











# More Insight on Deep Learning-Aided Cryptanalysis

Zhenzhen Bao<sup>1,6</sup>(), Jinyu Lu<sup>2,3</sup>(), Yiran Yao<sup>3</sup>(),  
and Liu Zhang<sup>3,4,5</sup>()

<sup>1</sup> Institute for Network Sciences and Cyberspace, BNRist,  
Tsinghua University, Beijing, China  
[zzbao@tsinghua.edu.cn](mailto:zzbao@tsinghua.edu.cn)

<sup>2</sup> College of Sciences, National University of Defense Technology,  
Changsha 410073, Hunan, China

<sup>3</sup> School of Physical and Mathematical Sciences,  
Nanyang Technological University, Singapore, Singapore  
[yiran005@e.ntu.edu.sg](mailto:yiran005@e.ntu.edu.sg)

<sup>4</sup> School of Cyber Engineering, Xidian University, Xi'an, China  
[liuzhang@stu.xidian.edu.cn](mailto:liuzhang@stu.xidian.edu.cn)

<sup>5</sup> State Key Laboratory of Cryptology, P.O.Box 5159, Beijing 100878, China

<sup>6</sup> Zhongguancun Laboratory, Beijing, China

**Abstract.** In CRYPTO 2019, Gohr showed that well-trained neural networks could perform cryptanalytic distinguishing tasks superior to differential distribution table (DDT)-based distinguishers. This suggests that the differential-neural distinguisher ( $\mathcal{ND}$ ) may use additional information besides pure ciphertext differences. However, the explicit knowledge beyond differential distribution is still unclear. In this work, we provide explicit rules that can be used alongside DDTs to enhance the effectiveness of distinguishers compared to pure DDT-based distinguishers. These rules are based on strong correlations between bit values in right pairs of XOR-differential propagation through addition modulo  $2^p$ . Interestingly, they can be closely linked to the earlier study of the multi-bit constraints and the recent study of the fixed-key differential probability. In contrast, combining these rules does not improve the  $\mathcal{ND}$ s' performance. This suggests that these rules or their equivalent form have already been exploited by  $\mathcal{ND}$ s, highlighting the power of neural networks in cryptanalysis.

In addition, we find that to enhance the differential-neural distinguisher's accuracy and the number of rounds, regulating the differential propagation is imperative. Introducing differences into the keys is typically believed to help eliminate differences in encryption states, resulting in stronger differential propagations. However, differential-neural attacks differ from traditional ones as they don't specify output differences or follow a single differential trail. This questions the usefulness of introducing differences in a key in differential-neural attacks and the resistance of SPECK against such attacks in the related-key setting. This work shows that the power of differential-neural cryptanalysis in the related-key setting can exceed that in the single-key setting by successfully conducting a 14-round key recovery attack on SPECK32/64.

**Keywords:** Neural Network · Interpretability · Modular Addition · Related-key · SPECK

## 1 Introduction

In 2019, Gohr [15] proposed differential-neural cryptanalysis, employing neural networks as superior distinguishers and exploiting them to perform efficient key recovery attacks. Impressively, the differential-neural distinguisher ( $\mathcal{ND}$ ) outperformed the traditional pure differential distinguishers using full differential distribution tables (DDT). However, interpreting these neural network-based distinguishers remains challenging, hindering the comprehension of the additional knowledge learned by differential-neural distinguishers.

Despite the intricate nature of neural network interpretability, researchers have made primary progress in understanding the differential-neural distinguisher’s inner workings. In EUROCRYPT 2021, Benamira *et al.* [7] proposed that Gohr’s neural distinguisher effectively approximates the cipher’s DDT during the learning phase. Moreover, the distinguisher relies on both the differential distribution of ciphertext pairs and that of the penultimate and antepenultimate rounds. Yet, the specific form of additional information remains undisclosed.

In AICrypt 2023, Gohr *et al.* [16] proved the differential-neural distinguisher for SIMON32/64 can use only differential features and achieve accuracy same as pure differential ones. Applying the same neural network to both SPECK and SIMON yields different conclusions: neural networks learned or did not learn features beyond full DDT. These intriguing findings motivate us to delve deeper into the neural network’s mechanisms, aiming to comprehend the specific features underpinning its conclusions for each cipher and to improve and exploit further the neural distinguishers should additional features be captured.

**Our Contributions.** In this work, we conclude that  $\mathcal{ND}$ s’ advantage over pure DDT-based distinguishers is in exploiting the differential distribution under the partially known value input to the last non-linear operation. Specifically,  $\mathcal{ND}$ s exploit the correlation between the ciphertexts’ partial value, ciphertext pair’s differences, and intermediate states’ differences. Furthermore, our work shows that differential-neural cryptanalysis in the related-key ( $\mathcal{RK}$ ) setting can attack more rounds than in the single-key setting, which was not apparent before. The concrete contributions include the following.

- **Improving full DDT-based distinguisher.** We observe that, apart from the information of differences, one knows the partial value of inputs, denoted by  $y$ , to the last modular addition of SPECK, leveraging by which one can improve DDT-based distinguishers. We show that the differential probability conditioned on a fixed value of  $y$  can differ from the average differential probability over all possible  $y$ . This insight enables more accurate classification based on the ciphertext pair’s differences and the ciphertexts’ partial value. The high-level idea is to consider conditional probabilities and specific cases where the fulfillment of the differential constraints can be predicted based on

**Table 1.** Summary of key recovery attacks on SPECK32/64

#R	Distinguisher	Configure	Time	Data	Succ. Rate	Key Space	Advantage	Ref.
13	$DD$	1+8+4	$2^{57}$	$2^{25}$	-	$2^{64}$	$2^7$	[12]
	$DD$	1+8+3	$2^{55.58}$	$2^{24.26}$	-	$2^{64}$	$2^{8.42}$	[14]
	$DD$	1+8+2+2	$2^{50.16}$	$2^{31.13}$	63%	$2^{64}$	$2^{13.84}$	[9]
	$\mathcal{ND}$	1+3+8+1	$2^{50.17}$	$2^{29}$	82%	$2^{63}$	$2^{13.83}$	[3]
	$\mathcal{ND}$	1+3+8+1	$2^{44.36}$	$2^{27}$	21%	$2^{63}$	$2^{18.64}$	[26]
	$\mathcal{RK}\text{-}\mathcal{ND}$	1+2+9+1	$2^{34.57}$	$2^{16}$	54.29%	$2^{50}$	$2^{15.43}$	Sect. 5.2
	$\mathcal{RK}\text{-}\mathcal{ND}$	1+2+9+1	$2^{31.79}$	$2^{10}$	43.33%	$2^{46}$	$2^{14.21}$	Sect. 5.2
14	$DD$	1+9+4	$2^{63}$	$2^{31}$	-	$2^{64}$	$2^1$	[12]
	$DD$	1+9+4	$2^{62.47}$	$2^{30.47}$	-	$2^{64}$	$2^{1.53}$	[24]
	$DD$	1+9+2+2	$2^{60.99}$	$2^{31.75}$	63%	$2^{64}$	$2^{3.01}$	[9]
	$DD$	2+9+3	$2^{60.58}$	$2^{30.26}$	76.00%	$2^{64}$	$2^{3.42}$	[14]
	$\mathcal{ND}$	1+3+8+2	$2^{60.36}$	$2^{27}$	21%	$2^{63}$	$2^{2.64}$	[26]
	$\mathcal{RK}\text{-}\mathcal{ND}$	1+3+9+1	<b><math>2^{35.59}</math></b>	<b><math>2^{16}</math></b>	75.71%	$2^{42}$	<b><math>2^{6.41}</math></b>	Sect. 5.2
	$\mathcal{RK}\text{-}\mathcal{ND}$	1+3+9+1	<b><math>2^{35.78}</math></b>	<b><math>2^{15}</math></b>	71.43%	$2^{41}$	<b><math>2^{5.22}</math></b>	Sect. 5.2
15	$DD$	1+10+4	$2^{63.39}$	$2^{30.39}$	-	$2^{64}$	$2^{0.61}$	[17]
	$DD$	1+10+2+2	$2^{62.25}$	$2^{30.39}$	-	$2^{64}$	$2^{1.75}$	[9]

–: Not available; “Advantage” denotes the time complexity advantage over a brute force attack.

the value of  $y$ . The results indicate that it is highly likely that  $\mathcal{ND}$ s rely on these specific cases to outperform pure DDT-based distinguishers.

- **Optimizing the performance and training process of  $\mathcal{ND}$ s.** Addressing the challenge of training high-round, especially 8-round,  $\mathcal{ND}$  of SPECK32/64, we introduce the Freezing Layer Method. By freezing all convolutional layers in a pre-trained 7-round  $\mathcal{ND}$ , we efficiently train an 8-round  $\mathcal{ND}$  using simple basic training with unaltered hyperparameters. This method matches Gohr’s accuracy but cuts training time and data.
- **Exploring differential-neural attacks in the related-key setting.** The conclusion that  $\mathcal{ND}$ s can efficiently capture features beyond full DDT encourages further exploration of  $\mathcal{ND}$ -based attacks. We observed that control over the differential propagation is vital for achieving effective high-round  $\mathcal{ND}$ s. Hence, we introduce related-key ( $\mathcal{RK}$ ) differences to slow down the diffusion of differences, aiding in training  $\mathcal{ND}$  for higher rounds. As a result, we achieve a 14-round key recovery attack on SPECK32/64 using related-key neural distinguishers ( $\mathcal{RK}\text{-}\mathcal{ND}$ s). Results are in Table 1. Furthermore, we constructed various distinguishers under various  $\mathcal{RK}$  differential trails and conducted comprehensive comparisons, reinforcing  $\mathcal{ND}$  explainability.

