Algebraic Cryptanalytic Extraction on Hard-Label Neural Networks

Zirui Chen[1], Shi Tang[2], Zhengchao Gao[2], Yongjia Su[2], Lingyue Qin[1], and Xiaoyang Dong[1]

[1] Tsinghua University, Beijing, P. R. China
chenzr25@mails.tsinghua.edu.cn, {qinly,xiaoyangdong}@tsinghua.edu.cn
[2] Shandong University, Qingdao, P. R. China
{shi.tang,chao_qwq,yongjia.su}@mail.sdu.edu.cn

Abstract. Although the state-of-the-art model extraction attack on the hard-label Fully-connected Neural Network (FCN) by Carlini *et al.* at EUROCRYPT 2025 has polynomial-time complexity in theory, its dual-point clustering relies on singular value decomposition (SVD) with a time complexity of $\mathcal{O}(n^2(d^{(k)})^3)$, resulting in huge runtime in practice. To address this computational bottleneck, this work transforms Carlini *et al.*'s geometric-view hard-label attack into an algebraic framework, and proposes two efficient clustering methods: Normal Rank Check (NRC) and Approximate Signature Vector (ASV). The NRC and ASV methods replace Carlini *et al.*'s heavy SVD-based rank checking with simple rank checking or inner-product operations, reducing the clustering complexity to $\mathcal{O}(n(d^{(k)})^3)$ on average. Furthermore, this paper presents the first model extraction attack against hard-label max-pooling Convolutional Neural Networks (CNNs) by combining the ASV method with the kernel-centric clustering scheme instead of the neuron-centric clustering, which fully exploits the property of weight sharing in convolutions and fills a cryptanalysis gap. Experiments on FCNs and the max-pooling LeNet-5 demonstrate that our NRC/ASV methods drastically cut clustering time, and improve the overall efficiency in the model extraction.

Keywords: Model Extraction · Hard-label · Convolutional Neural Networks · Algebraic Attack · Approximate Signature Vector

1 Introduction

Deep Neural Networks (DNNs) map inputs to outputs via a function defined by their architecture, weights, and biases. Extracting these weights and biases from a black-box DNN (given only its inputs and outputs) has been a long-standing challenge spanning three decades [3,14]. As early as 2005, Lowd and Meck introduced the first adversarial reverse engineering attack for this problem [23], with recent years seeing continued innovations from industry and academia [2,18,24,26,28]. At CRYPTO 2020, Carlini, Jagielski, and Mironov achieved the seminal parameter extraction attack on the ReLU-based fully connected neural

networks (FCNs) with a cryptographic differential attack [8]. This approach recovers the unsigned weights using a polynomial number of raw output queries, but relies on exponential time of guessing for sign recovery [8]. At EUROCRYPT 2024, Canales-Martínez *et al.* proposed the *neuron wiggle technique* to recover neuron signs in polynomial time [6]. Building on this, Foerster *et al.* developed an end-to-end attack for practical models by combining methods from Carlini *et al.* and Canales-Martínez *et al.* at NeurIPS 2024 [15]. At EUROCRYPT 2026, Liu *et al.* further extended parameter extraction to the deep layers with more novel techniques [22]. In 2026, the parameter extraction attacks are proposed against the Convolutional Neural Networks (CNNs) with max pooling [12,21] and average pooling [27]. Besides, various model extraction attacks have emerged for different networks, such as PReLU networks [10], non-linear activations networks [1,25], as well as in the physical leakage settings [2,16,13].

Another direction focuses on ReLU-based DNNs in the hard-label scenario, where attackers only access final classification labels (*e.g.*, "dog" or "car"). Yi Chen *et al.* first proposed an exponential-time extraction method at ASIACRYPT 2024 [9]. At EUROCRYPT 2025, Carlini *et al.* achieved a breakthrough by presenting a polynomial-time, polynomial-query attack by analyzing the geometric properties of decision boundaries [7]. Canales-Martínez later proposed the hard-label *SOE* to recover the signs at LATINCRYPT 2025 [5]. At CRYPTO 2026, Ito, Miura, and Todo proposed a cross-layer extraction strategy to handle always-active neurons [17]. Sun *et al.* proposed the first extraction attack on hard-label CNNs with average pooling [27], which was later improved by [20], while the attacks on max-pooling CNNs [11,21] require raw output queries, making the hard-label attacks against max-pooling CNNs an open problem.

Our Contributions. A major computational hurdle of Carlini *et al.*'s hard-label attack [7] arises during the signature recovery phase, which requires clustering the collected dual points into consistent groups. The authors explicitly acknowledge this limitation, noting that their proof-of-concept implementation avoids the full computational burden of the clustering in [7, Section 7.1]:

> "*Our proof-of-concept implementation disregards the runtime of the* n^2 *pairwise clustering* $t_{cluster}$ *needed to identify dual points that are mutually consistent. Although this method is known to work in principle (we validated it post-hoc) and runs in polynomial time, applying it naively (i.e., without many potential optimizations) to a real attack would involve a* $(1\ million)^2$ *time complexity, requiring weeks of computation.*"

Our contributions come from two aspects and (partial) black-box experiments:

– **Improved Carlini *et al.*'s Hard-label Attack:** We first convert Carlini *et al.*'s attack [7] from their geometric view to our algebraic view, which inspires two more efficient clustering methods named as *Normal Rank Check* (NRC) and *Approximate Signature Vector* (ASV) methods. The first method NRC computes the rank of only 4 normal vectors of the decision boundaries

of two dual points in layer k to determine the consistency, while Carlini *et al.* compute the rank of about $2d^{(k)}$ vectors of the same length [7].
The second method (*i.e.,* ASV) further enhances this efficiency: in unitary balanced networks, it theoretically bounds the expected directional distance between the derived ASV $\vec{v}$ and the corresponding neuron weight $\vec{\omega}$ by $1/d^{(k)}$. For FCNs with high weight dimension $d^{(k)}$, experiments have also confirmed that the directions of $\vec{v}$ and $\vec{\omega}$ are nearly parallel (identical or opposite), consistent with the theoretical results. Specifically, we first compute the ASV $\vec{v}$ for each dual point. If two dual points correspond to the same neuron, then their ASVs $\vec{v}_1$ and $\vec{v}_2$ will be nearly parallel to $\vec{\omega}$, simultaneously. Therefore, the directions of $\vec{v}_1$ and $\vec{v}_2$ will also be nearly parallel, which could be detected by a simple inner product in $\mathcal{O}(d^{(k)})$.
Suppose there are n dual points for layer k, Carlini *et al.*'s clustering [7] is to compute the rank of $2d^{(k)}$ vectors for any two dual points through SVD, which has a total time complexity of $\mathcal{O}(n^2(d^{(k)})^3)$. In contrast, the time complexities for both NRC or ASV methods are $\mathcal{O}(n(d^{(k)})^3 + n^2 d^{(k)})$ in the worst case, or $\mathcal{O}(n(d^{(k)})^3)$ in the average case. In our experiment with 2,000 dual points, the clustering time is $2^{6.09}$ seconds (about 1 minute) with our ASV method, while Carlini *et al.*'s method [7] requires about $2^{21.54}$ seconds (847 hours), *i.e.*, the time is significantly reduced by a factor of 2^{15}. More comparisons are given in Table 3. In practical attacks, the NRC and ASV methods could be combined to achieve efficiency and robustness.

– **Kernel-Centric ASV in Hard-label Attack on the Max-pooling CNN:** Based on the ASV method, we introduce the first hard-label attack on max-pooling CNNs, filling a cryptanalysis gap. We define the dual points (*i.e.*, dual RPCPs or dual PSPs) based on the critical points introduced by Zirui Chen *et al.* [11]. In Carlini *et al.*'s attack [7] on FCNs and Sun *et al.*'s attack [27] on CNNs, the dual points clustering methods are both *neuron-centric*, *i.e.*, two clustered dual points should correspond to the same neuron. When applying the *neuron-centric* method to the max-pooling CNN, one has to collect many dual RPCPs or dual PSPs to expect that at least two fall into the same neuron or max pooling. However, unlike FCNs, in layer k of a single-channel CNN, the convolutional kernel weight to be recovered is the same one for all neurons. Therefore, there is a natural question that *can we cluster different dual RPCPs and dual PSPs corresponding to different neurons or max poolings within the same layer k into the same category, and use them to determine the unique convolutional kernel weight?* This paper answers the question by proposing the *Kernel-Centric ASV* method with a *kernel-centric* (instead of Carlini *et al.*'s *neuron-centric*) clustering method to solve the unique convolutional kernel weight.

Black-box and Partial Black-box Experiments.

– **Black-box Signature Recovery on 64-64×4-10 FCN.** The experiment targets on 64-64×4-10 FCN, identical to the "Tiny" model from the source code provided by Carlini *et al.* [7], but trained on MNIST dataset. Due to

the high computational cost of Carlini *et al.*'s method [7], they only perform a proof-of-concept white-box signature recovery, *i.e.*, relying on white-box access to gradients or neuron indices across all three phases (dual point finding, clustering, and signature computing). For this network, we treat all the three phases as a black box and extract the signatures of the first two layers and 26 neurons of layer 3. The results and the comparison with the experiment [7] on the same platform are given in the first three rows of Table 1. Carlini *et al.*'s method needs 0.93 hours to extract the 1st layer, while our method only needs 0.04 hours, reducing the runtime by 94.6%. For the 2nd layer, about 300, 000 dual points are needed to cluster, and our method needs 4.22 hours, while Carlini *et al.*'s method cannot give a result within 1 month.

- **Black-box Extractions on Hard-label Max-pooling CNNs.** For the experiments on CNNs, we recover the full convolutional block of a (2+1) CNN and a (2+2) LeNet-5 with max pooling in the hard-label setting. The bottom four rows of Table 1 compare our results with Sun *et al.* (target average pooling) [27] and Chen *et al.* (requires raw outputs) [11].
- **Partial Black-box Signature Recovery on 3072-256×3-64-10 FCN.** For this experiment, only the *dual point finding* phase relies on white-box gradient access. The other two phases are still black-box. We use 100,000 dual points to recover the signatures of the first two layers in Table 2. Due to the high computational cost of Carlini *et al.*'s method [7], we estimate its runtime by multiplying the total required comparisons (obtained via white-box access) by the average comparison time measured on a subset of dual points.

The source codes for all the experiments can be found via

XXX

2 Preliminaries

All notations are summarized in Tab. 4 in Supp. A.1. A Deep Neural Network (DNN) with logit output is a function $\mathcal{F}_\theta$ parameterized by θ that takes inputs $\in \mathbb{R}^{d^{(1)}}$ and returns values $\in \mathbb{R}^{d^{(r+1)}}$. In the hard-label setting, DNN returns the most likely class label (*e.g.*, "cat" or "dog"), instead of the logit values.

Definition 1. (Hard-label DNN) *A hard-label DNN takes an input $X^{(1)} \in \mathbb{R}^{d^{(1)}}$, processes it through the DNN, and applies a function z on the resulting logit vector to returns a hard label:*

$$\mathcal{H}_\theta = (z \circ \mathcal{F}_\theta)(X^{(1)}) = z(\mathcal{F}_\theta(X^{(1)})) = \arg\max_i \mathcal{F}_\theta(X^{(1)})_i \tag{1}$$

Note that if there are ties, *i.e.*, multiple elements of $\mathcal{F}_\theta(X^{(1)})$ have the same maximum, the hard-label is the smallest index of these equal elements [7,9].

Table 1: Black-box Experiments on FCN and Max-pooling LeNet-5 in S1 Setting. All notations are summarized in Tab. 4 in Supp. A.1. and the architecture of LeNet-5 is given in Supp. E. The accuracy term $(\varepsilon, 0)$ is defined in Sect. 2.3, and $\max|\theta - \hat{\theta}|$ is the maximum extraction error of model parameters.

Models	**Architecture**† $d^{(k)}$–$d_f^{(k)}$–$d^{(k+1)}$	**Kernel**	**Method**	**Runtime** hours	**Queries**	$(\varepsilon, 0)$	$\max\|\theta - \hat{\theta}\|$
FCN 64-64×4-10	F_1: 64–64	–	ASV	0.05	$2^{24.30}$	$2^{-39.38}$	$2^{-44.20}$
		–	[7]	0.92	$2^{24.30}$	$2^{-39.14}$	$2^{-44.17}$
	F_2: 64–64	–	NRC	4.22	$2^{30.76}$	$2^{-36.76}$	$2^{-41.30}$
		–	[7]	≈1478.55♦	≈ $2^{30.76}$	–	–
	F_3: 64–64	–	NRC	6.58	$2^{32.18}$	26/64◇ recovered	$2^{-36.02}$
		–	[7]	≈1079.65♦	≈ $2^{30.76}$	25/64◇ recovered	–
(2+1) CNN	C_1: 1024–784–196 C_2: 196–100–25	(1,1,5,5) (1,1,5,5)	Kernel-Centric ASV	1.46	$2^{28.35}$	$2^{-31.41}$	$2^{-38.74}$
LeNet-5	C_1: 1024–4704–1176 C_2: 1176–1600–400	(1,6,5,5) (6,16,5,5)	Kernel-Centric ASV	12.65	$2^{30.96}$	$2^{-27.89}$	$2^{-36.18}$
LeNet-5 [27]♠	C_2: 1176–1600–400	(6,16,5,5)	Neuron-Centric Clustering	190.84‡	$2^{32.34}$	–	$2^{-17.75}$
LeNet-5 [11]♣	C_1: 1024–4704–1176 C_2: 1176–1600–400	(1,6,5,5) (6,16,5,5)	–	42.11‡	$2^{27.29}$	$2^{-23.38}$	$2^{-27.31}$

†: F_i is the FCN layer i; C_i is the CNN layer i, defined in Tab. 4.
♦: Due to the high computational cost of Carlini *et al.*'s method [7], the complexity of only F_1 layer extraction is given. For F_2 layer, we estimate their runtime as follows (also given in Sect. 6): We simulate the clustering process for all 300,000 dual points using white-box neuron indices, assuming each SVD-based consistency check is correct. This full-scale clustering requires a total of 2,653,799,211 comparison. We then evaluate the actual time cost on a smaller subset of 3,000 dual points, which takes 3,585.770 seconds to perform 1,790,440 comparison. According to Sect. 3.2, the estimated runtime for Carlini *et al.*'s method to cluster the entire dual points set is about $\frac{3{,}585.770\text{ s}}{1{,}790{,}440\times 3{,}600\text{ s/h}} \times 2{,}653{,}799{,}211 = 1{,}476$ hours. Together with the costs of dual point finding and signature computing, the extraction time for F_2 is about 1,478.55 hours. Similarly, for F_3, 300,000 dual points are clustered with about 1077 hours.
◇: In F_3, in order to make practical experimental comparison, only about 26 neurons' signatures are recovered.
♠: Sun *et al.* [27] target the LeNet-5 with *Average Pooling* in hard-label setting.
♣: Chen *et al.* [11] target LeNet-5 with *Max Pooling* in raw-output setting.
‡: The figures are borrowed from [27] and [11], respectively.

Table 2: Partial Black-box Experiments on *3072-FCN* in S1 Setting.

Models	**Architecture** $d^{(k)}$–$d^{(k+1)}$	**Method**	**Runtime** hours	**Queries**	$(\varepsilon, 0)$	$\max\|\theta - \hat{\theta}\|$
FCN 3072-256×3-64-10	F_1: 3072–256	ASV	6.36	$2^{22.77}$	$2^{-17.74}$	$2^{-24.56}$
		[7]	≈825845.75♦	≈ $2^{22.77}$	–	–
	F_2: 256–256	ASV	4.17	$2^{22.77}$	238/256 recovered	$2^{-23.19}$
		[7]	≈93568.23♦	≈ $2^{22.77}$	238/256 recovered	–

♦: Similarly, the runtime is estimated by multiplying the total comparisons (obtained via white-box access) by the average comparison time measured on a subset of dual points. For F_1 layer, the full-scale clustering for 100,000 dual points requires 1,296,384,990 comparison. We evaluate on 200 dual points, taking 42,371.518 seconds to perform 18,476 comparison. According to Sect. 3.2, the estimated runtime for Carlini *et al.*'s method to cluster is about $\frac{42{,}371.518\text{ s}}{18{,}476\times 3{,}600\text{ s/h}} \times 1{,}296{,}384{,}990 = 825{,}843$ hours. Similarly, for F_2 layer, 100,000 dual points are clustered with about 93567 hours.

2.1 Fully-connected Neural Network (FCN)

The FCN is a sequence of FCN layers, and each layer includes a linear function $f^{(k)} : \mathbb{R}^{d^{(k)}} \mapsto \mathbb{R}^{d^{(k+1)}}$ $(k \geq 1)$, and a component-wise ReLU function σ:

$$\mathcal{F}_\theta = f^{(r+1)} \circ \sigma \circ f^{(r)} \circ \sigma \circ \cdots f^{(2)} \circ \sigma \circ f^{(1)}, \tag{2}$$

where $f^{(k)} : \mathbb{R}^{d^{(k)}} \to \mathbb{R}^{d^{(k+1)}}$ is an affine transformation:

$$f^{(k)}(X) = A^{(k)}X^{(k)} + B^{(k)},$$

where $X^{(k)} \in \mathbb{R}^{d^{(k)}}$ is the input vector. The weight matrix $A^{(k)} \in \mathbb{R}^{d^{(k+1)} \times d^{(k)}}$ and the bias vector $B^{(k)} \in \mathbb{R}^{d^{(k+1)}}$ are composed of floating-point numbers, which are the model parameters. The unsigned weight $A^{(k)}$ and bias $B^{(k)}$ are named *signature* [6]. Given input $X^{(1)} \in \mathbb{R}^{d^{(1)}}$, the ReLU function σ in layer k is also interpreted as the matrix

$$I^{(k)} = \mathrm{diag}(\tau_0^{(k)}, \tau_1^{(k)}, \ldots, \tau_{d^{(k)}-1}^{(k)}), \tag{3}$$

where $\tau_i^{(k)} \in \{0, 1\},\ 0 \leq i < d^{(k)}$ are determined by the input $X^{(1)}$.

Definition 2. (Linear Neighborhood) *Given an input $x \in \mathbb{R}^{d^{(1)}}$, the matrices $A^{(k)},\ B^{(k)},\ I^{(k)},\ (1 \leq k \leq r+1)$ are all fixed. The linear neighborhood of x is defined as the subset $\chi \subset \mathbb{R}^{d^{(1)}}$, so that all $x' \in \chi$, the same $A^{(k)},\ B^{(k)},\ I^{(k)}$ will be applied to compute the output of the FCN.*

FCN has been proved to be piecewise linear function [8,6], *i.e.*, for x in the linear neighborhood χ of $X^{(1)}$, the FCN is reduced to

$$\begin{aligned} \mathcal{F}_\theta(x) &= A^{(r+1)}\Big(I^{(r)}\Big(A^{(r)}\cdots\Big(I^{(1)}\Big(A^{(1)}x + B^{(1)}\Big)\Big)\cdots + B^{(r)}\Big)\Big) + B^{(r+1)} \\ &= A^{(r+1)}I^{(r)}A^{(r)}\cdots I^{(2)}A^{(2)}I^{(1)}A^{(1)}x + \beta = \Gamma x + \beta. \end{aligned} \tag{4}$$

2.2 Convolutional Neural Network (CNN)

In $(m+n)$-deep CNN, the function $\mathcal{F}_\theta : \mathbb{R}^{d^{(1)}} \to \mathbb{R}^{d^{(m+n+2)}}$ is composed of two sequential blocks: the first block (Convolutional Block) consists of m CNN layers, and each layer includes a convolutional function $f_c^{(k)}$, a ReLU $\sigma_c^{(k)}$, and a pooling function $\rho^{(k)}$ $(1 \leq k \leq m)$; the second block consists of n FCN layers:

$$\mathcal{F}_\theta = \underbrace{f^{(n+1)} \circ \sigma^{(n)} \circ \cdots \circ \sigma^{(1)} \circ f^{(1)}}_{\text{Fully-connected Block}} \circ \underbrace{\rho^{(m)} \circ \sigma_c^{(m)} \circ f_c^{(m)} \circ \cdots \circ \rho^{(1)} \circ \sigma_c^{(1)} \circ f_c^{(1)}}_{\text{Convolutional Block}}. \tag{5}$$

At CRYPTO 2026, Zirui Chen *et al.* [11] gave an algebraic view of CNN. Given $X^{(1)} \in \mathbb{R}^{d^{(1)}}$, the CNN layer "$\rho^{(k)} \circ \sigma_c^{(k)} \circ f_c^{(k)}$" is translated into

$$\rho^{(k)} \circ \sigma_c^{(k)} \circ f_c^{(k)} = P^{(k)}I^{(k)}(A^{(k)}X^{(k)} + B^{(k)}), \tag{6}$$

where convolutional kernal matrix $A^{(k)}$ and the bias $B^{(k)}$ are the only trained parameters in CNN which will be extracted, $I^{(k)}$ is the same to Eq. (3), and $P^{(k)}$ is the so-call Boolean max-pooling matrix [11]. More details on those matrices are given in Supp. A.1, and a toy example is given Fig. 7 in Supp. A.2. CNN has been also proved to be piecewise linear [11], *i.e.*, for x in a linear neighborhood of $X^{(1)}$, the CNN layers and FCN layers of CNN "collapse" into an affine transformation. Assume that the function of the m CNN layers is $\mathcal{F}_1$, and the function of the n FCN layers is $\mathcal{F}_2$, then

$$\begin{aligned} y = \mathcal{F}_1(x) &= P^{(m)}(I^{(m)}(A^{(m)} \cdots (P^{(1)}(I^{(1)}(A^{(1)}x + B^{(1)}))) \cdots + B^{(m)})) \\ &= P^{(m)} I^{(m)} A^{(m)} \cdots P^{(2)} I^{(2)} A^{(2)} P^{(1)} I^{(1)} A^{(1)} x + \beta_1 = \Gamma_1 \cdot x + \beta_1. \end{aligned} \tag{7}$$

Together with Eq. (4), the full CNN "collapses" into

$$\mathcal{F}_\theta(x) = \Gamma_2(\Gamma_1 \cdot x + \beta_1) + \beta_2 = \Gamma_2 \cdot \Gamma_1 \cdot x + \Gamma_2\beta_1 + \beta_2. \tag{8}$$

Definition 3. (Layer Merging) *Given $(m+n)$-deep CNN network, suppose that we have complete knowledge of the first $k-1$ layers of the convolutional block, and we are currently recovering layer k. Let F_x^{k-1} and G_x^{k+1} represent, respectively, the fully recovered and non-recovered parts of the CNN.*

$$\mathcal{F}_\theta = \underbrace{f^{(n+1)} \circ \cdots \circ \sigma^{(1)} \circ f^{(1)} \circ \rho^{(m)} \circ \cdots \circ f_c^{(k+1)}}_{G_x^{k+1}} \circ \rho^{(k)} \circ \sigma_c^{(k)} \circ f_c^{(k)} \circ \underbrace{\rho^{(k-1)} \circ \cdots \circ f_c^{(1)}}_{F_x^{k-1}}.$$

Given $X^{(1)} \in \mathbb{R}^{d^{(1)}}$, G_x^{k+1} and F_x^{k-1} become $G_x^{k+1}(X^{(k+1)}) = G^{k+1}X^{(k+1)} + B^{k+1}$ and $F_x^{k-1}(X^{(1)}) = F^{k-1}X^{(1)} + B^{k-1}$, where the matrices $F^{k-1} \in \mathbb{R}^{d^{(k)} \times d^{(1)}}$, $B^{k-1} \in \mathbb{R}^{d^{(k)}}$, $G^{k+1} \in \mathbb{R}^{d^{(n+2)} \times d^{(k+1)}}$, $B^{k+1} \in \mathbb{R}^{d^{(n+2)}}$, respectively. The definition of layer merging for FCN is similarly defined.

