

# A Brain-Inspired Low-Dimensional Computing Classifier for Inference on Tiny Devices

Shijin Duan  
duan.s@northeastern.edu  
Northeastern University  
Boston, MA, USA

Xiaolin Xu  
x.xu@northeastern.edu  
Northeastern University  
Boston, MA, USA

Shaolei Ren  
sren@ece.ucr.edu  
UC Riverside  
Riverside, CA, USA

## ABSTRACT

By mimicking brain-like cognition and exploiting parallelism, hyperdimensional computing (HDC) classifiers have been emerging as a lightweight framework to achieve efficient on-device inference. Nonetheless, they have two fundamental drawbacks – heuristic training process and ultra-high dimension – which result in sub-optimal inference accuracy and large model sizes beyond the capability of tiny devices with stringent resource constraints. In this paper, we address these fundamental drawbacks and propose a low-dimensional computing (LDC) alternative. Specifically, by mapping our LDC classifier into an equivalent neural network, we optimize our model using a principled training approach. Most importantly, we can improve the inference accuracy while successfully reducing the ultra-high dimension of existing HDC models by orders of magnitude (e.g., 8000 vs. 4/64). We run experiments to evaluate our LDC classifier by considering different datasets for inference on tiny devices, and also implement different models on an FPGA platform for acceleration. The results highlight that our LDC classifier offers an overwhelming advantage over the existing brain-inspired HDC models and is particularly suitable for inference on tiny devices.

## 1 INTRODUCTION

Deploying deep neural networks (DNNs) for on-device inference, especially on tiny Internet of Things (IoT) devices with stringent resource constraints, presents a key challenge [13, 25, 36]. This is due in part to the fundamental limitation of DNNs that involve intensive mathematical operators and computing beyond the capability of many tiny devices [27].

More recently, inspired from the human-brain memorizing mechanism, hyperdimensional computing (HDC) for classification has been emerging as a lightweight machine learning framework targeting inference on resource-constrained edge devices [17, 20]. In a nutshell, HDC classifiers mimic the brain cognition process by representing an object as a vector (a.k.a. *hypervector*) with a very high dimension in the order of thousands or even more. They perform inference by comparing the similarities between the hypervector of a testing sample and a set of pre-built hypervectors representing different classes. Thus, with HDC, the conventional DNN inference process is essentially projected to parallelizable bit-wise operation in a hyperdimensional space. This offers several

key advantages to HDC over its DNN counterpart, including high energy efficiency and low latency, and hence makes HDC classifiers potentially promising for on-device inference [11, 15, 19]. As a consequence, the set of studies on optimizing HDC classifier performance in terms of inference accuracy, latency and/or energy consumption has quickly expanding [11, 14, 17, 22].

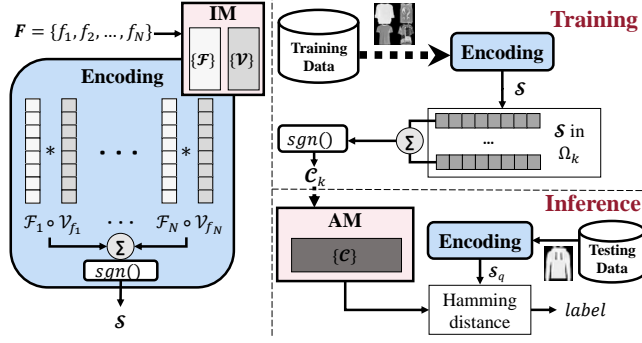
Nonetheless, state-of-the-art (SOTA) HDC classifiers suffer from fundamental drawbacks that prevent their successful deployment for inference on *tiny* devices. First and foremost, the *hyperdimensional* nature of HDC means that each value or feature is represented by a hypervector with several thousands of or even more bits, which can easily result in a prohibitively large model size beyond the limit of typical tiny devices. Even putting aside the large HDC model size, parallel processing of a huge number of bit-wise operations associated with hypervectors is barely feasible on tiny devices, thus significantly increasing the inference latency. Furthermore, the energy consumption by processing hypervectors can also be a deal breaker for inference on tiny devices. In addition, another crucial drawback of HDC classifiers is the low inference accuracy resulting from the lack of a principled training approach. Concretely, the training process of an HDC classifier is extremely simple – simply averaging over the hypervectors of labeled training samples to derive the corresponding class hypervectors. Although some heuristic techniques (e.g., re-training and regeneration [11, 18, 40]) have been recently added, the existing HDC training process still lacks rigorousness and heavily relies on a trial-and-error process without systematic guidance as in the realm of DNNs. In fact, even a well-defined loss function is lacking in the training of HDC classifiers.

In this paper, we address the fundamental challenges in the existing brain-inspired HDC classifiers and propose a new ultra-efficient classification model based on *low-dimensional computing* (LDC) for inference on resource-constrained tiny devices. Here, “low” is relative to the existing HDC model with a dimension of thousands or more.

First, we map the inference process of our LDC classifier into an equivalent neural network that includes a non-binary neural network followed by a binary neural network (BNN). Next, through the lens of this mapping, we optimize the neural network weights, from which we can extract low-dimensional vectors to represent object values and features for efficient inference. Most crucially, our LDC classifier eliminates the large hypervectors used in the existing HDC models, and utilizes optimized low-dimensional vectors with a much smaller size (e.g., 8000 vs. 4/64) to achieve even higher inference accuracy. Thus, compared to the existing SOTA HDC models, our LDC classifier can improve the inference accuracy and

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*tinyML Research Symposium '22, March 2022, San Jose, CA*  
© 2022 Copyright held by the owner/author(s).



**Figure 1: A brain-inspired HDC classifier. IM: item memory. AM: associative memory.**

meanwhile dramatically reduce the model size, inference latency, and energy consumption by orders of magnitude.