**Organization.** The paper’s structure is as follows: Sect. 2 provides preliminaries. Section 3 provides insights on the  $\mathcal{ND}$  explainability. Section 4 provides enhancements on the  $\mathcal{ND}$  training. Section 5 details of related-key differential-neural cryptanalysis. The conclusion is presented in Sect. 6.

## 2 Preliminary

### 2.1 Notations

Denote by  $C = (C_{n-1}, \dots, C_0)$  the binary vector of  $n$  bits, where  $C_i$  is the bit at position  $i$  and  $C_0$  is the least significant. Define  $n$  as the word size in bits and  $2n$  as the state size. Let  $(C_L^r, C_R^r)$  represent left and right state branches after  $r$  rounds, and  $k^r$  the  $r$ -round subkey. Bitwise XOR is denoted by  $\oplus$ , addition modulo  $2^n$  by  $\boxplus$ , bitwise AND by  $\odot$ , and bitwise right/left rotation by  $\ggg / \lll$ .

### 2.2 Brief Description of SPECK32/64

In 2013, the National Security Agency (NSA) proposed SPECK and SIMON block ciphers, aiming to ensure security on resource-constrained devices [5]. By 2018, both ciphers were standardized by ISO/IEC for air interface communication. The SPECK cipher uses a Feistel-like ARX design, emanating its non-linearity from modular addition and leveraging XOR and rotation for linear mixing. SPECK32/64 is the smallest SPECK variant [5]. Its round function, one of 22 rounds, takes a 16-bit subkey  $k^i$  and a state of two 16-bit words,  $(C_L^i, C_R^i)$ . Its key schedule reuses the round function to generate round keys. With  $K$  as a master key and  $k^i$  the  $i$ -th round key,  $K = (l^2, l^1, l^0, k^0)$ . The round function's details are in Fig. 1.

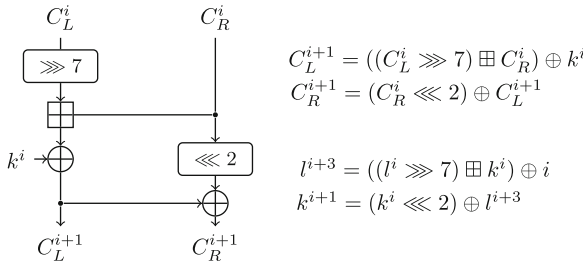


Fig. 1. The round function and key schedule algorithm of SPECK32/64

### 2.3 Overview of Differential-Neural Cryptanalysis

The differential-neural distinguisher operates as a supervised model, distinguishing whether ciphertext pairs originate from plaintext pairs with a defined input difference or from random pairs. Given  $m$  plaintext pairs  $\{(P_j, P'_j), j \in [0, m-1]\}$ , the corresponding ciphertext pairs  $\{(C_j, C'_j), j \in [0, m-1]\}$  constitute a sample (In [15],  $m = 1$ ). Each training sample is associated with a label  $Y$  defined as:

$$Y = \begin{cases} 1, & \text{if } P_j \oplus P'_j = \Delta, j \in [0, m-1] \\ 0, & \text{if } P_j \oplus P'_j \neq \Delta, j \in [0, m-1] \end{cases}$$

The  $\mathcal{ND}$  architecture from [15] uses the prevalent ResNet. It comprises an initial input block, several residual blocks, and a prediction output layer.

In [15], three training schemes are proposed: a) Basic training for short-round distinguishers. b) An enhanced method using the KEYAVERAGING simulation and an  $(r-1)$ -round distinguisher, achieving the optimal 7-round  $\mathcal{ND}$  for SPECK. c) A staged training approach evolving a pre-trained  $(r-1)$ -round distinguisher to an  $r$ -round one in stages, yielding the most extended  $\mathcal{ND}$  on SPECK, covering 8 rounds. In [15], Gohr also showed how to combine a neural distinguisher with a classical differential and use a Bayesian-optimized key-guessing strategy for key recovery. Later, in [16], the authors provide general guidelines for optimizing Gohr’s neural network and diverse optimization approaches across different ciphers, highlighting its efficacy and versatility. The authors also clarify which kind of ciphers the neural network can’t learn beyond differential features.


### 3 Explicitly Explain Knowledge Beyond Full DDT

Studies show differential-based neural distinguishers often outperform DDT-based ones in certain ciphers [3, 15, 16]. However, what specific knowledge these neural distinguishers learn beyond DDT remains elusive. Prior research suggests that these distinguishers rely on differential distributions in the last two rounds and differential-linear (DL) properties [7, 11]. In [15], a “Real Differences Experiment” was conducted to observe how well neural networks could detect real differences beyond DDT. The experiment used randomized ciphertext pairs with a blinding value  $R$  introduced to obscure information beyond the difference. Results showed that neural networks could detect real differences without explicit training, and ciphertext pairs have non-uniform distributions within their difference equivalence classes. But, using blinding values in the form  $R = aa$  (with  $a$  as any 16-bit word), the distinguishers failed (henceforth referred to as Gohr’s *aaaa*-blinding experiment). This underlines that the neural distinguishers aren’t exploiting the key schedule, and they can make finer distinctions than mere difference equivalence classes. These insights are crucial to explicitly explaining  $\mathcal{ND}$ ’s superior classification mechanism. Based on these studies, this section takes a further step towards fully interpreting the knowledge that an  $\mathcal{ND}$  has captured beyond full differential distribution.

We’ll initiate by locating the root of the performance improvement, then deduce the specific pattern that causes the improvement, and finally use this pattern to improve the pure DDT-based distinguisher.

#### 3.1 Locating Information Used by $\mathcal{ND}$ s of SPECK Beyond DDT

In the following, we start with a generalized definition of information that the differential-neural distinguisher might use.



**Fig. 2.** Definition of XOR information

**Table 2.** Experimental results detailing the information harnessed by  $\mathcal{N}\mathcal{D}$ s. Each set comprises both positive and negative samples. The notation  $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D})$  denotes ciphertext pairs derived from plaintext pairs with an input difference of  $(0040, 0000)$ , while *Random* signifies pairs generated from random values.  $\mathcal{R}_1$  refers to a random value.

Set	Positive Samples	Negative Samples	Acc.
1-1	$(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D})$	<i>Random</i>	0.7906
1-2	$(\mathcal{A}\mathcal{R}_1, \mathcal{B}\mathcal{R}_1, \mathcal{C}\mathcal{R}_1, \mathcal{D}\mathcal{R}_1)$	<i>Random</i>	0.7911
1-3	$(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D})$	$(\mathcal{A}\mathcal{R}_1, \mathcal{B}\mathcal{R}_1, \mathcal{C}\mathcal{R}_1, \mathcal{D}\mathcal{R}_1)$	<i>Fail</i>

**Generalized Definition of XOR Information.** In Gohr’s differential-neural distinguishers, given SPECK’s Feistel-like structure, samples are split into four words:  $\mathcal{A}, \mathcal{B}$  (forming the first ciphertext) and  $\mathcal{C}, \mathcal{D}$  (forming the second), as depicted in Fig. 2. In subsequent discussions, a symbol’s superscript denotes the number of encryption rounds. The absence of a superscript implies  $r$  rounds.

Traditional differential distinguishers focus solely on the difference of ciphertext pairs. Yet, as indicated in prior research [13, 22, 27], internal differentials can also be pivotal in cryptanalytic tasks.

We broaden the focus to include the XOR interactions among  $\mathcal{A}, \mathcal{B}, \mathcal{C}$ , and  $\mathcal{D}$ . For brevity, XOR combinations like  $\mathcal{A} \oplus \mathcal{B} \oplus \mathcal{C} \oplus \mathcal{D}$  are shortened to  $\mathcal{ABCD}$ . In other words, beyond the traditionally focused differences like  $\mathcal{AC}$  and  $\mathcal{BD}$ , we explore under-emphasized XORs such as  $\mathcal{AB}, \mathcal{CD}, \mathcal{AD}, \mathcal{BC}$ , and  $\mathcal{ABCD}$ . For clarity, we classify these XORs as: **Inter-XOR** ( $\mathcal{AC}, \mathcal{BD}$ ), **Intra-XOR** ( $\mathcal{AB}, \mathcal{CD}$ ), **Cross-XOR** ( $\mathcal{AD}, \mathcal{BC}$ ), and **Total-XOR** ( $\mathcal{ABCD}$ ).

In SPECK, Intra-XOR and Total-XOR relate to values and differences from the prior round. Specifically, Intra-XOR helps deduce the right-half values, and Total-XOR deduces the right-half differences of the preceding round.

**Is XOR Information the Sole Basis for Differential-Neural Distinguisher’s Decision Making?** Using a mechanical method to determine relations between information sets, it became evident that focusing solely on specified XOR information is natural for finding the source of the information that  $\mathcal{ND}$ s exploit beyond the difference information.

