

# HDTest: Differential Fuzz Testing of Brain-Inspired Hyperdimensional Computing

Dongning Ma<sup>‡</sup>, Jianmin Guo<sup>§</sup>, Yu Jiang<sup>§</sup>, Xun Jiao<sup>‡</sup>,  
<sup>‡</sup>Villanova University, <sup>§</sup>Tsinghua University

**Abstract**—Brain-inspired hyperdimensional computing (HDC) is an emerging computational paradigm that mimics brain cognition and leverages hyperdimensional vectors with fully distributed holographic representation and (pseudo)randomness. Compared to other machine learning (ML) methods such as deep neural networks (DNNs), HDC offers several advantages including high energy efficiency, low latency, and one-shot learning, making it a promising alternative candidate on a wide range of applications. However, the reliability and robustness of HDC models have not been explored yet. In this paper, we design, implement, and evaluate **HDTest** to test HDC model by automatically exposing unexpected or incorrect behaviors under rare inputs. The core idea of **HDTest** is based on guided differential fuzz testing. Guided by the distance between query hypervector and reference hypervector in HDC, **HDTest** continuously mutates original inputs to generate new inputs that can trigger incorrect behaviors of HDC model. Compared to traditional ML testing methods, **HDTest** does not need to manually label the original input. Using handwritten digit classification as an example, we show that **HDTest** can generate thousands of adversarial inputs with negligible perturbations that can successfully fool HDC models. On average, **HDTest** can generate around 400 adversarial inputs within one minute running on a commodity computer. Finally, by using the **HDTest**-generated inputs to retrain HDC models, we can strengthen the robustness of HDC models. To the best of our knowledge, this paper presents the first effort in systematically testing this emerging brain-inspired computational model.

## I. INTRODUCTION

Hyperdimensional computing (HDC) is an emerging computing scheme based on the working mechanism of brain that computes with deep and abstract patterns of neural activity instead of actual numbers. Compared with traditional ML algorithms such as DNN, HDC is more memory-centric, granting it advantages such as relatively smaller model size, less computation cost, and one-shot learning, making it a promising candidate in low-cost computing platforms [1]. Recently, HDC has demonstrated promising capability on various applications such as language classification [2], vision sensing [3], brain computer interfaces [4], gesture recognition [5], and DNA pattern matching [6]. However, despite the growing popularity of HDC, the reliability and robustness of HDC models have not been systematically explored yet.

The traditional approach to testing ML systems is to manually label a set of data and input them into the ML system to assess the accuracy. However, as the ML systems are scaling significantly and the input space is becoming more sophisticated, such manual approach is hardly feasible and scalable

anymore. On the other side, researchers have found that by adding even invisible perturbations to original inputs, ML systems can be “fooled” and produce wrong predictions [7].

Just like DNNs, HDC can also be vulnerable to small perturbations on inputs to produce wrong classification, as shown in Fig. I. Our attempt to enable automatic testing on HDC models faces the following challenges. First, unlike traditional ML systems such as DNNs with a well-defined mathematical formulation and relatively fixed architecture (specific layer types and network structure), HDC is less application-agnostic. The encoding of HDC is largely unique for different applications and relies on random indexing to project data onto vectors in a hyperdimensional space [2], adding difficulty to efficiently acquire adequate information to guide the testing process. As a result, adversarial generation techniques used in DNNs cannot be applied here since they rely on a set of well-defined mathematical optimization problems [8]. Second, the standard approach to testing ML systems is to gather and manually label as much real-world test data as possible [9], [10]. Google even used simulation to generate synthetic data [11]. However, such manual effort is not only largely short of scalability but also unguided as it does not consider the internal structure of ML systems, making it unable to cover more than a tiny fraction of all possible corner cases.