We implement our LDC classifier on an FPGA platform under stringent resource constraints. The results show that, in addition to the improved accuracy (92.72% vs. 87.38%) on MNIST, our LDC classifier has a model size of **6.48 KB**, inference latency of **3.99 microseconds**, and inference energy of **64 nanojoules**, which are 100+ times smaller than a SOTA HDC model; for cardiocography application, we achieve an accuracy of 90.50%, with a model size of **0.32 KB**, inference latency of **0.14 microseconds**, and inference energy of **0.945 nanojoules**. The overwhelming advantage over the existing HDC models makes our LDC classifier particularly appealing for inference on tiny devices.

## 2 BRAIN-INSPIRED HDC CLASSIFIERS

We provide an overview of HDC models as illustrated in Figure 1. Consider an input sample represented as a vector with  $N$  features  $F = \{f_1, f_2, \dots, f_N\}$ , where the value range for each feature is normalized and uniformly discretized into  $M$  values, i.e.,  $f_i \in \{1, \dots, M\}$  for  $i = 1, \dots, N$ . HDC encodes all the values, features, and samples as hyperdimensional bipolar vectors, e.g.,  $\mathcal{H} \in \{1, -1\}^{D=10,000}$ , which also equivalently correspond to binary 0/1 bits for hardware efficiency [11, 17]. In this paper, we also interchangeably use binary and bipolar if applicable. Note that the input vector  $F = \{f_1, f_2, \dots, f_N\}$  can represent raw features or extracted features (using, e.g., neural networks and random feature map [40]).

There are four types of hypervectors in HDC models: *value* hypervector  $\mathcal{V}$  (representing the value of a feature), *feature* hypervector  $\mathcal{F}$  (representing the index/position of a feature), *sample* hypervector  $\mathcal{S}$  (representing a training/testing sample), and *class* hypervector  $\mathcal{C}$  (representing a class).

To measure the similarity between two hypervectors, there are two commonly used distances – normalized Hamming and cosine, which are mutually equivalent as shown in the appendix. In this paper, we use the normalized Hamming distance defined as  $Hamm(\mathcal{H}_1, \mathcal{H}_2) = \frac{|\mathcal{H}_1 \neq \mathcal{H}_2|}{D}$ . If two hypervectors  $\mathcal{H}_1$  and  $\mathcal{H}_2$  have a normalized Hamming distance of 0.5, they are considered orthogonal. In a hyperdimensional space, two randomly-generated hypervectors are almost orthogonal [11, 21].

**Generating value and feature hypervectors:** In a typical HDC model [11, 20, 34], the value and feature hypervectors are randomly generated in advance and remain unchanged throughout the training and inference process. Most commonly,  $N$  feature hypervectors  $\mathcal{F}$  are randomly generated to keep mutual orthogonality (e.g., randomly sampling in the hyperdimensional space or rotating one random hypervector), whereas  $M$  value hypervectors  $\mathcal{V}$  are generated to preserve their value correlations (e.g., flipping a certain number of bits from one randomly-generated value hypervector) [11, 17, 20]. As a result, the Hamming distance between any two feature hypervectors is approximately 0.5, while the Hamming distance between two value hypervectors denoting normalized feature values of  $i, j \in \{1, \dots, M\}$  is  $Hamm(\mathcal{V}_i, \mathcal{V}_j) \approx \frac{|i-j|}{2(M-1)}$ .

**Encoding:** An input sample is encoded as a sample hypervector by fetching the pre-generated value and feature hypervectors from the item memory (IM).<sup>1</sup> Specifically, by combining each feature hypervector with its corresponding value hypervector, the encoding output for an input sample is given by

$$S = sgn \left( \sum_{i=1}^N \mathcal{F}_i \circ \mathcal{V}_{f_i} \right), \quad (1)$$

where  $f_i$  is the  $i$ -th feature value,  $\mathcal{V}_{f_i}$  is the corresponding value hypervector,  $\circ$  is the Hadamard product, and  $sgn(\cdot)$  is the sign function that binarizes the encoded sample hypervector. As a tiebreaker, we set  $sgn(0) = 1$ .

**Training:** Given  $K$  classes, the training process is to obtain  $K$  class hypervectors, each for one class. The basic training process is to simply average the sample hypervectors within a class:

$$C_k = sgn \left( \sum_{S \in \Omega_k} S \right) \quad (2)$$

where  $\Omega_k$  is the set of sample hypervectors in  $k$ -th class. All the class hypervectors are stored in the associative memory (AM).

More recently, to improve the accuracy, SOTA HDC models have also added re-training as part of the training process [11, 15]. Concretely, re-training fine tunes the class hypervectors  $\mathcal{C}$  derived in Equation (2): if a training sample is mis-classified, it will be given more weights in correct class hypervector and subtracted from the wrong class hypervector. Essentially, re-training will lead to an adjusted centroid for each class.

**Inference:** The testing input sample is first encoded in the same way as encoding a training sample. To be distinguished from the training sample hypervector, the testing sample hypervector is also referred to as a query hypervector. For inference, the query hypervector  $S_q$  is compared with all the class hypervectors fetched from the associative memory. The most similar one with the lowest Hamming distance indicates the classification result:

$$\arg \min_k Hamm(S_q, C_k). \quad (3)$$

<sup>1</sup>The study [40] applies random feature map for feature extraction [32] and then directly binarizes the extracted features as the encoded hypervector. Nonetheless, we can also use value hypervectors and feature hypervectors to encode the *non*-binary features extracted via random feature map, and this can preserve more information in the extracted features than direct binarization.




Figure 2: Value mapping. As an example, we use a 2-layer fully connected network as the *ValueBox*. FC: fully-connected, BN: batch normalization.

Due to the equivalence of Hamming distance and cosine similarity (appendix), the classification rule in Equation (3) is equivalent to

$$\arg \max_k S_q^T C_k, \quad (4)$$

which essentially transforms the bit-wise comparison to vector multiplication and is instrumental for establishing the equivalence between an HDC model and a neural network.