**Determine Relations Between Information Sets Mechanically.** Consider a pair of ciphertexts from a round-reduced SPECK, denoted as  $C_0 = (C_{0L}, C_{0R})$  and  $C_1 = (C_{1L}, C_{1R})$ . Each ciphertext splits into two parts, with  $C_{i,J} \in \mathbb{F}_2^b$  for  $i \in \{0, 1\}$  and  $J \in \{L, R\}$ . For SPECK32/64,  $b = 16$ . Let  $K$  be the last round key, with  $K \in \mathbb{F}_2^b$ . For each  $C_i$ , let  $M_{iL}$  and  $M_{iR}$  represent the state value immediately preceding the XOR with key  $K$  and before the XOR between the left and right branches for  $i \in \{0, 1\}$ . That is

$$C_{0L} = M_{0L} \oplus K, C_{0R} = M_{0L} \oplus M_{0R} \oplus K, C_{1L} = M_{1L} \oplus K, C_{1R} = M_{1L} \oplus M_{1R} \oplus K$$

The method to determine relations between information sets can be outlined in the following steps: Let  $\mathcal{R}_1$  and  $\mathcal{R}_2$  be two random values in  $\mathbb{F}_2^b$ .

1. **Setup:**

- (a) Set up a vector space  $\mathcal{V}$  over the field  $\mathbb{F}_2$  with dimension 7.
- (b) Define various basis vectors for  $\mathcal{V}$ , acting as linear masks whose non-zero bits indicate the variable selection from the following vector  $[M_{0L}, M_{0R}, M_{1L}, M_{1R}, K, \mathcal{R}_1, \mathcal{R}_2]$ . Concretely,

$$\begin{aligned} \Gamma_{M_{0L}} &= [1, 0, 0, 0, 0, 0, 0] & \Gamma_{M_{1L}} &= [0, 0, 1, 0, 0, 0, 0] \\ \Gamma_{M_{0R}} &= [0, 1, 0, 0, 0, 0, 0] & \Gamma_{M_{1R}} &= [0, 0, 0, 1, 0, 0, 0] \\ \Gamma_K &= [0, 0, 0, 0, 1, 0, 0] & \Gamma_{\mathcal{R}_1} &= [0, 0, 0, 0, 0, 1, 0] \\ & & \Gamma_{\mathcal{R}_2} &= [0, 0, 0, 0, 0, 0, 1] \end{aligned}$$

Accordingly,  $[C_{0L}, C_{0R}, C_{1L}, C_{1R}]$  can be obtained using the following masks:

$$\begin{aligned} \Gamma_{C_{0L}} &:= \Gamma_A = \Gamma_{M_{0L}} \oplus \Gamma_K &= [1, 0, 0, 0, 1, 0, 0], \\ \Gamma_{C_{0R}} &:= \Gamma_B = \Gamma_{M_{0L}} \oplus \Gamma_{M_{0R}} \oplus \Gamma_K &= [1, 1, 0, 0, 1, 0, 0], \\ \Gamma_{C_{1L}} &:= \Gamma_C = \Gamma_{M_{1L}} \oplus \Gamma_K &= [0, 0, 1, 0, 1, 0, 0], \\ \Gamma_{C_{1R}} &:= \Gamma_D = \Gamma_{M_{1L}} \oplus \Gamma_{M_{1R}} \oplus \Gamma_K &= [0, 0, 1, 1, 1, 0, 0]. \end{aligned}$$

Besides, we have  $\Gamma_{\mathcal{X}\mathcal{Y}} = \Gamma_{\mathcal{X}} \oplus \Gamma_{\mathcal{Y}}$  for  $\mathcal{X}, \mathcal{Y} \in \{\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{AC}, \mathcal{BD}, \mathcal{AB}, \mathcal{R}_1, \mathcal{R}_2\}$ .

2. **Subspace Generation:** Create the subspaces from given vectors and combinations:

- **Set-1-1:** span of  $\{\Gamma_A, \Gamma_B, \Gamma_C, \Gamma_D\}$ .
- **Set-1-2:** span of  $\{\Gamma_{AR_1}, \Gamma_{BR_1}, \Gamma_{CR_1}, \Gamma_{DR_1}\}$ .
- **Set-1-X:** span of  $\{\Gamma_{AC}, \Gamma_{BD}, \Gamma_{AB}\}$ .
- **Set-2-1:** span of  $\{\Gamma_{AR_1}, \Gamma_{BR_2}, \Gamma_{CR_1}, \Gamma_{DR_2}\}$ .
- **Set-2-2:** span of  $\{\Gamma_{AR_1}, \Gamma_{BR_2}, \Gamma_{CR_2}, \Gamma_{DR_1}\}$ .
- **Set-2-3:** span of  $\{\Gamma_{ABCD}\}$ .

**Table 3.** Experimental results  $\mathcal{N}\mathcal{D}$  leveraging select XOR information.

Set	Positive Samples	Negative Samples	Acc.
2-1	$(\mathcal{A}\mathcal{R}_1, \mathcal{B}\mathcal{R}_2, \mathcal{C}\mathcal{R}_1, \mathcal{D}\mathcal{R}_2)$	<i>Random</i>	0.7558
2-2	$(\mathcal{A}\mathcal{R}_1, \mathcal{B}\mathcal{R}_2, \mathcal{C}\mathcal{R}_2, \mathcal{D}\mathcal{R}_1)$	<i>Random</i>	0.6722
2-3	$(\mathcal{A}\mathcal{B}\mathcal{C}\mathcal{D}, \mathcal{A}\mathcal{B}\mathcal{C}\mathcal{D}, \mathcal{A}\mathcal{B}\mathcal{C}\mathcal{D}, \mathcal{A}\mathcal{B}\mathcal{C}\mathcal{D})$	<i>Random</i>	0.6721

Note that **Set-1-2** is the setting of Gohr’s *aaaa*-blinding experiment.

- Remove randomness:** In light of the observations from [15], where it’s determined that  $\mathcal{N}\mathcal{D}$ s in the single-key attack setting don’t leverage the key schedule, we can adapt the Speck32/64 key schedule to employ independent subkeys. This means we treat  $K$  along with  $\mathcal{R}_1$  and  $\mathcal{R}_2$  as random variables. Consequently, any vector that has a component of  $\Gamma_K$ , or  $\Gamma_{\mathcal{R}_1}$ , or  $\Gamma_{\mathcal{R}_2}$  is deemed random, and hence, devoid of information. For example,  $\Gamma_{C_{i,J}}$  has a linear component  $\Gamma_K$ , thus, a standalone  $C_{i,J}$  lacks information, where  $i \in \{0, 1\}$  and  $J \in \{L, R\}$ . Accordingly, we do as follows.
  - After creating each subspace, randomness is removed from each subspace according to whether a vector has a component from  $\Gamma_K$ , or  $\Gamma_{\mathcal{R}_1}$ , or  $\Gamma_{\mathcal{R}_2}$ . Without ambiguity, the sanitized sets are also denoted by **Set-i-j** for  $i \in \{1, 2\}$  and  $j \in \{1, 2, 3, X\}$ .
- Comparison:** The sanitized sets are then compared against each other to determine if one set equals or is a subset of the other.

The result shows that **Set-1-1** equals **Set-1-2** and **Set-1-X**, meaning that the combination of Inter-XOR and Intra-XOR is exactly what an information-theoretically optimal distinguisher accepting ciphertext pairs can use under the assumption that it does not use key-schedule.

As we proceed, we delve deeper to ascertain the specific XOR information that holds significance.

**Which of the XOR Information is Significant for Differential-Neural Distinguisher?** To isolate the pivotal XOR information, we conducted experiments where a differential-neural distinguisher was given access to only selected XOR data.

All our subsequent experiments were conducted on a 6-round SPECK32/64 with an input difference of (0040,0000), adhering to the configurations presented in Table 17 in [4]. The differential-neural distinguishers, trained as per Table 2 **Set.1-1** to **Set.1-3**, serve as baselines (**Set.1-2** and **Set.1-3** correspond to Gohr’s *aaaa*-blinding experiment)). In the sequel, we use **Set.i-j** to refer to the experimental setup, while **Set-i-j** represents the associated information set for the positive samples, where  $i \in \{1, 2\}$  and  $j \in \{1, 2, 3\}$ .

Defining  $\mathcal{R}_1$  and  $\mathcal{R}_2$  as two distinct random values, **Set.2-1** in Table 3 retains only Inter-XOR and Total-XOR, while **Set.2-2** keeps only Cross-XOR and Total-XOR. **Set.2-3**, on the other hand, exclusively considers Total-XOR. Firstly, our mechanical analysis of sanitized subspaces reveals the following relations:



- $\text{Set-2-1} \subset \text{Set-1-X}$ ,      –  $\text{Set-2-1} \not\subseteq \text{Set-2-2}$ ,      –  $\text{Set-2-3} \subset \text{Set-2-1}$ ,
- $\text{Set-2-2} \subset \text{Set-1-X}$ ,      –  $\text{Set-2-2} \not\subseteq \text{Set-2-1}$ ,      –  $\text{Set-2-3} \subset \text{Set-2-2}$ .

In Table 3 **Set.2-1**, the differential-neural distinguisher’s access is limited to Inter-XOR and Total-XOR – equivalent to what the DDT distinguisher utilizes. Its accuracy aligns closely with the 6-round DDT’s accuracy of 0.758, without any noticeable enhancement. This underscores the differential-neural distinguisher’s advantage over the DDT arising from its access to extra information. From this observation, we reinforce the subsequent conclusion.

