (2019).
- [4] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Gutttag. 2020. What is the state of neural network pruning? *Proceedings of machine learning and systems* 2 (2020), 129–146.
- [5] Davis Blalock and John Gutttag. 2021. Multiplying Matrices Without Multiplying. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 992–1004. <https://proceedings.mlr.press/v139/blalock21a.html>
- [6] Davis W Blalock and John V Gutttag. 2017. Bolt: Accelerated data mining with fast vector compression. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 727–735.
- [7] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. 2019. Once-for-all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791* (2019).
- [8] Ting Chen, Lala Li, and Yizhou Sun. 2020. Differentiable product quantization for end-to-end embedding compression. In *International Conference on Machine Learning*. PMLR, 1617–1626.
- [9] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. 2018. TVM: An automated end-to-end optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. 578–594.
- [10] Weihan Chen, Peisong Wang, and Jian Cheng. 2020. Towards Convolutional Neural Networks Compression via Global&Progressive Product Quantization.. In *BMVC*.
- [11] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. *Advances in neural information processing systems* 28 (2015).
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009)*, 20–25 June 2009, Miami, Florida, USA. IEEE Computer Society, 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [13] Misha Denil, Babak Shakibi, Laurent Dinh, Marc’Aurelio Ranzato, and Nando De Freitas. 2013. Predicting parameters in deep learning. *Advances in neural information processing systems* 26 (2013).
- [14] Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. 2014. Exploiting linear structure within convolutional networks for efficient evaluation. *Advances in neural information processing systems* 27 (2014).
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [16] Utsav Drolia, Katherine Guo, Jiaqi Tan, Rajeev Gandhi, and Priya Narasimhan. 2017. Cachier: Edge-caching for recognition applications. In *2017 IEEE 37th international conference on distributed computing systems (ICDCS)*. IEEE, 276–286.
- [17] Yunchao Gong, Liu Liu, Ming Yang, and Lubimir Bourdev. 2014. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115* (2014).
- [18] Google. 2019. *TensorFlow: An end-to-end open source machine learning platform*. <https://www.tensorflow.org/>
- [19] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. 2021. Knowledge distillation: A survey. *International Journal of Computer Vision* 129 (2021), 1789–1819.
- [20] R.M. Gray and D.L. Neuhoff. 1998. Quantization. *IEEE Transactions on Information Theory* 44, 6 (1998), 2325–2383. <https://doi.org/10.1109/18.720541>
- [21] Cong Guo, Chen Zhang, Jingwen Leng, Zihan Liu, Fan Yang, Yunxin Liu, Minyi Guo, and Yuhao Zhu. 2022. ANT: Exploiting Adaptive Numerical Data Type for Low-bit Deep Neural Network Quantization. In *55th IEEE/ACM International Symposium on Microarchitecture, MICRO 2022, Chicago, IL, USA, October 1-5, 2022*. IEEE, 1414–1433. <https://doi.org/10.1109/MICRO56248.2022.00095>
- [22] Peizhen Guo, Bo Hu, Rui Li, and Wenjun Hu. 2018. Foggycache: Cross-device approximate computation reuse. In *Proceedings of the 24th annual international conference on mobile computing and networking*. 19–34.
- [23] Peizhen Guo and Wenjun Hu. 2018. Potluck: Cross-application approximate deduplication for computation-intensive mobile applications. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*. 271–284.
- [24] Ruiqi Guo, Sanjiv Kumar, Krzysztof Choromanski, and David Simcha. 2015. Quantization based Fast Inner Product Search. *CoRR abs/1509.01469* (2015). [arXiv:1509.01469](http://arxiv.org/abs/1509.01469) <http://arxiv.org/abs/1509.01469>
- [25] Seungyeop Han, Haichen Shen, Matthai Philipose, Sharad Agarwal, Alec Wolman, and Arvind Krishnamurthy. 2016. Mcdnn: An approximation-based execution framework for deep stream processing under resource constraints. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. 123–136.
- [26] Kaiming He, Fang Wen, and Jian Sun. 2013. K-Means Hashing: An Affinity-Preserving Quantization Method for Learning Binary Compact Codes. In *2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, June 23-28, 2013*. IEEE Computer Society, 2938–2945. <https://doi.org/10.1109/CVPR.2013.378>
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [28] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. <https://doi.org/10.48550/ARXIV.1503.02531>
- [29] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2704–2713.
- [30] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. 2014. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866* (2014).
- [31] Himalaya Jain, Joaquin Zepeda, Patrick Pérez, and Rémi Gribonval. 2017. Subic: A supervised, structured binary code for image search. In *Proceedings of the IEEE international conference on computer vision*. 833–842.
- [32] Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical Reparameterization with Gumbel-Softmax. <https://doi.org/10.48550/ARXIV.1611.01144>

- [33] Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence* 33, 1 (2010), 117–128.
- [34] Yang Jiao, Liang Han, and Xin Long. 2020. Hanguang 800 NPU - The Ultimate AI Inference Solution for Data Centers. In *IEEE Hot Chips 32 Symposium, HCS 2020, Palo Alto, CA, USA, August 16-18, 2020*. IEEE, 1–29. <https://doi.org/10.1109/HCS49909.2020.9220619>
- [35] Norman P. Jouppi, Doe Hyun Yoon, Matthew Ashcraft, Mark Gottscho, Thomas B. Jablin, George Kurian, James Laudon, Sheng Li, Peter Ma, Xiaoyu Ma, Thomas Norrie, Nishant Patil, Sushma Prasad, Cliff Young, Zongwei Zhou, and David Patterson. 2021. Ten Lessons from Three Generations Shaped Google’s TPUv4i (ISCA ’21). IEEE Press, 1–14. <https://doi.org/10.1109/ISCA52012.2021.00010>
- [36] Benjamin Klein and Lior Wolf. 2017. In defense of product quantization. *arXiv preprint arXiv:1711.08589* 2, 3 (2017), 4.