### 3 THE DESIGN OF LDC CLASSIFIERS

In this section, we describe the LDC classifier design which exploits hardware-friendly association of low-dimensional vectors for efficient inference. Specifically, like in its HDC counterpart, our LDC classifier mimics brain cognition for hardware efficiency by representing features using vectors and performing inference via vector association [11, 20]. Nonetheless, the existing HDC models rely on randomly-generated hypervectors, which not only limits the accuracy but also results in a large inference latency and resource consumption beyond the capability of tiny devices. Our LDC classifier fundamentally differs from them – it uses vectors with orders-of-magnitude smaller dimensions and optimizes the vectors using a principled approach.

With a slight abuse of notations, we keep using  $\mathcal{F}$  and  $\mathcal{V}$  to denote feature and value vectors, respectively, in our LDC classifier wherever applicable. But, unlike in an HDC model that has the same hyperdimension of  $D$  for all hypervectors,  $\mathcal{F}$  and  $\mathcal{V}$  can have different and much lower dimensions of  $D_{\mathcal{F}}$  and  $D_{\mathcal{V}}$ , respectively.

#### 3.1 Mapping to a Neural Network

We now divide the LDC classification process into two parts – encoding and similarity checking – and map them to equivalent neural network operation.

**3.1.1 Encoding.** The encoding process shown in Equation (1) binarizes the summed bindings of value vectors and feature vectors. Instead of using random vectors as in the existing HDC models, we explicitly optimize the value and feature vectors by representing the encoder as a neural network and then using a principled training process. Next, we describe the equivalence between the encoder and its neural network counterpart.

**Value mapping:** As shown in Figure 2,  $\mathcal{V}_{f_i}$  represents a discretized feature value  $f_i \in \{1, \dots, M\}$  with a certain (bipolar) vector. For the ease of presentation, we refer to this mapping functionality as *ValueBox*. In an HDC model, *ValueBox* essentially assigns a random hypervector to a feature value. Alternatively, one may manually design a *ValueBox*, e.g., representing a value directly using its binary code (say,  $9 \rightarrow 1001$ ) [39]. In our LDC design, however,




Figure 3: Illustration of a *feature layer* that binds value and feature vectors.

we exploit the strong representation power of neural networks and use a trainable neural network for the *ValueBox*. For example, Figure 2 illustrates a simple fully-connected neural network with *tanh* activation to map a value  $f_i$  to a binary value vector  $\mathcal{V}_{f_i}$ . We will jointly train the *ValueBox* network together with subsequent operators to optimize the inference accuracy.

**Element-wise binding:** As shown in Equation (1), the encoder binds the feature and value vectors using a Hadamard product. Nonetheless, in our LDC classifier, we do not require  $D_{\mathcal{F}} = D_{\mathcal{V}}$ , which makes the Hadamard product inapplicable. Here, we set  $D_{\mathcal{F}}$  as an integer multiple of  $D_{\mathcal{V}}$ , i.e.,  $D_{\mathcal{F}}/D_{\mathcal{V}} = n$ , for  $n \in \mathbb{N}^+$ . As a result, a value vector  $\mathcal{V}_{f_i}$  can be stacked for  $n$  times in order to have the same dimension as its corresponding feature vector  $\mathcal{F}_i$  for Hadamard product. Equivalently, a feature vector can be evenly divided into  $n$  parts or sub-vectors, each aligned with the value vector for Hadamard product. Thus, the binding for the  $i$ -th feature vector and its corresponding value vector can be represented as

$$\begin{bmatrix} \mathcal{F}_i^1 \circ \mathcal{V}_{f_i} \\ \dots \\ \mathcal{F}_i^n \circ \mathcal{V}_{f_i} \end{bmatrix} = \begin{bmatrix} \text{diag}(\mathcal{F}_i^1) \\ \dots \\ \text{diag}(\mathcal{F}_i^n) \end{bmatrix} \mathcal{V}_{f_i}. \quad (5)$$

Through element-wise binding, the encoder outputs a sample vector given by

$$S = \text{sgn} \left( \sum_{i=1}^N \begin{bmatrix} \text{diag}(\mathcal{F}_i^1) \\ \dots \\ \text{diag}(\mathcal{F}_i^n) \end{bmatrix} \mathcal{V}_{f_i} \right). \quad (6)$$

Crucially, the element-wise binding used by the encoder is equivalently mapped to matrix multiplication in Equation (6). Therefore, it can be represented as a simple BNN whose architecture is shown in Figure 3 and referred to as a *feature layer*. Specifically, by stacking  $N$  value vectors  $\mathcal{V}_{f_i}$  for  $i = 1, \dots, N$ , the input to the feature layer is a vector with  $ND_{\mathcal{V}}$  elements. Also, the structurally sparse weight matrix  $\Theta$  in the feature layer is the collection of feature vectors  $\mathcal{F}_i$  for  $i = 1, \dots, N$ , with only  $ND_{\mathcal{F}}$  bipolar elements in total. By transforming the encoder into an equivalent BNN, we can leverage a principled training process (described in Section 3.2) to optimize the weight matrix  $\Theta$ , which in turn leads to optimized feature vectors  $\mathcal{F}_i$  with a small dimension instead of random hypervectors used by the existing HDC models.

**3.1.2 Similarity Checking.** In our LDC classifier, the classification process involves checking the similarities between a query hypervector and all class hypervectors. Specifically, as shown in Equations (3) and (4), the similarity check is equivalent to matrix calculation, which can also be mapped to the operation in a BNN. Hence, we transform Equation (4) into a *class layer* as shown in Figure 4. The input to the class layer is the sample vector, which is also the




Figure 4: Illustration of a class layer.

output of the preceding feature layer. The weight of the class layer is a  $D_{\mathcal{F}} \times K$  matrix that represents the collection of all the class vectors  $C_k$  for  $k = 1, \dots, K$ . The output of the class layer includes  $K$  products, and class  $k$  with the maximum product  $S^T C_k$  is chosen as the classification result.