Definition 4. (Neuron State and Critical Point [8]) *Let $\mathcal{V}(\eta; X^{(1)})$ denote the value that neuron η takes with $X^{(1)} \in \mathbb{R}^{d^{(1)}}$ before applying its corresponding ReLU. If $\mathcal{V}(\eta; X^{(1)}) > 0$ (respectively, $\mathcal{V}(\eta; X^{(1)}) < 0$), the neuron state of η is activated (respectively, inactivated). In FCN, if $\mathcal{V}(\eta; X^{(1)}) = 0$, the neuron state is critical, and $X^{(1)}$ is a critical point.*

Definition 5. (ReLU-Pooling Critical Point, RPCP [11]) *In CNN, a RPCP is an input $X^{(1)} \in \mathbb{R}^{d^{(1)}}$ that makes the input of a ReLU of $\sigma_c^{(k)}$ 0, and by varying the input in its tiny vicinity, the output of CNN changes non-linearly.*

Property 1. [11] If $X^{(1)}$ is a RPCP corresponding to the t-th neuron in layer k, then $A_t^{(k)}X^{(k)} + B_t^{(k)} = 0$, and the inputs of other neurons in the same Local Reception Field of Pooling (LRF-P) are smaller than 0.

Definition 6. (Pooling Switching Point, PSP [11]) *In CNN, a PSP is an input $X^{(1)} \in \mathbb{R}^{d^{(1)}}$ such that two inputs within the local receptive field of a pooling layer (LRF-P) achieve the same maximum value. By varying the inputs in the tiny vicinity of $X^{(1)}$, the output of CNN changes non-linearly.*

Property 2. [11] Given PSP $X^{(1)}$ for the i-th and j-th neuron in layer k, then $A_i^{(k)}X^{(k)} + b^{(k)} = A_j^{(k)}X^{(k)} + b^{(k)}$, and the inputs of other neurons in the same LRF-P are smaller than it.

2.3 Adversarial Goals and Assumptions

We make the following assumptions on the capabilities of the attacker:

- **Architecture knowledge.** The attacker has the full knowledge of the target neural network's architecture, while the parameters are treated as secret.
- **Full-domain inputs.** The attacker can query arbitrary inputs from $\mathbb{R}^{d^{(1)}}$.
- **Precise Computation.** The neural network is specified and evaluated using a sufficiently high precision floating-point arithmetic.
- **ReLU Activations.** All activation functions $\sigma^{(k)}$ are the ReLU function.
- **Hard-label Outputs.** Given input $X^{(1)}$, the model returns the label of the most likely class, which corresponds to the maximum of the raw output.

Adversarial Goal. The goal is to achieve an (ε, ξ)-functionally equivalent parameter extraction, rather than replicating the exact original parameters [8,7].

Definition 7. ((ε, ξ)-Functional Equivalence [8]) *Two neural networks $\mathcal{F}_\theta$ and $\mathcal{F}_{\hat{\theta}}$ are (ε, ξ)-functional equivalent on the input dataset S if*

$$Pr_{x \in S}(|\mathcal{F}_\theta(x) - \mathcal{F}_{\hat{\theta}}(x)| \leq \varepsilon) \geq 1 - \xi.$$

In our analysis, we specifically focus on the $(\varepsilon, 0)$-functional equivalence given by Carlini *et al.* [8] which represents an upper bound on the maximum error of the output, and adopt their error bounds propagation method to it.

3 Algebraic Analysis of Carlini *et al.*'s Hard-label Attack

In the hard-label setting, the only change of the model output that the adversary can observe is label switching. Hence, the boundary between pairs of class labels, which is called the *decision boundary*, can be detected. At EUROCRYPT 2025 [7], Carlini *et al.* observed that the decision boundary of a ReLU network is composed of piecewise-linear regions and it bends precisely when intersecting a *critical hyperplane* due to a single neuron's state change (see Fig. 1). This critical hyperplane consists of the critical points that are otherwise undetectable in this setting. Note that points on the two adjacent decision boundary patches share the same neuron states, except for the state of the neuron corresponding to the critical hyperplane. As shown in Fig. 1, in the input space, the intersection of the two $(d^{(1)} - 1)$-dimension decision boundary patches yields a $(d^{(1)} - 2)$-dimension subspace (called the *dual space*), which is entirely contained in the $(d^{(1)} - 1)$-dimension critical hyperplane (the yellow region).

3.1 Algebraic View on Carlini *et al.*'s Signature Recovery

Definition 8. (Decision Boundary) *The decision boundary consists of points $X^{(1)} \in \mathbb{R}^{d^{(1)}}$ so that two logit outputs of the FCN reach the same maximum value.*

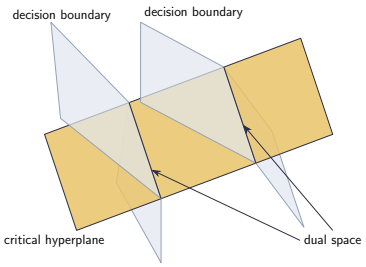


Fig. 1: Intersection of critical hyperplane and dual space with decision boundaries

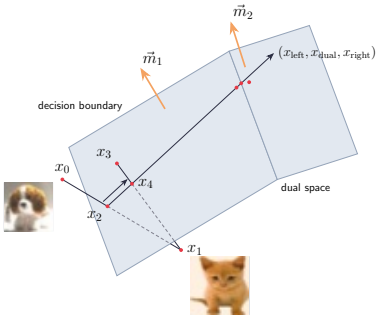


Fig. 2: Carlini *et al.*'s method for finding dual spaces [7]

Suppose the first $k-1$ layers are extracted. By Def. 3, for any input $X^{(1)}$, the forward mapping matrices F^{k-1}, B^{k-1} are known, yielding $X^{(k)} = F^{k-1}X^{(1)} + B^{k-1}$. If $X^{(1)}$ lies on the decision boundary where the network's the i-th and j-th outputs, $\mathcal{F}_i(X^{(1)})$ and $\mathcal{F}_j(X^{(1)})$, achieve the same maximum value, we have:

$$\mathcal{F}_i(X^{(1)}) - \mathcal{F}_j(X^{(1)}) = (G_i^{k+1} - G_j^{k+1})I^{(k)}(A^{(k)}X^{(k)} + B^{(k)}) + B_i^{k+1} - B_j^{k+1} = 0, \quad (9)$$

where G_i^{k+1} and B_i^{k+1} are the i-th rows of matrices G^{k+1} and B^{k+1}, respectively. Let $\vec{m} = (G_i^{k+1} - G_j^{k+1})I^{(k)}A^{(k)}F^{k-1} \in \mathbb{R}^{d^{(1)}}$, which is the normal vector of the decision boundary in the input space of the network.

Definition 9. (Critical Hyperplane) *For the t-th neuron in layer k, its critical hyperplane consists of this neuron's critical points $X^{(1)} \in \mathbb{R}^{d^{(1)}}$.*

If $X^{(1)}$ is a critical point corresponding to the t-th neuron in layer k, then

$$A_t^{(k)}X^{(k)} + B_t^{(k)} = 0. \quad (10)$$

At this point, the t-th diagonal entry of the activation matrix $I^{(k)}$ can be either 1 or 0 without altering the output, as the neuron always outputs 0. Consequently the corresponding activation matrices can be $I_+^{(k)}$ or $I_-^{(k)}$, differing only at the t-th diagonal entry (1 in $I_+^{(k)}$ and 0 in $I_-^{(k)}$). Then, we have:

$$G^{k+1}(I_+^{(k)} - I_-^{(k)})(A^{(k)}X^{(k)} + B^{(k)}) = 0. \quad (11)$$

Definition 10. (Dual Points) *Dual points are located at the intersection of a critical hyperplane and the decision boundary patches on both sides. Consequently, dual points fulfill the properties of both decision boundary points and critical hyperplane points (*i.e.*, Eqs. (9) and (10)).*

If $X^{(1)}$ is a dual point corresponding to the t-th neuron in layer k, and is located on the decision boundary between class i and class j, it satisfies:

$$(G_i^{k+1} - G_j^{k+1})(I_+^{(k)} - I_-^{(k)})(A^{(k)}X^{(k)} + B^{(k)}) = 0. \quad (12)$$

Definition 11. (Dual Space) *A dual space is the subspace spanned by dual points corresponding to the same neuron and the same decision boundary.*

Both a decision boundary and a critical hyperplane are governed by a single constraint (Eqs. (9) and (10) respectively), making them the $(d^{(1)}-1)$-dimension subspaces in the input space. Consequently, the dual space, constrained by both equations simultaneously, is $(d^{(1)}-2)$-dimension.

By employing the *Dual Points Finding* algorithm proposed by Carlini *et al.* in [7, Section 4], upon collecting a dual point x_{dual}, as shown in Fig. 2, we simultaneously obtain two adjacent points x_{left} and x_{right}. While x_{dual} lies precisely on the intersection of the decision boundary and the critical hyperplane, x_{left} and x_{right} reside on the two distinct decision boundary patches on either side of this critical hyperplane. From these points, we can extract the local normal vectors to the adjacent decision boundary patches at inputs x_{left} and x_{right}, denoted as $\vec{m}_\ell$ and $\vec{m}_r$, respectively as follows. According to [9,7], starting from a point on the decision boundary (*e.g.*, x_{left}), we take a small step α_0 along a reference basis direction $\vec{e_0}$ to step off the boundary, and then perform a binary search along another orthogonal basis direction $\vec{e_i}$ to find the distance α_i required to return to the decision boundary. By computing the ratio $\frac{\alpha_i}{\alpha_0}$, we obtain the relative coordinate weights of the normal vector. Iterating this procedure across all $d^{(1)}$ dimensions allows us to completely reconstruct the directions of $\vec{m}_\ell$ up to an unknown scalar factor, *i.e.*,

$$\widehat{\vec{m}}_\ell = (1, -\frac{\alpha_0}{\alpha_1}, \ldots, -\frac{\alpha_0}{\alpha_{d^{(1)}-1}}). \tag{13}$$

For arbitrary points $X^{(1)}$ and $X'^{(1)}$ sampled from the same dual space, their difference effectively eliminates the constant terms of Eq. (12):

$$(G_i^{k+1} - G_j^{k+1})(I_+^{(k)} - I_-^{(k)})A^{(k)}F^{k-1}(X^{(1)} - X'^{(1)}) = 0. \tag{14}$$

Note that the term $I_+^{(k)} - I_-^{(k)}$ yields a diagonal matrix with a single non-zero entry (value 1) at the (t,t)-th position, effectively acting as a selection operator. Let the row vector $G_i^{k+1} - G_j^{k+1} = (g_0, g_1, \ldots, g_{d^{(k+1)}-1}) \in \mathbb{R}^{d^{(k+1)}}$, then,

$$(G_i^{k+1} - G_j^{k+1})(I_+^{(k)} - I_-^{(k)})A^{(k)} = g_t A_t^{(k)}. \tag{15}$$

Let $\vec{\delta} = X^{(1)} - X'^{(1)}$. Then, Eq. (14) can be simplified as:

$$g_t A_t^{(k)} F^{k-1} \vec{\delta} = 0, \tag{16}$$

where $g_t A_t^{(k)}$ represents the weight vector of the target t-th neuron, scaled by an unknown constant g_t.

Remark. In the hard-label setting, the model extraction attack can still be regarded as a differential attack, where the difference between inputs from the same

dual space is studied as it propagates through the sequential cryptographic/DNN operations. For a given input difference $\vec{\delta}$, the corresponding output difference—specifically defined as the difference between the logit gaps $\mathcal{F}_i - \mathcal{F}_j$ in Eq. (9) of the two respective inputs—deterministically evaluates to 0.

Is a Single Dual Space Sufficient for One Neuron Weight Recovery? By sampling n (n is slightly larger than $d^{(k)}$) different dual points within the same dual space, compute $\vec{\delta}_j = X^{(1)} - X'^{(1)}$ $(0 \leq j \leq n-1)$, and construct a system of equations of Eq. (16): $(F^{k-1}\vec{\delta}_j)^{\mathsf{T}} g_t A_t^{(k)\mathsf{T}} = 0$. To obtain a unique non-trivial solution (up to a scalar) for this linear system, the $\{F^{k-1}\vec{\delta}_j\}_{0\leq j\leq n-1}$ must attain a rank of $d^{(k)} - 1$. Note that $F^{k-1}\vec{\delta}_j$ is the difference between $X^{(k)}$ and $X'^{(k)}$. Since both $X^{(k)}$ and $X'^{(k)}$ are sampled from the same dual space, they are simultaneously subject to the constraints of the decision boundary (Def. 8) and the critical hyperplane (Def. 9). Geometrically, this intersection forms at most a $(d^{(k)} - 2)$-dimension subspace (assuming all neurons in layer $k-1$ are activated, it is a $(d^{(k)} - 2)$-dimension subspace). Hence, the matrix formed by any subset of the differential vectors $\{X^{(k)} - X'^{(k)}\}$ derived from a single dual space has a rank of at most $d^{(k)} - 2$. Hence, we need at least two dual spaces corresponding to the same neuron to recover its weight vector.

Clustering Dual Points from Different Dual Spaces. Given a set of dual points collected thus far, the next step is to cluster them into disjoint groups, where each group corresponds to a distinct target neuron. To achieve this, Carlini *et al.* [7, Section 5.2] introduced the *ISCONSISTENT* algorithm, which verifies whether two dual spaces D_1 and D_2 correspond to the same neuron by analyzing the rank of their union. Let F^{k-1} and F'^{k-1} denote the corresponding local forward mapping matrices associated with D_1 and D_2, respectively. Suppose we sample differential vectors $\vec{\delta}_0, \ldots, \vec{\delta}_{s_1-1}$ from D_1, and $\vec{\delta}'_0, \ldots, \vec{\delta}'_{s_2-1}$ from D_2 (s_1 and s_2 are slightly larger than $d^{(k)}$). We stack each set of vectors: $S = \left(\vec{\delta}_0, \ldots, \vec{\delta}_{s_1-1}\right)^{\mathsf{T}}$, $S' = \left(\vec{\delta}'_0, \ldots, \vec{\delta}'_{s_2-1}\right)^{\mathsf{T}}$. The problem of determining consistency is thus reduced to checking whether the following homogeneous system of equations admits a non-trivial solution:

$$\begin{pmatrix} SF^{k-1\mathsf{T}} \\ S'F'^{k-1\mathsf{T}} \end{pmatrix} g_t A_t^{(k)\mathsf{T}} = 0. \tag{17}$$

Let $\Lambda_1^{(k-1)}$ and $\Lambda_2^{(k-1)}$ denote the sets of active neurons at layer $k-1$ for the dual points in D_1 and D_2, respectively. And let M_{SF} denote the coefficient matrix multiplying $g_t A_t^{(k)\mathsf{T}}$ in Eq. (17). As discovered by Carlini *et al.* [8, Section 4.3.2], the rank of the matrix is usually not full since the ReLU functions suppress negative outputs, *i.*e., some columns of M_{SF} are zero column vectors. Following the criterion established in [7, Section 5.2]:

- If D_1 and D_2 belong to the same neuron, $rank(M_{SF}) < |\Lambda_1^{(k-1)} \cup \Lambda_2^{(k-1)}|$.
- If D_1 and D_2 belong to different neurons, $rank(M_{SF}) = |\Lambda_1^{(k-1)} \cup \Lambda_2^{(k-1)}|$.

By evaluating this rank condition, all sampled dual points are clustered.

3.2 Limitations in Black-box Setting

As claimed by Carlini *et al.* in [7, Section 7.1], the clustering step presents a bottleneck in implementations. In contrast, the other phases of the attack were optimized and empirically viable in the black-box setting. According to [7], when using a 256-core server with GPU support to recover a 3072-256×3-64-10 FCN (referred to as *3072-FCN*) with four hidden layers, trained on CIFAR10, finding the dual points (accelerated by computing white-box gradients symbolically rather than via black-box binary search) and unifying the corresponding dual spaces (neglecting clustering) took approximately 16 hours, while the sign recovery phase was heavily parallelized, requiring only 8.5 hours. Consequently, the clustering becomes the bottleneck of the black-box model extraction attack. The procedure involves two main computational costs:

1. **Pairwise Comparisons:** Clustering n collected dual points requires $\mathcal{O}(n^2)$ independent invocations of the *ISCONSISTENT* algorithm.
2. **Rank Computation:** Each *ISCONSISTENT* call computes the rank of the $(s_1 + s_2) \times d^{(k)}$ projected matrix M_{SF} according to Eq. (17), where s_1 and s_2 are slightly larger than $d^{(k)}$. Consequently, performing Singular Value Decomposition (SVD) on this matrix costs $\mathcal{O}((d^{(k)})^3)$.

Multiplying these factors, the time complexity of Carlini *et al.*'s clustering [7] is:

$$\mathcal{O}(n^2(d^{(k)})^3). \tag{18}$$

For deeper layers, recovering near-dead neurons typically requires millions of dual points. Therefore, the black-box extraction becomes extremely time-consuming.

4 Improving Clustering by the Normal Vectors

4.1 Compute $\vec{G}_\ell$ and $\vec{G}_r$ up to a scalar

Given a dual point x corresponding to the t-th neuron in layer k. Recall Eq. (15) and denote the row vectors

$$\begin{aligned} \vec{G}_\ell &= (G_i^{k+1} - G_j^{k+1}) I_+^{(k)} A^{(k)} = \sum_{s \in \Lambda^{(k)}} g_s A_s^{(k)} + g_t A_t^{(k)}, \\ \vec{G}_r &= (G_i^{k+1} - G_j^{k+1}) I_-^{(k)} A^{(k)} = \sum_{s \in \Lambda^{(k)}} g_s A_s^{(k)}. \end{aligned} \tag{19}$$

where $\Lambda^{(k)}$ denotes the index set of active neurons in layer k for the dual point x. Then the difference between the row vectors $\vec{G}_\ell$ and $\vec{G}_r$ is the effective weight

vector of the target neuron t, scaled by a constant factor g_t, *i.e.*, $\vec{G}_\ell - \vec{G}_r = g_t A_t^{(k)}$. *Note that the vectors $\vec{G}_\ell$ and $\vec{G}_r$ also correspond to the normal vectors of the decision boundary patches on the left (activated) and right (inactivated) sides of the critical hyperplane, respectively.* Since we are operating in the hard-label setting, the exact magnitudes of the gradients $\vec{G}_\ell$ and $\vec{G}_r$ are inaccessible. Specifically, according to Def. 8 and Eq. (19), the normal vector to the decision boundary (*e.g.*, on the left patch) $\vec{m}_\ell = (G_i^{k+1} - G_j^{k+1}) I_+^{(k)} A^{(k)} F^{k-1}$, then

$$F^{k-1^\mathsf{T}} \vec{G}_\ell^\mathsf{T} = \vec{m}_\ell^\mathsf{T}. \tag{20}$$

This constructs a linear system with an unknown vector $\vec{G}_\ell^\mathsf{T}$. Recall from Eq. (13) that, the normal vectors $\vec{m}_\ell$ and $\vec{m}_r$ can be recovered up to a scalar multiple. Consequently, we can recover $\vec{G}_\ell^\mathsf{T}$ up to a scalar multiple from $\vec{m}_\ell^\mathsf{T}$ by solving this system. Let $\Lambda^{(k-1)}$ denote the index set of active neurons at layer $k-1$. For a vector $\vec{v}$, $\vec{v}[\Lambda^{(k-1)}]$ denotes the subvectors indexed by $\Lambda^{(k-1)}$. We will discuss how to solve the $\vec{G}_\ell^\mathsf{T}$ under 3 cases in the following section.