To address such challenges, we present **HDTest**, a highly-automated testing tool specifically designed for HDC based on differential fuzz testing. Fuzz testing is a software testing technique that strategically mutates inputs with the goal of generating faults or exceptions [12]. Assuming a greybox testing scenario, we leverage the unique property of HDC to guide and customize the fuzzing process. To avoid manual checking effort, we integrate fuzz testing with differential testing techniques [13], making it a highly automated and scalable testing solution. Our contributions are as follows:

- To the best of our knowledge, we present the first effort in systematically testing HDC models. Based on differential fuzz testing, **HDTest** iteratively mutates inputs to generate new inputs that can trigger incorrect behaviors of HDC models. Thanks to differential testing, **HDTest** does not require manual labeling effort.
- **HDTest** develops various mutation strategies to generate the inputs. **HDTest** leverages unique property of HDC and introduce the distance-guided fuzzing based on the distance between query vector and reference vector to improve fuzzing efficiency.
- Experimental results on MNIST dataset show that

Xun Jiao<sup>‡</sup> is the corresponding author




Fig. 1. An example of adversarial image of HDC by mutating some pixels in the image:(a) the original image as “8”; (b) the pixels mutated; (c) the mutated image wrongly predicted as “3”.

**HDTest** can generate thousands of adversarial inputs with invisible perturbations that can successfully fool the HDC models. On average, **HDTest** can generate 400 adversarial images within one minute running on commodity computer. We further show that the **HDTest**-generated inputs can be used to strengthen the robustness of HDC model.

## II. RELATED WORK

**HDC Model** While previous studies discussed the robustness of HDC with regard to hardware failures such memory errors [2], no known study has systematically look into the algorithmic robustness of HDC. Existing works on HDC focus mainly on two aspects: application of HDC and optimization of HDC processing. For application, Rahimi et al. used HDC on hand gesture recognition and achieve on average 97.8% accuracy, which surpasses support vector machine by 8.1% [14]. HDC was also used for language classification [2]. *VoiceHD* applied HDC to voice recognition [15]. Manabat et al., applied HDC to character recognition and conducted performance analysis [16]. For optimization of HDC processing, *HDC-IM* [17] proposed in-memory computing techniques for HDC scenarios based on Resistive RAM. There are also optimizations on HDC targeted at different computing platforms such as FPGA [18] and 3D IC [19].

**Fuzz testing** Fuzz testing has been broadly applied in the software testing community and has shown promising capability for detecting bugs or vulnerabilities in software systems [20], [21], [22], [23]. *American Fuzzy Lop* is a widely used fuzz testing platform that has been applied in industrial systems [24]. *RFuzz* uses coverage-directed fuzz testing in the hardware testing domain [25]. To discover errors occurring only for rare inputs in DNNs, *Tensorfuzz* uses approximate nearest neighbor algorithms and coverage-guided fuzzing (CGF) to enable property-based testing (PBT) [26]. Fuzz testing has broader application scenario and high versatility, particularly at systems that are not fully accessible, i.e., grey-box or even black-box systems [27], [28].

**Adversarial Generation in DNNs** Recently, adversarial deep learning [8], [29] have demonstrated that state-of-the-art DNN models can be fooled by crafted synthetic images by adding minimal perturbations to an existing image. Goodfellow et al. proposed a fast gradient sign method of generating adversarial examples with required gradient computed efficiently using backpropagation [8]. Nguyen et al. calculated the gradient of the posterior probability for a specific class (e.g., softmax output) with respect to the input image using backpropagation,