**3.1.3 End-to-end mapping.** By integrating different stages of the LDC classification pipeline, we now construct an end-to-end neural network shown in Figure 5. With both non-binary and binary weights, the neural network achieves the same function as our LDC classifier, and this equivalence allows a principled training approach to optimize the weights.

Specifically, in the equivalent neural network, each of the  $N$  feature values of an input sample is first fed to a *ValueBox*, which is a non-binary neural network and outputs bipolar value vectors. Note that a single *ValueBox* is shared by all features to keep the model size small, while our design can easily generalize to different *ValueBoxes* for different features at the expense of increasing the model size (in particular, the size of item memory). Then, the  $N$  value vectors enter a feature layer, which is a structurally sparse BNN as described in Section 3.1.1 and outputs a sample vector. Finally, the sample vector goes through a class layer, based on which we can decide the classification result.

**SOTA HDC models:** Besides the obvious drawback of ultra-high dimension, we can clearly see another major drawback by casting an existing HDC model into our equivalent neural network. Concretely, the *ValueBox* outputs and the feature layer weights corresponding to an HDC model are essentially *randomly* generated, and even the weights in the class layer (i.e.,  $K$  class hypervectors in an HDC model) are obtained by using simple averaging methods in conjunction with heuristic re-training [11, 17, 20]. Thus, the inference accuracy in the existing HDC models are highly sub-optimal.

## 3.2 Training LDC Classifiers

To address the fundamental drawbacks of existing HDC models and maximize the accuracy with a much smaller model size, we can optimize the weights in its equivalent neural network that has both non-binary weights (in the *ValueBox*) and binary weights (in the feature layer and class layer). Specifically, following SOTA training methods for BNNs [28], we use *Adam* with weight decay as the optimization method and consider softmax activation with *CrossEntropy* as the loss function. Due to the equivalence of Hamming distance and cosine similarity metrics for binary vectors, classification based on the largest softmax probability (or equivalently,  $\arg \max_k S_q^T C_k$ ) is equivalent to classification based on the minimum normalized Hamming distance. Additionally, while *CrossEntropy* is commonly used for classification tasks, we can also use other loss functions




Figure 5: Mapping an LDC classifier to a neural network.

such as hinge loss [12]. For training, we set the learning rate by starting with a large value (such as 0.001) with decaying linearly.

After training, the vectors used in our LDC classifier can be extracted as follows for testing.

① The value vectors  $\mathcal{V}$  can be extracted from the *ValueBox* by recording the output corresponding to each possible input. For example, the value vector ( $\mathcal{V}_f$ ) for a certain feature value  $f$  is  $ValueBox(f)$ . Due to  $D_{\mathcal{F}}/D_{\mathcal{V}} = n$  for Hadamard product, the value vector is stacked  $n$  times in the encoder to align with the dimension of feature vectors. Thus, the non-binary weights in the *ValueBox* are not utilized for inference and can be discarded once the value vectors are extracted.

② The feature vectors  $\mathcal{F}$  are extracted from the weight matrix in the feature layer. For the  $i$ -th feature, the feature vector  $\mathcal{F}_i$  can be extracted from the corresponding values in the  $i$ -th weight matrix block:

$$\begin{bmatrix} \mathcal{F}_{i,1}^1 & 0 & \cdots & 0 & \cdots & \mathcal{F}_{i,1}^n & 0 & \cdots & 0 \\ 0 & \mathcal{F}_{i,2}^1 & \cdots & 0 & \cdots & 0 & \mathcal{F}_{i,2}^n & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \cdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathcal{F}_{i,D_V}^1 & \cdots & 0 & 0 & \cdots & \mathcal{F}_{i,D_V}^n \end{bmatrix}$$

such that  $\mathcal{F}_i$  is composed of  $[\mathcal{F}_i^1, \mathcal{F}_i^2, \dots, \mathcal{F}_i^n]$  with dimension  $D_{\mathcal{F}}$ , where  $\mathcal{F}_i^j = [\mathcal{F}_{i,1}^j, \dots, \mathcal{F}_{i,D_V}^j]$  for  $j = 1, 2, \dots, n$ .

③ As shown in the right side of Figure 4, the class vectors  $C$  can be directly extracted from the weight matrix in *class layer*.

Like in Figure 1, the extracted value and feature vectors are stored in the item memory for encoding, and the class vectors are stored in the associative memory for similarity checking.

## 3.3 Inference

For inference, our LDC classifier still follows the encoding and similarity checking process as described in Section 2. Specifically, each feature value of an input sample is first mapped to a value vector, which is then combined together with the corresponding feature vector to form a query sample vector. The query vector is compared with the class vectors for similarity checking and yielding the classification results. In most BNNs, the fully-connected layers still use non-binary weights, which can slow down the inference process on tiny devices [39]. By contrast, although the offline training process involves non-binary weights in the *ValueBox* neural network, the inference process of our LDC classifier is fully binary by utilizing bit-wise operation and association for classification.

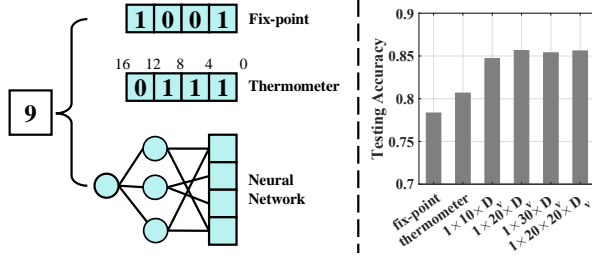


Figure 6: Different *ValueBoxes*. The testing accuracies are presented on the right side with Fashion-MNIST [37].