4.2 Normal Rank Check: Improved Rank Check by Normal Vectors

Case 1: $d^{(k)} = \dim(\vec{G}_\ell^\mathsf{T}) = \mathrm{rank}(F^{k-1}) = |\Lambda^{(k-1)}|$. In this case, the full vector $\vec{G}_\ell^\mathsf{T}$ in Eq. (20) is recovered up to a scalar. By normalizing the solution, we extract its unit normal row vector: $\vec{n}_\ell = \frac{\vec{G}_\ell}{\|\vec{G}_\ell\|}$. A similar procedure is applied to the right patch to compute $\vec{n}_r = \frac{\vec{G}_r}{\|\vec{G}_r\|}$. Consequently, according to Eq. (19), $g_t A_t^{(k)}$ can be expressed as a linear combination of the observed unit normal vectors:

$$g_t A_t^{(k)} = \|\vec{G}_\ell\| \cdot \vec{n}_\ell - \|\vec{G}_r\| \cdot \vec{n}_r, \tag{21}$$

which indicates that the target weight vector $g_t A_t^{(k)}$ lies within the 2-dimensional plane P spanned by $\vec{n}_\ell$ and $\vec{n}_r$, denoted as $P = span(\vec{n}_\ell, \vec{n}_r)$. Given another dual point belonging to the same neuron, similarly compute the normal vectors of its adjacent decision boundary patches, denoted as $\vec{n}'_\ell$ and $\vec{n}'_r$. Consequently, the corresponding target weight vector also satisfies $g_t A_t^{(k)} \in \mathrm{span}(\vec{n}'_\ell, \vec{n}'_r)$, which we denote as plane P'. If the two dual points correspond to the same neuron, $g_t A_t^{(k)}$ resides in both P and P', *i.e.*, these two planes must intersect along at least a 1-dimensional line (which corresponds to the direction of $A_t^{(k)}$). By constructing a combined matrix $N = (\vec{n}_\ell^\mathsf{T}, \vec{n}_r^\mathsf{T}, \vec{n}'^\mathsf{T}_\ell, \vec{n}'^\mathsf{T}_r)^\mathsf{T} \in \mathbb{R}^{4 \times d^{(k)}}$, the consistency between the two dual points can be evaluated via the rank of N:

- If the two dual points belong to the same neuron, then $\mathrm{rank}(N) < 4$;
- If the two dual points belong to different neurons, then $\mathrm{rank}(N) = 4$.

In practice, we employ Singular Value Decomposition (SVD) to compute the minimum singular value $\sigma_{\min}$ of N. Then the rank condition is converted into verifying whether $\sigma_{\min}$ is sufficiently close to zero. Given a small threshold τ, if $\sigma_{\min} < \tau$, the two dual points are assumed to belong to the same neuron.

Case 2: $d^{(k)} > \mathrm{rank}(F^{k-1}) = |\Lambda^{(k-1)}|$. In this case, we can recover $\vec{G}_\ell^\mathsf{T}[\Lambda^{(k-1)}]$ from Eq. (20) up to a scalar multiple. Given the active sets $\Lambda^{(k-1)}$ and $\Lambda'^{(k-1)}$ for dual points x and x', we compute and normalize its normal vectors on their respective active sets (*e.g.*, $\vec{n}_\ell = \frac{\vec{G}_\ell[\Lambda^{(k-1)}]}{\|\vec{G}_\ell[\Lambda^{(k-1)}]\|}$). When comparing the two dual points, we extract their shared active subset $\bar{\Lambda}^{(k-1)} = \Lambda^{(k-1)} \cap \Lambda'^{(k-1)}$, slice the raw vectors according to $\bar{\Lambda}^{(k-1)}$ and re-normalize them:

$$\vec{n}_\ell = \frac{\vec{G}_\ell[\bar{\Lambda}^{(k-1)}]}{\|\vec{G}_\ell[\bar{\Lambda}^{(k-1)}]\|}, \quad \vec{n}_r = \frac{\vec{G}_r[\bar{\Lambda}^{(k-1)}]}{\|\vec{G}_r[\bar{\Lambda}^{(k-1)}]\|}. \tag{22}$$

Then we perform the rank checking as Case 1. Note that the check requires $|\bar{\Lambda}^{(k-1)}| \geq 4$. Otherwise, the matrix $N \in \mathbb{R}^{4\times|\bar{\Lambda}^{(k-1)}|}$ necessarily has rank smaller than 4. We therefore exclude pairs with $|\bar{\Lambda}^{(k-1)}| < 4$ from the consistency check.

Case 3: $d^{(k)} > |\Lambda^{(k-1)}| > \mathrm{rank}(F^{k-1})$. For Eq. (20), if $\mathrm{rank}(F^{k-1}) < |\Lambda^{(k-1)}|$, it admits multiple solutions. We use the least-squares solver to obtain its minimum-norm solution, denoted by $\vec{h}$. Let $\vec{h}'$ be any other solution and define $\vec{z} = \vec{h}' - \vec{h}$. Since both vectors $\vec{h}$ and $\vec{h}'$ satisfy Eq. (20), we have

$$(F^{k-1})^\mathsf{T}\vec{z} = (F^{k-1})^\mathsf{T}(\vec{h}' - \vec{h}) = \vec{0}. \tag{23}$$

Hence, $\vec{z} \in \ker((F^{k-1})^\mathsf{T})$. For any scalar α, $\vec{h} + \alpha\vec{z}$ is also a solution. If $\vec{h}^\mathsf{T}\vec{z} \neq 0$, choosing $\alpha = -\frac{\vec{h}^\mathsf{T}\vec{z}}{\|\vec{z}\|^2}$ gives $\|\vec{h} + \alpha\vec{z}\|^2 = \|\vec{h}\|^2 - \frac{(\vec{h}^\mathsf{T}\vec{z})^2}{\|\vec{z}\|^2} < \|\vec{h}\|^2$, which contradicts the fact that $\vec{h}$ is the minimum-norm solution. Consequently, $\vec{h}^\mathsf{T}\vec{z} = 0$, *i.e.*, $\vec{h}$ is orthogonal to this kernel space.

For kernel space $\ker((F^{k-1})^\mathsf{T})$, let $q = |\Lambda^{(k-1)}| - \mathrm{rank}(F^{k-1})$ denote its dimension. Thus, q represents the number of directions that cannot be recovered from Eq. (20). Since the original $\vec{G}_\ell^\mathsf{T}$ is also a solution, the difference between $\vec{G}_\ell^\mathsf{T}$ and $\vec{h}$ must lie in this kernel space. Therefore, there exists $\vec{z}_0 \in \ker((F^{k-1})^\mathsf{T})$ such that $\vec{G}_\ell^\mathsf{T} = \vec{h} + \vec{z}_0$. Since $\vec{h}$ is orthogonal to $\vec{z}_0$, $\vec{h}$ is exactly the projection of $\vec{G}_\ell^\mathsf{T}$ obtained by discarding its kernel-space component $\vec{z}_0$. Therefore, when q is small and $\|\vec{z}_0\|/\|\vec{G}_\ell^\mathsf{T}\|$ is small, only a small number of directions and a small component of $\vec{G}_\ell^\mathsf{T}$ are lost. In this case, $\vec{h}$ still provides a good approximation to the direction of $\vec{G}_\ell^\mathsf{T}$ and can be used in the subsequent comparison.

Experiments on Consistency Verification Using the Minimum Singular Value $\sigma_{\min}$. Our experiment targets the same *3072-FCN* model used by Carlini *et al.* [7] and also given in Sect. 3.2. We sample 10,000 dual points across the input space. For each layer, we filter the points from the previous layers. Then evaluate two dual points by extracting their corresponding $\vec{n}_\ell$ and $\vec{n}_r$ from Eq. (22), and calculate the minimum singular value $\sigma_{\min}$ of $N = (\vec{n}_\ell^\mathsf{T}, \vec{n}_r^\mathsf{T}, \vec{n}_\ell'^\mathsf{T}, \vec{n}_r'^\mathsf{T})^\mathsf{T} \in \mathbb{R}^{4\times d^{(k)}}$.

As illustrated in Fig. 3, the $\sigma_{\min}$ from pairs of dual points corresponding to the *same* neuron in the target layer are all smaller than 10^{-8}, while that from

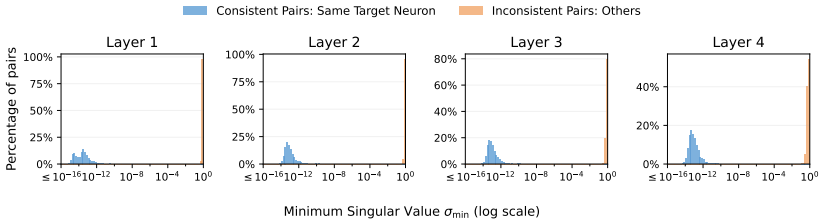


Fig. 3: Distribution of $\sigma_{\min}$ for same-target-neuron and other dual-point pairs

other pairs are close to 1. This separation on distribution of $\sigma_{\min}$ demonstrates the metric provides a reliable criterion for clustering dual points. To quantify the practical effectiveness of the consistency check, we evaluate its precision and recall during clustering:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}.$$

A "True Positive (TP)" occurs when both dual points corresponding to the same neuron in the target layer are identified as consistent. A "False Positive (FP)" arises when two dual points from different neurons, or the same neuron in another layer, are mistakenly clustered together. Conversely, a "False Negative (FN)" indicates two dual points belonging to the same target-layer neuron fail to be clustered together. "Precision" measures how reliably the consistency check avoids grouping unrelated dual points and maintains cluster purity, while "Recall" evaluates how effectively it recognizes true target-neuron pairs, thereby reducing missed consistent pairs.

- In Layer 1, every dual point meets Case 1 condition, and therefore the $\vec{G}_\ell^{\mathsf{T}}$ and $\vec{G}_r^{\mathsf{T}}$ are fully recovered. All dual points are clustered with 99.98% Precision and 100 % Recall.
- In Layer 2, the $\vec{G}_\ell^{\mathsf{T}}[\Lambda^{(k-1)}]$ and $\vec{G}_r^{\mathsf{T}}[\Lambda^{(k-1)}]$ are uniquely recoverable with no rank deficiency (Case 2). Only 0.0005% of the pairs meet $|\bar{\Lambda}^{(k-1)}| < 4$. Excluding these pairs, the consistency check achieves 99.96% Precision and 100% Recall.
- In Layer 3, let q denote the rank deficit $q = |\Lambda^{(k-1)} - \text{rank}(F^{k-1})|$. Among the remaining candidate points, only 1.44% have $q > 0$ (Case 3). We retain all candidates, using minimum-norm solutions for rank-deficient systems, and obtain 99.95% precision and 99.36% recall.
- In Layer 4, 14.23% have $q > 0$, Applying the same procedure without rank-deficiency filtering yields 100% precision and 88.55% recall.

4.3 Faster Consistency Checking: Approximate Signature Vector

This section introduces a more efficient method called *Approximate Signature Vector* (*ASV*) method using unit normal row vectors $\vec{n}_\ell, \vec{n}_r$ derived in Sect. 4.2

Definition 12. (Approximate Signature Vector, ASV) *Given $\vec{n}_\ell, \vec{n}_r$, the approximate signature vector (ASV) $\vec{v}$ is defined as:*

$$\vec{v} = \vec{n}_\ell - (\vec{n}_\ell \cdot \vec{n}_r^{\mathsf{T}})\vec{n}_r. \tag{24}$$

Let $\vec{\omega} = g_t A_t^{(k)}$. Then, $\vec{n}_\ell = \frac{1}{\|\vec{G}_\ell\|}(\vec{\omega} + \|\vec{G}_r\|\vec{n}_r)$. Since $\vec{n}_r\vec{n}_r^{\mathsf{T}} = 1$, Eq. (24) becomes

$$\vec{v} = \frac{1}{\|\vec{G}_\ell\|}(\vec{\omega} - (\vec{\omega} \cdot \vec{n_r}^{\mathsf{T}})\vec{n_r}). \tag{25}$$

Directional Distance between ASV $\vec{v}$ and Neuron Weight $\vec{\omega}$. Eq. (25) demonstrates that the ASV $\vec{v}$ deviates from the exact scaled weight $\vec{\omega}$ solely by the term $(\vec{\omega} \cdot \vec{n}_r^{\mathsf{T}})\vec{n}_r$. To measure the precision of this directional alignment, we employ the cosine similarity $\cos\varphi_{\vec{v},\vec{\omega}} = \frac{\vec{v}\cdot\vec{\omega}^{\mathsf{T}}}{\|\vec{v}\|\|\vec{\omega}\|}$, and introduce the directional distance $\mathsf{DGap}(\vec{v}, \vec{\omega})$ of $\vec{v}, \vec{\omega}$ by the following cosine distance

$$\mathsf{DGap}(\vec{v}, \vec{\omega}) = 1 - |\cos\varphi_{\vec{v},\vec{\omega}}| = 1 - \frac{|\vec{v} \cdot \vec{\omega}^{\mathsf{T}}|}{\|\vec{v}\| \cdot \|\vec{\omega}^{\mathsf{T}}\|}. \tag{26}$$

Following the assumption that the network is a unitary balanced network, which is also used in the analysis on the expected signal-to-noise ratio of *Neuron Wiggle* by Canales-Martínez *et al.* [6, Appendix A] at EUROCRYPT 2024, we propose the following proposition.

Proposition 1. *Assuming the network is a unitary balanced network, and each neuron weight is a randomly chosen unit vector in $d^{(k)}$-dimensional space. We have*

$$\mathbb{E}\big[\mathsf{DGap}(\vec{v}, \vec{\omega})\big] \leq \frac{1}{d^{(k)}}, \qquad \lim_{d^{(k)}\to\infty} \frac{1}{d^{(k)}} = 0. \tag{27}$$

Proof. See Supp. D. □

Proposition 1 shows that the expectation $\mathbb{E}[\mathsf{DGap}(\vec{v}, \vec{\omega})]$ approaches zero, and hence the ASV $\vec{v}$ and the neuron weight $\vec{w}$ are nearly parallel in direction, when the weight dimension increases.

White-box Experiment for Verification of $\mathsf{DGap}(\vec{v}, \vec{\omega}) \to 0$. For the *3072-FCN*, we randomly sample 10,000 dual points. For each point, compute its ASV $\vec{v}$ and DGap between $\vec{v}$ and the ground-truth weights of all neurons within the same layer. As shown in Fig. 4, the resulting distributions exhibit a clear separation: the cosine distance between an ASV and its corresponding target neuron $\vec{\omega}$ tightly clusters near 0, while the distances to other non-target neurons $\vec{\omega}'$ consistently concentrate approach 1.

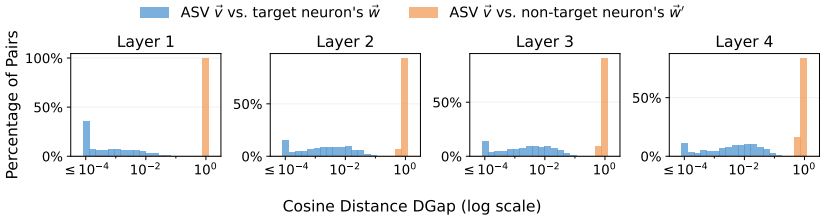


Fig. 4: Distribution of (DGap) between ASV ($\vec{v}$) and neuron weights ($\vec{w}$)

Consistency Check using ASVs. If two dual points belong to the same neuron, the cosine similarities between their ASVs (denoted as $\vec{v}_1$ and $\vec{v}_2$) and the ground-truth weight $\vec{\omega}$ both approach 0. According to the triangle inequality for spatial angles [4, Chapter I.1],

$$0 \leq \varphi_{\vec{v}_1,\vec{v}_2} \leq \varphi_{\vec{v}_1,\vec{\omega}} + \varphi_{\vec{v}_2,\vec{\omega}}. \tag{28}$$

Since both $\varphi_{\vec{v}_1,\vec{\omega}}$ and $\varphi_{\vec{v}_2,\vec{\omega}}$ approach 0, $\varphi_{\vec{v}_1,\vec{v}_2}$ will also approach 0.

Property 3. (Parallel ASVs) *If two dual points belong to the same neuron, their corresponding approximate signature vectors $\vec{v}_1$ and $\vec{v}_2$ will be nearly identical or opposite (both are nearly collinear with the underlying weight vector $\vec{\omega}$). Therefore, their consistency score is bounded by a small threshold τ:*

$$\textit{consistency score} = \mathsf{DGap}(\vec{v_1}, \vec{v_2}) = 1 - |\cos\varphi_{\vec{v}_1,\vec{v}_2}| = 1 - \frac{|\vec{v}_1\vec{v}_2^{\mathsf{T}}|}{\|\vec{v}_1\|\|\vec{v}_2\|} < \tau. \tag{29}$$

Immunity to Activation State Blindness. One potential concern in the hard-label setting is the inability to directly observe the activation states of the patches $X^{(1)}_{left}$ and $X^{(1)}_{right}$. Given a dual point, we can extract two distinct normal vectors from its adjacent patches as $\vec{n}_1$ and $\vec{n}_2$, but we cannot ascertain which corresponds to the activated normal $\vec{n}_\ell$ and which to the inactivated normal $\vec{n}_r$. To bypass this ambiguity, we compute two ASVs for each dual point:

$$\vec{v} = \vec{n}_1 - (\vec{n}_1 \cdot \vec{n}_2^{\mathsf{T}})\vec{n}_2, \quad \vec{v}' = \vec{n}_2 - (\vec{n}_2 \cdot \vec{n}_1^{\mathsf{T}})\vec{n}_1. \tag{30}$$

Given two distinct dual points, we obtain two pairs of ASVs: $(\vec{v}_1, \vec{v}_1')$ and $(\vec{v}_2, \vec{v}_2')$. To determine if these two dual points belong to the same neuron, we evaluate the consistency score $\mathsf{DGap}(\cdot,\cdot)$ across all possible cross-combinations. The final consistency score $\mathcal{S}$ between the two dual points is defined as the minimum value among these combinations:

$$\mathcal{S} = \min\Big\{\mathsf{DGap}(\vec{v}_1, \vec{v}_2), \mathsf{DGap}(\vec{v}_1, \vec{v}_2'), \mathsf{DGap}(\vec{v}_1', \vec{v}_2), \mathsf{DGap}(\vec{v}_1', \vec{v}_2')\Big\}. \tag{31}$$

Since the ASVs closely approximate the true target weight, the correct pairing will minimize the score $\mathcal{S}$, effectively resolving the activation state blindness.

Experiments on Consistency Verification. To verify the practical performance, we conducted an experiment on *3072-FCN*. As illustrated in Fig. 5, evaluating the consistency score $\mathcal{S}$ on pairs of dual points corresponding to the *same* neuron yields a significantly different distribution compared to pairs originating from *different* neurons. This separation on distribution demonstrates that the ASV-based metric provides a reliable criterion for clustering dual points.

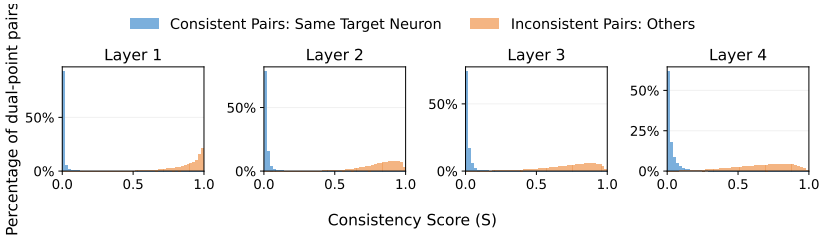


Fig. 5: Distribution of the Consistency Scores $\mathcal{S}$ of ASV pairs from the same versus different neurons across DNN layers

We observe a performance threshold regarding the shared active dimensions: when the number of shared active dimensions ($\bar{\Lambda}^{(k-1)}$ defined in Eq. (22) in Sect. 4.2) is small ($|\bar{\Lambda}^{(k-1)}| < 10$), the clustering precision is severely degraded (20.03% for Layer 3 and 25.81% for Layer 4). Conversely, when the shared dimensions $|\bar{\Lambda}^{(k-1)}|$ exceed this threshold, the precision leaps to 98.92% and 94.60%, respectively. Therefore, only the pairs of dual points with $|\bar{\Lambda}^{(k-1)}| \geq 10$ will be tested by the the ASV consistency check, and about 91.10% and 78.58% dual points belonging to Layers 3 and 4 are clustered after the ASV consistency check, and the remain dual points are discarded.

Remark on The Choice of the Small Threshold τ. As shown in Fig. 5, the small threshold τ serves as a separation boundary between the distributions of consistent and inconsistent pairs. An excessively large τ admits false positives, whereas an overly strict τ yields false negatives. In practice, τ can be determined empirically. We first experiment on a small subset of dual points. By computing the consistency score $\mathcal{S}$ for all pairs within this subset, we analyze their distribution. Similar to the phenomenon shown in Fig. 5, but without access to ground-truth neuron indices, the scores naturally form a distinct bimodal distribution. This natural separation allows us to select a threshold τ that cleanly divides the two groups. For example, the Fig. 5 indicates a sound $\tau = 0.2$. Then we apply this τ to cluster the full dataset of dual points.