**Conclusion 1.** *The differential-neural distinguisher  $\mathcal{ND}^{\text{SPECK}_{rR}}$ ’s superiority over  $\text{DD}^{\text{SPECK}_{rR}}$  is mainly due to its exploit of Intra-XOR and Cross-XOR.*

This conclusion naturally prompts a more intricate query: How does the differential-neural distinguisher effectively exploit Intra-XOR and Cross-XOR? Upon closer inspection, we can further dismiss the significance of Cross-XOR. Given that  $\text{Set-2-3} \subset \text{Set-2-2}$ , it’s evident that **Set-2-3** provides inherently less data than **Set-2-2**. While in **Set.2-2**, combining Total-XOR with either Intra-XOR or Cross-XOR results in a valid distinguisher, solely using Total-XOR in **Set.2-3** yields an accuracy identical to the distinguisher in **Set.2-2**. From this, we conclude that Cross-XOR on its own lacks significance. The differential-neural distinguisher likely uncovers new patterns by melding Inter-XOR with either Intra-XOR or Cross-XOR. This line of reasoning culminates in the following conclusion.

**Conclusion 2.** *Unlike Inter-XOR, neither Intra-XOR nor Cross-XOR independently offers useful information. The differential-neural distinguisher relies on combinations of Inter-XOR with either Intra-XOR or Cross-XOR.*

*Remark 1 (On  $\mathcal{ND}$  exploiting the key schedule).* Gohr’s study in [15] indicates that  $\mathcal{ND}$ s, in a single-key attack on SPECK, do not exploit the key schedule. It naturally raises the question: Do  $\mathcal{ND}$ s behave similarly in related-key scenarios? Motivated by this, we conduct comparison experiments similar to Gohr’s *aaaa*-blinding experiment (comparing  $\mathcal{RK}\text{-}\mathcal{ND}$ s in **Set.1-1** and **Set.2-1**), investigating whether  $\mathcal{RK}\text{-}\mathcal{ND}$ s use the same ciphertext equivalence classes as the single-key  $\mathcal{ND}$ s by [15]. In Sect. 5.1, we delve deep into our  $\mathcal{RK}\text{-}\mathcal{ND}$ s and present an interesting observation reinforcing our following  $\mathcal{ND}$  explainability in Sect. 3.2.

### 3.2 Explicitly Rules to Exploit the Information Beyond Full DDT: From a Cryptanalytic Perspective

This section delves into the exact patterns harnessed by the differential-neural distinguisher. Our exploration commences with an intriguing observation from **Experiment A**, as described in [7]. The experiment unfolds as follows:

1. For each 5-round ciphertext pair difference,  $\delta$ , which results in extreme scores surpassing 0.9 (indicative of a good score) and exhibiting a high frequency of occurrence:

- (a) Generate a set of  $10^4$  random 32-bit numbers.
- (b) Utilize the difference  $\delta$  to construct a dataset encompassing  $10^4$  data pairs, each bearing the difference  $\delta$ .
- (c) Feed the dataset to the differential-neural distinguisher and count the predicted labels.

While DDT-based distinguishers would predict **Experiment A**'s entire data as positive, the differential  $\mathcal{ND}$  does not. For  $\mathcal{ND}$ , the proportion of each difference is consistently at 0.75 (refer to Table 19 in the full version [4]), suggesting that the  $\mathcal{ND}$  employs criteria beyond simple differential probability in its classifications. The consistent proportion of 0.75 also implies a discernible pattern linked to two specific bits. If a ciphertext pair aligns with this bi-bit pattern, it's classified as negative, regardless of high output difference probabilities. This observation prompts an investigation into the potential two-bit pattern, motivating us to look into properties of the addition modulo  $2^n$  ( $\boxplus$ ) from a cryptanalytic perspective.

**Enhancing DDT-Based Distinguishers via Conditional Probabilities.**

In the  $r$ -round SPECK32/64, denote the input and output differences of the last  $\boxplus$  by  $(\alpha, \beta, \gamma)$ , and their respective values by  $(x \ y \ z)$  and  $(x' \ y' \ z')$ . For each output pairs  $((C_L, C_R), (C'_L, C'_R))$ , one knows the following information:  $\gamma = C_L \oplus C'_L$ ,  $\beta = (C_L \oplus C_R \oplus C'_L \oplus C'_R) \ggg 2$ , and  $y = (C_L \oplus C_R) \ggg 2$ . Namely, apart from knowing two differences (*i.e.*,  $\beta$  and  $\gamma$ ), one knows a value (*i.e.*,  $y$ ) around the last  $\boxplus$ . Besides, the input difference  $\alpha$  is unknown but might be biased among positive samples and thus is predictable. Concretely, attributes of the information around the last  $\boxplus$  are as follows:

$\alpha$	unknown but biased	$x$	unknown and balanced
$\beta$	known	$y$	known
$\gamma$	known	$z$	unknown and balanced

The knowledge of  $y$ , which is one of two inputs of the last  $\boxplus$ , provides additional information apart from the differences. The concrete analysis is as follows.

When conditioned on a fixed  $y$ , the differential probability can differ from the average probability over all possible  $y$ . For a valid differential propagation  $(\alpha, \beta \mapsto \gamma)$  through  $\boxplus$ , consider each bit position  $i$  where  $0 \leq i < n - 1$ : If  $\text{eq}(\alpha, \beta, \gamma)_i = 1$ , the difference propagation at the  $(i + 1)$ -th position is deterministic, as elucidated in [20]; Conversely, for  $\text{eq}(\alpha, \beta, \gamma)_i = 0$ , the  $(i + 1)$ -th bit's difference propagation is probabilistic; for a given  $(i + 1)$ -th bit differences to be fulfilled, the input values at the  $i$ -th position (namely,  $x_i, y_i, c_i$  – the carry's  $i$ -th bit) must satisfy a certain linear constraint, detailed in Observation 1.

**Observation 1 ([10]).** *Let  $\delta = (\alpha, \beta \mapsto \gamma)$  be a possible XOR-differential through addition modulo  $2^n$  ( $\boxplus$ ). Let  $(x, y)$  and  $(x \oplus \alpha, y \oplus \beta)$  be a conforming*

**Table 4.** Necessary and sufficient conditions for a one-bit difference from Observation 1

Case No.	Difference	Constraint on values	Known
Cxy <sub>(i+1,i)</sub>	$\begin{cases} \text{eq}(\alpha, \beta, \gamma)_i = 0, \\ \alpha_i \oplus \beta_i = 0. \end{cases}$	$\text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \alpha_i = x_i \oplus y_i$	None
Cxc <sub>(i+1,i)</sub>	$\begin{cases} \text{eq}(\alpha, \beta, \gamma)_i = 0, \\ \alpha_i \oplus \beta_i = 1, \\ \alpha_i \oplus \text{xor}(\alpha, \beta, \gamma)_i = 0. \end{cases}$	$\text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \alpha_i = x_i \oplus c_i$	None
Cyc <sub>(i+1,i)</sub>	$\begin{cases} \text{eq}(\alpha, \beta, \gamma)_i = 0, \\ \alpha_i \oplus \beta_i = 1, \\ \alpha_i \oplus \text{xor}(\alpha, \beta, \gamma)_i = 1. \end{cases}$	$\text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \beta_i = y_i \oplus c_i$	$y_i \oplus c_i$

The column titled “Known” indicates whether the fulfilment of the condition might be known in SPECK’s last  $\boxplus$ .

pair of  $\delta$ ,  $x$  and  $y$  should satisfy the follows. For  $0 \leq i < n - 1$ , if  $\text{eq}(\alpha, \beta, \gamma)_i = 0$

$$\left. \begin{aligned} x_i \oplus y_i &= \text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \alpha_i, && \text{if } \alpha_i \oplus \beta_i = 0, \\ x_i \oplus c_i &= \text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \alpha_i, && \text{if } \alpha_i \oplus \text{xor}(\alpha, \beta, \gamma)_i = 0, \\ y_i \oplus c_i &= \text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \beta_i, && \text{if } \alpha_i \oplus \text{xor}(\alpha, \beta, \gamma)_i = 1, \end{aligned} \right\} \text{if } \alpha_i \oplus \beta_i = 1,$$

where  $c_i$  is the  $i$ -th carry bit,  $x \boxplus y = z$ ,  $\text{eq}(a, b, d) = (\neg a \oplus b) \wedge (\neg a \oplus d)$  (i.e.,  $\text{eq}(a, b, d) = 1$  if and only if  $a = b = d$ ), and  $\text{xor}(a, b, d) = a \oplus b \oplus d$ .

In other words, at bit positions  $i$  and  $i + 1$ , a valid difference tuple  $(\alpha_{i+1,i}, \beta_{i+1,i}, \gamma_{i+1,i})$  that satisfies  $\text{eq}(\alpha_i, \beta_i, \gamma_i) = 0$  imposes a 1-bit linear constraint on the tuple  $(x_i, y_i, c_i)$ . As  $c_i$  is determined by lower bits, the freedom for conforming to the constraint comes exclusively from the  $i$ -th bits of  $x$  and  $y$ , independent of constraints at other bit positions. Accordingly, the constraints on  $(x_i, y_i)$ ,  $(x_i, c_i)$ , or  $(y_i, c_i)$  as listed in Observation 1 are necessary and sufficient. Therefore, when the constraint at a bit position is fulfilled, the conditional probability  $\tilde{p}$  of a differential whose unconditional probability is  $p$  should be calculated as  $2 \cdot p$ ; when unfulfilled, it is 0. In comparison, the conditional probability for random pairs is still at most  $2^{-n}$ . Hence, leveraging conditional probability for classification amplifies the advantage.