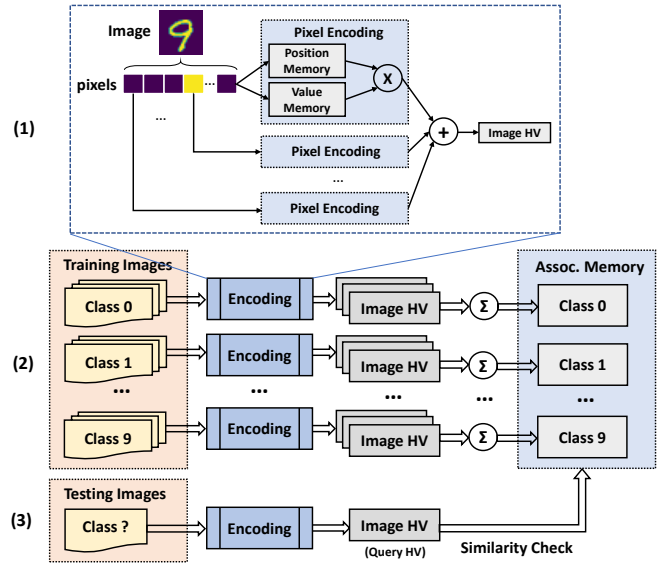


Fig. 2. HDC for image classification including: (1) **Encoding**, (2) **Training** and (3) **Testing**. **Encoding** section maps, or encodes an image into a representative image HV. **Training** section sums the image HVs from same class to train the Associative Memory (AM). **Testing** section performs similarity check with the image HV from an unlabelled image with the trained AM to predict its class.

and then used gradient to increase a chosen unit’s activation to obtain adversarial images [29]. While these inspiring methods are effective in generating adversarial examples for DNN, they cannot be applied to HDC because HDC is not built upon solving a mathematical optimization problem to find hyperparameters. The detail of HDC modeling process is described in the next Section.

**Our work** Despite the growing popularity of HDC, there is no known study on systematic HDC testing. This paper presents the first effort in such paradigm by customizing fuzz testing method for HDC models. Further, unique properties of HDC are explored to guide the testing process and the fuzz testing is enhanced with differential testing to enable manual labeling-free testing.

## III. HDC MODEL DEVELOPMENT

We develop an HDC model for image classification as our testing target. We develop three key phases: **Encoding**, **Training** and **Testing** as illustrated in Fig. 2.

### A. Encoding

The encoding phase is to use HDC arithmetic to encode an image into a hypervector called “Image HV”. Hypervector (HV) is the fundamental building block of HDC. HVs are high-dimensional, holographic, and (pseudo-)random with independent and identically distributed (i.i.d.) components. In HDC, HV supports three types of arithmetic operation for encoding: addition ( $\oplus$ ), multiplication ( $\otimes$ ) and permutation ( $\rho$ ), where ( $\oplus$ ) and ( $\otimes$ ) are element-wise addition and multiplication of two operand HVs while ( $\rho$ ) is to perform cyclic shifting of one operand HV. Multiplication and permutation will produce HVs

that are orthogonal to the original operand HVs while addition will preserve 50% of each original operand HVs [14].

As shown in Fig. 2 – Encoding, in order to encode one image into the representing HV, there are three steps. The first step is to decompose and flat the image into an array of pixels. The indices of the pixels in the array reflect the position of the pixel in the original image while the values of the pixels reflect the greyscale level of each pixel. For the MNIST dataset we use in this paper, since the image size is  $28 \times 28$  and the pixel range is 0 to 255 in greyscale, we flat a single image into an array with 784 elements with values ranging from 0 to 255.

The second step is to construct HVs representing each pixel from the index and value information provided by the image array. We randomly generate two memories of HVs: the position HV memory and the value HV memory based on the size and pixel value range of the image. The position HV memory accommodates  $28 \times 28 = 784$  HVs, each representing a pixel’s position in the original image. The value HV memory accommodates 255 HVs, each representing a pixel’s greyscale value. All the HVs generated are bipolar, i.e, with elements either “1” or “-1”. For each pixel from the image array, we look up the position HV and value HV from the two memories and use multiplication operation to combine them in order to construct the pixel’s representing HV.