## 4 PERFORMANCE EVALUATION

In this section, we evaluate the performance of our LDC classifier on different datasets and highlight that it offers an overwhelming advantage over the SOTA HDC models: better accuracy and orders-of-magnitude smaller dimension.

### 4.1 Experimental Setup

Like in the existing HDC research [11, 17], we select four application scenarios for inference on tiny devices: computer vision (MNIST [24] and Fashion-MNIST [37]), human activity (UCIHAR [1]), voice recognition (ISOLET [7]), and cardiocography (CTG [8]). Each feature value is normalized to  $[0, 255]$  and quantized to an 8-bit integer. The configurations for each dataset is shown in Table 1. We compare LDC with the SOTA HDC model with re-training [15], and the basic HDC model without re-training [11]. As reported in [15], the HDC accuracy has significant reduction when the hypervector dimension is lower than 8,000. Thus, we use  $D = 8,000$  for both SOTA and basic HDC. For training, all the experiments are executed in Python with Tesla V100 GPU. For the inference, we also build a hardware acceleration platform on a Zynq UltraScale+ ZU3EG FPGA embedded on the Ultra96 evaluation board.

### 4.2 Hyperparameter Selection

**4.2.1 Value mapping.** We test different neural networks for the *ValueBox* and also compare them against two manual designs (i.e., fix-point encoding and thermometer encoding [39]), as illustrated in Figure 6). We set  $D_V = 8$  to make fix-point encoding and thermometer encoding capable of representing 256 values. Considering the Fashion-MNIST dataset, we see that our neural networks can discover better *ValueBoxes* than manual designs to achieve higher accuracy. On the other hand, it has no significant differences by varying the network size, such as  $1 \times 10 \times D_V$  vs.  $1 \times 20 \times D_V$ . In our experiments, we will use  $1 \times 20 \times D_V$  for the neural network in the *ValueBox*.

**4.2.2 Dimensions.** The dimensions  $D_V$  and  $D_F$  used in our LDC classifier are important hyperparameters. As a case study, we use Fashion-MNIST to evaluate the impact of  $D_V$  and  $D_F$  on the accuracy. By setting different  $n = D_F/D_V$ , we also vary  $D_V = 2, 4, 8, 16$ , as presented in Figure 7. In general, a higher  $D_V$  retains richer information about the input, thus yielding a higher accuracy. Nonetheless, even with  $D_V = 2$  or  $D_V = 4$ , we can still achieve a reasonably




Figure 7: The inference accuracy with different  $D_V$  and  $D_F$  on Fashion-MNIST. Initial learning rate is  $1 \times 10^{-3}$ , weight decay is  $1 \times 10^{-4}$ , and learning rate will decay by half every 5 iterations.

Table 1: Configuration of our LDC classifier for each dataset. The batch size is 64.

Dataset	N	# of (train, test, class)	$D_V, D_F$	LR <sup>1</sup>	WD <sup>2</sup>
MNIST	784	(60000, 10000, 10)	4, 64	0.0001	0
Fashion-MNIST	784	(60000, 10000, 10)	4, 64	0.0002	0.00001
UCIHAR	561	(7352, 2947, 6)	4, 128	0.001	0.0001
ISOLET	617	(6328, 1559, 26)	4, 128	0.005	0.0001
CTG	21	(1701, 425, 3)	4, 64	0.008	0.0001

<sup>1</sup> LR = Learning Rate <sup>2</sup> WD = Weight Decay

Table 2: Inference accuracy (% in  $mean \pm std$ ).

Classifier	MNIST	Fashion-MNIST	UCIHAR	ISOLET	CTG
LDC	92.72 $\pm$ 0.18	85.47 $\pm$ 0.29	92.56 $\pm$ 0.41	91.33 $\pm$ 0.50	90.50 $\pm$ 0.46
Basic HDC	79.35 $\pm$ 0.03	69.11 $\pm$ 0.03	89.17 $\pm$ 0.16	85.90 $\pm$ 0.25	71.35 $\pm$ 0.76
SOTA HDC [15]	87.38 $\pm$ 0.21	79.24 $\pm$ 0.52	90.31 $\pm$ 0.06	90.66 $\pm$ 0.31	89.51 $\pm$ 0.43

good accuracy by increasing  $D_F$ . In all cases, our dimensions are orders-of-magnitude smaller than the dimensions of 8,000 or higher in existing HDC models.

### 4.3 Inference Accuracy


In Table 2, We evaluate the inference accuracy that LDC can achieve compared to both basic and SOTA HDC models. The result shows that LDC can outperform existing HDC models that use random value/feature vectors and heuristically generate class vectors. While retraining improves the inference accuracy against the basic HDC, the hypervector dimension must be as large as 8,000 to prevent accuracy degradation. By contrast, our LDC classifier reduces the dimension significantly, without introducing any extra cost during the inference process.

### 4.4 Hardware Acceleration

**4.4.1 Hardware design.** For hardware implementation, bipolar values  $\{1, -1\}$  are represented as binary values  $\{0, 1\}$ , respectively. The multiplication and accumulation on bipolar values are realized using *XOR* and *popcount* which is constructed using tree adders, on the binary representation [26].

**Table 3: Hardware acceleration on inference with two selected datasets. The system frequency is 200MHz.**

Dataset	Name	Accuracy (%)	Platform	Model Size (KB)	Latency ( $\mu$ s)	(LUT, BRAM, DSP)	Energy (nJ)
MNIST	LDC	92.72	Zynq UltraScale+	<b>6.48</b>	<b>3.99</b>	(745, 5, 1)	<b>64</b>
	SOTA HDC	87.38	Zynq UltraScale+	1050	499	(768, 178, 1)	36926
	FINN [35]	95.83	Zynq-7000	600	240	(5155, 16, -)	96000
CTG	LDC	90.50	Zynq UltraScale+	<b>0.32</b>	<b>0.14</b>	(345, 3, 1)	<b>0.945</b>
	SOTA HDC	89.51	Zynq UltraScale+	280	16.88	(362, 9, 1)	169
	Compressed HDC [2]	$\sim$ 82	Odroid XU3	45.1	90	NA	6300




**Figure 8: The hardware acceleration for LDC on FPGA. IM: Item Memory. AM: Associative Memory.**

The existing HDC models typically exploit full parallelism for acceleration. For example, given  $N$  features, the FPGA prepares  $N$  identical hypervector multiplication blocks to encode all the features simultaneously, incurring a high resource expenditure (e.g., over  $10^5$  LUTs, 100 BRAMs, and 800 DSPs in SOTA FPGA acceleration [19]). Nonetheless, this design does not fit into tiny devices, for which we must limit the resource utilization.