To clarify when the fulfilment of the constraints at the last  $\boxplus$  can be effectively predicted, we catalog cases from Observation 1 in Table 4, naming them Cxy<sub>(i+1,i)</sub>, Cxc<sub>(i+1,i)</sub>, and Cyc<sub>(i+1,i)</sub>. As above analyzed, in SPECK32/64’s last  $\boxplus$ , among the tuple  $(x, y, c)$  (with  $c = z \oplus x \oplus y$  and unknown  $z$ ), only  $y$  is known. Hence, exploiting knowledge of  $y$  requires examining bit positions with differential constraints fulfilling Cyc<sub>(i+1,i)</sub> in Table 4.

In the Cyc<sub>(i+1,i)</sub> case, the constraint is on  $y_i \oplus c_i$ . While  $c_i$  may seem unknown, it is determined by lower bits:  $c_i = x_{i-1}y_{i-1} \oplus (x_{i-1} \oplus y_{i-1})c_{i-1}$ . The knowledge on  $c_i$  might be inferred if the  $(i - 1)$ -th bit differences meet the condition  $\text{eq}(\alpha_{i-1}, \beta_{i-1}, \gamma_{i-1}) = 0$ , as per Observation 1. For example, when

**Table 5.** Cases for deducing the  $i$ -th carry bit  $c_i$

Case No.	Difference	Value	Known
Cy0c0 <sub>(i,i-1)</sub>		$y_{i-1} = 0, c_{i-1} = 0$	$c_i = 0$
Cylc1 <sub>(i,i-1)</sub>		$y_{i-1} = 1, c_{i-1} = 1$	$c_i = 1$
Cxy0 <sub>(i,i-1)</sub>	Cxy <sub>(i,i-1)</sub> and $\text{xor}(\alpha, \beta, \gamma)_i \oplus \alpha_{i-1} = 0$	$x_{i-1} \oplus y_{i-1} = 0$	$c_i = y_{i-1}$
Cxy1 <sub>(i,i-1)</sub>	Cxy <sub>(i,i-1)</sub> and $\text{xor}(\alpha, \beta, \gamma)_i \oplus \alpha_{i-1} = 1$	$x_{i-1} \oplus y_{i-1} = 1$	$c_i = c_{i-1}$
Cxc0 <sub>(i,i-1)</sub>	Cxc <sub>(i,i-1)</sub> and $\text{xor}(\alpha, \beta, \gamma)_i \oplus \alpha_{i-1} = 0$	$x_{i-1} \oplus c_{i-1} = 0$	$c_i = c_{i-1}$
Cxc1 <sub>(i,i-1)</sub>	Cxc <sub>(i,i-1)</sub> and $\text{xor}(\alpha, \beta, \gamma)_i \oplus \alpha_{i-1} = 1$	$x_{i-1} \oplus c_{i-1} = 1$	$c_i = y_{i-1}$
Cyc0 <sub>(i,i-1)</sub>	Cyc <sub>(i,i-1)</sub> and $\text{xor}(\alpha, \beta, \gamma)_i \oplus \beta_{i-1} = 0$	$y_{i-1} \oplus c_{i-1} = 0$	$c_i = y_{i-1}$
Cyc1 <sub>(i,i-1)</sub>	Cyc <sub>(i,i-1)</sub> and $\text{xor}(\alpha, \beta, \gamma)_i \oplus \beta_{i-1} = 1$	$y_{i-1} \oplus c_{i-1} = 1$	$c_i = x_{i-1}$

**Table 6.** Cases where the knowledge on  $y$  can be used to check the fulfilment of the differential constraints

Case No.	Difference	Known
C1	Cyc <sub>(0,-1)</sub>	$\text{xor}(\alpha, \beta, \gamma)_1 \oplus \beta_0 = y_0$
C2	Cyc <sub>(2,1)</sub> and Cy0 <sub>(1,0)</sub>	$\text{xor}(\alpha, \beta, \gamma)_2 \oplus \beta_1 = y_1$
C3	Cyc <sub>(i+1,i)</sub> and (Cxy0 <sub>(i,i-1)</sub> or Cxc1 <sub>(i,i-1)</sub> or Cyc0 <sub>(i,i-1)</sub> )	$\text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \beta_i = y_i \oplus y_{i-1}$
C4	Cyc <sub>(i+1,i)</sub> and (Cxy1 <sub>(i,i-1)</sub> or Cxc0 <sub>(i,i-1)</sub> ) and (Cxy0 <sub>(i-1,i-2)</sub> or Cxc1 <sub>(i-1,i-2)</sub> or Cyc0 <sub>(i-1,i-2)</sub> )	$\text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \beta_i = y_i \oplus y_{i-2}$

$$\left\{ \begin{array}{l} (\alpha_i, \beta_i, \gamma_i) = (0, 1, 0), \\ (\alpha_{i-1}, \beta_{i-1}, \gamma_{i-1}) = (1, 1, 0) \end{array} \right. , \text{ one knows that } \left\{ \begin{array}{l} \text{eq}(\alpha, \beta, \gamma)_{i-1} = 0, \\ \alpha_{i-1} \oplus \beta_{i-1} = 0, \\ \text{xor}(\alpha, \beta, \gamma)_i \oplus \alpha_{i-1} = 0. \end{array} \right.$$

From Table 4, one has  $x_{i-1} \oplus y_{i-1} = 0$ . Thus,  $c_i = x_{i-1}y_{i-1} \oplus (x_{i-1} \oplus y_{i-1})c_{i-1} = y_{i-1}$ . Therefore,  $y_i \oplus c_i = y_i \oplus y_{i-1}$ . As a consequence, one can predict the fulfilment of constraint in case Cyc<sub>(i+1,i)</sub> by observing whether  $y_i \oplus y_{i-1} = \text{xor}(\alpha, \beta, \gamma)_{i+1} \oplus \beta_i$ . Table 5 lists more cases where  $c_i$  might be known.

Incorporating observations from Table 4 and Table 5, one gets Table 6, which lists various cases where the knowledge of  $y$  can be used to determine the satisfaction of differential constraints.

Note that apart from the general cases (C3 and C4) at the  $i$ -th bit, special cases (C1 and C2) emerge at the two least significant bits due to the carry bit  $c_0$  being 0. For example,

1. at the 0th bit position, observing  $\beta_0 = 0$  and  $\gamma_0 = 1$  determines  $\alpha_0 = 1$  based on Alg. 3 in [4]. From case Cyc<sub>(i+1,i)</sub> in Table 4 and given  $c_0 = 0$ , one knows that  $\text{xor}(\alpha, \beta, \gamma)_1 \oplus \beta_0 = y_0 \oplus c_0 = y_0$ ;
2. at the 1st bit position,  $c_1 = x_0y_0 \oplus (x_0 \oplus y_0)c_0 = x_0y_0$ . Given an observed  $y_0 = 0$ , one knows  $c_1 = 0$ . Consequently, in case Cyc<sub>(2,1)</sub> and  $y_0 = 0$ , one knows  $\text{xor}(\alpha, \beta, \gamma)_2 \oplus \beta_1 = y_1 \oplus c_1 = y_1$ ;

3. in general case C3, based on Table 5,  $c_i$  is determined as  $y_{i-1}$ , leading to the use of  $y_i \oplus y_{i-1}$ ;
4. in general case C4, applying Table 5 to the  $(i - 1, i - 2)$ -th bit position, it is inferred that  $c_i = c_{i-1} = y_{i-2}$ , leading to the use of  $y_i \oplus y_{i-2}$ ;
5. for cases where  $c_{i-1} = c_{i-2}$ , one can further observe differences at the  $(i - 2)$ -th bit position and continues deducting  $c_{i-2}$  by observing bit differences at the  $i - 3$  position.

Table 7 lists some concrete examples of differential patterns where the observation of  $y$  enables prediction of whether differential constraints are met.

*Remark 2.* These constraints on values for valid differential propagation resonate with established concepts. Specifically, insights derived from Table 6 align with findings on multi-bit constraints from [18, 19], quasi-differential trails in [8], and extended differential-linear approximations in [11]. Table 7 exhibits the correspondence between examples of cases in Table 6 and these established concepts. For instance, given a differential propagation  $(\alpha_{i+1,i,i-1}, \beta_{i+1,i,i-1} \mapsto \gamma_{i+1,i,i-1}) = (*01, *11 \mapsto *00)$  (for  $0 < i < n - 1$ ),

1. using the 1.5-bit constraints concept and the finite state machines representing the differential properties of modular addition from [18, 19], one can get a new constraint and refine the propagation  $(-x, -xx \mapsto ---)$  to  $(-x, ->x \mapsto ---)$  (where the notations  $\{-, x, >, <, =, !\}$  are explained below Table 7); more generally, C3 cases correspond to the 1.5-bit constraints  $\{>, <, =, !\}$  in [18, 19];
2. using the quasi-differential trail concept from [8], the differential trail  $(001, 011 \mapsto 000)$  comprises a non-trivial quasi-differential trail with a mask of  $(000, 011 \mapsto 000)$ . The non-trivial quasi-differential trail has correlation  $-2^{-1}$  (i.e., additional weight of 0). Consequently, the “fixed- $y$ ” probability of this differential trail is  $(1 - (-1)^{y_i \oplus y_{i-1}}) \cdot 2^{-1}$ , i.e., the probability equals 1 when  $y_i \oplus y_{i-1} = 1$  and 0 in the opposite case;
3. using the extended differential-linear connectivity table (EDLCT) concept from [11], assessing the constraint  $y_i \oplus y_{i-1} = \alpha_{i+1} \oplus \beta_{i+1} \oplus \gamma_{i+1} \oplus \beta_i$  aligns with gauging the bias of the linear approximation  $(x_{i+1} \oplus x'_{i+1}) \oplus (y_{i+1} \oplus y'_{i+1}) \oplus (z_{i+1} \oplus z'_{i+1}) \oplus (y_i \oplus y'_i) \oplus (y_i \oplus y_{i-1})$  that corresponds to selecting bits  $[x_{i+1}, y_{i+1}, z_{i+1}, y_{i-1}]$  and  $[x'_{i+1}, y'_{i+1}, z'_{i+1}, y'_i]$ .

