

Late Breaking Results: Scalable and Efficient Hyperdimensional Computing for Network Intrusion Detection

Junyao Wang[†], Hanning Chen[†], Mariam Issa[†], Sitao Huang[§], Mohsen Imani[†]

[†] Department of Computer Science, University of California, Irvine, CA, United States

[§] Department of Electrical Engineering and Computer Science, University of California, Irvine, CA, United States
 {junyaow4, hanningc, mariamai, sitaoh, m.imani}@uci.edu

Abstract—Cybersecurity has emerged as a critical challenge for the industry. With the large complexity of the security landscape, sophisticated and costly deep learning models often fail to provide timely detection of cyber threats on edge devices. Brain-inspired hyperdimensional computing (HDC) has been introduced as a promising solution to address this issue. However, existing HDC approaches use static encoders and require very high dimensionality and hundreds of training iterations to achieve reasonable accuracy. This results in a serious loss of learning efficiency and causes huge latency for detecting attacks. In this paper, we propose CyberHD, an innovative HDC learning framework that identifies and regenerates insignificant dimensions to capture complicated patterns of cyber threats with remarkably lower dimensionality. Additionally, the holographic distribution of patterns in high dimensional space provides CyberHD with notably high robustness against hardware errors.

I. INTRODUCTION

The rapid development of information technology has introduced increasingly sophisticated cyber threats. As one of the most widely deployed security devices, network intrusion detection systems (NIDS) is designed to monitor network traffic and identify suspicious activity. As shown in Fig. 1(a), when traditional firewalls fail to intercept intruders, NIDS are aimed at providing timely detection and alerts to prevent the spread of infection through local area networks. With an increasingly gigantic amount of network traffic nowadays, along with the constant evolution of cyber attacks, a more intelligent, robust, and fully automated NIDS framework is of absolute necessity.

Hyperdimensional Computing (HDC) is considered a promising solution for NIDS for its (i) high computational efficiency ensuring real-time attack detection, (ii) holographic pattern distributions offering ultra-robustness against failures, and (iii) lightweight hardware implementations allowing efficient execution on edge devices [1]. As shown in Fig. 1(b), closely mimicking information representation and memorization functionalities of human brains, HDC encodes low-dimensional inputs to hypervectors with more than 10^4 elements to perform various learning tasks. In this way, HDC conducts highly parallelizable operations and achieves high-quality results with significantly faster convergence and higher efficiency.

Besides its lightweight and robust nature, utilizing encoded data points in hyperspace, HDC exhibits outstanding capabilities to distinguish various sophisticated attack patterns. Unfortunately, existing HDC algorithms use pre-generated encoders lacking the capability to distinguish the importance of each dimension, and hence require extremely high dimensionality to outperform DNNs. Consequently, the learning efficiency and scalability are compromised due to large numbers of unnecessary computations, especially for network intrusion detection tasks analyzing billions of network traffic instances. This work aims at addressing this issue by proposing CyberHD, a novel HDC learning framework utilizing a dynamic HDC encoding technique that identifies and regenerates insignificant dimensions. It ensures a highly effective dimensionality and achieves the desired accuracy with significantly higher efficiency.

II. RELATED WORK

A recently proposed HDC framework, NeuralHD [2], demonstrated a dynamic encoding approach by eliminating dimensions with minor impacts on classifying patterns. Our work optimizes and parallelizes this dynamic HDC framework with matrix operations and applies it

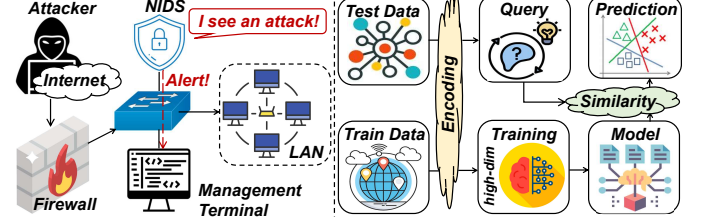


Fig. 1. An Overview of NIDS and HDC Classification




Fig. 2. An Overview of CyberHD Workflow




Fig. 3. Comparing Accuracy of CyberHD with State-of-the-art Algorithms to the domain of cyber security for the first time, aiming to deliver real-time effective attack detection on resource-constrained devices.

