SHEARer: Highly-Efficient Hyperdimensional Computing by Software-Hardware Enabled Multifold AppRoximation

Behnam Khaleghi, Sahand Salamat, Anthony Thomas, Fatemeh Asgarinejad, Yeseong Kim, and Tajana Rosing
Computer Science and Engineering Department, UC San Diego, La Jolla, CA 92093, USA
{bkhaleghi, sasalama, althomas, fasgarinejad, yek048, tajana}@ucsd.edu

ABSTRACT

Hyperdimensional computing (HD) is an emerging paradigm for machine learning based on the evidence that the brain computes on high-dimensional, distributed, representations of data. The main operation of HD is encoding, which transfers the input data to hyperspace by mapping each input feature to a hypervector, followed by a bundling procedure that adds up the hypervectors to realize the encoding hypervector. The operations of HD are simple and highly parallelizable, but the large number of operations hampers the efficiency of HD in embedded domain. In this paper, we propose SHEARer, an algorithm-hardware co-optimization to improve the performance and energy consumption of HD computing. We gain insight from a prudent scheme of approximating the hypervectors that, thanks to error resiliency of HD, has minimal impact on accuracy while provides high prospect for hardware optimization. Unlike previous works that generate the encoding hypervectors in full precision and then and then perform *ex-post* quantization, we compute the encoding hypervectors in an approximate manner that saves resources yet affords high accuracy. We also propose a novel FPGA architecture that achieves striking performance through massive parallelism with low power consumption. Moreover, we develop a software framework that enables training HD models by emulating the proposed approximate encodings. The FPGA implementation of SHEARer achieves an average throughput boost of 104,904× (15.7×) and energy savings of up to 56,044× (301×) compared to state-of-the-art encoding methods implemented on Raspberry Pi 3 (GeForce GTX 1080 Ti) using practical machine learning datasets.

ACM Reference Format:

Behnam Khaleghi, Sahand Salamat, Anthony Thomas, Fatemeh Asgarinejad, Yeseong Kim, and Tajana Rosing. 2020. SHEARer: Highly-Efficient Hyperdimensional Computing by Software-Hardware Enabled Multifold AppRoximation. In *ISLPED '20: ACM/IEEE International Symposium on Low Power Electronics and Design*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3370748.3406587>

1 INTRODUCTION

Networked sensors with native computing power – otherwise known as the “internet of things” (IoT) – are a rapidly growing source of data. Applications based on IoT devices typically use machine learning (ML) algorithms to generate useful insights from data. While modern ML techniques – in particular deep neural networks (DNNs) – can produce state-of-the-art results, they often entail substantial memory and compute requirements which may exceed the resources available on light-weight edge devices. Thus, there is a pressing need to develop novel ML techniques which provide accuracy and flexibility while meeting the tight resource constraints imposed by edge-sensing devices.

Hyperdimensional computing – HD for short – is an emerging paradigm for machine learning based on evidence from the neuroscience community that the brain “computes” on high-dimensional, distributed, representations of data [1–3]. In HD, the primitive units of computation

are high-dimensional vectors of length d_{hv} sampled randomly from the uniform distribution over the binary cube $\{\pm 1\}^{d_{hv}}$. Typical values of d_{hv} are in the range 5-10,000. Because of their high-dimensionality, any randomly chosen pair of points will be approximately orthogonal (i.e., their inner product will be approximately zero). A useful consequence of this is that sets can be encoded simply by summing (or “bundling”) together their constituent vectors. For any collection of vectors P, Q, V their element-wise sum $S = P + Q + V$ is, in expectation, closer to P, Q and V than any other randomly chosen vector in the space.

Given HD representations of data, this provides a simple classification scheme: we simply take the data points corresponding to a particular class and superimpose them into a single representation for the set. Then, given a new piece of data with unknown class label, we compute the similarity with the hypervectors representing each class and return the label corresponding to the most similar one. More formally, suppose we are given a set of labeled data $X = \{(x_i, y_i)\}_{i=1}^N$ where $x \in \mathbb{R}^{d_{io}}$ corresponds to an observation in low-dimensional space and $y \in C$ is a categorical variable indicating the class to which a particular x belongs. HD classification proceeds by generating a set of “class hypervectors” which represent the training data corresponding to each class. Then, given a piece of data for which we do not know the correct label – the “query” – we simply compute the similarity between the query and each class hypervector and return the label of the most similar one. This process is illustrated in Figure 1.

Suppose we wish to generate the class hypervector corresponding to some class $k \in C$. The prototype can be generated simply by superimposing (also called “bundling”) the HD-encoded representation of the training data corresponding to that particular class [1, 4]:

$$C_k = \sum_{i \text{ s.t. } y_i=k} \text{enc}(x_i) \quad (1)$$

where $\text{enc} : \mathbb{R}^{d_{io}} \rightarrow \{\pm 1\}^{d_{hv}}$ is some *encoding function* which maps a low-dimensional signal to a binary HD representation. Then, given some piece of “query” data x_q for which we *do not* know the correct label we simply return the predicted label as:

$$k^* = \underset{k \in C}{\operatorname{argmax}} \delta(\text{enc}(x_q), C_k) \quad (2)$$

where δ is an appropriate similarity metric, e.g. inner-product/cosine distance – appropriate for integer or real valued encoding schemes – and the hamming distance – appropriate for binary HD representations. This phase is commonly referred to in literature as “associative search”. Despite the simplicity of this “learning” scheme, HD computing has been successfully applied to a number of practical problems in the literature ranging from optimizing the performance of web-browsers [5], to DNA sequence alignment [6], bio-signal processing [7], robotics [8, 9], and privacy preserving learning [10, 11].