As noted in [7],  $\mathcal{ND}$ s rely on differential-linear (DL) properties. We note that pure DL properties do not provide additional information beyond full DDT; the differential-linear distribution can be directly derived from the full differential distribution. It is the *extended* differential-linear distribution [11] (which includes the selection of ciphertext values apart from differences) that contains additional information.

To directly exploit these observations for an  $r$ -round SPECK32/64, a preliminary is to effectively predict the input difference  $\alpha$  at the last  $\boxplus$ , which equals  $((\delta_R^{r-2})^{\lll 2} \oplus \delta_R^{r-1})^{\ggg 7}$ . Given the known  $\delta_R^{r-1}$  from  $r$ -round outputs, the focus shifts to predicting  $(\delta_R^{r-2})^{\lll 2}$ . Notably, for  $r \leq 7$  and input difference  $(0040, 0000)$ , some bits of  $(\delta_R^{r-2})^{\lll 2}$  exhibit bias, as detailed in Table 8, enabling predictions of  $\alpha$  for positive samples.

**Table 7.** Concrete examples of differential patterns where one can predict the fulfilment of the differential constraints by observing the value of  $y$ 

Case	Observation 1			Multi-bit Constraints in [18, 19]		Quasi-differential in [8]		Extended DLCT in [11]	
	No.	Difference local map	Value	Observations	org	new	diff		mask (add. w)
C1	$\alpha_{1,0}$	*1	$x_{1,0}$	**	$y_0 =$	-x	-x	01 00	[ $x_1, y_1, z_1$ ], [ $x'_1, y'_1, z'_1, y'_0$ ]
	$\beta_{1,0}$	*0	$y_{1,0}$	**	$\alpha_1 \oplus \beta_1 \oplus \gamma_1 \oplus 0$	--	-0	00 01 (+2 <sup>0</sup> )	
	$\gamma_{1,0}$	*1	$z_{1,0}$	**		-x	-x	01 00	
C2	$\alpha_{2,1,0}$	*1*	$x_{2,1,0}$	***	$y_1 =$	-x?	-x?	010 000	[ $x_2, y_2, z_2, y_0$ ], [ $x'_2, y'_2, z'_2, y'_1$ ]
	$\beta_{2,1,0}$	*0*	$y_{2,1,0}$	**0	$\alpha_2 \oplus \beta_2 \oplus \gamma_2 \oplus 0$	--	-00	000 011 (+2 <sup>-1</sup> )	
	$\gamma_{2,1,0}$	*1*	$z_{2,1,0}$	***		-x?	-x?	010 000	
C3	$\alpha_{i+1,i,i-1}$	*01	$x_{i+1,i,i-1}$	***	$y_i \oplus y_{i-1} =$	-x	--x	001 000	[ $x_{i+1}, y_{i+1}, z_{i+1}, y_{i-1}$ ], [ $x'_{i+1}, y'_{i+1}, z'_{i+1}, y'_i$ ]
	$\beta_{i+1,i,i-1}$	*11	$y_{i+1,i,i-1}$	**	$\alpha_{i+1} \oplus \beta_{i+1} \oplus \gamma_{i+1} \oplus 1$	--x	->x	011 011 (-2 <sup>0</sup> )	
	$\gamma_{i+1,i,i-1}$	*00	$z_{i+1,i,i-1}$	***		---	---	000 000	
C3	$\alpha_{i+1,i,i-1}$	*11	$x_{i+1,i,i-1}$	***	$y_i \oplus y_{i-1} =$	--x	--x	011 000	[ $x_{i+1}, y_{i+1}, z_{i+1}, y_{i-1}$ ], [ $x'_{i+1}, y'_{i+1}, z'_{i+1}, y'_i$ ]
	$\beta_{i+1,i,i-1}$	*00	$y_{i+1,i,i-1}$	**	$\alpha_{i+1} \oplus \beta_{i+1} \oplus \gamma_{i+1} \oplus 0$	---	--	000 011 (+2 <sup>0</sup> )	
	$\gamma_{i+1,i,i-1}$	*11	$z_{i+1,i,i-1}$	***		--x	--x	011 000	
C4	$\alpha_{i+1,i,i-1,i-2}$	*111	$x_{i+1,i,i-1,i-2}$	****	$y_i \oplus y_{i-2} =$	-xxx	-xxx	0111 0000	[ $x_{i+1}, y_{i+1}, z_{i+1}, y_{i-2}$ ], [ $x'_{i+1}, y'_{i+1}, z'_{i+1}, y'_i$ ]
	$\beta_{i+1,i,i-1,i-2}$	*010	$y_{i+1,i,i-1,i-2}$	****	$\alpha_{i+1} \oplus \beta_{i+1} \oplus \gamma_{i+1} \oplus 0$	--x-	- $\frac{2}{0}$ x-	0010 0101 (+2 <sup>0</sup> )	
	$\gamma_{i+1,i,i-1,i-2}$	*101	$z_{i+1,i,i-1,i-2}$	****		-x-x	-x-x	0101 0000	

0:  $y_i = y'_i = 0$     1:  $y_i = y'_i = 1$     -:  $y_i = y'_i$     x:  $y_i \neq y'_i$      $\frac{2}{0}$ : uncommon 2.5-bit constraint "28000014"  
 \*:  $y'_i = y_i = y_{i-1}$     !:  $y'_i = y_i \neq y_{i-1}$     <:  $y'_i \neq y_i = y_{i-1}$     >:  $y'_i \neq y_i \neq y_{i-1}$

**Table 8.** Bit bias towards ‘0’ of  $(\delta_R^{r-2})^{\lll 2}$  for  $4 \leq r \leq 7$ , where the input difference of the plaintext is (0040, 0000). A positive (resp. negative) value indicates a bias towards ‘0’ (resp. ‘1’).

Position	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
$(\delta_2^r)^{\lll 2}$	0.4689	0.4377	0.3752	0.2498	-0.0002	-0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	-0.5000	0.5000	-0.4922	0.4844
$(\delta_3^r)^{\lll 2}$	0.3749	0.2809	0.1247	-0.1241	0.0100	0.4687	0.4451	0.4062	0.3435	0.2498	-0.1251	-0.0000	-0.4922	0.4844	-0.4608	0.4297
$(\delta_4^r)^{\lll 2}$	0.0926	-0.0347	0.0004	-0.0617	0.0028	-0.3709	0.3012	0.2046	-0.0578	-0.0004	0.0312	-0.0009	0.4451	0.4059	-0.2931	0.2035
$(\delta_5^r)^{\lll 2}$	-0.0002	-0.0016	-0.0002	-0.0238	0.0002	0.1531	-0.0243	-0.0001	-0.0034	-0.0011	-0.0076	-0.0001	0.2700	0.1772	-0.0597	-0.0061
$(\delta_6^r)^{\lll 2}$	0.0001	-0.0006	0.0005	0.0103	0.0002	-0.0002	-0.0009	-0.0003	-0.0004	-0.0002	-0.0033	0.0007	-0.0438	-0.0048	-0.0007	-0.0010

*A Simple Procedure to Improve the DDT-Based Distinguisher.* To improve a DDT-based distinguisher for an  $r$ -round SPECK32/64 using its  $\text{DDT}_{(0040, 0000)}$ , we proceed as follows, resulting in distinguishers named  $\mathcal{YD}^{\text{SPECK}_{r,R}}$ :

1. Compute the bias (towards 0) of each bit of  $(\delta_R^{r-2})^{\lll 2}$ ,
2. Predict bit values for  $(\delta_R^{r-2})^{\lll 2}$  based on their biases: assign a value of 0 if bias  $\geq 0$  and 1 otherwise,
3. Define the absolute bias of the  $i$ -bit of  $((\delta_R^{r-2})^{\lll 2})^{\ggg 7}$  as  $\epsilon_\alpha(i)$ ,
4. For each output pair of  $r$ -round SPECK32/64, use Alg. 1 to predict its classification.

*Results of Improving the DDT-Based Distinguisher.* Table 9 presents the performance of  $\mathcal{YD}^{\text{SPECK}_{r,R}}$  distinguishers, derived from the described enhancement of  $\text{DD}^{\text{SPECK}_{r,R}}$ . For rounds  $4 \leq r \leq 7$ ,  $\mathcal{YD}^{\text{SPECK}_{r,R}}$  typically shows improvement.