4.4 The Steps and Complexities of NRC and ASV Clustering

Both NRC and ASV share the following steps:

- **Step 1. Filtering and Normal Extraction.** Given a collection of dual points, first filter out the points of previous layers. For each remaining point, compute $\vec{n}_\ell$ and $\vec{n}_r$ of its adjacent decision boundary patches. For ASV method, calculate the ASV $\vec{v} = \vec{n}_\ell - (\vec{n}_\ell \cdot \vec{n}_r^{\mathsf{T}})\vec{n}_r$.
- **Step 2. Seed Selection and Testing.** Randomly select a unclustered dual point from the candidate set, and evaluate it with all other points using NRC or ASV methods.
- **Step 3. Cluster and Remove.** All dual points that yield a minimum singular value $\sigma_{\min} < \tau_{\text{NRC}}$ for NRC or consistency score $\mathcal{S} < \tau_{\text{ASV}}$ for ASV are grouped into the same cluster as the seed. Once clustered, these points are strictly removed from the candidate pool. We then iteratively repeat Step 2, selecting a new seed from the shrinking pool, until all dual points are successfully partitioned.

Complexity. The NRC and ASV methods have similar complexities:

- **Complexity of Step 1:** Computing the normal vectors requires solving local linear equations, taking $\mathcal{O}((d^{(k)})^3)$ operations per dual point. For n dual points, this preprocessing strictly bounds to $\mathcal{O}(n(d^{(k)})^3)$.
- **Complexity of Step 2 and 3:** For NRC, the cost of applying SVD to 4 vectors is $\mathcal{O}(d^{(k)})$. For ASV, the vector inner product costs $\mathcal{O}(d^{(k)})$.Therefore, both of these time complexities are strictly lower than the rank check given by Carlini *et al.* in Eq. (17), which has a time complexity of $\mathcal{O}((d^{(k)})^3)$.
  For NRC and ASV, the worst-case matching scales as $\mathcal{O}(n^2 d^{(k)})$. In average case, we assume the n points are uniformly distributed across $d^{(k)}$ hidden neurons, each neuron corresponds to a cluster of approximately $n/d^{(k)}$ points. The clustering procedure effectively proceeds in $d^{(k)}$ rounds with $\sum_{i=0}^{d^{(k)}-1} \left(n - i\frac{n}{d^{(k)}}\right) = \mathcal{O}(nd^{(k)})$ pairwise comparisons. Since each comparison takes $\mathcal{O}(d^{(k)})$ operations for NRC or ASV, this lowers the average-case time complexity to $\mathcal{O}(nd^{(k)} \cdot d^{(k)}) = \mathcal{O}(n(d^{(k)})^2)$.

Consequently, our method successfully reduces the clustering complexity from Carlini *et al.*'s original $\mathcal{O}(n^2(d^{(k)})^3)$ [7] to $\mathcal{O}(n(d^{(k)})^3 + n(d^{(k)})^2) \approx \mathcal{O}(n(d^{(k)})^3)$ in average case or $\mathcal{O}(n(d^{(k)})^3 + n^2 d^{(k)})$ in worst case by NRC or ASV .

4.5 Comparison and Combination of the NRC and ASV Methods

Both the NRC and ASV methods derive $\vec{n}_\ell$ and $\vec{n}_r$ from normal vectors already recorded during the *Dual Points Finding* phase [7], neither method incurs additional model queries during clustering. Carlini *et al.*'s method [7] adopts heavy SVD and NRC employs a very lightweight SVD. While, ASV method relies on efficient vector inner products rather than complex SVDs.

Dual Points Sharing the Same Decision Boundary. ASV method resolves a critical vulnerability in both the NRC and Carlini *et al.*'s [7] methods: if two

dual points share the same decision boundary patch, they inherently share the same normal vector (*e.g.*, $\vec{n}_r = \vec{n}'_\ell$). Although these two points belong to different neurons, the shared normal vector reduces the matrix rank, causing the consistency check to fail. For NRC and Carlini *et al.*'s method [7], we must identify if any pair of their normal vectors is nearly identical. In contrast, the computed ASVs directly approximate the target weight vectors in direction, it inherently eliminates the need for these redundant pairwise checks. Consequently, the ASV-based clustering is faster than NRC.

Conditions of ASV and the Combination of NRC and ASV. According to Proposition 1, the $\mathsf{DGap}(\vec{v}, \vec{\omega})$ of the ASV method becomes larger when $d^{(k)}$ decreases, which may lead to ineffective distinguishing between consistent dual point pairs and inconsistent pairs. Also ASV becomes ineffective when the the number of shared active neurons is extremely small (*e.g.*, $|\bar{A}^{(k-1)}| < 10$). Depending on the specific network architecture, extreme cases may arise when the network is relatively narrow.

For example, in our experiment on the relatively narrow network, a 64-64×4-10 FCN, after trained on the MNIST dataset, the average DGap between ASV $\vec{v}$ and its underlying neuron weight $\vec{w}$ rise to around 0.055-0.066 in Layers 2 and 3, compared to around 0.008 for the wide *3072-FCN* network in Table 5 in Supp. C, leading to a higher false positive rate for the ASV method. Favorable conditions arise at Layer 1, where ASVs are computed directly from the input-space normals $\vec{m}_\ell$ and $\vec{m}_r$, avoiding errors from recovering normals through F^{k-1}, and at Layer 4, where there are no deeper hidden neurons whose dual points could be mistakenly clustered into the target layer.

In general, ASV is particularly suitable for networks with high-dimensional weight spaces according to Proposition 1. For example, targeting the *3072-FCN*, ASV clusters 100,000 dual points and successfully recovers 100% and 92% of the target neurons for the first two layers (achieving partial black-box signature recovery in Table 2), identical to absolutely correct results using white-box clustering in Carlini *et al.*'s proof-of-concept experiment [7]. The NRC method does not rely on the high-dimensional weight vectors and provides a robust consistency check applicable to a wide range of narrow or wide architectures.

In narrow network like 64-64×4-10, NRC maintains a precision of over 99%. Therefore, in practical experiment on 64-64×4-10, a combined method of ASV-NRC is applied, *i.e.*, the Layer 1 and 4 are handled with ASV method and Layer 2 and 3 are handled with NRC method.

Experimental Comparisons between NRC, ASV, and Carlini *et al.*'s Methods. The experiment is performed on the *3072-FCN*. We evaluate the methods on sets of 200, 1000, and 2000 dual points. These sets comprise dual points of different neurons across 4 layers. We record the runtime and accuracy, including precision and recall. Note that, due to the substantial runtime overhead of Carlini *et al.*'s method, we only run their algorithm up to 200 dual points. For larger sets, their clustering performances are estimated by multiplying the time

of one pairwise comparison with the number of pairwise comparisons, denoted by "≈" as shown in Table 3.

Table 3: Comparison of clustering methods for layer 1. T denotes the computational runtime in seconds; n represents the number of accumulated dual points.

n	Approach	$T\,(s)$	Comparisons	Precision	Recall
200	Carlini's	$2^{15.37}$	$2^{14.17}$	100%	100%
	NRC	$2^{1.83}$	$2^{14.17}$	100%	100%
	ASV	$2^{-0.48}$	$2^{14.17}$	100%	100%
1000	Carlini's	$\approx 2^{19.75}$	$\approx 2^{18.55}$	-	-
	NRC	$2^{6.15}$	$2^{18.55}$	100%	100%
	ASV	$2^{3.96}$	$2^{18.55}$	100%	99.54%
2000	Carlini's	$\approx 2^{21.54}$	$\approx 2^{20.34}$	-	-
	NRC	$2^{8.05}$	$2^{20.35}$	100%	100%
	ASV	$2^{6.09}$	$2^{20.34}$	100%	99.90%

Remarks: Refining Clusters after NRC or ASV Clustering.

1. *False Positive.* Although the dual points are partitioned into distinct clusters, each cluster may contain a minor fraction of false positives. This case also happens to Carlini *et al.* [7], and they apply a secondary filter that evaluates triples of supposedly consistent dual points, discarding the entire triplet if an inconsistency is detected. This verification strategy can also be integrated into both the NRC and ASV methods. Moreover, we propose another trick: since the number of candidate points assigned to each cluster is relatively small, we can identify potential false positives by inspecting the residuals of the initial least-squares solution to Eq. (17) when recovering the signatures. We then discard the rows with anomalous residuals and recompute $g_t A_t^{(k)}$ using the remaining equations.
2. *False Negative.* In a small fraction of cases, a few dual point pairs from the same target-layer neuron exhibit relatively high minimum singular values $\sigma_{\min}$ or consistency scores $\mathcal{S}$, rendering them indistinguishable from different-neuron pairs under these metrics. We collect more dual points to ensure a sufficient number of valid samples within each target neuron cluster. *E.g.*, in Layer 4 of *3072-FCN*, we observe that both the NRC and ASV require about 1.2 times as many dual points as the original clustering of Carlini *et al.*'s method to yield a comparable number of clustered candidate points.

5 Kernel-Centric ASV in Hard-label CNN

5.1 *Dual RPCP* and *Dual PSP*

In the hard-label CNN, similar to Eq. (9), the input $X^{(1)}$ corresponding to a point on the decision boundary satisfies

$$\mathcal{F}_i(X^{(1)}) - \mathcal{F}_j(X^{(1)}) = (G_i^{k+1} - G_j^{k+1})P^{(k)}I^{(k)}(A^{(k)}X^{(k)} + B^{(k)}) + B_i^{k+1} - B_j^{k+1} = 0. \tag{32}$$

Then, we define the *Dual RPCP* and *Dual PSP* in a similar way.

Definition 13. (Dual ReLU-Pooling Critical Point, Dual RPCP) *Dual RPCP is in the intersection of a critical hyperplane of RPCP and the decision boundary patches on both sides. So, they meet Eq.* (32) *and Pro. 1.*

Like Def. 11 in FCNs, the Dual RPCP space is defined as the subspace spanned by Dual RPCPs of the same neuron and the same decision boundary. Suppose that a Dual RPCP $X^{(1)}$ corresponds to the t-th neuron in the k-th layer, covered by the u-th pooling window. Since the value of the neuron t is 0, the corresponding activation matrix can be either $I_+^{(k)}$ or $I_-^{(k)}$, differing only at the t-th diagonal entry (*e.g.*, 1 in $I_+^{(k)}$ and 0 in $I_-^{(k)}$). By subtracting the corresponding Eq. (32), we have

$$(G_i^{k+1} - G_j^{k+1})P^{(k)}(I_+^{(k)} - I_-^{(k)})(A^{(k)}X^{(k)} + B^{(k)}) = 0. \tag{33}$$

For any Dual RPCP $X^{(1)}$ and $X'^{(1)}$ sampled from the same Dual RPCP Space, by subtracting the corresponding Eq. (33), we have

$$(G_i^{k+1} - G_j^{k+1})P^{(k)}(I_+^{(k)} - I_-^{(k)})A^{(k)}F^{k-1}(X^{(1)} - X'^{(1)}) = 0. \tag{34}$$

According to Eq. (55) in Supp. A.1, if the neuron t is selected by the u-th pooling window, the u-th row vector $P_u^{(k)}$ of the pooling matrix $P^{(k)}$ contains a "1" at its t-th position and "0"s elsewhere. The term $I_+^{(k)} - I_-^{(k)}$ yields a diagonal matrix with a single non-zero entry (value 1) at the entry (t, t), effectively acting as a selection operator. Consequently, the product $P^{(k)}(I_+^{(k)} - I_-^{(k)})$ extracts the t-th column of the pooling matrix $P^{(k)}$. Since neuron t is covered by exactly one pooling window, this t-th column contains a single "1" at its u-th row. Let the row vector $G_i^{k+1} - G_j^{k+1} = (g_1, g_2, \cdots, g_{d^{(k+1)}})$. Then, we have:

$$(G_i^{k+1} - G_j^{k+1})P^{(k)}(I_+^{(k)} - I_-^{(k)})A^{(k)} = g_u A_t^{(k)}. \tag{35}$$

Let $\vec{\delta} = X^{(1)} - X'^{(1)}$, Eq. (34) can be simplified as: $g_u A_t^{(k)} F^{k-1}\vec{\delta} = 0$. Sample m directions $S = (\vec{\delta}_0, \vec{\delta}_1, \cdots, \vec{\delta}_{m-1})^\mathsf{T}$ from the same Dual RPCP Space to construct a system and recover $g_u A_t^{(k)}$. Specifically,

$$S{F^{k-1}}^\mathsf{T} g_u {A_t^{(k)}}^\mathsf{T} = 0. \tag{36}$$

Definition 14. (Dual Pooling Switching Point, Dual PSP) *The Dual PSP is in the intersection of a critical hyperplane of PSP and the decision boundary patches on both sides. Thus, they fulfill Eq. (32) and Pro. 2.*

Similarly, Dual PSP Space is defined as the subspace spanned by Dual PSPs corresponding to the same two competing neurons and the same decision boundary. Assume that the u-th ($0 \leq u \leq d^{(k+1)}-1$) pooling window includes neurons p and q ($p < q$) whose input values are $Y^{(k)}[p]$ and $Y^{(k)}[q]$ respectively, and that they share the same maximum value. According to Eq. (55) in Supp. A.1, each row $P_u^{(k)}$ of the pooling matrix $P^{(k)}$ contains a single "1" to indicate the selection of the maximum value for the u-th pooling window. Specifically, if the p-th neuron is selected, the row is $(0\cdots0, 1_p, 0\cdots0)$. Since $Y^{(k)}[p] = Y^{(k)}[q]$, the max pooling can arbitrarily select either neuron p or neuron q without altering its output. We can represent these two equivalent selecting operations using two distinct pooling matrices, $P_+^{(k)}$ and $P_-^{(k)}$, which differ only in their u-th row. Specifically, in $P_+^{(k)}$, this row is $(0\cdots0, 1_p, 0\cdots0, 0_q, 0\cdots0)$, while in $P_-^{(k)}$, the same row becomes $(0\cdots0, 0_p, 0\cdots0, 1_q, 0\cdots0)$. Consequently, the pooling matrix can be either $P_+^{(k)}$ or $P_-^{(k)}$. By substituting $P_+^{(k)}$ and $P_-^{(k)}$ into Eq. (32) respectively and subtracting the two resulting equations, we have

$$(G_i^{k+1} - G_j^{k+1})(P_+^{(k)} - P_-^{(k)})I^{(k)}(A^{(k)}X^{(k)} + b^{(k)}) = 0. \tag{37}$$

For any Dual PSP $X^{(1)}$ and $X'^{(1)}$ sampled from the same Dual PSP space, like Eq. (14) in DNN, subtracting the corresponding Eq. (37), we have

$$(G_i^{k+1} - G_j^{k+1})(P_+^{(k)} - P_-^{(k)})I^{(k)}A^{(k)}F^{k-1}(X^{(1)} - X'^{(1)}) = 0. \tag{38}$$

Note that the u-th row of $P_+^{(k)} - P_-^{(k)}$ is $(0\cdots0, 1_p, 0\cdots0, -1_q, 0\cdots0)$, and the other rows are zero vectors; let $G_i^{k+1} - G_j^{k+1} = (g_1, g_2, \cdots, g_{d^{(k+1)}})$, then

$$(G_i^{k+1} - G_j^{k+1})(P_+^{(k)} - P_-^{(k)})I^{(k)}A^{(k)} = g_u(A_p^{(k)} - A_q^{(k)}), \tag{39}$$

Let $\vec{\delta} = X^{(1)} - X'^{(1)}$, Eq. (38) can be simplified as: $g_u(A_p^{(k)} - A_q^{(k)})F^{k-1}\vec{\delta} = 0$. Sample n distinct directions $S = (\vec{\delta}_0, \vec{\delta}_1, \cdots, \vec{\delta}_{n-1})^\mathsf{T}$ from the same Dual PSP space to construct a system of equations and recover $g_u(A_p^{(k)} - A_q^{(k)})$:

$$SF^{k-1\mathsf{T}} \cdot g_u(A_p^{(k)} - A_q^{(k)}) = 0. \tag{40}$$

We refer to Dual RPCPs and Dual PSPs as *Dual Points*, and denote their corresponding Dual RPCP and Dual PSP Spaces simply as the *Dual Spaces*.

5.2 Kernel-Centric ASV Method

Limitations of *Neuron-Centric Clustering* for CNN. Recall from Eq. (17) in the *Clustering step* of Sect. 3.1 that the clustering principle is to ensure Eq.

(17) admits a unique non-trivial solution $g_t A_t^{(k)\mathsf{T}}$. Similarly clustering strategy applied to CNN is to cluster dual points and build similar system of Eq. (17) from Eq. (36) for Dual RPCPs or from Eq. (40) for Dual PSPs, and to ensure the system admits a unique non-trivial solution. Hence, one must collect at least two dual points belonging to the same category to make the system solvable, *i.e.*, the *neuron-centric clustering*. Specifically, it requires at least two Dual RPCPs associated with the same t-th neuron to recover $g_u A_t^{(k)}$ of Eq. (36), or at least two Dual PSPs corresponding to the same competing pair of neurons (p, q) within the same u-th pooling window to recover $g_u(A_p^{(k)} - A_q^{(k)})$ of Eq. (40).

Taking a toy layer with all symbols given in Fig. 6 of Supp. A as an example[3]: there are $d_f^{(k)} = 36$ distinct types of Dual RPCPs corresponding to 36 neurons. For Dual PSPs, since each 2×2 pooling window contains $\binom{4}{2} = 6$ possible competing neuron pairs, there are in total $d^{(k+1)} \times \binom{h_m^{(k)} \cdot w_m^{(k)}}{2} = 54$ distinct types of Dual PSPs in the k-th layer. If the clustering process is *neuron-centric*—treating each of these $36 + 54 = 90$ categories as completely independent—then according to the Pigeonhole Principle, in the worst case, one might collect at least 91 dual points to guarantee the construction of one solvable linear system from Eq. (36) or Eq. (40) to solve $A_t^{(k)}$ or $A_p^{(k)} - A_q^{(k)}$. Although all 91 dual points share the *same* convolutional kernel, isolated linear systems fail to jointly utilize them trivially, thereby wasting most dual points. Note that, the extraction processes introduced by Carlini *et al.* [7] on FCN and Sun *et al.* [27] on CNN are both *neuron-centric*. This severe inefficiency naturally raises a critical question: *Can we establish a "kernel-centric" clustering and recovery framework that exploit the weight sharing property in CNN, rather than treating each neuron independently?*

Weight Sharing Property in CNN. According to Eq. (49) in Supp. A.1, $A_i^{(k)}$ is a sparse vector which only contains non-zero entries at the indices where the convolutional kernel is applied. This sparsity allows us to explicitly extract the positional indices of the convolutional kernel's local receptive field (LRF-C)—defined as the specific sub-matrix of the input matrix $IN^{(k)}$, which is multiplied by the same kernel matrix $C^{(k)}$ according to Supp. A.1. Denote the vector form of $C^{(k)}$ as $\vec{C}^{(k)} = (C_0^{(k)}, C_1^{(k)}, \ldots, C_{h_c^{(k)}-1}^{(k)})^{\mathsf{T}}$. Thus, the length of $\vec{C}^{(k)}$ is $\ell_c^{(k)} = h_c^{(k)} w_c^{(k)}$. Assume the stride $s_c^{(k)} = 1$, and let $\mathcal{K}_t^{(k)} \subset \{0, 1, \ldots, d^{(k)} - 1\}$ denote the set of indices belonging to the LRF-C of the t-th neuron $(0 \le t \le d_f^{(k)} - 1)$, then

$$\mathcal{K}_t^{(k)} := \left\{ w_{in}^{(k)} \cdot \left\lfloor \frac{t}{w_o^{(k)}} \right\rfloor + t \bmod w_o^{(k)} + w_{in}^{(k)} \cdot \left\lfloor \frac{s}{w_c^{(k)}} \right\rfloor + s \bmod w_c^{(k)},\ 0 \le s < \ell_c \right\}. \tag{41}$$

For Dual RPCPs, recall from Eq. (36) that $g_u A_t^{(k)} F^{k-1} \vec{\delta} = 0$, the set $\mathcal{K}_t^{(k)}$ picks out the nonzero entries of $A_t^{(k)}$ according to Eq. (49). Then,

$$g_u A_t^{(k)} (F^{k-1} \vec{\delta}) = g_u \vec{C}^{(k)} (F^{k-1} \vec{\delta})[\mathcal{K}_t^{(k)}] = 0. \tag{42}$$

[3] The readers are suggested to read Supp. A.2 to better understand all the symbols.