III. METHODOLOGY

As shown in Fig. 2, CyberHD starts with encoding data points into high-dimensional space with existing encoding methods (A) depending on the data type. CyberHD then leverage matrix operations to train the encoded data (B) in a highly-parallel way. We then normalize each class hypervector (D) and calculate the variance of each dimension over all classes to identify dimensions with minimum impacts on the classification accuracy (F). CyberHD then drops insignificant dimensions from our model depending on the regeneration rate (G). Finally, CyberHD updates the base vectors on the selected dimensions and retrains the model (H).

HDC Learning: After encoding training data onto hyperspace (A), to reduce model saturation, we bundle encoded data by scaling a proper weight to each of them depending on how much new information they bring to class hypervectors. For instance, for a new encoded training sample $\mathcal{H}$, we update the model base on its cosine similarities with all class hypervectors, i.e. $\delta_i = \frac{\mathcal{H} \cdot \vec{C}_i}{\|\mathcal{H}\| \cdot \|\vec{C}_i\|}$, where $\mathcal{H} \cdot \vec{C}_i$ is the dot product between $\mathcal{H}$ and a class hypervector $\vec{C}_i$. If $\vec{H}$ has the highest cosine similarity with class l' while it actually has label l , the model updates as $\vec{C}_l \leftarrow \vec{C}_l + \eta(1 - \delta_l) \times \vec{H}$ and $\vec{C}_{l'} \leftarrow \vec{C}_{l'} - \eta(1 - \delta_{l'}) \times \vec{H}$, where η is a learning rate. A large δ_l indicates the input data point is common or already exists in the




Fig. 4. (log scale) Comparing Training and Inference Efficiency of CyberHD with SOTA Algorithms

model and updates the model by adding a very small portion of the encoded query ($1 - \delta_l \approx 0$). In contrast, a small δ_l indicates a noticeably new pattern that is uncommon or does not already exist in the model so we update the model with a large factor ($1 - \delta_l \approx 1$).

Insignificant dimensions: HDC represents each class with a class hypervector encoding patterns of that class (C). In inference, a query vector (Q) is assigned to the class where it has the highest cosine similarity (I). An effective classifier achieves the desired accuracy with a strong capability to distinguish patterns so that query vectors have very differentiated cosine similarities to each class. In contrast, dimensions with similar values over all classes store common information and play minimal roles in differentiating patterns. CyberHD identifies and drops such dimensions by calculating variances over all classes. After computing variances (F) over the normalized model (E), CyberHD selects $\mathcal{R}\%$ of dimensions with the lowest variance to drop (G) depending on a regeneration rate $\mathcal{R}$.

Dimension Regeneration: To further improve classification accuracy, CyberHD regenerates the dimensions selected to drop (H), so that the new dimensions can potentially better differentiate patterns. For cybersecurity datasets, considering the non-linear relationship between features, we utilize an encoder inspired by the Radial Basis Function (RBF) [3]. During regeneration, CyberHD replaces the base vector of the selected dimension in the encoding module with another randomly generated vector from Gaussian Distribution.

IV. EXPERIMENTAL RESULT

We evaluate CyberHD with CPU (Intel Core i9-12900) and FPGA (Xilinx Alveo U50) on popular cyber-security datasets: NSL-KDD [4], UNSW-15 [5], CIC-IDS-2017 [6], and CIC-IDS-2018 [7]. We compare CyberHD with state-of-the-art (SOTA) DNNs [8] and SVMs [9]. We also compare CyberHD with SOTA HDCs [1] without the capability to regenerate dimensions (baselineHDC), and report results in two dimensionality: physical dimensionality ($\mathcal{D} = 0.5k$) of CyberHD, and effective dimensionality ($\mathcal{D}^* = 4k$) as CyberHD. The effective dimensionality is defined as the sum of the physical dimensionality with regenerated dimensions throughout the retraining.

Accuracy: As shown in Fig. 3, CyberHD provides comparable accuracy to SOTA DNNs, while on average a 1.63% higher accuracy than SVMs. CyberHD also shows on average 4.28% higher accuracy than baselineHDC ($\mathcal{D} = 0.5k$). Additionally, CyberHD delivers comparable accuracy to baselineHDC using the same dimensionality as the effective dimension ($\mathcal{D}^* = 4k$), indicating that CyberHD is capable of providing comparably high classification accuracy as the SOTA HDC while using $8.0\times$ lower physical dimensionality.