The third step is to establish the HV representing the entire image. After constructing all the HVs representing each pixel, the final HV representing the image is established by adding all the pixel HVs together:  $ImgHV = Pixel_0HV \oplus Pixel_1HV \oplus \dots \oplus Pixel_{783}HV$ . So far, we encode one image into a representing image HV that is ready for HDC training and testing. However, addition can destroy the bipolar distribution of the HV (some elements can be numbers other than “1” or “-1”), thus the image HV is bipolarized again by Eq. 1.

$$HV[i] = \begin{cases} -1 & HV[i] < 0 \\ 1 & HV[i] > 0 \\ RandomSelect(1, -1) & HV[i] = 0 \end{cases} \quad (1)$$

## B. Training

The training phase is to iteratively incorporate the information (or features) contained in each image in the training set into an associative memory (AM) with corresponding label. AM stores a group of HVs, each representing a class. In other words, training is the process of building the AM by adding up all the training images’ HVs belong to one class together. In the beginning, we initialize every element inside the associative memory  $AM$  with 0. The dimension of the HVs inside  $AM$  is also consistent with the images’ HVs. Then, for each image  $t_i$  with class label  $c_i$  in the training set  $T$ , we first encode the image into its representing HV  $HV_i$  and then we add  $HV_i$  into the corresponding class HV inside the associative memory  $AM[c_i]$ . After one epoch when all the training images are added into  $AM$ , we finally bipolarize the HVs inside  $AM$  again using Eq.1 and the  $AM$  is constructed and ready for testing and evaluation.

## C. Testing

The testing phase is to evaluate the trained associative memory using the testing dataset. First, for each testing images  $t_i$  in the testing set  $T$  is encoded into its representing query HV  $HV_i$  using the same encoding mechanism in training. We then calculate the cosine similarity  $Sim$  between  $HV_i$  and every class’ hypervector  $AM[j]$  in the associative memory:  $Sim = Cosim(HV_i, AM[j]) = \frac{HV_i \cdot AM[j]}{\|HV_i\| \|AM[j]\|}$ . The class with the maximum similarity with  $HV_i$  subsequently becomes the prediction result  $c'_i$ . We then compare the true label  $c_i$  with the prediction  $c'_i$  for each image to determine if the prediction is correct and to evaluate the accuracy of the HDC model.

## IV. HDTEST FRAMEWORK

The overview of **HDTest** is illustrated in Fig. 3. **HDTest** takes the original input image  $t$  without necessarily knowing the label of it. **HDTest** then applies mutation algorithms on the original input  $t$  to generate new input  $t'$ . Both the generated input and the original input are then sent to the HDC classifier for prediction. We then check if the two predicted labels are different, and if yes, this indicates a successful generation of an adversarial input. Otherwise, **HDTest** will continue repeating the fuzzing process until success. The details are as following.

---

### Algorithm 1 HDTest Mutation Algorithm

---

**Input** *inputs*: unlabeled input images for testing.  
*HDC*: the HDC model under test.  
*strategies*: strategies for mutation listed in Table I.  
*iter\_times*: maximum allowed iteration times for a seed.

**Output**  $S$ : (set of) adversarial (input) images.

```

1:  $S = []$ 
2: for  $t$  in inputs do
3: \ * Prediction of the original input image. * /
4:    $y = HDC(t)$ 
5:   while iter_times not reached do
6:      $seeds = strategies(t)$ 
7:      $y\_seeds = HDC(seeds)$ 
8:     if  $y'$  in  $y\_seeds!$  =  $y$  then
9: \ * Indicating successful mis-predicted image generated. * /
10:     $S.append(y')$ 
11:    break
12:   else
13: \ * Continue fuzzing using only the fittest seeds. * /
14:     $seeds = seeds.fittest()$ 
15:   end if
16: end while
17: end for
```

---

The mutation algorithm of **HDTest** is shown in Alg. 1. Its objective is to generate adversarial images by applying mutation strategies to change, or add perturbation on the image. For each input in the unlabelled input image dataset, first, **HDTest** uses the HDC model to get the predicted label of the original input image (Line 4) as a reference label. Then, **HDTest** applies mutation strategies on the original image to generate different mutated images as seeds (Line 6). Again, **HDTest** feeds the seeds into HDC and obtains the corresponding label of the seeds as query labels (Line 7). By