The primary appeal of HD computing lies in its amenability to implementation in modern hardware accelerators. Because the HD representations are simply long Boolean vectors, they can be processed extremely efficiently in highly parallel platforms like GPUs, FPGAs and PIM architectures. The principal challenge of HD computing – and the focus of this paper – lies in designing good encoding schemes which (1) represent the data in a format suitable for learning and (2) are efficient to implement in hardware. In general, encoding is the most expensive stage in the HD learning pipeline – in some cases taking up to 10×

This work is licensed under a Creative Commons Attribution International 4.0 License.

Figure 1: Encoding and training in HD.

longer than training or prediction [12]. Existing encoding methods generate hypervectors in full integer-precision and then *ex-post* quantizing to $\{\pm 1\}$. While this accelerates the associative search phase, it does not address encoding which is the primary source of inefficiency.

In this work, we propose novel techniques to compute the encodings in an approximate manner that saves a substantial amount of resources with an insignificant impact on accuracy. Of independent interest is our novel FPGA implementation that achieves striking performance through massive parallelism with low power consumption. Approximate encodings entail models to be trained in a similar approximate fashion. Thus we also develop a software emulation to enable users to train desired HD models. Our software framework enables users to explore the tradeoff between the degree of approximation, accuracy, and resource utilization (hence power consumption) by generating a pre-compiled library that correlates approximation schemes and FPGA resource utilization and power consumption. We show that our techniques achieve a performance boost of $104,904\times$ ($15.7\times$) and energy savings of up to $56,044\times$ ($301\times$) compared to state-of-the-art encoding methods implemented on Raspberry Pi 3 (GeForce GTX 1080 Ti).

2 BACKGROUND AND MOTIVATION

2.1 HD Encoding Algorithms

The literature has proposed a number of encoding methods for the multitude of data types which arise in practical learning settings. We here focus on a method from [1, 4, 13] which we refer to as “ID-vector” based encoding. This encoding method is widely used (see e.g. [7, 13–15]) and works well on both discrete and continuous data. We focus the discussion on continuous data as discrete data is a simple extension.

Suppose we wish to encode some set of vectors $\mathcal{X} = \{x_i\}_{i=1}^N$ where x_i is supported on some compact subset of $\mathbb{R}^{d_{iv}}$. To begin, we first quantize the domain of each feature into a set of L discrete values $\mathcal{L} = \{l_i\}_{i=1}^L$ and assign each $l_i \in \mathcal{L}$ a codeword $L_i \in \{\pm 1\}^{d_{hv}}$. To preserve the *ordinal relationship* between the quantizer bins (the l_i), we wish the similarity between the codewords L_i, L_j to be inversely proportional to distance between the corresponding quantization bins; e.g. $\delta(L_i, L_j) \propto |l_i - l_j|^{-1}$. To enforce this property we generate the codeword L_1 corresponding to the minimal quantizer bin l_1 by sampling randomly from $\{\pm 1\}^{d_{hv}}$. The codeword for the second bin is generated by flipping $\frac{d_{hv}}{2L}$ random coordinates in L_1 . The codeword for the third bin is generated analogously from L_2 and so on. Thus, the codewords for the minimal and maximal bins are orthogonal and $\delta(L_i, L_j)$ decays as $|j - i|$ increases. This scheme is appropriate for quantizers with linearly spaced bins – however, it can be extended to variable bin-width quantizers.

To complete the description of encoding, let $q(x_i)$ be a function which returns the appropriate codeword L for a component $x_i \in \mathbf{x}$. Then encoding proceeds as follows:

$$\mathbf{X} = \sum_{j=i}^{d_{iv}} q(x_i) \otimes \mathbf{P}_i \quad (3)$$

Figure 2: (a) Adder-tree and (b) counter-based implementation of popcount. $\oplus$ denotes add operation.

Where $\mathbf{P}_i$ is a “position hypervector” which encodes the index of the feature value (e.g. $i \in \{1, \dots, d_{iv}\}$) and $\otimes$ is a “binding” operation which is typically taken to be XOR.

2.2 Motivation

The basic operations of HD are simple but are also numerous due to large number of dimensions. Prior work has proposed various algorithmic and hardware innovations to tackle the computational challenges of HD. Acceleration in hardware has typically focused on FPGAs [16–18] or ASIC-ish accelerators [19, 20]. FPGA-based implementations provide high parallelism and bit-level granularity of operations that significantly improves the effective utilization of resources and performance. Also, FPGAs are advantageous over specialized ASICs by allowing easy customization of model parameters such as lengths of hypervectors (d_{hv}) and input-vectors (d_{iv}) and the quantization levels. This flexibility is crucial since learning applications are heterogeneous in practice. Therefore, we here focus on an FPGA based implementation but emphasize our techniques are generic and can be integrated with ASIC- [19] and processor-based [20] implementations.

As noted in the preceding section, the element-wise sum is a critical operation in the encoding pipeline. Thus, popcount operations play a major role in determining the efficiency of HD computing. Figure 2(a) shows a popular tree-based implementation of popcount that adds d_{iv} binary bits (note that we can replace ‘–1’s by 0 in the hardware). Each six-input look-up table (LUT-6) of conventional FPGAs consists of two LUT-5. Hence, we can implement the first stage of the tree using $\frac{d_{iv}}{3}$ of three-port one-bit adders. Each subsequent stage comprises two-port k -bit adders where k increases by one at each stage, while the number of adders per stage decreases by a factor of $\frac{1}{2}$. An n -bit adder requires n LUT-6. Thus, the number of LUT-6 for a d_{iv} -input popcount can be formulated as Equation (4).

$$n_{\text{LUT6}}(\text{adder-tree}) = \sum_{i=1}^{\log d_{iv}} \frac{d_{iv}}{3} \times \frac{i}{2^{i-1}} \approx \frac{4}{3} d_{iv} \quad (4)$$

HD operations can be parallelized at the granularity of a single coordinate in each hypervector: all dimensions of the encoding hypervector and associative search can be computed in parallel. Nonetheless, Equation (4) reveals that we need ~ 820 LUTs for a single popcount module of a typical benchmark (*isolet* [21] with 617-feature inputs). This limits the upper-bound of encoding dimensions we can generate at each cycle.