Efficiency: For fairness, we compare the training and inference efficiency of the SOTA DNN, SVM, baselineHDC ($\mathcal{D}^* = 4k$), and CyberHD ($\mathcal{D} = 0.5k$) as they achieve comparable accuracy as shown in Fig. 3. Since cybersecurity datasets generally include millions of samples, SVM algorithms take an extraordinarily long time for both training and inference. Additionally, as shown in Fig. 4, CyberHD delivers on average $2.47\times$ faster training than SOTA DNN. CyberHD also provides on average $1.85\times$ faster training and $15.29\times$ faster inference compared to the baselineHDC ($\mathcal{D}^* = 4k$).

Cross Platform and Quantization Evaluation: TABLE I shows the impact of the hypervectors' dimensions and bitwidths on the

TABLE I
IMPACT OF BITWIDTH ON CPU'S AND FPGA'S ENERGY EFFICIENCY

	32 bits	16 bits	8 bits	4 bits	2 bits	1 bit
Effective $\mathcal{D}$	1.2k	2.1k	3.6k	5.6k	7.5k	8.8k
CPU	$6.6\times$	$4.0\times$	$2.4\times$	$1.5\times$	$1.2\times$	$1.0\times$
FPGA	$16\times$	$24\times$	$34\times$	$31\times$	$28\times$	$26\times$

training efficiency, normalized to the efficiency of 1-bit CPU implementation. HDC models on CPUs achieve higher efficiency with low dimensionality and high element bitwidth due to the limited parallelism on CPU. CPUs demonstrate more strength for high bitwidth data due to their high frequency and powerful arithmetic logic unit (ALU). We also implement CyberHD on FPGA with different hypervector dimensions and element bitwidths for comparison. FPGA shows excellent energy efficiency improvement compared to CPU due to its high parallelism and low power consumption. On the Xilinx Alveo U50 FPGA board, the power consumption of the CyberHD accelerator is less than 20 W under 200 MHz frequency.

Robustness Against Hardware Failures: As shown in Fig. 5, in DNNs, random bit flip results in significant quality loss as corruptions on most significant bits can cause major weight changes. In contrast, CyberHD provides notably higher robustness against noise due to its redundant and holographic distribution. Additionally, all dimensions equally contribute to storing information so that failure on partial data will not result in the loss of entire information. CyberHD delivers the maximum robustness using hypervectors in 1-bit precision, which is on average $12.90\times$ higher than the robustness of the DNN. An increase in precision lowers the robustness of CyberHD since random flips on more significant bits will cause more accuracy loss.

Hardware Error	1.0%	2.0%	5.0%	10.0%	15.0%
DNN	3.9%	10.7%	17.8%	32.1%	41.2%
CyberHD 1 bit	0.0%	0.0%	1.0%	3.1%	4.1%
CyberHD 2 bits	1.9%	2.3%	4.5%	7.9%	10.4%
CyberHD 4 bits	2.3%	4.7%	8.4%	13.1%	17.3%
CyberHD 8 bits	3.6%	7.9%	13.7%	18.3%	22.9%

Fig. 5. Comparing Robustness of CyberHD with DNNs

V. ACKNOWLEDGEMENTS

This work was supported in part by National Science Foundation #2127780, Semiconductor Research Corporation (SRC), Office of Naval Research, grants #N00014-21-1-2225 and #N00014-22-1-2067, the Air Force Office of Scientific Research under award #FA9550-22-1-0253, and generous gifts from Xilinx and Cisco.

REFERENCES

- [1] Abbas Rahimi et al. A robust and energy-efficient classifier using brain-inspired hyperdimensional computing. In *ISLPED*, 2016.
- [2] Zhuowen Zou et al. Scalable edge-based hyperdimensional learning system with brain-like neural adaptation. In *SC*, 2021.
- [3] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 2007.
- [4] Mahbod Tavallaei et al. A detailed analysis of the KDD CUP 99 data set. In *CISDA*. IEEE, 2009.
- [5] Nour Moustafa and Jill Slay. UNSW-NB15: a comprehensive data set for network intrusion detection systems. In *MilCIS*. IEEE, 2015.
- [6] Iman Sharafaldin et al. Toward generating a new intrusion detection dataset and intrusion traffic characterization. *ICISSP*, 2018.
- [7] Joffrey Leevy et al. A survey and analysis of intrusion detection models based on CSE-CIC-IDS2018 big data. *Journal of Big Data*, 2020.
- [8] Hind Taud et al. Multilayer perceptron (mlp). In *Geomatic approaches for modeling land change scenarios*. Springer, 2018.
- [9] Marti A. Hearst et al. Support vector machines. *Intelligent Systems and their applications*, 1998.