Fig. 3. Overview of **HDTest**

comparing the query labels with reference labels, **HDTest** is able to know if there are any discrepancies which indicate successful generation of an adversarial image. The adversarial images are added to the set  $S$  and **HDTest** proceeds to the next image (Line 10). If all the seeds are still predicted the same as the original input image, **HDTest** will update the seeds and repeat the process, until a successful adversarial image is generated, or the maximum allowed  $iter\_times$  is reached (Line 14).

To maximize the efficiency of the fuzz testing process, **HDTest** uses distance-guided mutation for seeds update based on the distance between query HV and reference HV. That is, during the mutation process, only the top- $N$  fittest seeds can survive (In our experiments,  $N = 3$ ). The fitness of seeds are defined as:  $fitness = 1 - Cosim(AM[y], HDC(seed))$ , where  $AM[y]$  is the reference vector in the associative memory and  $HDC(seed)$  is the query HV of the seed, encoded by the HDC model. Higher fitness means lower similarity between the HV of the seed and the original input image’s HV, indicating higher possibility to generate an adversarial image. Experimental results show that using such guided testing can generate adversarial inputs faster than unguided testing by 12% on average.

To ensure the added perturbations are within an “invisible” range, we set a threshold for the distance metric during fuzzing (e.g.,  $L_2 < 1$ ). When generated images are beyond this limit, it is regarded as unacceptable and then discarded. This constraint can be modified by the user to achieve customized and adaptive performance control when using **HDTest**.

We introduce multiple mutation strategies for **HDTest**’s mutation algorithm as listed in Table I. The mutation strategies can be used independently or jointly to implement **HDTest** with different mutation strategies. In this paper, we focus on four mutation strategies of implementation: **HDTest** with (1) **gauss**, (2) **rand**, (3) **row & col\_rand**, and (4) **shift**. The first two strategies, **gauss** and **rand** add Gaussian and random noise onto the image, resulting in holographic mutation on the image. Comparatively, **row & col\_rand** only mutates a specific row or column. It is worthwhile to note that **shift** is different from all the other strategies. **Shift** does not modify the pixels’ values of the image, but just rearranges the pixel locations.

## V. EXPERIMENTAL RESULTS

### A. Experimental Setup

We use the MNIST database of handwritten digits [30] as our dataset for training and testing the HDC model at an

TABLE I  
MUTATION STRATEGIES OF **HDTest**.

Name	Description
row_rand	randomly mutate all pixels in one single row
col_rand	randomly mutate all pixels in one single column
rand	apply random noise over the entire image
gauss	apply gaussian noise over the entire image
shift	apply horizontal or vertical shifting to the image

accuracy around 90%. We use different metrics to evaluate **HDTest**: normalized  $L_1$  distance and normalized  $L_2$  distance between the mutated image and original image, the average fuzzing iterations and the execution time to successfully generate 1000 adversarial images. A smaller distance between the generated image and the original image is preferred by **HDTest** because it means less perturbation added to the image. The average fuzzing iterations indicate the speed of fuzz testing to obtain adversarial images. One fuzzing iteration refers to a cycle of seed generation, prediction discrepancy check and seed update, as denoted by Line 5 – 16 in Alg. 1. A lower number of iterations means **HDTest** spends less computation cost and time in generating adversarial images. It is calculated after running **HDTest** over a set of images, by  $Avg.\#iterations = \frac{\#total\_iterations}{\#images}$ . To reflect the speed of fuzz testing, we also record the execution time of **HDTest** on generating 1,000 adversarial images on a commodity computer with AMD Ryzen™5 3600 (6C12T, 3.6GHz) processor with 16GB memory.