Denote an empty set $\mathcal{K}_{\emptyset}^{(k)} = \emptyset$, then Eq. (42) can be rewritten as:

$$g_u A_t^{(k)}(F^{k-1}\vec{\delta}) = g_u \vec{C}^{(k)}((F^{k-1}\vec{\delta})[\mathcal{K}_t^{(k)}] - (F^{k-1}\vec{\delta})[\mathcal{K}_{\emptyset}^{(k)}]) = 0. \tag{43}$$

For Dual PSPs, recall from Eq. (40) that $g_u(A_p^{(k)} - A_q^{(k)})F^{k-1}\vec{\delta} = 0$, the set $\mathcal{K}_p^{(k)}$ and $\mathcal{K}_q^{(k)}$ picks out the nonzero entries of $A_p^{(k)}$ and $A_q^{(k)}$ respectively. Then,

$$g_u(A_p^{(k)} - A_q^{(k)})(F^{k-1}\vec{\delta}) = g_u \vec{C}^{(k)} \cdot ((F^{k-1}\vec{\delta})[\mathcal{K}_p^{(k)}] - (F^{k-1}\vec{\delta})[\mathcal{K}_q^{(k)}]) = 0. \tag{44}$$

Therefore, Eq. (43) for Dual RPCPs and Eq. (44) for Dual PSPs share the same unknown vectors $g_u\vec{C}^{(k)}$ due the weight sharing property of CNN. From a collection of dual points, if we can recognize Dual RPCPs and Dual PSPs, and identify the critical neuron t for Dual RPCPs and the competing neurons p and q for Dual PSPs, then the $\mathcal{K}_t^{(k)}$, $\mathcal{K}_p^{(k)}$, and $\mathcal{K}_q^{(k)}$ will be determined. Hence, Eqs. (43) and (44) can be combined with the same unknown $g_u\vec{C}^{(k)}$, which are solved using the SVD method to recover $\vec{C}^{(k)}$ up to a scalar multiple.

The ASV Method to Identify the Dual RPCPs and Dual PSPs. Chen *et al.* introduced the identification method for RPCPs and PSPs in a raw-output CNN [11], based on the structural sparsity of $\vec{\omega} = A_t^{(k)}$ (or $\vec{\omega} = A_p^{(k)} - A_q^{(k)}$), according to Eq. (49) in Supp. A.1. In hard-label CNN, we introduce the ASV method to identify different dual points.

1. For Dual RPCPs, denote row vectors

$$\vec{G}_\ell = (G_i^{k+1} - G_j^{k+1})P^{(k)}I_+^{(k)}A^{(k)}, \quad \vec{G}_r = (G_i^{k+1} - G_j^{k+1})P^{(k)}I_-^{(k)}A^{(k)}. \tag{45}$$

   Recall from Eq. (35) that, the scaled row vector $g_u A_t^{(k)}$ of the convolutional matrix $A^{(k)}$, is the difference between $\vec{G}_\ell$ and $\vec{G}_r$, *i.e.*, $g_u A_t^{(k)} = \vec{G}_\ell - \vec{G}_r$.
2. Similarly, for Dual PSPs, denote row vectors

$$\vec{G}_\ell = (G_i^{k+1} - G_j^{k+1})P_+^{(k)}I^{(k)}A^{(k)}, \quad \vec{G}_r = (G_i^{k+1} - G_j^{k+1})P_-^{(k)}I^{(k)}A^{(k)}. \tag{46}$$

   Recall from Eq. (39) that, the difference of the scaled row vectors $g_u(A_p^{(k)} - A_q^{(k)})$, is the difference between $\vec{G}_\ell$ and $\vec{G}_r$, *i.e.*, $g_u(A_p^{(k)} - A_q^{(k)}) = \vec{G}_\ell - \vec{G}_r$.

For any dual point in a hard-label CNN (including both Dual RPCPs and Dual PSPs), compute the unit normal vectors of the adjacent decision boundary patches as $\vec{n}_\ell = \frac{\vec{G}_\ell}{\|\vec{G}_\ell\|}$ and $\vec{n}_r = \frac{\vec{G}_r}{\|\vec{G}_r\|}$. The ASV in the hard-label CNN is then similarly defined by Eq. (24). Let $\vec{\omega} = \vec{G}_\ell - \vec{G}_r$ be the neuron weight ($g_u A_t^{(k)}$ for Dual RPCPs, or $g_u(A_p^{(k)} - A_q^{(k)})$ for Dual PSPs). Then, Eq. (24) becomes $\vec{v} = \frac{1}{\|\vec{G}_\ell\|}(\vec{\omega} - (\vec{\omega} \cdot \vec{n}_r^{\mathsf{T}})\vec{n}_r)$. Similarly, we measure the directional distance $\mathsf{DGap}(\vec{v},\vec{\omega})$ between $\vec{v}$ and $\vec{\omega}$ by Eq. (26). According to our experiments given in Supp. F.1, the $\mathsf{DGap}(\vec{v},\vec{\omega})$ is close to 0. Moreover, $\vec{v}$ strictly preserves the structural sparsity due to the localized connectivity of CNNs. This allows us

to verify dual points and identify the corresponding neurons by examining the consistency of their sparsity. An experiment to verify the sparsity alignment between $\vec{v}$ and $\vec{\omega}$ is given in Supp. G.

With known architecture of a CNN, we could iteratively construct $A_t^{(k)}$ for each $0 \leq t \leq d_f^{(k)} - 1$ according to Eq. (49) in Supp. A.1, thereby determining the locations of zero and non-zero values. Then, for Dual RPCPs, verify if $A_t^{(k)}$ and $\vec{v}$ are of the same sparsity consistency:

1. **Non-zero-Entry Consistency**: At any index where $A_t^{(k)}$ is non-zero, the corresponding element in the ASV $\vec{v}$ should be either non-zero or NaN.
2. **Zero-Entry Consistency**: At any index where $A_t^{(k)}$ is zero, the corresponding element in $\vec{v}$ should be either 0 or NaN.

Similarly, for Dual PSPs, build $A_p^{(k)}$ and $A_q^{(k)}$ with Eq. (49) by enumerating all possible competing neuron pairs (p, q), and verify whether $A_p^{(k)} - A_q^{(k)}$ and $\vec{v}$ share the same sparsity consistency:

1. **Union Non-zero Consistency**: At any index where either $A_p^{(k)}$ or $A_q^{(k)}$ is non-zero, the corresponding element in $\vec{v}$ should be either non-zero or NaN.
2. **Common Zero Consistency**: At any index where both $A_p^{(k)}$ and $A_q^{(k)}$ are zero, the corresponding element in $\vec{v}$ should be either 0 or NaN.

In practice, the consistency checks are performed for Dual RPCPs before proceeding to Dual PSPs. At last, a dual point is identified as Dual RPCP or Dual PSP, together with the deduced $\mathcal{K}_t^{(k)}$, or $\mathcal{K}_p^{(k)}$, $\mathcal{K}_q^{(k)}$. Hence, Eqs. (43) and (44) can be combined and solved to determine $\vec{C}^{(k)}$ up to a scalar multiple.

6 Black-box Experiments On Signature Recovery

Description of the Implementation. In practice, our signature recovery attack[4] proceeds layer by layer. For FCNs, when targeting a specific layer, we assume all preceding layers are extracted and directly *use their white-box parameters*. The reason is that when replaying Carlini *et al.*'s sign recovery [7] on 64-64×4-10 network, it does not completely and correctly recover all the signs in the deep layers. This issue also leaves an open question on how to invent a more robust sign recovery method.

For CNNs, we use the actual empirical parameters we recovered from the preceding layers during our experiments. Note that for CNNs, sign recovery is inherently straightforward, as described in [11, Sect. 4.3 and 5.3]. We provide the sign recovery process in hard-label CNNs in Supp. B.

The signature recovery for each layer includes 3 phases: dual point finding, clustering, signature computing. In Carlini *et al.*'s proof-of-concept experiment

[4] Our paper focuses on the signature recovery phase, the sign recovery method is the same to Carlini *et al.* [7].

[7], they use white-box gradient access for both dual point finding and signature computing phases, and rely on white-box neuron indices to cluster the dual points. In contrast, our signature recoveries on the 64-64×4-10 FCN and all CNNs are black-box, while only the dual point finding for the *3072-FCN* uses white-box gradient access (thus partial black-box). Most experiments are conducted on a server equipped with an NVIDIA RTX 5080 GPU (16 GB memory).

Hard-label FCN Extraction By the ASV-NRC Combined Clustering. The FCN with a 64-64×4-10 architecture is trained on MNIST dataset. As shown in Table 1 in Sect. 1, to extract the 1st layer, about 3000 dual points are needed to cluster. Carlini *et al.*'s method [7] requires about 0.93 hours to extract the 1st layer. In contrast, the extraction time based on our ASV method is only 0.05 hours. The difference of time complexities is even more pronounced when the number of dual points increases. For example, extracting the 2nd layer requires clustering roughly $300,000$ dual points. Our NRC method completes this in 4.22 hours, while Carlini *et al.*'s method [7] cannot give a result within a reasonable time (over 1 month). We estimate their runtime as follows: We simulate the clustering process for all $300,000$ dual points using white-box neuron indices, assuming each SVD-based consistency check is correct. This full-scale clustering requires a total of $2,653,799,211$ comparison. We then evaluate the actual time cost on a smaller subset of $3,000$ dual points, which takes $3,585.770$ seconds to perform $1,790,440$ comparison. According to Sect. 3.2, the estimated runtime for Carlini *et al.*'s method to cluster the entire dual points set is about $\frac{3{,}585.770\text{ s}}{1{,}790{,}440\times 3{,}600\text{ s/h}} \times 2,653,799,211 = 1,476$ hours. In Table 1 and 2, the runtime of Carlini *et al.*'s method is similarly estimated for the 64-64×4-10 FCN's layer F_3 and *3072-FCN*'s layers F_1 and F_2.

Black-box Hard-label CNN Extraction with Kernel-Centric ASV. The LeNet-5, first proposed by LeCun *et al.* in 1998 [19], is a foundational work in CNNs. Modern implementations typically update its classic (2+2) architecture by incorporating ReLU activations and max-pooling layers. Our practical attacks target a (2+1) CNN, and this modern LeNet-5 (specified in Supp. E) in hard-label setting with fully black-box access. The results are given in Table 1.

7 Conclusion

This paper resolves a computation bottleneck in Carlini *et al.*'s hard-label attacks on DNNs, *i.e.*, the dual points clustering step, by proposing the *Normal Rank Check (NRC)* and *Approximate Signature Vector (ASV)* to replace Carlini *et al.*'s heavy SVD-driven clustering, drastically lowering the computational overhead. Besides, for hard-label max-pooling CNNs, we present a highly efficient *kernel-centric* methodology, namely *Kernel-Centric ASV*, to overcome the limitations of previous *neuron-centric* extraction strategies. Experimental results on FCNs and max-pooling LeNet-5 confirm that our method achieves high clustering precision and accelerates the extraction by several orders of magnitude.

References

1. Roderick Asselineau, Patrick Derbez, Pierre-Alain Fouque, and Brice Minaud. Cryptanalytic extraction of deep neural networks with non-linear activations. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part VII*, volume 16806 of *Lecture Notes in Computer Science*, pages 36–66. Springer, 2026.
2. Lejla Batina, Shivam Bhasin, Dirmanto Jap, and Stjepan Picek. {CSI}{NN}: Reverse engineering of neural network architectures through electromagnetic side channel. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 515–532, 2019.
3. Avrim Blum and Ronald Rivest. Training a 3-node neural network is np-complete. *Advances in neural information processing systems*, 1, 1988.
4. Martin R Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319. Springer Science & Business Media, 2013.
5. Isaac A. Canales-Martínez and David Santos. Extracting some layers of deep neural networks in the hard-label setting. In Daniel Escudero and Ivan Damgård, editors, *Progress in Cryptology - LATINCRYPT 2025 - 9th International Conference on Cryptology and Information Security in Latin America, Medellín, Colombia, October 1-3, 2025, Proceedings*, volume 16129 of *Lecture Notes in Computer Science*, pages 399–421. Springer, 2025.
6. Isaac Andrés Canales-Martinez, Jorge Chávez-Saab, Anna Hambitzer, Francisco Rodríguez-Henríquez, Nitin Satpute, and Adi Shamir. Polynomial time cryptanalytic extraction of neural network models. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part III*, volume 14653 of *Lecture Notes in Computer Science*, pages 3–33. Springer, 2024.
7. Nicholas Carlini, Jorge Chávez-Saab, Anna Hambitzer, Francisco Rodríguez-Henríquez, and Adi Shamir. Polynomial time cryptanalytic extraction of deep neural networks in the hard-label setting. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part I*, Lecture Notes in Computer Science, pages 364–396. Springer, 2025.
8. Nicholas Carlini, Matthew Jagielski, and Ilya Mironov. Cryptanalytic extraction of neural network models. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, Lecture Notes in Computer Science, pages 189–218. Springer, 2020.
9. Yi Chen, Xiaoyang Dong, Jian Guo, Yantian Shen, Anyu Wang, and Xiaoyun Wang. Hard-label cryptanalytic extraction of neural network models. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part VIII*, Lecture Notes in Computer Science, pages 207–236. Springer, 2024.
10. Yi Chen, Xiaoyang Dong, Ruijie Ma, Yantian Shen, Anyu Wang, Hongbo Yu, and Xiaoyun Wang. Delving into cryptanalytic extraction of prelu neural networks. In

International Conference on the Theory and Application of Cryptology and Information Security, pages 576–607. Springer, 2025.
11. Zirui Chen, Shi Tang, Zhengchao Gao, Yongjia Su, Lingyue Qin, and Xiaoyang Dong. Algebraic attack on convolutional neural network with max pooling. *IACR Cryptol. ePrint Arch.*, 2026:241, 2026.
12. Zirui Chen, Shi Tang, Zhengchao Gao, Yongjia Su, Lingyue Qin, and Xiaoyang Dong. Algebraic attack on convolutional neural networks with max pooling. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part VII*, volume 16806 of *Lecture Notes in Computer Science*, pages 3–35. Springer, 2026.
13. Benoit Coqueret, Mathieu Carbone, Olivier Sentieys, and Gabriel Zaid. A divide-and-conquer strategy for hard-label extraction of deep neural networks via side-channel attacks. Cryptology ePrint Archive, Paper 2024/1870, 2024.
14. Charles Fefferman et al. Reconstructing a neural net from its output. *Revista Matemática Iberoamericana*, 10(3):507–556, 1994.
15. Hanna Foerster, Robert Mullins, Ilia Shumailov, and Jamie Hayes. Beyond slow signs in high-fidelity model extraction. *Advances in Neural Information Processing Systems*, 37:19496–19522, 2024.
16. Péter Horváth, Dirk Lauret, Zhuoran Liu, and Lejla Batina. Sok: Neural network extraction through physical side channels. In *USENIX Security Symposium*, 2024.
17. Akira Ito, Takayuki Miura, and Yosuke Todo. Is the hard-label cryptanalytic model extraction really polynomial? In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part VII*, volume 16806 of *Lecture Notes in Computer Science*, pages 67–98. Springer, 2026.
18. Matthew Jagielski, Nicholas Carlini, David Berthelot, Alex Kurakin, and Nicolas Papernot. High accuracy and high fidelity extraction of neural networks. In *29th USENIX security symposium (USENIX Security 20)*, pages 1345–1362, 2020.
19. Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
20. Chun Li, Zheng Gong, Di Li, Liping Zhuang, Yufeng Tang, Yin Lv, and Xingfu Yan. End-to-end polynomial-time cryptanalytic extraction of convolutional neural networks in the hard-label setting. Cryptology ePrint Archive, Paper 2026/902, 2026.
21. Haolin Liu, Adrien Siproudhis, Christina Boura, and Thomas Peyrin. Model extraction of convolutional neural networks with max-pooling. *Cryptology ePrint Archive*, 2026.
22. Haolin Liu, Adrien Siproudhis, Samuel Experton, Peter Lorenz, Christina Boura, and Thomas Peyrin. Navigating the deep: End-to-end extraction on deep neural networks. In Joan Daemen and Emmanuel Thomé, editors, *Advances in Cryptology - EUROCRYPT 2026 - 45th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Rome, Italy, May 10-14, 2026, Proceedings, Part VI*, volume 16546 of *Lecture Notes in Computer Science*, pages 482–512. Springer, 2026.
23. Daniel Lowd and Christopher Meek. Adversarial learning. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 641–647, 2005.
24. Daryna Oliynyk, Rudolf Mayer, and Andreas Rauber. I know what you trained last summer: A survey on stealing machine learning models and defences. *ACM Computing Surveys*, 55(14s):1–41, 2023.

25. Xiaokang Qi, Hao Lei, Longxiang Wei, Xiaohan Sun, and Meiqin Wang. Cryptanalytic extraction of neural networks with various activation functions. *Cryptology ePrint Archive*, 2026.
26. David Rolnick and Konrad Kording. Reverse-engineering deep relu networks. In *International conference on machine learning*, pages 8178–8187. PMLR, 2020.
27. Xiaohan Sun, Hao Lei, Longxiang Wei, Xiaokang Qi, Kai Hu, Meiqin Wang, and Wei Wang. Cryptanalytic extraction of convolutional neural networks. Cryptology ePrint Archive, Paper 2026/139, 2026.
28. Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction {APIs}. In *25th USENIX security symposium (USENIX Security 16)*, pages 601–618, 2016.

Supplementary Material

A Detailed Notations of CNN

A.1 Basic Definitions and Notations

The basic definitions and notations is summarized in Table 4.

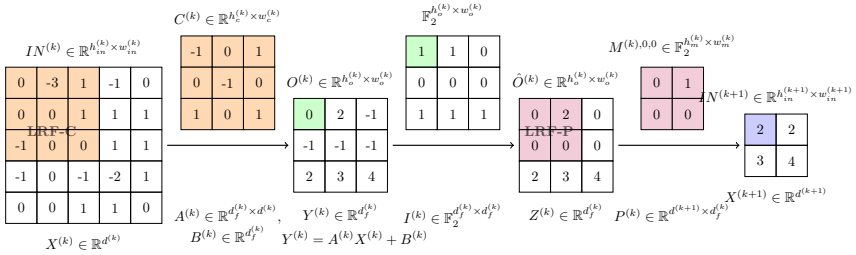


Fig. 6: Symbolic View of a CNN Layer

The CNN layer includes a Convolutional Function $f_c^{(k)}$, a nonlinear activation function $\sigma_c^{(k)}$ (ReLU), and pooling function $\rho^{(k)}$ (Max Pooling).

Convolutional Function. The CNN's inputs are formulated as matrices. The convolution operation of the convolutional function $f_c^{(k)}$ employs a sliding window mechanism, where a small matrix of learnable parameters (the *Convolutional Kernel Matrix*) traverses the input from left to right, top to bottom. At each position, the sum of element-wise products between the *Convolutional Kernel Matrix* and the overlapping local input patch (*Local Receptive Field of convolutional function*, abbreviated as *LRF-C*), is computed to generate a single output value. A shared bias is then added to the result at each position to constitute the output matrix (*Feature Map*) of the convolutional function.

In the algebraic view in Figure 6, the input matrix of $f_c^{(k)}$ is denoted as $IN^{(k)}$ with dimension $h_{in}^{(k)} \times w_{in}^{(k)}$, so $d^{(k)} = h_{in}^{(k)} \cdot w_{in}^{(k)}$. The convolutional kernel matrix is defined by the shape $(\mathsf{c}_{in}^{(k)}, \mathsf{c}_{out}^{(k)}, h_c^{(k)}, w_c^{(k)})$, where $\mathsf{c}_{in}^{(k)}$ and $\mathsf{c}_{out}^{(k)}$ denote the number of input and output channels, and $h_c^{(k)}, w_c^{(k)}$ denote the height and width of the convolutional kernel matrix $C^{(k)}$, respectively.