To save resources, [18] and [19] suggest using counters to implement the popcount of each dimension, as shown in Figure 2 (b). It seems a compact implementation but in practice it is less efficient than an adder-tree implementation: a counter needs “ $\log d_{iv}$ ” LUTs and has a latency of d_{iv} cycles, while adder-trees require $O(\frac{4}{3} d_{iv})$ LUTs per dimension with a per-dimension throughput of one cycle. Therefore, the typical adder-tree is $\frac{3}{4} \log d_{iv} \times$ more performance-efficient.

Figure 3: Our proposed approximate encoding techniques. MAJ and $\oplus$ denote majority and addition, respectively.

Work in [16] and [17] quantize the dimensions of encoding and class hypervectors which eliminates DSP modules (or large number of cascaded LUTs) that are conventionally used for the associative search stage, as through quantization the inner product for cosine similarity will be replaced by popcount or lower-bit multiplications. The resulting improvement is minor because the quantization is applied *after* full-bit encoding, so the encoding remains complex. Furthermore, it can be shown that the associative search stage requires considerably less resources and has a limited impact on total improvement (e.g., for the aforementioned *isolet*, associative search contributes to $\sim 23\%$ of resources). Thus, we target the popcount portion that contributes to the major part of resources. Indeed, *ex-post* quantizing of encoding hypervectors can be orthogonal to our technique for further improvement.

3 PROPOSED METHOD: SHEARer

3.1 Approximate Encoding

In the previous section, we explained prior work that applies quantization after obtaining the encoding hypervector in full bit-width. As noted there, while this approach is simple it only accelerates the associative search phase and does not improve encoding - which is often the principal bottleneck. Because the HD representation of data entails substantial redundancy and information is uniformly distributed over a large number of bits, it is robust to bit-level errors: flipping 10% of hypervectors' bits shows virtually zero accuracy drop, while 30% bit-error impairs the accuracy by a mere 4% [22]. We leverage such resilience to improve the resource utilization through approximate encoding, as shown in Figure 3. In the following, we discuss each technique in greater detail and estimate its resource usage.

(1) Local majority. From Equation (4) we can observe that the number of resources (in terms of LUT-6) of the exact adder-tree to see that the complexity encoding each dimension linearly depends on the number of data features, d_{iv} . We, therefore, aim to reduce the number of inputs to the primary adder-tree by sub-sampling using the majority function so as to shrink the tree inputs while (approximately) extracting the information contained in the input. Note that, here, 'inputs' are the binary dimensions of the level hypervectors (see Figure 1 ② and Figure 2). As shown in Figure 3(a), each LUT-6 is configured to return the majority of its six input bits. When three out of six inputs are 0/1, we break the tie by designating *all* LUTs that perform majority functions of a specific encoding dimension to deterministically output 0 or 1. We specify this randomly for every dimension (i.e., an entire adder-tree) but it remains fixed for a model during the training and inference. We choose groups of six bits as a single LUT-6 can vote for up to six inputs. Using smaller majority groups diminishes the resource saving, especially taking the majorities adds extra LUTs. Moreover, following the Shannon decomposition, implementing a ' $k+1$ '-input LUT requires two k -input LUTs (and a two-input multiplexer). Thus, the number of LUTs for majority groups larger than six inputs grows exponentially.

There are $\frac{d_{iv}}{6}$ MAJ LUTs in the first stage of Figure 3(a), hence the number of inputs for the subsequent adder-tree reduces to $\frac{d_{iv}}{6}$. From

Equation (4) we also know that a k -input adder-tree requires $\frac{4}{3}k$ LUT-6. Thus, the design of Figure 3(a) consumes

$$\underbrace{\frac{d_{iv}}{6}}_{\text{MAJ LUT-6}} + \underbrace{\frac{4}{3} \frac{d_{iv}}{6}}_{\text{adder-tree}} = \frac{7}{18} d_{iv} \text{ LUT-6} \quad (5)$$

which uses $1 - \frac{7/18}{4/3} = 70.8\%$ less LUT resources than an exact adder-tree.

In [17], the authors report an average accuracy loss of 1.6% by *post-hoc* quantizing the encodings to binary. Thus, one might think of repeating the majority functions in the subsequent stages to obtain final one-bit encoding dimensions. Using local majority functions is efficient, but degrades the encoding quality as majority is not associative. In particular, the MAJ LUTs add another layer of approximation by breaking ties. Thus, a so-called MAJ-tree causes considerable accuracy loss. Therefore, in our cascaded-MAJ design in Figure 3(b), we limit the MAJ stages to the first two stages. Our cascaded-MAJ utilizes:

$$\underbrace{\frac{d_{iv}}{6}}_{\text{1st stage MAJs}} + \underbrace{\frac{d_{iv}/6}{6}}_{\text{2nd stage MAJs}} + \underbrace{\frac{4}{3} \frac{d_{iv}/6}{6}}_{\text{adder-tree}} = \frac{25}{108} d_{iv} \text{ LUT-6} \quad (6)$$

which saves $1 - \frac{25/108}{4/3} = 82.6\%$ resources compared to exact encoding. We emphasize that a cascaded all-MAJ popcount needs $\approx \sum_{i=1}^1 \frac{1}{6^i} = 0.2 d_{iv}$ LUTs, which saves 85.0% of LUTs. So the two-stage MAJ implementation with 82.6% resource saving is nearly optimal because the first two stages of the exact tree were consuming the most resources.