- [37] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Technical Report.
- [38] Yun Li, Chen Zhang, Shihao Han, Li Lina Zhang, Baoqun Yin, Yunxin Liu, and Mengwei Xu. 2021. Boosting mobile CNN inference through semantic memory. In *Proceedings of the 29th ACM International Conference on Multimedia*. 2362–2371.
- [39] Rendong Liang, Ting Cao, Jicheng Wen, Manni Wang, Yang Wang, Jianhua Zou, and Yunxin Liu. 2022. Romou: Rapidly generate high-performance tensor kernels for mobile gpus. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*. 487–500.
- [40] Stuart P. Lloyd. 1982. Least squares quantization in PCM. *IEEE Trans. Inf. Theory* 28, 2 (1982), 129–136. <https://doi.org/10.1109/TIT.1982.1056489>
- [41] MACE. 2020. <https://github.com/XiaoMi/mace>.
- [42] Avner Magen and Anastasios Zouzias. 2011. Low rank matrix-valued Chernoff bounds and approximate matrix multiplication. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 1422–1436.
- [43] Yusuke Matsui, Yusuke Uchida, Hervé Jégou, and Shin’ichi Satoh. 2018. A survey of product quantization. *ITE Transactions on Media Technology and Applications* 6, 1 (2018), 2–10.
- [44] Calvin McCarter and Nicholas Dronen. 2022. Look-ups are not (yet) all you need for deep learning inference. *arXiv preprint arXiv:2207.05808* (2022).
- [45] Jeffrey L McKinstry, Steven K Esser, Rathinakumar Appuswamy, Deepika Bablani, John V Arthur, Izzet B Yildiz, and Dharmendra S Modha. 2018. Discovering low-precision networks close to full-precision networks for efficient embedded inference. *arXiv preprint arXiv:1809.04191* (2018).
- [46] Microsoft. 2019. *ONNX Runtime*. <https://github.com/microsoft/onnxruntime>
- [47] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*.
- [48] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. 2020. Designing network design spaces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10428–10436.
- [49] Akshay Krishna Ramanathan, Gurpreet S. Kalsi, Srivatsa Srinivasa, Tarun Makesh Chandran, Kamlesh R. Pillai, Om Ji Omer, Vijaykrishnan Narayanan, and Sreenivas Subramoney. 2020. Look-Up Table based Energy Efficient Processing in Cache Support for Neural Network Acceleration. In *53rd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2020, Athens, Greece, October 17-21, 2020*. IEEE, 88–101. <https://doi.org/10.1109/MICRO50266.2020.00020>
- [50] Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. 2011. The German Traffic Sign Recognition Benchmark: A multi-class classification competition. In *The 2011 International Joint Conference on Neural Networks*. 1453–1460. <https://doi.org/10.1109/IJCNN.2011.6033395>
- [51] Pierre Stock, Armand Joulin, Rémi Gribonval, Benjamin Graham, and Hervé Jégou. 2020. And the Bit Goes Down: Revisiting the Quantization of Neural Networks. In *ICLR 2020-Eighth International Conference on Learning Representations*. 1–11.
- [52] Yaman Umuroglu, Yash Akhauri, Nicholas James Fraser, and Michaela Blott. 2020. LogicNets: Co-designed neural networks and circuits for extreme-throughput applications. In *2020 30th International Conference on Field-Programmable Logic and Applications (FPL)*. IEEE, 291–297.
- [53] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *CoRR* abs/1804.07461 (2018). [arXiv:1804.07461](http://arxiv.org/abs/1804.07461) <http://arxiv.org/abs/1804.07461>
- [54] Longguang Wang, Xiaoyu Dong, Yingqian Wang, Li Liu, Wei An, and Yulan Guo. 2022. Learnable lookup table for neural network quantization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12423–12433.
- [55] Manni Wang, Shaohua Ding, Ting Cao, Yunxin Liu, and Fengyuan Xu. 2021. Asymo: scalable and efficient deep-learning inference on asymmetric mobile cpus. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 215–228.
- [56] Pete Warden. 2018. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv preprint arXiv:1804.03209* (2018).
- [57] Ofri Wechsler, Michael Behar, and Bharat Daga. 2019. Spring Hill (NNP-I 1000) Intel’s Data Center Inference Chip. In *2019 IEEE Hot Chips 31 Symposium (HCS)*. 1–12. <https://doi.org/10.1109/HOTCHIPS.2019.8875671>
- [58] Tan Yu, Junsong Yuan, Chen Fang, and Hailin Jin. 2018. Product quantization network for fast image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 186–201.
- [59] Cao Yue, M Long, J Wang, Zhu Han, and Q Wen. 2016. Deep quantization network for efficient image retrieval. In *Proc. 13th AAAI Conf. Artif. Intell.* 3457–3463.
- [60] Zhifei Zhang, Yang Song, and Hairong Qi. 2017. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5810–5818.
- [61] Shuchang Zhou, Zekun Ni, Xinyu Zhou, He Wen, Yuxin Wu, and Yuheng Zou. 2016. DoReFa-Net: Training Low Bitwidth Convolutional Neural Networks with Low Bitwidth Gradients. *CoRR* abs/1606.06160 (2016). [arXiv:1606.06160](http://arxiv.org/abs/1606.06160) <http://arxiv.org/abs/1606.06160>
- [62] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160* (2016).
- [63] Yuhao Zhu, Anand Samajdar, Matthew Mattina, and Paul Whatmough. 2018. Euphrates: Algorithm-SoC Co-Design for Low-Power Mobile Continuous Vision. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 547–560.