For briefness, we will only describe the simple case of single input and output channel, *i.e.*, $\mathsf{c}_{in}^{(k)} = \mathsf{c}_{out}^{(k)} = 1$. For the case of multiple channels, please refer to [11]. Given a convolution stride $s_c^{(k)}$, the $h_o^{(k)} \times w_o^{(k)}$ output matrix or the *feature map* of $f_c^{(k)}$ is denoted as $O^{(k)}$, where

$$h_o^{(k)} = \left\lfloor \frac{h_{in}^{(k)} - h_c^{(k)}}{s_c^{(k)}} \right\rfloor + 1, \quad w_o^{(k)} = \left\lfloor \frac{w_{in}^{(k)} - w_c^{(k)}}{s_c^{(k)}} \right\rfloor + 1. \tag{47}$$

Table 4: Basic Definitions and Notations

FCN: Fully Connected Neural Network (FCN).
CNN: Convolutional Neural Network (CNN).
ReLU: Rectified Linear Unit.
FCN layer: Including a linear function $f^{(k)}$ called fully-connected function, and a nonlinear function σ (component-wise ReLU function).
CNN layer: Including Convolutional Function $f_c^{(k)}$, nonlinear activation function $\sigma_c^{(k)}$ (ReLU), and pooling function $\rho^{(k)}$ (Max Pooling).
$(\mathbf{m}+\mathbf{n})$-Deep Convolutional Neural Network (CNN): A CNN with m CNN layers and n FCN layers, as well as an output layer with a fully-connected linear function.
$(\mathbf{c_{in}^{(k)}}, \mathbf{c_{out}^{(k)}}, \mathbf{h_c^{(k)}}, \mathbf{w_c^{(k)}})$: Convolutional Function with $\mathtt{c}_{in}^{(k)}$ input and $\mathtt{c}_{out}^{(k)}$ output channels, and the size of kernel matrices is $h_c^{(k)} \times w_c^{(k)}$.
$IN^{(k)}$: input matrix of $f_c^{(k)}$, $IN^{(k)} \in \mathbb{R}^{h_{in}^k \times w_{in}^{(k)}}$
$C^{(k)}$: convolutional kernel matrix of layer k, $C^{(k)} \in \mathbb{R}^{h_c^k \times w_c^{(k)}}$
$O^{(k)}$: $h_o^{(k)} \times w_o^{(k)}$ output matrix or the feature map of $f_c^{(k)}$, $O^{(k)} \in \mathbb{R}^{h_o^{(k)} \times w_o^{(k)}}$
$I^{(k)}$: diagonal activation indicator matrix of layer k, $I^{(k)} \in \mathbb{F}_2^{d_f^{(k)} \times d_f^{(k)}}$
$\hat{O}^{(k)}$: output matrix of the k-th non-linear activation function $\sigma_c^{(k)}$, $\hat{O}^{(k)} \in \mathbb{R}^{h_o^{(k)} \times w_o^{(k)}}$
$M^{(k),i,j}$: Boolean selection matrix for the (i,j)-th local receptive field of max pooling function in layer k, $M^{(k),i,j} \in \mathbb{F}_2^{h_m^{(k)} \times w_m^{(k)}}$
$P^{(k)}$: Boolean pooling matrix of layer k, $P^{(k)} \in \mathbb{R}^{d^{(k+1)} \times d_f^{(k)}}$
$X^{(k)}, Y^{(k)}, Z^{(k)}$: Input vectors of the operations $f_c^{(k)}$, $\sigma_c^{(k)}$, $\rho^{(k)}$ in the k-th CNN layer, $k \geq 1$. $X^{(k)} \in \mathbb{R}^{d^{(k)}}, Y^{(k)}, Z^{(k)} \in \mathbb{R}^{d_f^{(k)}}$, since $\sigma_c^{(k)}$ does not change the dimension. $X^{(k)}[i]$ is the i-th $(0 \leq i < d^{(k)})$ entry of $X^{(k)}$.
$A^{(k)}$: Matrix whose entry in i-th row and j-th is $A_{i,j}^{(k)}$, i, j start from 0.
RPCP, PSP, FCP: ReLU-Pooling Critical Point, Pooling Switching Point, Critical points in the Fully-connected Block, respectively.
LRF-C, Local Receptive Field of Convolutional Function: A sub-matrix of $IN^{(k)}$, which is multiplied by *Convolutional Kernel Matrix* $C^{(k)}$ to produce each entry in $Y^{(k)}$. The LRF-C will be different for different elements of $Y^{(k)}$. Let $Y^{(k)}[t]$ correspond to the t-th $(0 \leq t \leq d_f^{(k)} - 1)$ LRF-C.
LRF-P, Local Receptive Field of Pooling function: A sub-matrix of $Z^{(k)}$, which is processed by a pooling function to produce each entry of $X^{(k+1)}$. The LRF-P will be different for different elements of $X^{(k+1)}$. Let $X^{(k+1)}[t]$ correspond to the t-th $(0 \leq t \leq d^{(k+1)} - 1)$ LRF-P.
$s_c^{(k)}$: The convolution stride, which is by default set to 1 in this paper.

For briefness, we let the convolution stride $s_c^{(k)} = 1$. The input of $f_c^{(k)}$ can be a $h_{in}^{(k)} \times w_{in}^{(k)}$ $IN^{(k)}$ or a $d^{(k)}$-dimension vector $X^{(k)}$. Denote $IN^{(k)} = [a_{i,j}^{(k)}]_{h_{in}^{(k)} \times w_{in}^{(k)}}$ and the i-th ($0 \le i \le h_{in}^{(k)} - 1$) row $IN_i^{(k)} = (a_{i,0}^{(k)}, a_{i,1}^{(k)}, \cdots, a_{i,w_{in}^{(k)}-1}^{(k)})$. Then the vector $X^{(k)} = (IN_0^{(k)}, IN_1^{(k)}, \cdots, IN_{h_{in}^{(k)}-1}^{(k)})^{\mathsf{T}}$.

Denote the kernel matrix $C^{(k)} = [c_{i,j}^{(k)}]_{h_c^{(k)} \times w_c^{(k)}}$, where $C_i^{(k)} = (c_{i,0}^{(k)}, c_{i,1}^{(k)}, \cdots, c_{i,w_c^{(k)}-1}^{(k)})$. Then convolution operation computes the output matrix $O^{(k)}$ as

$$O_{i,j}^{(k)} = \sum_{p=0}^{h_c^{(k)}-1} \sum_{q=0}^{w_c^{(k)}-1} IN_{i+p,j+q}^{(k)} C_{p,q}^{(k)} + b^{(k)}, \tag{48}$$

where the bias $b^{(k)}$ is the same for all the outputs. In vector form, the output matrix $O^{(k)}$ will be a ($d_f^{(k)} = h_o^{(k)} \cdot w_o^{(k)}$)-dimension vector $Y^{(k)} = (O_0^{(k)}, \cdots, O_{h_o^{(k)}-1}^{(k)})^{\mathsf{T}}$.

Definition 15. (Convolutional Matrix) *Let $\mathsf{c}_{in}^{(k)} = \mathsf{c}_{out}^{(k)} = s_c^{(k)} = 1$. In the k-th convolutional function, given a convolution kernal matrix $C^{(k)}$, the function $f_c^{(k)} : \mathbb{R}^{d^{(k)}} \to \mathbb{R}^{d_f^{(k)}}$, is defined as $f_c^{(k)}(X^{(k)}) = A^{(k)}X^{(k)} + B^{(k)}$, where $X^{(k)} \in \mathbb{R}^{d^{(k)}}$ is the input vector. $A^{(k)} = (A_0^{(k)\mathsf{T}}, A_1^{(k)\mathsf{T}}, \cdots, A_{d_f^{(k)}-1}^{(k)\mathsf{T}})^{\mathsf{T}} \in \mathbb{R}^{d_f^{(k)} \times d^{(k)}}$, whose i-th row ($0 \le i \le d_f^{(k)} - 1$) is:*

$$A_i^{(k)} = (\underbrace{0, \cdots, 0}_{w_{in}^{(k)} \cdot \left\lfloor \frac{i}{w_o^{(k)}} \right\rfloor}, \underbrace{\underbrace{0, \cdots, 0}_{i \bmod w_o^{(k)}}, C_0^{(k)}, \underbrace{0, \ldots, 0}_{padding}}_{w_{in}^{(k)}}, \cdots, \underbrace{\underbrace{0, \cdots, 0}_{i \bmod w_o^{(k)}}, C_{h_c^{(k)}-1}^{(k)}, \underbrace{0, \ldots, 0}_{padding}}_{w_{in}^{(k)}}, \underbrace{0, \cdots, 0}_{padding}), \tag{49}$$

and $B^{(k)} = (b^{(k)}, b^{(k)}, \cdots, b^{(k)})^T \in \mathbb{R}^{d_f^{(k)}}$ (the bias for each neuron is the same).

Nonlinear Activation Function. The k-th activation function $\sigma_c^{(k)}$ of a CNN consists of a set of parallel $d_f^{(k)}$ nonlinear activation functions $\mathrm{ReLU}(x_i) = \max\{x_i, 0\}$ for $i = 0, \cdots, d_f^{(k)} - 1$.

The input of $\sigma_c^{(k)}$ function is the $h_o^{(k)} \times w_o^{(k)}$ matrix $O^{(k)}$, and we denote its output as $\hat{O}^{(k)}$ with the same dimension. In vector form, the input is denoted as $Y^{(k)} = (y_0, y_1, \cdots, y_{d_f^{(k)}-1})^{\mathsf{T}}$ and then $\sigma_c^{(k)}$ function can be expressed as a matrix $I^{(k)}$ multiplied by $Y^{(k)}$, *i.e.*, $Z^{(k)} = (\hat{O}_0^{(k)}, \cdots, \hat{O}_{h_o^{(k)}-1}^{(k)})^{\mathsf{T}} = \sigma_c^{(k)}(Y) = I^{(k)}Y^{(k)}$, where

$$I^{(k)} = \begin{pmatrix} \tau_0^{(k)} & 0 & \cdots & 0 \\ 0 & \tau_1^{(k)} & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & \tau_{d_f^{(k)}-1}^{(k)} \end{pmatrix}, \ 0 \le i \le d_f^{(k)} - 1, \ \tau_i^{(k)} = \begin{cases} 1, & \text{if} \quad y_i > 0, \\ 0, & \text{if} \quad y_i \le 0. \end{cases} \tag{50}$$

Pooling Function. In the pooling function, a small fixed-size window (analogous to a convolutional kernel, and referred to as a pooling kernel) slides across a feature map (input matrix $\hat{O}^{(k)}$) with a predefined stride. The region of the input matrix covered by the pooling window is also referred to as the *Local Receptive Field* of Pooling (LRF-P for short). For each window position, it applies a pooling function (usually use the max or average function) to compress the contents in the LRF-P into a single value, producing a dimension reduced output matrix (denoted as $IN^{(k+1)}$). Unlike convolutional kernels, pooling kernels are parameter-free (no learnable weights and biases) — their behavior is determined solely by size, stride, and compressing rule.

In the pooling function $\rho^{(k)}$, we assume that the max pooling kernel is applied to an LRF-P of size $h_m^{(k)} \times w_m^{(k)}$ with stride $s_\rho^{(k)}$. The input is $h_o^{(k)} \times w_o^{(k)}$ matrix $\hat{O}^{(k)}$ in Figure 6. The output is denoted as a $h_\rho^{(k)} \times w_\rho^{(k)}$ matrix $IN^{(k+1)}$, where:

$$h_\rho^{(k)} = \left\lfloor \frac{h_o^{(k)} - h_m^{(k)}}{s_\rho^{(k)}} \right\rfloor + 1, \quad w_\rho^{(k)} = \left\lfloor \frac{w_o^{(k)} - w_m^{(k)}}{s_\rho^{(k)}} \right\rfloor + 1.$$

In particular, when the pooling kernel size is $h_m^{(k)} \times w_m^{(k)} = 2 \times 2$ and the stride is $s_\rho^{(k)} = 2$, the windows partition the input into disjoint LRF-Ps, and the calculation of the output size reduces to:

$$h_\rho^{(k)} = \left\lfloor \frac{h_o^{(k)}}{2} \right\rfloor, \quad w_\rho^{(k)} = \left\lfloor \frac{w_o^{(k)}}{2} \right\rfloor.$$

The entry of the i-th row and j-th column ($0 \le i \le h_\rho^{(k)} - 1$, $0 \le j \le w_\rho^{(k)} - 1$) of the output matrix $IN^{(k+1)}$ takes the maximum value of the elements in the $(i \cdot w_\rho^{(k)} + j)$-th LRF-P of $\hat{O}^{(k)}$, whose indexes are denoted as $\Omega_{i,j} = \{(i \cdot s_\rho^{(k)} + h,\ j \cdot s_\rho^{(k)} + w),\ 0 \le h \le h_m^{(k)} - 1, 0 \le w \le w_m^{(k)} - 1\}$. Then,

$$IN_{i,j}^{(k+1)} = \max\{\hat{O}_{\Omega_{i,j}}^{(k)}\}. \tag{51}$$

When the input matrix $\hat{O}^{(k)}$ is given, and the max pooling operation can be expressed as multiplying a series of Boolean matrices $M^{(k),i,j} \in \mathbb{F}_2^{h_m^{(k)} \times w_m^{(k)}}$ ($0 \le i \le h_\rho^{(k)} - 1$, $0 \le j \le w_\rho^{(k)} - 1$) by different $h_m^{(k)} \times w_m^{(k)}$ LRF-P of $\hat{O}^{(k)}$,

$$IN_{i,j}^{(k+1)} = \sum_{0 \le h \le h_m^{(k)} - 1, 0 \le w \le w_m^{(k)} - 1} M_{\{h,w\}}^{(k),i,j} \cdot \hat{O}_{\{i \cdot s_\rho^{(k)} + h, j \cdot s_\rho^{(k)} + w\}}^{(k)}, \tag{52}$$

In the $(i \cdot w_\rho^{(k)} + j)$-th LRF-P of $\hat{O}^{(k)}$, let $\Lambda_{i,j}$ be the set of all index pairs that achieve the maximum value:

$$\Lambda_{i,j} = \left\{ (h,w) \mid \hat{O}_{\{i \cdot s_\rho^{(k)} + h, j \cdot s_\rho^{(k)} + w\}}^{(k)} = \max\{\hat{O}_{\Omega_{i,j}}^{(k)}\} \right\}. \tag{53}$$

To ensure exactly one neuron is selected from the LRF-P, let $(h^*, w^*) \in \Lambda_{i,j}$ be an arbitrarily chosen index pair. The elements of $M^{(k),i,j}$ are defined as:

$$M^{(k),i,j}_{\{h,w\}} = \begin{cases} 1, & \text{if } (h,w) = (h^*, w^*), \\ 0, & \text{otherwise,} \end{cases} \tag{54}$$

with $0 \le h \le h_m^{(k)} - 1$, $0 \le w \le w_m^{(k)} - 1$. Note that any arbitrary choice of $(h^*, w^*) \in \Lambda_{i,j}$ yields a functionally equivalent matrix $M^{(k),i,j}$ that produces the exact same pooling output $IN_{i,j}^{(k+1)}$.

In vector form, the input and output vector of pooling function $\rho^{(k)}$ are $Z^{(k)} = (\hat{O}_0^{(k)}, \cdots, \hat{O}^{(k)}_{h_o^{(k)}-1})^\mathsf{T} \in \mathbb{R}^{d_f^{(k)} = h_o^{(k)} \cdot w_o^{(k)}}$ and $X^{(k+1)} = (IN_0^{(k+1)}, IN_1^{(k+1)}, \cdots, IN^{(k+1)}_{h_\rho^{(k)}-1})^\mathsf{T} \in \mathbb{R}^{d^{(k+1)}}$, $d^{(k+1)} = h_\rho^{(k)} \cdot w_\rho^{(k)}$.

Definition 16. (Boolean Pooling Matrix) *There exists a $d^{(k+1)} \times d_f^{(k)}$ Boolean pooling matrix $P^{(k)}$ so that $\rho^{(k)} : \mathbb{R}^{d_f^{(k)}} \to \mathbb{R}^{d^{(k+1)}}$ satisfying $X^{(k+1)} = \rho^{(k)}(Z^{(k)}) = P^{(k)} \cdot Z^{(k)}$. Assuming the LRF-Ps are disjoint and square (*i.e.*, $s_\rho^{(k)} = w_m^{(k)} = h_m^{(k)}$), then the max pooling operation on the r-th LRF-P can be defined as the r-th row ($0 \le r \le d^{(k+1)} - 1$) of $P^{(k)}$:*

$$P_r^{(k)} = (\underbrace{\underbrace{0, \cdots, 0}_{w_o^{(k)} \cdot i \cdot s_\rho^{(k)}}, \underbrace{0, \cdots, 0}_{j \cdot s_\rho^{(k)}}, M_0^{(k),i,j}, \underbrace{0, \ldots, 0}_{padding}}_{w_o^{(k)}}, \cdots, \underbrace{\underbrace{0, \cdots, 0}_{j \cdot s_\rho^{(k)}}, M^{(k),i,j}_{h_m^{(k)}-1}, \underbrace{0, \ldots, 0}_{padding}, \underbrace{0, \cdots, 0}_{padding}}_{w_o^{(k)}}), \tag{55}$$

where $i = \left\lfloor \frac{r}{w_\rho^{(k)}} \right\rfloor$, $j = r \bmod w_\rho^{(k)}$.

Obviously, similar to the matrix of $I^{(k)}$ for $\sigma_c^{(k)}$ function, the pooling matrix $P^{(k)}$ varies with the input $X^{(1)}$ due to Eq. (54). For example, the input matrix $\hat{O}^{(k)}$ has dimensions $h_o^{(k)} \times w_o^{(k)} = 4 \times 4$, i.e., $d_f^{(k)} = 16$, the max pooling kernel is of $h_m^{(k)} \times w_m^{(k)} = 2 \times 2$ with stride $s_\rho^{(k)} = 2$. The output matrix $IN^{(k+1)}$ will have dimension $h_\rho^{(k)} \times w_\rho^{(k)} = 2 \times 2$, i.e., $d^{(k+1)} = 4$. The $d^{(k+1)} \times d_f^{(k)} = 4 \times 16$ pooling matrix $P^{(k)} = \begin{pmatrix} P_{0,0}^{(k)} & \mathbf{0}_{2\times 8} \\ \mathbf{0}_{2\times 8} & P_{1,1}^{(k)} \end{pmatrix}$, where $\mathbf{0}_{2\times 8}$ is 2×8 zero matrix, and

$$P_{0,0}^{(k)} = \begin{pmatrix} M_{0,0}^{(k),0,0} & M_{0,1}^{(k),0,0} & 0 & 0 & M_{1,0}^{(k),0,0} & M_{1,1}^{(k),0,0} & 0 & 0 \\ 0 & 0 & M_{0,0}^{(k),0,1} & M_{0,1}^{(k),0,1} & 0 & 0 & M_{1,0}^{(k),0,1} & M_{1,1}^{(k),0,1} \end{pmatrix},$$

$$P_{1,1}^{(k)} = \begin{pmatrix} M_{0,0}^{(k),1,0} & M_{0,1}^{(k),1,0} & 0 & 0 & M_{1,0}^{(k),1,0} & M_{1,1}^{(k),1,0} & 0 & 0 \\ 0 & 0 & M_{0,0}^{(k),1,1} & M_{0,1}^{(k),1,1} & 0 & 0 & M_{1,0}^{(k),1,1} & M_{1,1}^{(k),1,1} \end{pmatrix}.$$

Definition 17. (Fully-connected Function) *The k-th fully-connected function within the fully-connected block of the CNN is a function $f^{(k)} : \mathbb{R}^{d^{(k)}} \to \mathbb{R}^{d^{(k+1)}}$ given by an affine transformation:*

$$f^{(k)}(X) = A^{(k)}X^{(k)} + B^{(k)},$$

where $X^{(k)} \in \mathbb{R}^{d^{(k)}}$ represents the input vector. The weight matrix $A^{(k)} \in \mathbb{R}^{d^{(k+1)} \times d^{(k)}}$ and the bias vector $B^{(k)} \in \mathbb{R}^{d^{(k+1)}}$ are composed of floating-point numbers.

The matrix for the nonlinear function σ (ReLU functions) of the fully-connected block can be similarly defined by Eq. (50).

Given $X^{(1)} = x \in \mathbb{R}^{d^{(1)}}$, the CNN layers and FCN layers of CNN "collapse" into an affine transformation. Assume that the function of the m CNN layers is $\mathcal{F}_1$, and the function of the n FCN layers is $\mathcal{F}_2$, then

$$\begin{aligned} y = \mathcal{F}_1(x) &= P^{(m)}(I^{(m)}(A^{(m)} \cdots (P^{(1)}(I^{(1)}(A^{(1)}x + B^{(1)}))) \cdots + B^{(m)})) \\ &= P^{(m)}I^{(m)}A^{(m)} \cdots P^{(2)}I^{(2)}A^{(2)}P^{(1)}I^{(1)}A^{(1)}x + \beta_1 \\ &= \Gamma_1 \cdot x + \beta_1 \end{aligned} \tag{56}$$

According to [6, Section 3.1], given y, $\mathcal{F}_2$ collapses into

$$\mathcal{F}_2(y) = \Gamma_2 \cdot y + \beta_2. \tag{57}$$

Consequently, the full CNN "collapses" into

$$\mathcal{F}(x) = \Gamma_2(\Gamma_1 \cdot x + \beta_1) + \beta_2 = \Gamma_2 \cdot \Gamma_1 \cdot x + \Gamma_2\beta_1 + \beta_2. \tag{58}$$

Definition 18. (Linear Neighborhood of CNN) *Given an input $x \in \mathbb{R}^{d^{(1)}}$ with the corresponding output of the CNN computed by Eq. (58), the linear neighborhood of x is defined as the set*

$$\{u \in \mathbb{R}^{d^{(1)}} | \mathcal{F}(u) = \Gamma_2 \cdot \Gamma_1 \cdot u + \Gamma_2\beta_1 + \beta_2\},$$

i.e., the same affine transformation is used for u and x to compute the output of the CNN.

If we make a change of Δ to the input x, and $x+\Delta$ remains within the linear neighborhood of x, we can observe the corresponding change of the output

$$\mathcal{F}(x+\Delta)-\mathcal{F}(x) = \Gamma_2 \cdot \Gamma_1 \cdot (x+\Delta)+\Gamma_2\beta_1+\beta_2-(\Gamma_2 \cdot \Gamma_1 \cdot x+\Gamma_2\beta_1+\beta_2) = \Gamma_2 \cdot \Gamma_1 \cdot \Delta.$$

Definition 19. (Layer Merging) *We focus on the convolutional block of the $(m+n)$ network. We assume that we have complete knowledge of the first $k-1$ layers of the convolutional block, and we are currently recovering layer k. Let*

F_x^{k-1} and G_x^{k+1} represent, respectively, the fully recovered and non-recovered parts of the CNN.

$$\mathcal{F} = \underbrace{f^{(n+1)} \circ \cdots \circ \sigma^{(1)} \circ f^{(1)} \circ \rho^{(m)} \circ \cdots \circ f_c^{(k+1)}}_{G_x^{k+1}} \circ \rho^{(k)} \circ \sigma_c^{(k)} \circ f_c^{(k)} \circ \underbrace{\rho^{(k-1)} \circ \cdots \circ f_c^{(1)}}_{F_x^{k-1}}.$$

Given $X^{(1)} \in \mathbb{R}^{d^{(1)}}$, G_x^{k+1} and F_x^{k-1} become $G_x^{k+1}(X^{(k+1)}) = G^{k+1}X^{(k+1)} + B^{k+1}$ and $F_x^{k-1}(X^{(1)}) = F^{k-1}X^{(1)} + B^{k-1}$, where the matrices $F^{k-1} \in \mathbb{R}^{d^{(k)} \times d^{(1)}}$, $B^{k-1} \in \mathbb{R}^{d^{(k)}}$, $G^{k+1} \in \mathbb{R}^{d^{(n+2)} \times d^{(k+1)}}$, $B^{k+1} \in \mathbb{R}^{d^{(n+2)}}$, respectively.