(2) Input overfeeding. In Figure 2(a) we can observe that each LUT-5 pair of the first stage computes $\overline{s_1 s_0} = hv_i + hv_{i+1} + hv_{i+2}$. Since only three of their inputs are used, these LUTs left underutilized. With one more input, the output range will be $[0-3]$, which requires three bits (outputs) to represent, so we cannot add more than three bits using two LUT-5s. However, instead of using the LUT-5s to carry out regular addition, we can supply a pair of LUT-5s with five inputs to perform quantized/truncated addition. For actual outputs (sum of five bits) of 0 or 1, the LUT-5 pair would produce 00 (zero); for 2 or 3 they produce 01 (one), and for 4 or 5 they produce 10 (two). That is one LUT-5 computes the actual carry out of the five bits, and the other computes MSB of the sum. To ensure the synthesis tool infers a single LUT-6 for each pair, we can directly instantiate LUT primitives. As a LUT-6 comprises a pair of LUT-5, the number of resources of Figure 3(c) is:

$$\sum_{i=1}^{\log d_{iv}} \frac{d_{iv}}{5} \times \frac{i}{2^{i-1}} \approx \frac{4}{5} d_{iv} \text{ LUT-6} \quad (7)$$

The first stage comprises $\frac{d_{iv}}{5}$ LUT-6s, and each subsequent stage contains i -bit adders while their count decreases by $\frac{1}{2} \times$ at each stage. Total number of LUTs is reduced by $1 - \frac{4/5}{4/3} = 40\%$ (same as input over-using ratio). The saving is less than the local majority approach but we expect higher accuracy as this approximation is intuitively more moderate.

(3) Truncated nodes. Out of $\frac{4}{3}d_{iv}$ LUTs used in an exact adder-tree, d_{iv} (75%) are used in the intermediate adder units. More precisely, following $\frac{i}{2^i}$ ratio (see Equation (4)), stages 1–4 of the adder contribute to 25%, 25%, 18.75%, and 12.5% of total resources, respectively. Note that the number of adder units halves at each stage but the area of each one increases linearly. We avoid a blowup of adder sizes by truncating the least significant bit (LSB) of each adder. As demonstrated in Figure 3(d), the LSB of the second stage (which is supposed to have three-bit output) is discarded. Thus, instead of using two LUT-6s to compute $\overline{s_2s_1s_0} = \overline{a_1a_0 + b_1b_0}$, we can use two LUT-5s (equivalent to one LUT-6) to obtain $\overline{s_2s_1} = \overline{a_1a_0 + b_1b_0}$, where one LUT-5 computes s_2 and the other produces s_1 using four inputs a_0, a_1, b_0 , and b_1 . Truncating the output of the second stage consequently decreases the output bit-width of the third stage by one bit as its inputs became two bits. Thus, we can apply the LSB truncating to the third stage to implement it using two LUT-5s, as well. We can apply the same procedure in all the consecutive nodes and implement them by only two LUT-5s. The output of the first stage is already two bits so we do not modify its original implementation.

We apply truncating to first stages, especially from the left side of Equation 4 we can perceive the first five stages that contribute to ~90% of the adder-tree resources. Otherwise, the accuracy loss becomes severe. Equation (8) formulates the resource usage of the adder-tree in which the first k stages are implemented using 2-bit adders shown in Figure 3(d) (including the stage one, which uses the exact mode).

$$\underbrace{\sum_{i=1}^k \frac{d_{iv}}{3} \frac{1}{2^{i-1}}}_{\text{the first } k \text{ stages}} + \underbrace{\sum_{i=k+1}^{\log d_{iv}} \frac{d_{iv}}{3} \frac{i+1-k}{2^{i-1}}}_{\text{subsequent stages}} \simeq \frac{d_{iv}}{3} \left(2 + \frac{4}{2^k}\right) \text{ LUT-6} \quad (8)$$

We can see that for $k = 1$, i.e. when none of intermediate stages are truncated, the equation returns $\frac{4}{3}d_{iv}$ (equal to exact adder-tree). $k = 2, 3, 4$ achieves 25%, 37.5%, and 43.75% resource saving, respectively.

3.2 SHEARer Architecture

Recall from Figure 1 that the HD encoding procedure needs to convert all input features to equivalent level hypervectors, bind them with the associated ID hypervector, and bundle (e.g. sum) the resulting hypervectors to generate the encoding. FPGAs, however, contain limited logic and on-chip memory blocks (a.k.a BRAMs) to provide high performance with affordable power. Previous work break down encoding into multiple cycles whereby at each cycle they process d_{seg} dimensions [15, 16, 23]. When processing dimensions $n \cdot d_{seg}$ to $(n+1) \cdot d_{seg}$, those architectures fetch the same dimensions of all $\mathcal{L}$ level hypervectors. Each of d_{seg} adder-trees are augmented with $\mathcal{L}$ -to-1 multiplexers in all of their d_{iv} input ports, where the $k^{\text{th}} \leq d_{seg}$ adder-tree's multiplexers are connected to k^{th} dimension of the fetched level hypervectors, and the (quantized) value of associated feature selects the right level dimension to pass. Therefore only $d_{seg} \cdot \mathcal{L}$ bits are need to be fetched at each cycle. However, it requires $d_{iv} \cdot d_{seg}$ multiplexers. For a modest $\mathcal{L} = 16$ which translates to 16-input multiplexers each of which needs four LUTs, the multiplexers will consume total $4 \cdot d_{iv} \cdot d_{seg}$ LUTs, while the adder-trees occupy $d_{seg} \cdot \frac{4}{3}d_{iv}$ (remember a d_{iv} input adder-tree uses $\frac{4}{3}d_{iv}$ LUTs). This means that the augmented multiplexers occupy $3 \times$ LUTs of the adder area. In our approximate encoding, this ratio would be even larger as we trim the exact adder. Thus, multiplexer-based implementation overshadows the gain of approximating the adders.