### B. Mutation Strategies Analysis

Table II presents the  $L_1/L_2$  distance, as well as execution time of generating 1000 adversarial images. Generally, **HDTest** can generate 1000 adversarial images around 100s - 200s. Among them, **shift** is the fastest strategy with only 88.4s (because it only changes the pixel locations, or more exactly, indices), while **rand** is the slowest with 228s, because it requires much higher number of iterations for generation. We provide sample original images, mutated pixels, and generated images of **HDTest** in Fig. 4–6, under **gauss**, **rand**, and **shift** respectively.

For the amount of perturbations, **rand** shows smallest  $L_1$  and  $L_2$  distance of 0.58 and 0.09, respectively. However, its average fuzzing iterations is 9X higher than that of **gauss**. Additionally, while **gauss** shows the lowest average fuzzing iterations of 1.46, its distance metrics are around 5X higher than that of **rand**. **row & col\_rand** does not perform well, with 9.45  $L_1$  and 0.65  $L_2$  distance and 7.94 average fuzzing iterations, which is outperformed by the **gauss** in general, therefore we decide to exclude it from further demonstration.

We can also observe that **shift** exhibits significantly higher distance and average fuzzing iterations. This however, does not indicate **shift** has an overall bad performance. Shifting adjusts their geographical location in the image. This can lead to disastrous metrics since all the pixels are mutated and thus we deem the distance metrics are thus not meaningful




Fig. 4. Sample adversarial images produced by **HDTest** with “gauss”.




Fig. 5. Sample adversarial images produced by **HDTest** with “rand”.

in reflecting the effectiveness of **shift** strategy. The average fuzzing iteration of **shift** is 4.25, which means **HDTest** on average **shifts** 4.25 pixels, either horizontally or vertically, to generate an adversarial image.

An interesting observation is that the difficulty of generating adversarial inputs tend to vary for different samples. While some samples require extensive mutations with large perturbations added to obtain an adversarial image such as the third image in Fig. 4 and the first image in Fig. 5, some are easy to generate adversarial images with only few pixel mutations such as the last image in Fig. 4 and the second image in




Fig. 6. Sample adversarial images produced by **HDTest** with “shift”.

TABLE II  
 $L_1$  AND  $L_2$  DISTANCE, AVERAGE FUZZING ITERATIONS AND RUNTIME OF **HDTest** UNDER DIFFERENT MUTATION STRATEGIES.

Metric		gauss	rand	row & col_rand	shift*
Avg. Norm. Dist.	$L_1$	2.91	0.58	9.45	10.19*
	$L_2$	0.38	0.09	0.65	0.68*
Avg. #Iter.		1.46	12.18	7.94	4.25
Time Per-1K Gen. Img. (s)		173.0	228.3	114.2	88.4

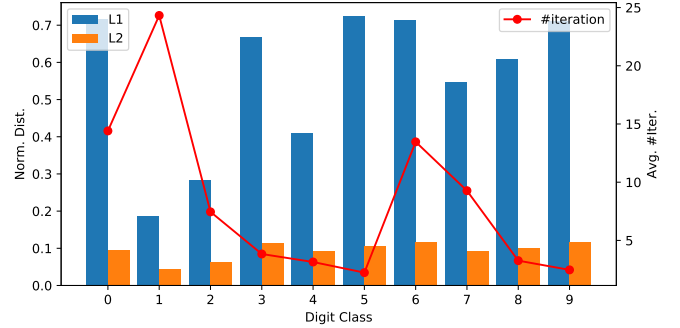


Fig. 7. Per-class normalized  $L_1$  and  $L_2$  distances and average pixel mutations to generate an adversarial image.

Fig. 5, which we refer to as vulnerable cases. Such vulnerable cases bring potential security loopholes since they can result in incorrect behaviors of HDC models by only minor and even negligible perturbations. Therefore, such images should be emphasized when defending attacks to the HDC systems, and **HDTest** is able to pinpoint and highlight them.