A.2 A Toy Example of CNN

To facilitate a clearer understanding of the algebraic notations introduced in Section A.1, this appendix provides a concrete, step-by-step toy example. We explicitly demonstrate the complete computational pipeline — including Convolution, ReLU activation, and Max Pooling — for a single-input and single-output channel scenario (*i.e.*, $\mathsf{c}_{in}^{(k)} = \mathsf{c}_{out}^{(k)} = 1$).

Figure 7 visualizes the spatial mapping mechanism of this process. In the following subsections, we formulate this exact process algebraically by completely expanding the equivalent transformation matrices without any truncation.

Setup and Vector Flattening. Consider an input matrix $IN^{(k)} \in \mathbb{R}^{h_{in}^{(k)} \times w_{in}^{(k)}} = \mathbb{R}^{5\times5}$. We apply a Convolutional Function $f_c^{(k)}$ with a kernel $C^{(k)} \in \mathbb{R}^{h_c^{(k)} \times w_c^{(k)}} = \mathbb{R}^{2\times2}$, a stride of $s_c^{(k)} = 1$, and valid padding, yielding an intermediate output matrix $O^{(k)}$. According to the definitions in Section A.1, the spatial dimensions of $O^{(k)}$ are calculated as follows:

$$h_o^{(k)} = \left\lfloor \frac{h_{in}^{(k)} - h_c^{(k)}}{s_c^{(k)}} \right\rfloor + 1 = \left\lfloor \frac{5-2}{1} \right\rfloor + 1 = 4, \tag{59}$$

$$w_o^{(k)} = \left\lfloor \frac{w_{in}^{(k)} - w_c^{(k)}}{s_c^{(k)}} \right\rfloor + 1 = \left\lfloor \frac{5-2}{1} \right\rfloor + 1 = 4. \tag{60}$$

Consequently, the resulting intermediate output matrix is $O^{(k)} \in \mathbb{R}^{4\times4}$. Subsequently, a ReLU activation function $\sigma_c^{(k)}$ is applied to $O^{(k)}$, generating an activated output matrix $\hat{O}^{(k)} \in \mathbb{R}^{4\times4}$.

Following the activation, a max pooling function $\rho^{(k)}$ with a window size of $h_m^{(k)} \times w_m^{(k)} = 2\times2$ and a stride of $s_\rho^{(k)} = 2$ processes $\hat{O}^{(k)}$. The spatial dimensions of the final output matrix $IN^{(k+1)}$ for the subsequent layer are determined by:

$$h_\rho^{(k)} = \left\lfloor \frac{h_o^{(k)} - h_m^{(k)}}{s_\rho^{(k)}} \right\rfloor + 1 = \left\lfloor \frac{4-2}{2} \right\rfloor + 1 = 2, \tag{61}$$

$$w_\rho^{(k)} = \left\lfloor \frac{w_o^{(k)} - w_m^{(k)}}{s_\rho^{(k)}} \right\rfloor + 1 = \left\lfloor \frac{4-2}{2} \right\rfloor + 1 = 2. \tag{62}$$

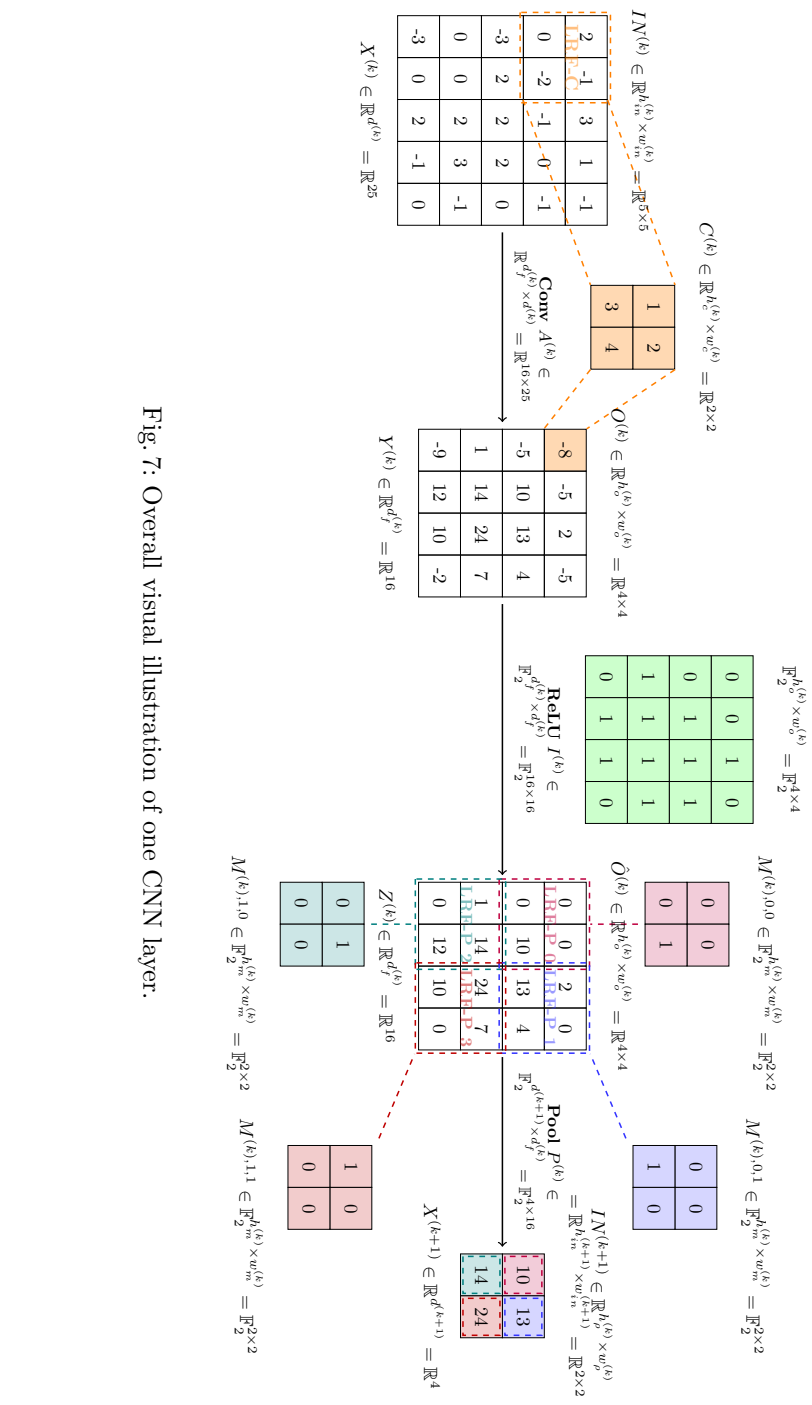


Fig. 7: Overall visual illustration of one CNN layer.

This yields the matrix $IN^{(k+1)} \in \mathbb{R}^{2\times 2}$.

To represent these operations algebraically, we flatten the spatial matrices into vectors using row-major ordering. For the input matrix $IN^{(k)}$, its i-th row $(0 \le i \le h_{in}^{(k)} - 1)$ is denoted as $IN_i^{(k)} = (a_{i,0}^{(k)}, a_{i,1}^{(k)}, \ldots, a_{i,w_{in}^{(k)}-1}^{(k)})$. Based on the specific numerical values from Figure 7, we instantiate the first and the last row vectors as follows:

$$\begin{aligned} IN_0^{(k)} &= (2, -1, 3, 1, -1), \\ &\vdots \\ IN_4^{(k)} &= (-3, 0, 2, -1, 0). \end{aligned} \tag{63}$$

The flattened input vector $X^{(k)} \in \mathbb{R}^{d^{(k)}} = \mathbb{R}^{h_{in}^{(k)} \cdot w_{in}^{(k)}} = \mathbb{R}^{25}$ is constructed by concatenating these row vectors as $X^{(k)} = (IN_0^{(k)}, IN_1^{(k)}, \ldots, IN_4^{(k)})^\mathsf{T}$, which explicitly evaluates to:

$$X^{(k)} = [2, -1, 3, 1, -1, 0, -2, -1, 0, -1, -3, 2, 2, 2, 0, 0, 0, 2, 3, -1, -3, 0, 2, -1, 0]^\mathsf{T}. \tag{64}$$

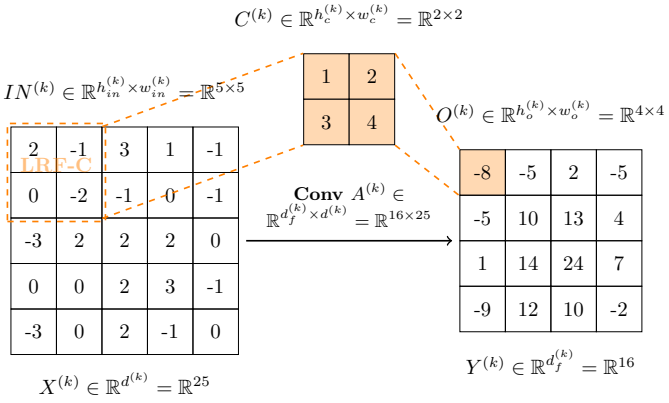


Fig. 8: Local view of the convolution operation $f_c^{(k)}$

The Convolutional Matrix ($A^{(k)}$). Let the 2×2 convolutional kernel matrix be $C^{(k)} = \left[\begin{smallmatrix} 1 & 2 \\ 3 & 4 \end{smallmatrix}\right]$. For simplicity in this example, we assume the bias is $b^{(k)} = 0$, yielding an all-zero bias vector $B^{(k)} = \mathbf{0}$. Therefore, the convolution operation translates into a linear matrix multiplication $Y^{(k)} = A^{(k)}X^{(k)}$, where $Y^{(k)} \in \mathbb{R}^{d_f^{(k)}} = \mathbb{R}^{h_o^{(k)} \cdot w_o^{(k)}} = \mathbb{R}^{16}$ is the flattened output vector and $A^{(k)} \in \mathbb{R}^{d_f^{(k)} \times d^{(k)}} = \mathbb{R}^{16\times 25}$ is the convolutional matrix.

As illustrated by the orange dashed lines in Figure 8, the first element of the feature map $O_{0,0}^{(k)}$ (which corresponds to $Y_0^{(k)}$) is computed by the sum of the element-wise multiplications between $C^{(k)}$ and the top-left LRF-C. Formally, this is calculated as:

$$\begin{aligned} O_{0,0}^{(k)} &= \sum_{p=0}^{h_c^{(k)}-1} \sum_{q=0}^{w_c^{(k)}-1} IN_{0+p,0+q}^{(k)} C_{p,q}^{(k)} + b^{(k)} \\ &= IN_{0,0}^{(k)} C_{0,0}^{(k)} + IN_{0,1}^{(k)} C_{0,1}^{(k)} + IN_{1,0}^{(k)} C_{1,0}^{(k)} + IN_{1,1}^{(k)} C_{1,1}^{(k)} + b^{(k)} \\ &= (2)(1) + (-1)(2) + (0)(3) + (-2)(4) + 0 = -8. \end{aligned} \tag{65}$$

According to Eq. (49), the first row of the convolutional matrix ($i = 0$) is systematically constructed by substituting the kernel rows into the padded structure. Given $w_{in}^{(k)} = 5$ and $w_o^{(k)} = 4$, the substitution and expansion explicitly demonstrate how the kernel parameters are positioned within $A_0^{(k)}$:

$$\begin{aligned} A_0^{(k)} = (\ & \underbrace{0, \cdots, 0}_{w_{in}^{(k)} \cdot \lfloor 0/w_o^{(k)} \rfloor}, \underbrace{\underbrace{0, \cdots, 0}_{0 \bmod w_o^{(k)}}, C_0^{(k)}, \underbrace{0, \ldots, 0}_{padding}}_{w_{in}^{(k)}}, \underbrace{\underbrace{0, \cdots, 0}_{0 \bmod w_o^{(k)}}, C_1^{(k)}, \underbrace{0, \ldots, 0}_{padding}, \underbrace{0, \cdots, 0}_{padding}}_{w_{in}^{(k)}}) \\ = (& \underbrace{}_{0}, \underbrace{\underbrace{}_{0}, \mathbf{1}, \mathbf{2}, \underbrace{0, 0, 0}_{3}}_{5}, \underbrace{\underbrace{}_{0}, \mathbf{3}, \mathbf{4}, \underbrace{0, 0, 0}_{3}}_{5}, \underbrace{0, \cdots, 0}_{15}) \\ = (& 1, 2, 0, 0, 0, 3, 4, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0). \end{aligned} \tag{66}$$

By representing the exhaustive sliding window operations over the 2D input matrix, the entire equivalent matrix $A^{(k)}$ is expanded below. We highlight the non-zero elements corresponding to the convolutional kernel in each row in bold:

$$A^{(k)} = \begin{bmatrix} \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & \mathbf{3} & \mathbf{4} \end{bmatrix} \tag{67}$$

Performing the multiplication yields the pre-activation vector $Y^{(k)}$:

$$Y^{(k)} = A^{(k)} X^{(k)} = [-8, -5, 2, -5, -5, 10, 13, 4, 1, 14, 24, 7, -9, 12, 10, -2]^{\mathsf{T}}. \tag{68}$$

The ReLU Activation Matrix ($I^{(k)}$) Fig. 9 shows the operation of ReLU activation. Algebraically, it is represented as $Z^{(k)} = I^{(k)}Y^{(k)}$, where $I^{(k)} \in \mathbb{F}_2^{d_f^{(k)} \times d_f^{(k)}} = \mathbb{F}_2^{16\times 16}$ is a diagonal matrix. According to Eq. (50), the i-th diagonal entry of $I^{(k)}$ is set to 1 if the corresponding element $Y_i^{(k)}$ is strictly positive, and 0 otherwise. Based on the numerical values of $Y^{(k)}$ calculated previously, the fully

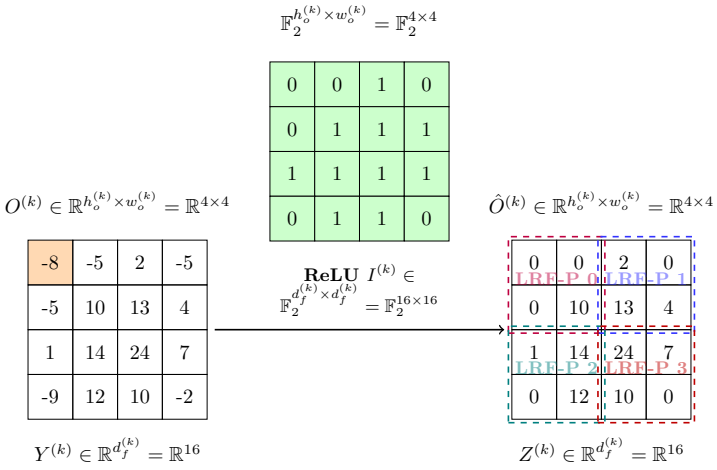


Fig. 9: Local view of the ReLU activation $\sigma_c^{(k)}$

expanded activation matrix is constructed as follows. We highlight the entire main diagonal in bold:

$$I^{(k)} = \begin{bmatrix} \mathbf{0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \mathbf{0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \mathbf{0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{0} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{0} \end{bmatrix} \tag{69}$$

Applying this masking matrix to $Y^{(k)}$ zeros out its non-positive components, yielding the activated vector $Z^{(k)}$:

$$Z^{(k)} = I^{(k)} Y^{(k)} = [0, 0, 2, 0, 0, 10, 13, 4, 1, 14, 24, 7, 0, 12, 10, 0]^{\mathsf{T}}. \tag{70}$$

The Max Pooling Matrix ($P^{(k)}$) In the pooling function $\rho^{(k)}$, the 2×2 max pooling operation with stride 2 partitions the 4×4 intermediate feature map $\hat{O}^{(k)}$ into 4 disjoint Local Receptive Fields of Pooling (LRF-Ps). For each LRF-P, the pooling function selects the maximum value. Fig. 10 illustrates the

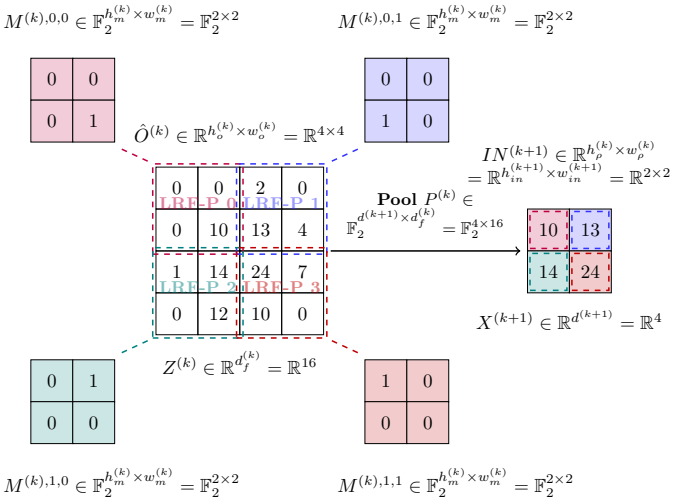


Fig. 10: Local view of the max pooling $\rho^{(k)}$

operation of max pooling. According to Eq. (51), the first element of the pooled output is the maximum value within the $(i \cdot w_\rho^{(k)} + j) = (0 \cdot w_\rho^{(k)} + 0) = 0$-th LRF-P denoted as $\Omega_{0,0}$:

$$IN_{0,0}^{(k+1)} = \max\{\hat{O}_{\Omega_{0,0}}^{(k)}\} = \max(\hat{O}_{0,0}^{(k)}, \hat{O}_{0,1}^{(k)}, \hat{O}_{1,0}^{(k)}, \hat{O}_{1,1}^{(k)}) = \max(0, 0, 0, 10) = 10. \tag{71}$$

Since the maximum value "10" in $\Omega_{0,0}$ is located at local index $(h^* = 1, w^* = 1)$, the Boolean matrix $M^{(k),0,0}$ assigns 1 exclusively to this position, making its row vectors $M_0^{(k),0,0} = (0, 0)$ and $M_1^{(k),0,0} = (0, 1)$. According to Eq. (55), the first row $(r = 0)$ of the pooling matrix is constructed by mapping these Boolean rows of $M^{(k),0,0}$ into the padded structure. Given $w_o^{(k)} = 4$ and $s_\rho^{(k)} = 2$, the substitution explicitly demonstrates how the selected bits are positioned within

$P_0^{(k)}$:

$$
\begin{aligned}
P_0^{(k)} &= (\underbrace{0,\cdots,0}_{w_o^{(k)}\cdot 0\cdot s_\rho^{(k)}},\underbrace{\underbrace{0,\cdots,0}_{0\cdot s_\rho^{(k)}},M_0^{(k),0,0},\underbrace{0,\ldots,0}_{padding}}_{w_o^{(k)}},\underbrace{\underbrace{0,\cdots,0}_{0\cdot s_\rho^{(k)}},M_1^{(k),0,0},\underbrace{0,\ldots,0}_{padding}}_{w_o^{(k)}},\underbrace{0,\cdots,0}_{padding})\\
&= (\underbrace{}_{0},\underbrace{\underbrace{}_{0},\mathbf{0},\mathbf{0},\underbrace{0,0}_{2},\underbrace{}_{0},\mathbf{0},\mathbf{1},\underbrace{0,0}_{2}}_{4},\underbrace{\underbrace{}_{0},\underbrace{0,\cdots,0}_{8}}_{4})\\
&= (0,0,0,0,0,1,0,0,0,0,0,0,0,0,0,0).
\end{aligned}
\tag{72}
$$

By extending this construction to all four LRF-Ps, the entire equivalent matrix $P^{(k)}$ extracts exactly one element per row. We highlight these selected positions in bold:

$$
P^{(k)} = \begin{bmatrix}
0\,0\,0\,0\,0\,\mathbf{1}\,0\,0\,0\,0\,0\,0\,0\,0\,0\,0\\
0\,0\,0\,0\,0\,0\,\mathbf{1}\,0\,0\,0\,0\,0\,0\,0\,0\,0\\
0\,0\,0\,0\,0\,0\,0\,0\,0\,\mathbf{1}\,0\,0\,0\,0\,0\,0\\
0\,0\,0\,0\,0\,0\,0\,0\,0\,0\,\mathbf{1}\,0\,0\,0\,0\,0
\end{bmatrix}
\tag{73}
$$

The final output vector for the entire layer is therefore:

$$
X^{(k+1)} = P^{(k)}Z^{(k)} = [10,13,14,24]^{\mathsf{T}}. \tag{74}
$$

B Sign Recovery in Hard-label CNN

According to [11, Sect. 4.3 and 5.3], for raw-output CNNs, sign recovery is inherently straightforward. The situation and the recovery procedure remain identical for hard-label CNNs.

Sign Recovery from Dual RPCPs. For Dual RPCPs, after recovering $\hat{A}^{(k)}$ and $\hat{B}^{(k)}$ up to an unknown sign, we compute the input vector of the neurons as $\hat{A}^{(k)}(F^{k-1}X + B^{k-1}) + \hat{B}^{(k)}$, where X is the Dual RPCP of the t-th neuron of layer k. Assume that neuron t belongs to the u-th pooling window Ω_u. According to Prop. 1, all the inputs of the neurons within Ω_u should be ≤ 0. Then,

- if all these inputs are ≤ 0, the sign of the extracted kernel matrix is correct;
- if there exists a positive input, the extracted kernel matrix has an inverted sign, and we multiply the matrices $\hat{A}^{(k)}$ and $\hat{B}^{(k)}$ by "-1".