To address this issue, we propose a novel FPGA implementation that relies on on-chip memories rather than adding extra resources. Figure 4 illustrates the overview of SHEARer architecture. At each cycle, we partially process $\mathcal{F}$ (out of d_{iv}) input features, where $\mathcal{F} \leq d_{iv}$. Our implementation is BRAM-oriented, so each (quantized) feature translates to an address from which the corresponding level hypervector can be read. This entails a dedicated memory block group for each of $\mathcal{F}$ features being processed simultaneously. The number of BRAMs

Figure 4: A scheme of SHEARer datapath.

in a group (group size) is equal to $\frac{\mathcal{L} \cdot d_{hv}}{C_{bram}}$ as there are $\mathcal{L}$ different level hypervectors of length d_{hv} bits, for a memory capacity of C_{bram} bits. Therefore, the number of features $\mathcal{F}$ that can be partially processed in a cycle is limited to $\mathcal{F} < 2 \frac{\text{total BRAMs}}{\text{group size}}$. The coefficient 2 is because the BRAMs have two ports from which we can independently read (that is why in Figure 4 two pixels share the same BRAM group). The address translator ("level to address" in Figure 4) activates only the right BRAM row of the group, so the other BRAMs do not dissipate dynamic power. Depending on its configuration, each memory block can deliver up to d_{mem} bits, as indicated in the figure. We can process more dimensions per cycle by increasing the size of memory groups but then $\mathcal{F}$ – the number of features being processed per cycle – reduces proportionally.

Each of d_{mem} fetched level hypervector bit is XORed with the corresponding bit of the ID (position) hypervector. As detailed in Section 2.1, each feature *index* is associated with an ID hypervector, which is a randomly chosen (but fixed) hypervector of length d_{hv} . We thus require $\frac{d_{hv} \cdot d_{iv}}{C_{bram}}$ additional BRAM blocks to store ID hypervectors which further limits the parallelism by exhausting BRAMs. To resolve this, we only store a single ID hypervector (seed ID) and generate the other ones by rotating the seed ID, i.e. ID of index k can be obtained by rotating the ID of index 1 (seed ID) by $k-1$. This does not affect the HD accuracy as the resulting ID hypervectors are still *iid* and approximately orthogonal. For the first feature, we need to read d_{mem} bits, while for the subsequent $\mathcal{F} - 1$ features we need one more bit as each ID has $d_{mem} - 1$ common bits with its predecessor. Therefore we need a data-width of $d_{mem} + \mathcal{F} - 1$ for ID memory, meaning that we need $1 + \frac{\mathcal{F}}{d_{mem}}$ memory blocks of the seed ID hypervector. That is, the seed ID fits in a single BRAM but the required data-width demands more memory blocks. However, it is still significantly smaller than the case of storing all different IDs in BRAM blocks and either releases BRAMs for processing more features, or power gates the unused BRAMs. Moreover, using seed ID BRAM also saves dynamic power as $d_{mem} + \mathcal{F} - 1$ bits are read (compared to $d_{mem} \times \mathcal{F}$ of storing different IDs). It is also noteworthy that at each cycle the first d_{mem} bits read from the ID memory are passed to the first feature of the features currently being processed (i.e., feature 1, $\mathcal{F} + 1, 2\mathcal{F} + 1, \dots$). Similarly, bits 2 to $d_{mem} + 1$ of the fetched ID are passed to the second feature, and so on. Thus, connecting ID BRAMs to processing logic needs a fixed (multiplexer-free) routing.

After XORing the fetched level hypervectors with the ID hypervectors, each of the d_{mem} approximate adder-trees add $\mathcal{F}$ binary bits, so the input size of all adders is $\mathcal{F}$. Since the result is only the sum of the first $\mathcal{F}$ features, we use a buffer to store partial sums. In the next cycle, the procedure repeats for the next group of features, i.e., features $\mathcal{F} + 1$ to $2\mathcal{F}$. Therefore, SHEARer produces d_{mem} encoding dimensions in $\frac{d_{iv}}{\mathcal{F}}$ cycles, hence the entire encoding hypervector is generated in $\left\lceil \frac{d_{hv}}{d_{mem}} \right\rceil \times \left\lceil \frac{d_{iv}}{\mathcal{F}} \right\rceil$ cycles.