### C. Per-class Analysis

We then perform a per-class analysis on  $L_1/L_2$  distance and fuzzing iterations as shown in Fig. 7. We can observe that for some classes such as “1”, the average fuzzing iteration is drastically higher, which means that it is relatively difficult to generate adversarial examples. For other classes such as “9”, the average fuzzing iteration is lower. This is reasonable because all the other digits except for “7” are visually dissimilar from “1” while “9” has quite a few similarities such as “8” and “3”. Moreover, we do not observe apparent relation between fuzzing iteration and distance metrics. For example, the average fuzzing iterations of “1” and “6” are both high, yet the average distance metrics of “6” is noticeably smaller. In summary, smaller average distance and/or lower fuzzing iterations means that **HDTest** performs better on these digits, or vice versa.

### D. Case Study on Defense Against Adversarial Attacks

Adversarial attack is one major threat to ML systems [31]. The attacker mutates images by adding perturbations and generate adversarial images to fool the model to predict wrong labels. As a case study of **HDTest**, we propose a defense mechanism to strengthen the robustness of HDC models to such adversarial attacks, as shown in Fig. 8. First, for a specific HDC model, we run **HDTest** to generate 1000 adversarial




Fig. 8. Defend attacks using **HDTest**-generated images: (1): Attack image generation; (2): HDC Retraining; (3): Attack HDC again using unseen images.

images. We then randomly split such 1000 images into two subsets. We feed the first subset of adversarial images to the HDC model with correct labels to retrain the HDC model, i.e., updating the reference HVs. Next, we use the second subset of adversarial images to attack the HDC model. Our experimental results show that, after retraining, the rate of successful attack rate drops more than 20%.

### E. Discussion

HDC is a fast-growing emerging area that has seen increasing applications in various domains. While optimizing prediction accuracy of existing HDC models is out of scope for this paper, we note intensive ongoing research in this area focusing on advancing HV encoding, model structure, and training mechanism (e.g., retraining) [32]. Complementary and orthogonal to the above efforts, we focus on the problem of how to test HDC models and introduce a customized highly-automated and scalable testing tool for HDC. While we only consider one dataset in this paper due to the fact that HDC is still exploring its use in image classification domain, **HDTest** can be naturally extended to other HDC model structures because it considers a general greybox assumption with only HV distance information.

## VI. CONCLUSION

Brain-inspired HDC is an emerging computational paradigm but is gaining increasing capability in various application domains. This paper presents the first effort in HDC testing by introducing **HDTest**, a highly-automated and scalable testing approach without manual labeling. Based on differential fuzz testing, **HDTest** iteratively mutates inputs to generate adversarial inputs to expose incorrect behaviors of HDC models. We develop multiple mutation strategies in **HDTest**, and evaluate **HDTest** on MNIST dataset. Experimental results show that **HDTest** is able to generate around 400 adversarial images within one minute. As a case study, we use the **HDTest**-generated inputs to retrain HDC models, which can strengthen the robustness of HDC models. While HDC performance still expects more advancements both in theoretical and implementation aspects, this paper aims to raise community attention to the reliability and robustness aspects of this emerging technique.