Sign Recovery from Dual PSPs. For Dual PSPs, we reconstruct the convolutional matrix $\hat{A}^{(k)}$, and once we extract $\hat{p}$ and $\hat{q}$, the involved pooling window Ω_u $(0 \leq u < d^{(k+1)})$ corresponding to the PSP could be deduced. According to Prop. 2, the outputs of the p-th and q-th neurons should be the maximum value in Ω_u. Then, with $\hat{A}^{(k)}$,

- if $\hat{A}_p^{(k)}(F^{k-1}X+B^{k-1}) = \hat{A}_q^{(k)}(F^{k-1}X+B^{k-1}) > \hat{A}_\ell^{(k)}(F^{k-1}X+B^{k-1})$, $\ell \in \Omega_u \setminus \{p,q\}$, the sign of the extracted kernel matrix is correct;
- if $\hat{A}_p^{(k)}(F^{k-1}X+B^{k-1}) = \hat{A}_q^{(k)}(F^{k-1}X+B^{k-1}) < \hat{A}_\ell^{(k)}(F^{k-1}X+B^{k-1})$, $\ell \in \Omega_u \setminus \{p,q\}$, the extracted kernel matrix has an inverted sign, and we multiply the vector $\hat{A}_p^{(k)}$ by "-1".

C White-box Experiments for verification of $\mathsf{DGap}(\vec{v},\vec{\omega})$ on 64-64×4-10 FCN and *3072-FCN*.

For both the 64-64×4-10 FCN and the *3072-FCN*, we randomly sample 10,000 dual points. For each point, we compute the DGap between its ASV $\vec{v}$ in the corresponding layer and the underlying neuron weight $\vec{w}$.

Table 5: Layer-wise statistics of average DGap for the 64-64×4-10 FCN and 3072-FCN.

Layer	**64-64×4-10 FCN**		**3072-FCN**	
	Valid Points	**Average DGap**	**Valid Points**	**Average DGap**
0	3,248	0.0161	3,582	0.0032
1	2,678	0.0559	2,389	0.0083
2	1,498	0.0669	2,695	0.0087
3	1,926	0.0610	1,181	0.0173
Total	**9,350**	**0.0449**	**9,847**	**0.0076**

D Proof of Proposition 1 in Unitary Balanced Networks

This appendix provides a proof of Proposition 1 in unitary balanced networks, where each neuron weight is a randomly chosen unit vector in $d^{(k)}$-dimensional space. This assumption follows the back-of-the-envelope analysis in [6, Appendix A].

To compute the cosine distance $\mathsf{DGap}(\vec{v},\vec{\omega}) = 1 - |\cos\varphi_{\vec{v},\vec{\omega}}| = 1 - \frac{|\vec{v}\cdot\vec{\omega}^\mathsf{T}|}{\|\vec{v}\|\cdot\|\vec{\omega}^\mathsf{T}\|}$, we first compute

$$\vec{v}\cdot\vec{\omega}^\mathsf{T} = \frac{1}{\|\vec{G}_\ell\|}\left(\vec{\omega} - (\vec{\omega}\cdot\vec{n}_r^\mathsf{T})\vec{n}_r\right)\cdot\vec{\omega}^\mathsf{T} = \frac{1}{\|\vec{G}_\ell\|}\left(\|\vec{\omega}\|^2 - (\vec{\omega}\cdot\vec{n}_r^\mathsf{T})^2\right), \tag{75}$$

$$\|\vec{v}\| = \sqrt{\vec{v}\cdot\vec{v}^\mathsf{T}} = \sqrt{\frac{1}{\|\vec{G}_\ell\|^2}\left(\|\vec{\omega}\|^2 - (\vec{\omega}\cdot\vec{n}_r^\mathsf{T})^2\right)} = \frac{1}{\|\vec{G}_\ell\|}\sqrt{\|\vec{\omega}\|^2 - (\vec{\omega}\cdot\vec{n}_r^\mathsf{T})^2}. \tag{76}$$

Then,

$$\cos\varphi_{\vec{v},\vec{\omega}} = \frac{\|\vec{\omega}\|^2 - (\vec{\omega}\cdot\vec{n}_r^\mathsf{T})^2}{\sqrt{\|\vec{\omega}\|^2 - (\vec{\omega}\cdot\vec{n}_r^\mathsf{T})^2}\cdot\|\vec{\omega}\|} = \sqrt{1 - \left(\frac{\vec{\omega}\cdot\vec{n}_r^\mathsf{T}}{\|\vec{\omega}\|}\right)^2}. \tag{77}$$

According to Eq. (19), we have $\vec{n}_r = \frac{\vec{G}_r}{\|\vec{G}_r\|} = \frac{1}{\|\vec{G}_r\|}\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}$. Hence,

$$\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}} = g_t A_t^{(k)}\cdot\frac{1}{\|\vec{G}_r\|}\left(\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\right)^{\mathsf{T}} = \frac{g_t}{\|\vec{G}_r\|}\sum_{s\in\Lambda^{(k)}} g_s\left(A_t^{(k)}\cdot A_s^{(k)\mathsf{T}}\right). \tag{78}$$

Let $\varphi_{t,s}$ denote the angle between $A_t^{(k)}$ and $A_s^{(k)}$, *i.e.*, $A_t^{(k)}\cdot A_s^{(k)\mathsf{T}} = \|A_t^{(k)}\|\cdot\|A_s^{(k)}\|\cos\varphi_{t,s}$, then $\frac{\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|}$ can be rewritten as:

$$\frac{\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|} = \frac{g_t\|A_t^{(k)}\|}{\|\vec{\omega}\|}\frac{\sum_{s\in\Lambda^{(k)}} g_s\|A_s^{(k)}\|\cos\varphi_{t,s}}{\|\vec{G}_r\|} = \frac{\sum_{s\in\Lambda^{(k)}} g_s\|A_s^{(k)}\|\cos\varphi_{t,s}}{\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\|}. \tag{79}$$

Note that in unitary balanced networks, each weight (row vectors) satisfies $\|A_s^{(k)}\| = 1$. We have:

$$\mathbb{E}[A_t^{(k)}] = \vec{0}, \qquad \mathbb{E}[A_t^{(k)\mathsf{T}}A_t^{(k)}] = \frac{I_{d^{(k)}}}{d^{(k)}}, \tag{80}$$

where $I_{d^{(k)}}$ is the $d^{(k)}\times d^{(k)}$ identity matrix. We can explicitly verify these moments component by component:

- **First Moment:** For the target weight $A_t^{(k)}$, denote its coordinates as $(a_1, a_2, \ldots, a_{d^{(k)}})$. For any a_i, the spherical reflection symmetry implies that a_i and $-a_i$ have the same distribution. Thus, $\mathbb{E}[a_i] = 0$ for all i, yielding $\mathbb{E}[A_t^{(k)}] = \vec{0}$.
- **Second Moment:** The outer product matrix $A_t^{(k)\mathsf{T}}A_t^{(k)}$ has entries $a_i a_j$. For the off-diagonal entries ($i \neq j$), the reflection symmetry across coordinate planes ensures that $\mathbb{E}[a_i a_j] = 0$. For the diagonal entries ($i = j$), we consider the deterministic unit length constraint:

$$\sum_{i=1}^{d^{(k)}} a_i^2 = 1 \implies \sum_{i=1}^{d^{(k)}} \mathbb{E}[a_i^2] = 1.$$

  Because the distribution is isotropic, all coordinate axes are statistically equivalent, meaning $\mathbb{E}[a_1^2] = \mathbb{E}[x_a^2] = \cdots = \mathbb{E}[a_{d^{(k)}}^2]$. Consequently, each diagonal entry evaluates to $\mathbb{E}[a_i^2] = 1/d^{(k)}$. Combining the zero off-diagonal elements and the $1/d^{(k)}$ diagonal elements directly yields the covariance matrix $I_{d^{(k)}}/d^{(k)}$.

Evaluation on the Error Term. According to Eq. (79), the error term $\frac{\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|}$ can be written as:

$$\frac{\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|} = \frac{\sum_{s\in\Lambda^{(k)}} g_s\|A_s^{(k)}\|\cos\varphi_{t,s}}{\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\|}.$$

We evaluate its expectation by first conditioning on all non-target weights $A_s^{(k)}$ ($s \neq t$). The target weight $A_t^{(k)}$ is statistically independent of other weight

vectors $A_s^{(k)}$ for $s \in \Lambda^{(k)}$. Since $\cos\varphi_{t,s} = A_t^{(k)} \cdot {A_s^{(k)}}^{\mathsf{T}}$, we have the conditional expectation:

$$\mathbb{E}_{A_t^{(k)}}[\cos\varphi_{t,s}] = \mathbb{E}[A_t^{(k)}] \cdot {A_s^{(k)}}^{\mathsf{T}} = 0. \tag{81}$$

Note that $t \notin \Lambda^{(k)}$. We adopt a local idealization in which $\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\|$ and g_s are deterministic. The conditional expectation of the entire ratio vanishes:

$$\mathbb{E}_{A_t^{(k)}}\left[\frac{\sum_{s\in\Lambda^{(k)}} g_s \cos\varphi_{t,s}}{\left\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\right\|}\right] = 0. \tag{82}$$

For any $s, q \in \Lambda^{(k)}$ (including $s = q$), the expectation of the product of two cosine terms leverages the covariance matrix of $A_t^{(k)}$:

$$\mathbb{E}_{A_t^{(k)}}[\cos\varphi_{t,s}\cos\varphi_{t,q}] = A_s^{(k)}\mathbb{E}[{A_t^{(k)}}^{\mathsf{T}} A_t^{(k)}]{A_q^{(k)}}^{\mathsf{T}} = \frac{A_s^{(k)} \cdot {A_q^{(k)}}^{\mathsf{T}}}{d^{(k)}}. \tag{83}$$

The numerator $\left(\sum g_s \cos\varphi_{t,s}\right)^2$ naturally produces a double sum over all pairs (s, q), including the cross-terms ($s \neq q$). We apply the covariance identity derived above to every pair in the double sum. This substitution transforms the expectation of the numerator into a sum of inner products $\left(A_s^{(k)} \cdot {A_q^{(k)}}^{\mathsf{T}}\right)$, perfectly mirroring the algebraic expansion of the denominator's squared norm. Explicitly, evaluating the expectation of the squared ratio yields:

$$\begin{aligned}\mathbb{E}_{A_t^{(k)}}\left[\left(\frac{\sum_{s\in\Lambda^{(k)}} g_s \cos\varphi_{t,s}}{\left\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\right\|}\right)^2\right] &= \frac{\sum_{s,q\in\Lambda^{(k)}} g_s g_q \,\mathbb{E}_{A_t^{(k)}}[\cos\varphi_{t,s}\cos\varphi_{t,q}]}{\left\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\right\|^2} \\ &= \frac{1}{d^{(k)}} \frac{\sum_{s,q\in\Lambda^{(k)}} g_s g_q \left(A_s^{(k)} \cdot {A_q^{(k)}}^{\mathsf{T}}\right)}{\left\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\right\|^2}.\end{aligned} \tag{84}$$

The double summation in the numerator is precisely the expanded form of the squared norm in the denominator:

$$\sum_{s,q\in\Lambda^{(k)}} g_s g_q \left(A_s^{(k)} \cdot {A_q^{(k)}}^{\mathsf{T}}\right) = \left\|\sum_{s\in\Lambda^{(k)}} g_s A_s^{(k)}\right\|^2. \tag{85}$$

Consequently, the terms perfectly cancel, evaluating to exactly $1/d^{(k)}$. Applying the definitions of $\vec{\omega}$ and $\vec{n}_r$, we conclude:

$$\mathbb{E}\left[\left(\frac{\vec{\omega} \cdot \vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|}\right)^2\right] = \frac{1}{d^{(k)}}. \tag{86}$$

Finally, we map this projection error back to the Directional Gap (DGap). According to Eq. (26) and Eq. (77), the exact relationship between DGap and the normalized projection term is given by:

$$\mathsf{DGap}(\vec{v},\vec{\omega}) = 1 - \sqrt{1 - \left(\frac{\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|}\right)^2}. \tag{87}$$

Since $1-\sqrt{1-z} \le z$ for $z \in [0,1]$, we can strictly bound the directional gap by the squared error term:

$$\mathsf{DGap}(\vec{v},\vec{\omega}) \le \left(\frac{\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|}\right)^2. \tag{88}$$

Taking the expectation of both sides and substituting our derived second moment yields a rigorous upper bound on the average directional error:

$$\mathbb{E}\big[\mathsf{DGap}(\vec{v},\vec{\omega})\big] \le \mathbb{E}\left[\left(\frac{\vec{\omega}\cdot\vec{n}_r^{\mathsf{T}}}{\|\vec{\omega}\|}\right)^2\right] = \frac{1}{d^{(k)}}. \tag{89}$$

This formally demonstrates that the expected directional gap gracefully converges to zero as the layer width increases:

$$\lim_{d^{(k)}\to\infty} \mathbb{E}\big[\mathsf{DGap}(\vec{v},\vec{\omega})\big] = 0.$$

E Detailed Model Structure Information in Experiments

(2+1)-Deep CNN:

1. CNN layer 1 (C_1):
   - **Convolutional function** $f_c^{(1)}$: 32×32 Input matrix $IN^{(1)}$, convolved with $(1,1,5,5)$ convolution kernel, with stride $s_c^{(1)}=1$, padding $pad^{(1)}=0$, produces 28×28 output matrix $O^{(1)}$, followed by ReLU activation.
   - **Max Pooling function** $\rho^{(1)}$: 2×2 max pooling kernel, with stride $s_\rho^{(1)}=2$, produces 14×14 output matrix $IN^{(2)}$.
2. CNN layer 2 (C_2):
   - **Convolutional function** $f_c^{(2)}$: 14×14 Input matrix $IN^{(2)}$, convolved with $(1,1,5,5)$ convolution kernel, with stride $s_c^{(2)}=1$, padding $pad^{(2)}=0$, produces 10×10 output matrix $O^{(2)}$, followed by ReLU activation.
   - **Max Pooling function** $\rho^{(2)}$: 2×2 max pooling kernel, with stride $s_\rho^{(2)}=2$, produces 5×5 output matrix $IN^{(3)}$.
3. FCN layer 1 (F_1):
   - Flatten the matrix $IN^{(3)}$ and get the 25-dimensional vector $X^{(1)}$.
   - **Fully-connected function** $f^{(1)}$: $d^{(1)}=25$ input neurons and $d^{(2)}=18$ output neurons, followed by ReLU activation.
4. Last layer $f^{(2)}$: $d^{(2)}=18$ input neurons and $d^{(3)}=10$ output neurons.

(2+2)-Deep LeNet-5:

1. CNN layer 1 (C_1):
   - **Convolutional function** $f_c^{(1)}$: 32×32 Input matrix $IN^{(1)}$, convolved with $(6, 1, 5, 5)$ convolution kernel, with stride $s_c^{(1)} = 1$, padding $pad^{(1)} = 0$, produces 6 output matrices $O^{(1)}$ of size 28×28 , followed by ReLU activation.
   - **Max Pooling function** $\rho^{(1)}$: 2×2 max pooling kernel, with stride $s_\rho^{(1)} = 2$, produces 6 output matrices $IN^{(2)}$ of size 14×14 .
2. CNN layer 2 (C_2):
   - **Convolutional function** $f_c^{(2)}$: 6 Input matrix $IN^{(2)}$ of size 14×14, convolved with $(6, 16, 5, 5)$ convolution kernel, with stride $s_c^{(2)} = 1$, padding $pad^{(2)} = 0$, produces 16 output matrix $O^{(2)}$ of size 10×10, followed by ReLU activation.
   - **Max Pooling function** $\rho^{(2)}$: 2×2 max pooling kernel, with stride $s_\rho^{(2)} = 2$, produces 16 output matrix $IN^{(3)}$ of size 5×5.
3. FCN layer 1 (F_1):
   - Flatten the matrix $IN^{(3)}$ and get the 400-dimensional vector $X^{(1)}$.
   - **Fully-connected function** $f^{(1)}$: $d^{(1)} = 400$ input neurons and $d^{(2)} = 120$ output neurons, followed by ReLU activation.
4. FCN layer 2 (F_2):
   - **Fully-connected function** $f^{(2)}$: $d^{(2)} = 120$ input neurons and $d^{(3)} = 84$ output neurons, followed by ReLU activation.
5. Last layer $f^{(3)}$: $d^{(3)} = 84$ input neurons and $d^{(3)} = 10$ output neurons.

F The Experiments on the DGap among the ASV and the weight for CNN

F.1 White-box Experiment for Verification of $\mathsf{DGap}(\vec{v}, \vec{\omega}) \to 0$ in Hard-Label CNN

We experiment on a (2+1)-deep CNN with detail configurations given in Supp. E. We randomly sample 1000 valid dual points across the input space and compute their corresponding ASVs $\vec{v}$. For each $\vec{v}$, we then calculate the DGap with respect to all row vectors $A_t^{(k)}$, as well as all possible pairwise differences $A_p^{(k)} - A_q^{(k)}$ (neurons p and q are in the same pooling window) of the convolutional matrix $A^{(k)}$. As shown in Fig. 11, the distributions exhibit a clear separation: the directional distances between ASVs and their corresponding target row vectors $A_t^{(k)}$ (or $A_p^{(k)} - A_q^{(k)}$) tightly cluster near 0, while the distances to all other non-target row vectors consistently concentrate around 1.

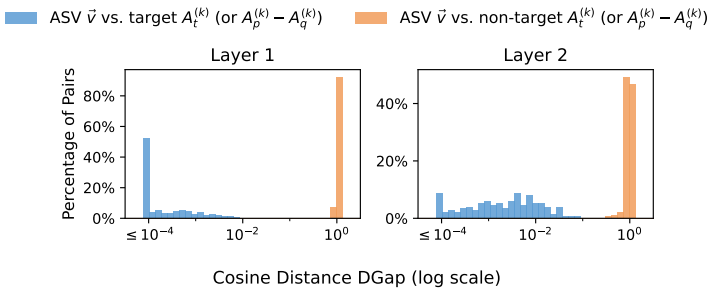


Fig. 11: Distribution of the Cosine Distance (DGap) between ASV $\vec{v}$ and all possible pairs of row vectors across CNN layers

F.2 Experimental Verification of ASV-based Consistency Check

Targeting a $(2+1)$ CNN, we evaluate the consistency scores across pairs of dual points. As shown in Fig. 12, dual points from the *same* target vector $A_t^{(k)}$ (or $A_p^{(k)} - A_q^{(k)}$) yield a significantly different distribution compared to pairs from *different* targets. This separation in distribution demonstrates that the ASV-based consistency score provides a reliable criterion for clustering dual points.

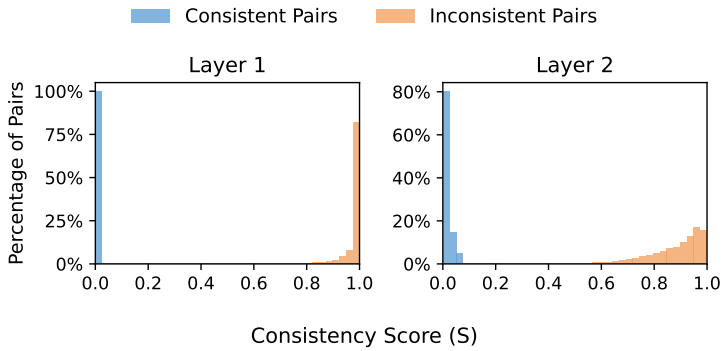


Fig. 12: Distribution of Consistency Scores for pairs sharing the same target vector versus those from different target vectors across CNN layers.

G Experiment Verification of Sparsity Alignment between $\vec{v}$ and $\vec{\omega}$ in CNNs

Targeting a $(2+1)$ CNN, we collect a set of 1000 dual points and compute their ASVs $\vec{v}$. If $\vec{v}$ can serve as a reliable proxy for the ground-truth vectors $\vec{\omega}$ in identifying Dual RPCPs and Dual PSPs, the positions where $\vec{\omega}$ is exactly 0 should also remain 0 in $\vec{v}$. This sparsity alignment guarantees that $\vec{v}$ preserves the structural identity of either the neuron weight vector $(A_t^{(k)})$ or the difference vector $(A_p^{(k)} - A_q^{(k)})$. To quantify this zero-position fidelity, we extract the absolute values of our ASVs $\vec{v}$ at the indices where $\vec{\omega}$ are 0. Our statistical analysis across all 1000 dual points evaluates the distribution of these theoretically inactive entries. As summarized in Table 6, the element-wise noise is small enough to recognize the sparse structure across both convolutional layers.

Table 6: Statistics of element-wise noise at the theoretical zero-positions.

Layer	Theoretical Value	Valid Duals	max	min	mean
C_1	zero	807	$2^{-4.28}$	$2^{-47.97}$	$2^{-11.05}$
	non-zero		$2^{-1.01}$	$2^{-18.99}$	$2^{-2.80}$
C_2	zero	140	$2^{-3.36}$	$2^{-44.25}$	$2^{-7.81}$
	non-zero		$2^{-0.53}$	$2^{-13.35}$	$2^{-2.58}$