We design a pipeline structure for feature encoding, rather than in parallel, to fit into tiny devices. The structure is demonstrated in Figure 8. Only one vector multiplication block and several BRAMs are required. Although the encoding time will increase to  $N + 1$  cycles, the latency is still on a microsecond scale. Subsequently, one adder is utilized to accumulate the multiplied vectors, followed by a threshold ( $\tau = N/2$ ) comparator to binarize the encoding output. For checking similarity, we also use a pipeline structure for comparison on all class vectors. Then, a tree adder along the vector dimension is used for Hamming distance calculation. Finally, we transfer the Hamming distances back to the CPU to execute the  $\arg \min()$  function for classification.<sup>2</sup>

<sup>2</sup>As an alternative, the  $\operatorname{argmin}()$  function can also be integrated on FPGA with a small overhead.



**Figure 9: Bit error robustness of different models for Fashion-MNIST dataset. The accuracy is averaged over five different runs for each case.**

**4.4.2 Results.** In Table 3, we show the efficiency results. As we focus on tiny devices, we limit the resource utilization, e.g.,  $< 1$ MB model size and  $< 10$ k LUT utilization. To evaluate the existing HDC models, for fair comparison, we also implement the SOTA HDC classifier [15] with  $D = 8,000$  using our acceleration designs. Moreover, we choose two other lightweight models for comparison: a compressed HDC model [2] that uses a small vector but has non-binary weights and per-feature *ValueBoxes* founded using evolutionary search; and SFC-fix with FINN [35] that applies a 3-layer binary MLP on FPGA. For these models, we only report their results available for our considered datasets. Like in the literature [15, 19], the results are measured for model inference only, excluding data transmission between the FPGA and CPU.

For the MNIST dataset, by reducing the dimension of SOTA HDC models by 125 times, the model size of LDC is only 6.48 KB. Further, the low dimension can benefit the resource utilization and execution time. For the SOTA HDC model with  $D = 8,000$ , the number of BRAMs increases greatly to store all hypervectors, but other resources such as LUT and DSP do not increase dramatically due to sequential execution to account for tiny devices; consequently, the latency increases by 125 times. For energy consumption, the result shows that our LDC classifier has the lowest, because of the very low resource utilization and short latency. For the CTG dataset, the results are even more impressive as shown in Table 3.

### 4.5 Robustness Against Random Bit Errors

While inefficient, a natural byproduct of the hyperdimensionality of HDC classifiers is the robustness against random hardware bit errors [29, 31]. By contrast, our LDC classifier reduces the dimension by orders of magnitude and hence might become less robust against random bit errors. Here, we show in Figure 9 the robustness analysis for both LDC and HDC classifiers. Specifically, we inject

bit errors in the associative memory under various bit error rates, and assess the inference accuracy caused by the injected bit errors. We see that our LDC classifier has the highest accuracy and also achieves comparable robustness on a par with the HDC counterparts. The counter-intuitive results can be explained by the fact that, although our LDC classifier significantly reduces dimensions, the information is still spread uniformly within each compact vector (i.e., different bits are equally important), thus exhibiting robustness against random bit errors.

Note finally that we can also run multiple LDC classifiers and apply the majority rule for better robustness. Due to the orders-of-magnitude dimension reduction, the total size of multiple LDC classifiers is still much less than that of a single HDC classifier.

## 5 RELATED WORK

The set of applications with HDC classifiers have been proliferating, include language classification [16], image classification [5, 11], emotion recognition based on physiological signals [6], distributed fault isolation in power plants [23], gesture recognition for wearable devices [3], seizure onset detection [4], and robot navigation [30].

Nonetheless, the training process for existing HDC classifiers is mostly heuristics-driven and often relies on random hypervectors for encoding, resulting in low inference accuracy [11]. In many cases, even a well-defined loss function is lacking for training HDC classifiers. More recently, [10] learns a neural network model based on non-binary hypervector representations, and [9] uses generalized learning vector quantization (GLVQ) to optimize the class vectors. But, they both consider hyperdimensional representations based on random hypervectors, thus resulting in an overly large model size as well as suboptimal accuracy.

In addition, the ultra-high dimension is another fundamental drawback of HDC models. Simply reducing the dimension without fundamentally changing the HDC design can dramatically decrease the inference accuracy. A recent study [2] reduces the dimension but considers non-binary HDC that is unfriendly to hardware acceleration. Moreover, it uses different sets of value vectors for different features, and hence results in a large model size (Table 3). Last but not least, its training process is still heuristic-driven as in the existing HDC models. By sharp contrast, our LDC classifier is optimized based on a principled training approach guided by a loss function, and offers an overwhelming advantage over the existing HDC models, in terms of accuracy, model size, inference latency, and energy consumption.

Our LDC classifier is relevant to but also differs significantly from BNNs which use binary weights to speed up inference [33, 38, 39]. To avoid information loss, non-binary weights are still utilized in the early stage of typical BNNs [33, 38] which may not be supported by tiny devices, whereas the inference of our LDC classifier is fully binary and follows a brain-like cognition process. More recently, [39] manually designs the value mapping from raw features to binary features. In our design, we use a neural network to automatically learn the mapping which, as shown in Figure 6, outperforms manual designs in terms of accuracy.