## REFERENCES

- [1] L. Ge *et al.*, "Classification using hyperdimensional computing: A review," *IEEE Circuits and Systems Magazine*, 2020.
- [2] A. Rahimi *et al.*, "A robust and energy-efficient classifier using brain-inspired hyperdimensional computing," in *ISLPED*, 2016.
- [3] M. Hersche, E. M. Rella, A. Di Mauro, L. Benini, and A. Rahimi, "Integrating event-based dynamic vision sensors with sparse hyperdimensional computing: a low-power accelerator with online learning capability," in *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design*, 2020, pp. 169–174.
- [4] A. Rahimi *et al.*, "Hyperdimensional computing for blind and one-shot classification of eeg error-related potentials," *Mobile Networks and Applications*, 2017.
- [5] A. Moin *et al.*, "An emg gesture recognition system with flexible high-density sensors and brain-inspired high-dimensional classifier," in *ISCAS*, 2018.
- [6] Y. Kim *et al.*, "Geniehd: efficient dna pattern matching accelerator using hyperdimensional computing," in *DATE*, 2020.
- [7] C. Szegedy *et al.*, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [8] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [9] O. Russakovsky *et al.*, "Imagenet large scale visual recognition challenge," *IJCV*, 2015.
- [10] S. Wang *et al.*, "Exploring causes and effects of automated vehicle disengagement using statistical modeling and classification tree based on field test data," *Accident Analysis & Prevention*, 2019.
- [11] A. C. Madrigal, "Inside waymo's secret world for training self-driving cars," *The Atlantic*, 2017.
- [12] J. Liang *et al.*, "Deepfuzzer: Accelerated deep greybox fuzzing," *IEEE TDSC*, 2019.
- [13] W. M. McKeeman, "Differential testing for software," *Digital Technical Journal*, 1998.
- [14] A. Rahimi *et al.*, "Hyperdimensional biosignal processing: A case study for emg-based hand gesture recognition," in *ICRC*, 2016.
- [15] M. Imani *et al.*, "Voicehd: Hyperdimensional computing for efficient speech recognition," in *2017 IEEE International Conference on Rebooting Computing (ICRC)*. IEEE, 2017.
- [16] A. X. Manabat *et al.*, "Performance analysis of hyperdimensional computing for character recognition," in *ISMAC*, 2019.
- [17] J. Liu *et al.*, "Hdc-im: Hyperdimensional computing in-memory architecture based on rram," in *ICECS*, 2019.
- [18] M. Schmuck *et al.*, "Hardware optimizations of dense binary hyperdimensional computing: Rematerialization of hypervectors, binarized bundling, and combinational associative memory," *JETC*, 2019.
- [19] T. F. Wu *et al.*, "Brain-inspired computing exploiting carbon nanotube fets and resistive ram: Hyperdimensional computing case study," in *ISSCC*, 2018.
- [20] Y. Fu *et al.*, "Evmfuzzer: detect evm vulnerabilities via fuzz testing," in *ESEC/FSE*, 2019.
- [21] Y. Chen *et al.*, "Enfuzz: Ensemble fuzzing with seed synchronization among diverse fuzzers," in *28th {USENIX} Security Symposium*, 2019.
- [22] J. Gao *et al.*, "Vulseeker-pro: enhanced semantic learning based binary vulnerability seeker with emulation," in *ESEC/FSE*, 2018.
- [23] J. Guo *et al.*, "Dlfuzz: Differential fuzzing testing of deep learning systems," in *ESEC/FSE*, 2018.
- [24] M. Zalewski, "American fuzzy lop," 2014.
- [25] K. Laeufer *et al.*, "Rfuzz: coverage-directed fuzz testing of rtl on fpgas," in *ICCAD*, 2018.
- [26] A. Odena *et al.*, "Tensorfuzz: Debugging neural networks with coverage-guided fuzzing," in *ICML*, 2019.
- [27] C. Lemieux *et al.*, "Fairfuzz: A targeted mutation strategy for increasing greybox fuzz testing coverage," in *ASE*, 2018.
- [28] P. Godefroid, "Random testing for security: blackbox vs. whitebox fuzzing," in *2nd international workshop on Random testing*, 2007.
- [29] A. Nguyen *et al.*, "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images," in *CVPR*, 2015.
- [30] Y. LeCun *et al.*, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, 1998.
- [31] N. Thakur *et al.*, "Evaluating a simple retraining strategy as a defense against adversarial attacks," *arXiv preprint arXiv:2007.09916*, 2020.
- [32] M. Nazemi *et al.*, "Synergiclearning: Neural network-based feature extraction for highly-accurate hyperdimensional learning," *ICCAD*, 2020.