To make these tangible, in the Xilinx FPGAs we use for experiments, d_{mem} is 64 and $C_{bram} = 512_{row} \times 64_{col}$. We also noticed that 16 level hypervectors gives the same accuracy of having more, so we set $\mathcal{L} = 16$. We also select the hypervector lengths to be a multiple of 512. Taking the

previously mentioned language recognition benchmark *isolet* [21] as an example, we observed that $d_{hv} = 2,560$ provides acceptable accuracy (see Section 4 for more details). For this benchmark we thus need group size of $\lceil \frac{16 \times 2560}{512 \times 64} \rceil = 2$ BRAMs, where each group covers two input features. The FPGA we use has a total 445 BRAMs and can make at most $\lfloor \frac{445}{2} \rfloor = 222$ groups, capable of processing 444 features per cycle. Therefore, we divide 617 input features of the benchmark into two repeating cycles using 310 BRAMs (155 BRAM groups) to process the first 310 features in the first cycle, and the rest 307 cycles in the second cycle, generating $d_{mem} = 64$ encoding dimensions per 2 cycles. All 64 adder-trees have 310 1-bit inputs. An encoding takes $2560 \text{ dim} \times \frac{2 \text{ cycles}}{64 \text{ dim}} = 80$ cycles. Note that on-chip BRAMs have one-cycle latency and the off-chip memory latency is buried in the computation pipeline.

3.3 Software Layer

Because of approximation, the encoding and hence the class hypervectors become different than training with exact encoding. Therefore we need to train the model using the same approximate encoding(s), since the associative search just looks for the *similarity* (rather than exactness) of an approximately encoded hypervector with trained class hypervectors. Our FPGA implementation is tailored for inference, so we carry out the training step on CPU. We developed an efficient SIMD vectorized Python implementation to emulate the exact and approximate techniques in software. Emulation of the proposed techniques is straightforward. For instance, for the local majority approximation (Figure 3(a)), instead of adding up all d_{iv} hypervectors, we divide them to groups of six hypervectors, add up all six hypervectors of each group, and compare if each resultant dimension is larger than 3. We also break the ties in software by generating a constant vector dictating how the ties of each dimension should be served. This acts as the MAJ LUTs of the first stage. Thereafter, we simply add up all these temporary hypervectors to realize the subsequent exact adders. This guarantees matching of software output with approximate hardware, while achieving a fast implementation by avoiding cycle-level imitating of hardware.

In addition to d_{iv} that is the dataset’s attribute, d_{hv} , α , and training epochs are the other parameters of software implementation. α is the learning rate of HD. As explained in Section 1, HD bundles all encoding hypervectors belonging to the same-label data to create the initial class hypervectors. In the subsequent epochs iterations, HD updates the class hypervectors by observing if the model correctly predicts the training data. If the model mispredicts an encoded query $\mathcal{H}^l$ of label l as class C^l , HD updates as shown by Equation (9). If learning rate α is not provided, SHEARer finds the best α through bisectioning for a certain number of iterations.

$$C^l = C^l + \alpha \cdot \mathcal{H}^l \quad C^l = C^l - \alpha \cdot \mathcal{H}^l \quad (9)$$

We supply the software implementation of SHEARer with the number of BRAM and LUT resources of the target FPGA to estimate the architectural parameters according to Section 3.2 as well as using the resource utilization formulated in Section 3.1. We have also implemented the exact and approximate adder-trees of different input sizes and interpolated their power consumption – which is linear w.r.t. the adder size – for different switching activities. Therefore, we calculate the average signal activity observed by the adders according to the values of temporary-generated binding hypervectors (level XOR ID). We similarly estimate the toggle rate of BRAMs according to consecutive d_{mem} bits read from BRAMs. As alluded earlier, we do not replicate the hardware implementation in software; we just need to determine each fetched level hypervector belongs to which BRAM group (based on the index of feature), so we can keep track of toggle rates. By using the signal information with an offline look-up table created for activity-power, along with the information of instantiated resources calculated as mentioned, SHEARer is able to estimate the power consumption of an application targeted for a specific device.

Table 1: Baseline (exact) implementation results.

Parameter ↓	Benchmark →	speech	activity	face	digit
Input features (d_{iv})		617	561	608	784
Hypervector length (d_{hv})		2,560	3,072	6,144	2,048
Baseline accuracy		93.18%	93.91%	95.47%	89.07%

Table 2: LUT count for a 512-input adder-tree.

	exact	MAJ	MAJ-2	over-feed	truncate
Synthesis	638	183	116	383	340
Equation	675	195	116	405	343
Error	5.8%	6.6%	0.0%	5.7%	0.9%

Table 3: Relative accuracies SHEARer approximate encodings.

	exact	MAJ	MAJ-2	over-feed	trunc-3	trunc-4
speech	93.2%	−0.7%	−2.3%	−0.8%	−0.9%	−1.9%
activity	93.9%	−0.8%	−1.2%	−1.3%	−1.1%	−1.0%
face	95.5%	−1.8%	−3.3%	−1.7%	−1.6%	−1.9%
digit	89.1%	−0.8%	−0.3%	−1.7%	0.1%	−0.1%
average		−1.0%	−1.8%	−1.4%	−0.9%	−1.2%
LUT saving	0	71.1%	82.8%	40.0%	37.5%	43.8%

4 EXPERIMENTAL RESULTS

(1) General Setup. We have implemented the SHEARer architecture using Vivado High-Level Synthesis Design Suite on Xilinx Kintex-7 FPGA KC705 Evaluation Kit which embraces a XC7K325T device with 203,800 LUT-6 and 445 36 Kb BRAM memory blocks. By pipelining the adder-tree stages we achieve a clock frequency of 200 MHz. We compare the performance and energy results with the high-end NVIDIA GeForce GTX 1080 Ti GPU, and Raspberry Pi 3 embedded processor. We optimize the CUDA implementation by packing the hypervectors as 32-bit integers, so a single logical XOR operation can bind 32 dimensions. We use speech [21], activity [24], and MNIST hand-written digit recognition as well as a face detection dataset [25] as benchmarks. Table 1 summarizes the length of hypervectors and accuracy of each dataset in the exact mode. For a fair comparison, we obtained the accuracies using $d_{hv} = 10,000$, then reduced d_{hv} until the accuracies remain within 0.5% of the original values. This avoids over-saturated hypervectors and better reveals the accuracy drop due to approximation.