## 6 CONCLUSION

In this paper, we address the fundamental drawbacks of existing HDC models and propose an LDC alternative. We map our LDC classifier into an equivalent neural network, and optimize the model using a principled training approach. We run experiments to evaluate our LDC classifier, and also implement different models on an FPGA platform for acceleration. The results highlight that our LDC classifier has an overwhelming advantage over the existing HDC models and is particularly suitable for inference on tiny devices.

## APPENDIX

We show the equivalence between the (normalized) Hamming distance and *cosine* similarity on two bipolar vectors. The normalized Hamming distance and *cosine* similarity are defined as:

$$Ham(\mathcal{H}_1, \mathcal{H}_2) = \frac{|\mathcal{H}_1 \neq \mathcal{H}_2|}{D}, \text{ and } cosine(\mathcal{H}_1, \mathcal{H}_2) = \frac{\mathcal{H}_1^T \mathcal{H}_2}{\|\mathcal{H}_1\| \|\mathcal{H}_2\|},$$

where  $|\mathcal{H}_1 \neq \mathcal{H}_2|$  denotes the number of different bits in  $\mathcal{H}_1$  and  $\mathcal{H}_2$ ,  $\|\mathcal{H}_1\|$  and  $\|\mathcal{H}_2\|$  denote the  $l_2$  norms of  $\mathcal{H}_1$  and  $\mathcal{H}_2$ , respectively. Since vectors are in bipolar  $\{1, -1\}$ , we have  $\mathcal{H}_1^T \mathcal{H}_2 = (|\mathcal{H}_1 = \mathcal{H}_2| - |\mathcal{H}_1 \neq \mathcal{H}_2|)$ . Plus the fact that  $\|\mathcal{H}_1\| \|\mathcal{H}_2\| = D$  and  $(|\mathcal{H}_1 \neq \mathcal{H}_2| + |\mathcal{H}_1 = \mathcal{H}_2|) = D$ , we can conclude  $cosine(\mathcal{H}_1, \mathcal{H}_2) = \frac{D - 2|\mathcal{H}_1 \neq \mathcal{H}_2|}{D} = 1 - 2Ham(\mathcal{H}_1, \mathcal{H}_2)$ .

## REFERENCES

- [1] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, Jorge Luis Reyes-Ortiz, et al. 2013. A public domain dataset for human activity recognition using smartphones. In *Esann*, Vol. 3. 3.
- [2] Toygun Basaklar, Yigit Tuncel, Shruti Yadav Narayana, Suat Gumussoy, and Umit Y Ogras. 2021. Hypervector Design for Efficient Hyperdimensional Computing on Edge Devices. *arXiv preprint arXiv:2103.06709* (2021).
- [3] Simone Benatti, Fabio Montagna, Victor Kartsch, Abbas Rahimi, Davide Rossi, and Luca Benini. 2019. Online Learning and Classification of EMG-Based Gestures on a Parallel Ultra-Low Power Platform Using Hyperdimensional Computing. *IEEE Transactions on Biomedical Circuits and Systems* 13, 3 (2019), 516–528.
- [4] Alessio Burrello, Kaspar Schindler, Luca Benini, and Abbas Rahimi. 2020. Hyperdimensional Computing With Local Binary Patterns: One-Shot Learning of Seizure Onset and Identification of Ictogenic Brain Regions Using Short-Time iEEG Recordings. *IEEE Transactions on Biomedical Engineering* 67, 2 (2020), 601–613.
- [5] Cheng-Yang Chang, Yu-Chuan Chuang, and An-Yeu (Andy) Wu. 2020. Task-Projected Hyperdimensional Computing for Multi-task Learning. In *Artificial Intelligence Applications and Innovations*.
- [6] En-Jui Chang, Abbas Rahimi, Luca Benini, and An-Yeu Andy Wu. 2019. Hyperdimensional Computing-based Multimodality Emotion Recognition with Physiological Signals. In *IEEE International Conference on Artificial Intelligence Circuits and Systems*.
- [7] Ron Cole and Mark Fanty. 1994. ISOLET. UCI Machine Learning Repository.
- [8] J P Marques de Sa, J Bernardes, and D Ayres de Campos. 2010. Cardiotocography. UCI Machine Learning Repository.
- [9] Cameron Diao, Denis Kleyko, Jan M. Rabaey, and Bruno A. Olshausen. 2021. Generalized Learning Vector Quantization for Classification in Randomized Neural Networks and Hyperdimensional Computing. *arXiv:2106.09821 [cs.LG]*
- [10] Ashwinkumar Ganesan, Hang Gao, Sunil Gandhi, Edward Raff, Tim Oates, James Holt, and Mark McLean. 2021. Learning with Holographic Reduced Representations. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.). <https://openreview.net/forum?id=RX6PrpXP>
- [11] Lulu Ge and Keshab K. Parhi. 2020. Classification Using Hyperdimensional Computing: A Review. *IEEE Circuits and Systems Magazine* 20, 2 (2020), 30–47.
- [12] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- [13] Song Han, Huizi Mao, and William J. Dally. 2016. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. In *ICLR*.
- [14] Longbo Huang, Xin Liu, and Xiaohong Hao. 2014. The Power of Online Learning in Stochastic Network Optimization. In *The 2014 ACM International Conference on Measurement and Modeling of Computer Systems* (Austin, Texas, USA) (SIGMETRICS '14). Association for Computing Machinery, New York, NY, USA, 153–165. <https://doi.org/10.1145/2591971.2591990>