---

**Algorithm 1:** A simple procedure to improve the DDT-based distinguisher:  $\mathcal{YD}^{\text{SPECK}_{r,R}}$

---

1. Get the differential probability  $p$  of  $(0040, 0000) \mapsto (C_L \oplus C'_L, C_R \oplus C'_R)$  by looking up the table  $\text{DDT}_{(0040, 0000)}[(C_L \oplus C'_L, C_R \oplus C'_R)]$
  2. Compute the following information around the last  $\boxplus$  from  $((C_L, C_R), (C'_L, C'_R))$ :
    - (a)  $\gamma \leftarrow C_L \oplus C'_L$ ,  $\beta \leftarrow (C_L \oplus C_R \oplus C'_L \oplus C'_R) \ggg^2$ ,
    - (b)  $\alpha \leftarrow ((\delta_R^{r-2}) \lll^2 \oplus \beta) \ggg^7$ ,  $y \leftarrow (C_L \oplus C_R) \ggg^2$ .
  3. For bit position 0, if  $\epsilon_\alpha(1) > \tau$  and  $\epsilon_\alpha(0) > \tau$ , do:
    - (a) If  $\text{Cyc}_{(0,-1)}$ , do:  $p \leftarrow (1 + (-1)^{\text{xor}(\alpha, \beta, \gamma)1 \oplus \beta_0 \oplus y_0}) \cdot p$ .
  4. For bit position 1, if  $\epsilon_\alpha(2) > \tau$  and  $\epsilon_\alpha(1) > \tau$ , do:
    - (a) If  $\text{Cyc}_{(2,1)}$  and  $y_0 = 0$ , do:  $p \leftarrow (1 + (-1)^{\text{xor}(\alpha, \beta, \gamma)2 \oplus \beta_1 \oplus y_1}) \cdot p$ .
  5. For each bit position  $i$  ( $1 < i < n - 1$ ), if  $\epsilon_\alpha(i + 1) > \tau$  and  $\epsilon_\alpha(i) > \tau$  and  $\epsilon_\alpha(i - 1) > \tau$ , do:
    - (a) If  $\text{Cyc}_{(i+1,i)}$  and  $(\text{Cxy}0_{(i,i-1)}$  or  $\text{Cxc}1_{(i,i-1)}$  or  $\text{Cyc}0_{(i,i-1)})$ , do:
 
$$p \leftarrow (1 + (-1)^{\text{xor}(\alpha, \beta, \gamma)_{i+1 \oplus \beta_i \oplus y_i \oplus y_{i-1}}}) \cdot p.$$
    - (b) If  $\text{Cyc}_{(i+1,i)}$  and  $(\text{Cxy}1_{(i,i-1)}$  or  $\text{Cxc}0_{(i,i-1)})$  and  $(\text{Cxy}0_{(i-1,i-2)}$  or  $\text{Cxc}1_{(i-1,i-2)}$  or  $\text{Cyc}0_{(i-1,i-2)})$  and  $\epsilon_\alpha(i - 2) > \tau$ , do:
 
$$p \leftarrow (1 + (-1)^{\text{xor}(\alpha, \beta, \gamma)_{i+1 \oplus \beta_i \oplus y_i \oplus y_{i-2}}}) \cdot p.$$
  6. If  $p > 2^{-n}$ , predict  $Z \leftarrow 1$ ; else predict  $Z \leftarrow 0$ .
- 

In contrast, when applying a similar method to adjust the  $\mathcal{N}\mathcal{D}^{\text{SPECK}_{r,R}}$  score  $Z$  (converting score  $Z$  to probability  $p$  using  $p = Z/(1 - Z) \cdot 2^{-n}$ ), the accuracy does not get improved. It is unchanged for  $\mathcal{N}\mathcal{D}^{\text{SPECK}_{4R}}$  and marginally degrades for rounds  $5 \leq r \leq 7$  since the threshold  $\tau$  is set less than 0.5. This suggests that the additional information useful in improving DDT-based distinguishers does not help improve  $\mathcal{N}\mathcal{D}$ 's; thus, the  $\mathcal{N}\mathcal{D}$ 's might have maximally utilized this information already. Thus, we conclude as follows.

**Conclusion 3.** *By utilizing conditional differential distributions when the input and/or output values of the last nonlinear operation are observable, a distinguisher can surpass pure DDT-based counterparts. Accordingly, if these conditional distributions differ greatly from the averaged differential distribution, and the satisfaction of the conditions is either observable or effectively predictable, then  $r$ -round  $\mathcal{N}\mathcal{D}$ s can outperform  $r$ -round DDT-based distinguishers.*

For SPECK, one of the two inputs of the last non-linear operation ( $\boxplus$ ) is observable. If conditioned on this input, the conditional differential distribution can diverge significantly from the averaged one. Therefore, an optimal distinguisher can obviously outperform a pure DDT-based counterpart. A similar analysis applies to SIMON. In SIMON, the values that go through the last nonlinear operation are fully observable. Consequently, it is interpretable that in the case of SIMON, an  $r$ -round  $\mathcal{N}\mathcal{D}$  can achieve an accuracy close to the  $(r - 1)$ -round DDT [3].

This conclusion can be further supported by the following experimental result: In a modified  $r$ -round SPECK32/64 where the last key XORing is omitted, revealing both  $z$  and  $y$  (equating to full awareness of the satisfaction of the last round's differential constraints given a predictable input difference  $\alpha$ ), a well-trained  $r$ -round  $\mathcal{ND}$  achieves an accuracy close to the  $(r - 1)$ -round  $\mathcal{DD}$ . Interestingly, subsequent observations on  $\mathcal{RK-ND}$ s reinforce our conclusion, while the conclusion itself aids in interpreting those observations.

### 3.3 Distinguishers Using Systematic Computation of Conditional Differential Probability Under Known $y$

The simple process in Algorithm 1 is fast, but it requires evaluating the bias of each bit of the difference on the right branch of round  $r - 2$  to estimate the input difference  $\alpha$  for the last modular addition. The differential probability can only be adjusted if the estimated bias of the corresponding bit of  $\alpha$  exceeds a certain threshold. As a result, it does not make the most of the information in  $y$ . Therefore, we further designed a process, described in Algorithm 2, to systematically calculate the differential probability conditioned under the known value of  $y$  and predict based on the  $(r - 1)$ -round DDT<sup>1</sup>.

In essence, the systematic process involves using  $\beta$ ,  $\gamma$ , and  $y$  to determine all possible  $\alpha$ s and the conditional differential probabilities of the last round. It combines this information with the probabilities of the previous  $(r - 1)$  rounds to calculate the conditional differential probability for  $r$  rounds under the known value of  $y$ . Finally, it uses the systematically computed conditional probability for prediction.

More concretely, in the process, we have the following procedures:

1. **Precomputation:** We generate three  $b$ -bit conditional DDTs, denoted as  $\mathbf{A}_0$ ,  $\mathbf{A}_{\text{next}}$ , and  $\mathbf{A}_{\text{next}}^c$ , of the single modular addition operation  $\boxplus$ . These resemble Dinur's  $b$ -bit filter in [12]:
  - (a)  $\mathbf{A}_0$  tells all valid  $b$ -bit values of  $\alpha$  with their associated probability  $pr$  for given  $b$ -bit inputs  $\beta$ ,  $\gamma$ , and  $y$  at the first  $b$  least significant bits (LSB) where the first carry bits are zeros.
  - (b)  $\mathbf{A}_{\text{next}}$  tells all valid 1-bit values of  $\alpha_{\text{next}}$  with their associated probability  $pr$  for given  $b$ -bit inputs  $\beta$ ,  $\gamma$ ,  $y$ , and  $(b - 1)$ -bit  $\alpha$  at intermediate consecutive  $b$  bits where the LSB of carry is undetermined.
  - (c)  $\mathbf{A}_{\text{next}}^c$  is similar to  $\mathbf{A}_{\text{next}}$  but serves scenarios with known carry LSBs.
2. **Initialization:** From a received ciphertext pair, we derive the output difference  $\gamma$ , input difference  $\beta$ , and input value  $y$ ; initialize the to-be-calculated probability  $p$  and the last round's probability factor  $q$  with 0 and 1.
3. **Generate candidate LSB  $b$ -bit of  $\alpha$ :**
  - (a) Using table  $\mathbf{A}_0$ , we obtain candidates for the LSB  $b$ -bit of  $\alpha$  based on the LSB  $b$ -bit of  $\beta$ ,  $\gamma$ , and  $y$ , update  $q$  with the associated  $pr$ .

<sup>1</sup> Please refer to [1] for the implementation codes and experimental results.



- (b) For each valid LSB  $b$ -bit of  $\alpha$ , we invoke ‘ComputeCarryNextBit’ to determine the carry bits wherever possible according to Table 5.
4. **Iterative Calculation:** For each valid LSB  $b$ -bit of  $\alpha$ ,
- (a) Starting from the  $(b - 1)$ -th bit, we invoke ‘ComputeAlphaPrNextBit’ to sequentially determine  $\alpha$ ’s later bits and the respective augmentation of the probability factor to  $q$ ; alongside, we use ‘ComputeCarryNextBit’ to determine the carry bits wherever possible, preparing to be used to derive later bits of  $\alpha$  in case of Cyc or be used to look up  $\mathbf{A}_{\text{next}}^c$ .