(2) Resource Utilization. To validate the efficiency of the proposed approximation techniques, besides holistic performance and energy comparisons, we examine them by synthesizing a 512-input adder-tree. Table 2 represents the LUT utilization of the adder implemented in exact and approximate modes. MAJ, MAJ-2, over-feed and truncate refer to the designs of Figure 3(a)–(d). We see that our equations in Section 3.1 have a modest average error of 3.8%. For instance, synthesis results indicate MAJ (MAJ-2) saves 71.3% (81.8%) LUTs, which is very close to the predicted 71.1% (82.8%).

(3) Accuracy. Table 3 compares the accuracies of the proposed encodings relative to the exact encoding. LUT saving, which is dataset-independent, is represented again for the comparison purpose. “trunc-3” and “trunc-4” stand for truncated encoding (Figure 3(d)) where, respectively, three and four intermediate stages are truncated. Overall, MAJ encoding (Figure 3(a)) achieves an acceptable accuracy with significant resource saving, though it is not always the highest-accurate one. For instance, in the face detection benchmark, the over-feed and 3-stage truncated encodings offer slightly better accuracy. More interestingly, in the digit recognition dataset, trunc-4 shows a negligible −0.1% accuracy drop while trunc-3 even improves the accuracy by 0.1%. This is because emulating the approximation takes a long time for the digit dataset, so we limited SHEARer to try five different learning rates (α) and repeat the entire training for five times (epochs = 50) so the result might be slightly skewed. For the other datasets we conducted the training for 25 times each with 50 epochs to average out the variance.

(4) Performance. Figure 5 compares the throughput of SHEARer FPGA implementation with Raspberry Pi and Nvidia GPU. SHEARer implementation is BRAM-bound, so all the exact and approximate implementations yield the same performance. In Section 3.2 we elaborated that the speech dataset requires two cycles per $d_{mem} = 64$ dimensions.

Figure 5: Throughput of SHEARer versus Raspberry Pi 3 and Nvidia GTX 1080 Ti. Y-axis is logarithmic scale.

Figure 6: Energy (Joule) consumption of SHEARer, Raspberry Pi and GPU for 10 million inference. Y-axis is logarithmic.

We can similarly show that activity and digit datasets also need two cycles per 64 dimensions, while digit requires three cycles as its level hypervectors are larger ($d_{ho} = 6,144$) and occupy more BRAMs. In the worst scenario, SHEARer improves the throughput by 58,333 $\times$ and 6.7 $\times$ compared to Raspberry Pi and GPU implementation. On average SHEARer provides a throughput of 104,904 $\times$ and 15.7 $\times$ as compared to Raspberry Pi and GPU, respectively. The substantial improvements arise from that SHEARer adds up $\frac{d_{ho}}{2} \times 64$ (e.g., $\sim 25,000$) numbers per cycle while also performs the binding (XOR operations) on the fly. However, Raspberry Pi executes sequentially and also its cache cannot fit all the class hypervectors with non-binary dimensions. We assume the dataset is available in the *off-chip* memory (DRAM) of FPGA. Otherwise, per-input *latency* increases but throughput remains the same as off-chip memory latency is buried in the computation cycles.

(5) Energy Consumption. Figure 6 compares the energy consumption of the exact and approximate implementations with Raspberry Pi and GPU. We have scaled the energy to 10 million inferences for the sake of illustration (Y-axis is logarithmic). We used Hioki 3334 power meter and NVIDIA system management interface to measure the power consumption of Raspberry Pi and GPU, respectively. We used Xilinx Power Estimator (XPE) to estimate the FPGA power consumption. The average power of Raspberry Pi for all datasets hovers around 3.10 Watt, while this is ~ 120 Watt for the GPU. In FPGA implementation, powers showed more variation as the number of active LUTs and BRAMs differ between applications. E.g., the face dataset with two-stage majority encoding (MAJ-2) consumes 3.11 Watt, while the digit recognition in the exact mode consumes 10.80 Watt. The smaller power consumption of face is mainly because of smaller off-chip data transfer as face has the largest hypervector length and takes 288 cycles to process an entire input, while for digit it takes 64 cycles. On average, our *exact* encoding decreases the energy consumption by 45,988 $\times$ and 247 $\times$ (average of all datasets) as compared to Raspberry Pi and GPU implementations. MAJ-2 encoding of SHEARer consumes the minimum energy, which throttles the energy consumption by 56,044 $\times$ and 301 $\times$ compared to Raspberry Pi and GPU, respectively. Note that power improvement of the approximate encodings is not proportional to their resource (LUT) utilization as BRAM power remains the same for all encodings.

5 CONCLUSION

We leveraged the error resiliency of HD to develop different approximate encodings with varied accuracy and resource utilization attributes. With a modest 1.0% accuracy drop, our approximate encoding reduces the LUT utilization by 71.1%. By effectively utilizing FPGA on-chip

BRAMs, we proposed a highly efficient implementation that outperforms an optimized GPU implementation over 15 $\times$, and surpasses Raspberry Pi by over five orders of magnitude. Our FPGA implementation consumes a moderate power: as low as 3.11 Watt for a face detection dataset using approximate encoding, and a maximum of 10.8 Watt on a digit recognition dataset when using exact encoding. Eventually, our implementation reduces the energy consumption by 247 $\times$ (45,988 $\times$) compared to GPU and Raspberry Pi in exact encoding, which further improves by a factor of 1.22 $\times$ using approximate encoding.