- [15] Mohsen Imani, Samuel Bosch, Sohun Datta, Sharadhi Ramakrishna, Sahand Salamat, Jan M Rabaey, and Tajana Rosing. 2019. QuantHD: A quantization framework for hyperdimensional computing. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2019).
- [16] Mohsen Imani, John Hwang, Tajana Rosing, Abbas Rahimi, and Jan M Rabaey. 2017. Low-Power Sparse Hyperdimensional Encoder for Language Recognition. *IEEE Design Test* 34, 6 (2017), 94–101.
- [17] Mohsen Imani, John Messerly, Fan Wu, Wang Pi, and Tajana Rosing. 2019. A Binary Learning Framework for Hyperdimensional Computing. In *DATE*.
- [18] Mohsen Imani, Xunzhao Yin, John Messerly, Saransh Gupta, Michael Niemier, Xiaobo Sharon Hu, and Tajana Rosing. 2019. Searchd: A memory-centric hyperdimensional computing with stochastic training. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39, 10 (2019), 2422–2433.
- [19] Mohsen Imani, Zhuowen Zou, Samuel Bosch, Sanjay Anantha Rao, Sahand Salamat, Venkatesh Kumar, Yeseong Kim, and Tajana Rosing. 2021. Revisiting hyperdimensional learning for fpga and low-power architectures. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 221–234.
- [20] Pentti Kanerva. 2009. Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive computation* 1, 2 (2009), 139–159.
- [21] Geethan Karunaratne, Manuel Le Gallo, Giovanni Cherubini, Luca Benini, Abbas Rahimi, and Abu Sebastian. 2020. In-memory Hyperdimensional Computing. *Nature Electronics* (Jun 2020).
- [22] Denis Kleyko, Mansour Kheffache, E. Paxon Frady, Urban Wiklund, and Evgeny Osipov. 2019. Density Encoding Enables Resource-Efficient Randomly Connected Neural Networks. *arXiv preprint arXiv:1909.09153* (2019).
- [23] Denis Kleyko, Evgeny Osipov, Nikolaos Papakonstantinou, and Valeriy Vyatkin. 2018. Hyperdimensional Computing in Industrial Systems: The Use-case of Distributed Fault Isolation in a Power Plant. *IEEE Access* 6 (May 2018), 30766–30777.
- [24] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324. <https://doi.org/10.1109/5.726791>
- [25] Chaojian Li, Zhongzhi Yu, Yonggan Fu, Yonggan Zhang, Yang Zhao, Haoran You, Qixuan Yu, Yue Wang, Cong Hao, and Yingyan Lin. 2021. HW-NAS-Bench: Hardware-Aware Neural Architecture Search Benchmark. In *ICLR*. [https://openreview.net/forum?id=\\_0kaDkv3dVf](https://openreview.net/forum?id=_0kaDkv3dVf)
- [26] Shuang Liang, Shouyi Yin, Leibo Liu, Wayne Luk, and Shaojun Wei. 2018. FP-BNN: Binarized neural network on FPGA. *Neurocomputing* 275 (2018), 1072–1086.
- [27] Ji Lin, Wei-Ming Chen, Yujun Lin, John Cohn, Chuang Gan, and Song Han. 2020. MCUNet: Tiny Deep Learning on IoT Devices. In *NeurIPS*.
- [28] Zechun Liu, Zhiqiang Shen, Shichao Li, Koen Helwegen, Dong Huang, and Kwang-Ting Cheng. 2021. How Do Adam and Training Strategies Help BNNs Optimization? *arXiv preprint arXiv:2106.11309* (2021).
- [29] Dongning Ma, Jianmin Guo, Yu Jiang, and Xun Jiao. 2021. HDTest: Differential Fuzz Testing of Brain-Inspired Hyperdimensional Computing. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*. 391–396. <https://doi.org/10.1109/DAC18074.2021.9586169>
- [30] Anton Mitrokhin, P Sutor, Cornelia Fermüller, and Yiannis Aloimonos. 2019. Learning Sensorimotor Control with Neuromorphic Sensors: Toward Hyperdimensional Active Perception. *Science Robotics* 4, 30 (2019).
- [31] P. Podual, Y. Ni, Y. Kim, K. Ni, R. Kuma, R. Cammarota, and M. Imani. 2021. Hyperdimensional Self-Learning Systems Robust to Technology Noise and Bit-Flip Attacks. In *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*.
- [32] Ali Rahimi and Benjamin Recht. 2007. Random Features for Large-Scale Kernel Machines. In *Proceedings of the 20th International Conference on Neural Information Processing Systems (Vancouver, British Columbia, Canada) (NIPS'07)*. Curran Associates Inc., Red Hook, NY, USA, 1177–1184.
- [33] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. 2016. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks. In *ECCV*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.).
- [34] Sahand Salamat, Mohsen Imani, Behnam Khaleghi, and Tajana Rosing. 2019. F5-HD: Fast Flexible FPGA-Based Framework for Refreshing Hyperdimensional Computing. In *FPGA*.
- [35] Yaman Umuroglu, Nicholas J Fraser, Giulio Gambardella, Michaela Blott, Philip Leong, Magnus Jahre, and Kees Vissers. 2017. Finn: A framework for fast, scalable binarized neural network inference. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. 65–74.
- [36] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. 2016. Learning Structured Sparsity in Deep Neural Networks. In *NIPS*.
- [37] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. *Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms*. arXiv:cs.LG/1708.07747 [cs.LG]
- [38] Chunyu Yuan and Sos S. Agaian. 2021. A comprehensive review of Binary Neural Network. *CoRR* abs/2110.06804 (2021).
- [39] Yichi Zhang, Junhao Pan, Xinheng Liu, Hongzheng Chen, Deming Chen, and Zhiru Zhang. 2021. FracBNN: Accurate and FPGA-efficient binary neural networks with fractional activations. In *The 2021 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. 171–182.
- [40] Zhuowen Zou, Yeseong Kim, Farhad Imani, Haleh Alimohamadi, Rosario Cammarota, and Mohsen Imani. 2021. Scalable Edge-Based Hyperdimensional Learning System with Brain-like Neural Adaptation. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (St. Louis, Missouri) (SC'21)*. Article 38, 15 pages.