ACKNOWLEDGEMENTS

This work was supported in part by CRISP, one of six centers in JUMP, an SRC program sponsored by DARPA, in part by SRC Global Research Collaboration (GRC) grant, DARPA HyDDENN grant, and NSF grants #1911095 and #2003279.

REFERENCES

- [1] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors," *Cognitive computation*, vol. 1, no. 2, pp. 139–159, 2009.
- [2] G. C. Turner, M. Bazhenov, and G. Laurent, "Olfactory representations by drosophila mushroom body neurons," *Journal of Neurophysiology*, vol. 99, no. 2, pp. 734–746, 2008.
- [3] R. I. Wilson, "Early olfactory processing in drosophila: mechanisms and principles," *Annual Review of Neuroscience*, vol. 36, pp. 217–241, 2013.
- [4] T. A. Plate, "Holographic reduced representations," *IEEE Transactions on Neural networks*, vol. 6, no. 3, pp. 623–641, 1995.
- [5] M. Wan, A. Jönsson, C. Wang, L. Li, and Y. Yang, "Web user clustering and web prefetching using random indexing with weight functions," *Knowledge and information systems*, vol. 33, no. 1, pp. 89–115, 2012.
- [6] Y. Kim, M. Imani, N. Moshiri, and T. Rosing, "Geniehd: efficient dna pattern matching accelerator using hyperdimensional computing," in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 115–120, IEEE, 2020.
- [7] F. Asgarinejad, A. Thomas, and T. Rosing, "Detection of epileptic seizures from surface eeg using hyperdimensional computing," in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2020.
- [8] A. Mitrokhin, P. Sutor, C. Fermüller, and Y. Aloimonos, "Learning sensorimotor control with neuromorphic sensors: Toward hyperdimensional active perception," *Science Robotics*, vol. 4, no. 30, p. eaaw6736, 2019.
- [9] P. Neubert, S. Schubert, and P. Protzel, "An introduction to hyperdimensional computing for robotics," *KI-Künstliche Intelligenz*, vol. 33, no. 4, pp. 319–330, 2019.
- [10] M. Imani, Y. Kim, S. Riaz, J. Messerly, P. Liu, F. Koushanfar, and T. Rosing, "A framework for collaborative learning in secure high-dimensional space," in *2019 IEEE 12th International Conference on Cloud Computing (CLOUD)*, pp. 435–446, IEEE, 2019.
- [11] B. Khaleghi, M. Imani, and T. Rosing, "Prive-hd: Privacy-preserved hyperdimensional computing," in *Proceedings of the 57th Annual Design Automation Conference*, 2020.
- [12] M. Imani, J. Morris, J. Messerly, H. Shu, Y. Deng, and T. Rosing, "Bric: Locality-based encoding for energy-efficient brain-inspired hyperdimensional computing," in *Proceedings of the 56th Annual Design Automation Conference 2019*, p. 52, ACM, 2019.
- [13] A. Rahimi, P. Kanerva, and J. M. Rabaey, "A robust and energy-efficient classifier using brain-inspired hyperdimensional computing," in *Proceedings of the 2016 International Symposium on Low Power Electronics and Design*, pp. 64–69, 2016.
- [14] M. Imani, D. Kong, A. Rahimi, and T. Rosing, "Voicehd: Hyperdimensional computing for efficient speech recognition," in *2017 IEEE International Conference on Rebooting Computing (ICRC)*, pp. 1–8, IEEE, 2017.
- [15] M. Imani *et al.*, "Sparsehd: Algorithm-hardware co-optimization for efficient high-dimensional computing," in *IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pp. 190–198, IEEE, 2019.
- [16] S. Salamat, M. Imani, B. Khaleghi, and T. Rosing, "F5-hd: Fast flexible fpga-based framework for refreshing hyperdimensional computing," in *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pp. 53–62, 2019.
- [17] M. Imani, S. Bosch, S. Datta, S. Ramakrishna, S. Salamat, J. M. Rabaey, and T. Rosing, "Quanthd: A quantization framework for hyperdimensional computing," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2019.
- [18] M. Schmuck, L. Benini, and A. Rahimi, "Hardware optimizations of dense binary hyperdimensional computing: Rematerialization of hypervectors, binarized bundling, and combinational associative memory," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 15, no. 4, pp. 1–25, 2019.
- [19] M. Imani, J. Messerly, F. Wu, W. Pi, and T. Rosing, "A binary learning framework for hyperdimensional computing," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 126–131, IEEE, 2019.
- [20] S. Datta, R. A. Antonio, A. R. Ison, and J. M. Rabaey, "A programmable hyperdimensional processor architecture for human-centric iot," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 9, no. 3, pp. 439–452, 2019.
- [21] "Uci machine learning repository." <http://archive.ics.uci.edu/ml/datasets/ISOLET>.
- [22] M. Imani, A. Rahimi, D. Kong, T. Rosing, and J. M. Rabaey, "Exploring hyperdimensional associative memory," in *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 445–456, IEEE, 2017.
- [23] M. Imani, S. Salamat, S. Gupta, J. Huang, and T. Rosing, "Fach: Fpga-based acceleration of hyperdimensional computing by reducing computational complexity," in *Proceedings of the 24th Asia and South Pacific Design Automation Conference*, pp. 493–498, 2019.
- [24] "Uci machine learning repository." <https://archive.ics.uci.edu/ml/datasets/human+activity+recognition+using+smartphones>.
- [25] G. Griffin, A. Holub, and P. Perona, "Caltech-256 object category dataset," 2007.