Within procedure ‘ComputeAlphaPrNextBit’:

- (a) Once  $\alpha$  is fully assigned, we calculate the output difference of the penultimate round and use it to look up the  $(r - 1)$ -round DDT. The resultant value, upon multiplied by the last round’s probability factor  $q$ , yields a contribution term to the final probability  $p$ .
- (b) At an intermediate bit position  $i$ , equal three input/output bits differences facilitate the direct determination of the subsequent  $\alpha$  bit.
- (c) When input/output bits differences at position  $(i + 1, i)$  conforms to the Cyc $_{(i+1,i)}$  condition with an determined value for  $c_i$ , the subsequent  $\alpha$  bit is deduced using  $y_i \oplus c_i$ . After determining  $\alpha_{i+1}$ , we invoke ‘ComputeCarryNextBit’ to determine the carry bit  $c_{i+1}$  wherever possible.
- (d) Otherwise (in the absence of conformity or a determined  $c_i$  value),  $\alpha_{i+1}$  is enumerated using either  $\mathbf{A}_{\text{next}}$  or  $\mathbf{A}_{\text{next}}^c$ , depending on whether the carry bit before  $b$  bits of the  $(i + 1)$ -th bit is determined.
- (e) After obtaining  $\alpha_{i+1}$  and its probability  $pr$ , we continue to determine the next  $\alpha$  bit, updating the probability factor by multiplying  $pr$  to  $q$ .

The resulting procedure is slower than the simple one; however, the resulting distinguishers, named “ $\mathcal{AD}_{\mathbf{YD}}$ ”, have accuracy exceeds not only that of the distinguishers  $\mathcal{DD}$ s but also the neural distinguishers  $\mathcal{ND}$ s, comparable to the  $r - 1$ -round DDT-based key-averaging distinguishers  $\mathcal{AD}_{\mathbf{KD}}$ s [2] (refer to Table 9 and Table 20 in [4]), indicating an exemplary accuracy for  $\mathcal{ND}$ s.

### 3.4 Discussion on $\mathcal{ND}$ ’s Advantages

Based on the above observations and experiments, we can conclude that  $\mathcal{ND}$ ’s advantage over pure differential-based distinguishers comes from exploiting the conditional differential distribution under the partially known value from ciphertexts input to the last non-linear operation. More specifically,  $\mathcal{ND}$ s exploited the correlation between the ciphertexts’ partial value, the ciphertext pair’s differences, and the intermediate states’ differences. Specifically, when some of the last-round nonlinear operations’ inputs and outputs are known (*i.e.*, not XORed with independently randomized key bits), a distinguisher can achieve higher distinguishing accuracy than an  $r$ -round pure differential-based distinguisher.

These findings apply not only to the SPECK but also to other block ciphers, such as SIMON and GIFT (refer to Appendix D.1 in [4]), and demonstrate the ability of neural networks to capture and utilize complex relationships between

---

**Algorithm 2:** Known- $y$  differential distinguishers:  $\mathcal{AD}_{\text{YD}}^{\text{SPECK}_{r,R}}$ 


---

1.  $b \leftarrow 6$  // for practical reason, we consider 6-bit conditional DDT of  $\boxplus$
2.  $\mathbf{A}_0, \mathbf{A}_{\text{next}}, \mathbf{A}_{\text{next}}^c \leftarrow \text{GenMultiBitsConditionalDDTs}(b)$
3.  $p \leftarrow 0.0, q \leftarrow 1.0$
4. Compute the following around the last  $\boxplus$  from  $((C_L, C_R), (C'_L, C'_R))$ :
  - (a)  $\gamma \leftarrow C_L \oplus C'_L, \beta \leftarrow (C_L \oplus C_R \oplus C'_L \oplus C'_R) \ggg^2, y \leftarrow (C_L \oplus C_R) \ggg^2$ .
5.  $\alpha \leftarrow \mathbf{0}, c \leftarrow \mathbf{0}$
6.  $\beta_b \leftarrow$  LSB  $b$  bits of  $\beta, \gamma_b \leftarrow$  LSB  $b$  bits of  $\gamma, y_b \leftarrow$  LSB  $b$  bits of  $y$
7. For  $(\alpha_b, pr) \in \mathbf{A}_0[\beta_b, \gamma_b, y_b]$ 
  - (a)  $\alpha \leftarrow \alpha_b$
  - (b) For  $i$  in  $\{0, 1, \dots, b-2\}$ :  $\text{ComputeCarryNextBit}(c, \alpha, \beta, \gamma, y, i)$
  - (c)  $\text{ComputeAlphaPrNextBit}(c, \alpha, \beta, \gamma, y, b-1, q \times pr, p)$
8. If  $p > 2^{-n}$ , predict  $Z \leftarrow 1$ ; else predict  $Z \leftarrow 0$ .

$\text{ComputeAlphaPrNextBit}(c, \alpha, \beta, \gamma, y, i, q, p)$  // update  $c_{i+1}, \alpha_{i+1}, p$  in-place

1. If  $i = \text{WordSize} - 1$ :  $p \leftarrow p + q \times \mathcal{DD}^{\text{SPECK}_{r-1R}}(\alpha \lll^7 \parallel \beta)$ ; return
2. If  $\text{eq}(\alpha_i, \beta_i, \gamma_i)$ :
  - (a)  $\alpha_{i+1} \leftarrow \beta_{i+1} \oplus \gamma_{i+1} \oplus \beta_i$ ;  $\text{ComputeCarryNextBit}(c, \alpha, \beta, \gamma, y, i)$ ;
  - (b)  $\text{ComputeAlphaPrNextBit}(c, \alpha, \beta, \gamma, y, i+1, q \cdot 1, p)$ ; return
3. Else if  $\text{Cyc}_{(i+1,i)}$  and  $c_i \neq \perp$ :
  - (a)  $\alpha_{i+1} \leftarrow \beta_{i+1} \oplus \gamma_{i+1} \oplus \beta_i \oplus y_i \oplus c_i$ ;  $\text{ComputeCarryNextBit}(c, \alpha, \beta, \gamma, y, i)$
  - (b)  $\text{ComputeAlphaPrNextBit}(\alpha, \beta, \gamma, y, i+1, q \cdot 1, p)$ ; return
4. Else:
  - (a)  $\beta_b \leftarrow \beta_{\{i+1, \dots, i+2-b\}}, \gamma_b \leftarrow \gamma_{\{i+1, \dots, i+2-b\}}, y_b \leftarrow y_{\{i+1, \dots, i+2-b\}},$   
 $\alpha_b \leftarrow \alpha_{\{i, \dots, i+2-b\}}$
  - (b) If  $c_{i+2-b} \neq \perp$ : For  $(\alpha_{i+1}, pr) \leftarrow \mathbf{A}_{\text{next}}^c[\beta_b, \gamma_b, y_b, \alpha_b, c_{i+2-b}]$ 
    - $\text{ComputeCarryNextBit}(c, \alpha, \beta, \gamma, y, i)$
    - $\text{ComputeAlphaPrNextBit}(\alpha, \beta, \gamma, y, i+1, q \cdot pr, p)$
  - (c) If  $c_{i+2-b} = \perp$ : For  $(\alpha_{i+1}, pr) \leftarrow \mathbf{A}_{\text{next}}[\beta_b, \gamma_b, y_b, \alpha_b]$ 
    - $\text{ComputeCarryNextBit}(c, \alpha, \beta, \gamma, y, i)$
    - $\text{ComputeAlphaPrNextBit}(\alpha, \beta, \gamma, y, i+1, q \cdot pr, p)$

$\text{ComputeCarryNextBit}(c, \alpha, \beta, \gamma, y, i)$  // update  $c_{i+1}$  in-place

1. If  $y_i = 0$  and  $c_i = 0$ :  $c_{i+1} \leftarrow 0$
2. Else if  $y_i = 1$  and  $c_i = 1$ :  $c_{i+1} \leftarrow 1$
3. Else if  $\text{Cxy0}_{(i+1,i)}$  or  $\text{Cxc1}_{(i+1,i)}$  or  $\text{Cyc0}_{(i+1,i)}$ :  $c_{i+1} \leftarrow y_i$
4. Else if  $\text{Cxy1}_{(i+1,i)}$  or  $\text{Cxc0}_{(i+1,i)}$ :  $c_{i+1} \leftarrow c_i$
5. Else:  $c_{i+1} \leftarrow \perp$ . //  $\perp$  means unknown

$\text{GenMultiBitsConditionalDDTs}(b)$

1.  $\mathbf{A}_0 \leftarrow$  Generate  $b$ -bit conditional DDT of  $\boxplus$ , each entry is indexed by  $(b$ -bit  $\beta, b$ -bit  $\gamma, b$ -bit  $y$ ), the values are  $(b$ -bit  $\alpha, \text{non-zero } pr)$ . //  $\mathbf{A}_0$  will be used for the first  $b$  bits since one knows that both LSB carry bits are 0.
  2.  $\mathbf{A}_{\text{next}} \leftarrow$  Generate  $b$ -bit conditional DDT of  $\boxplus$ , each entry is indexed by  $(b$ -bit  $\beta, b$ -bit  $\gamma, b$ -bit  $y, (b-1)$ -bit  $\alpha$ ), the values are  $(1$ -bit  $\alpha_{\text{next}}, \text{non-zero } pr)$ . //  $\mathbf{A}_{\text{next}}$  will be used for the intermediate bits when LSB carry  $c$  is unknown.
  3.  $\mathbf{A}_{\text{next}}^c \leftarrow$  Generate  $b$ -bit conditional DDT of  $\boxplus$ , each entry is indexed by  $(b$ -bit  $\beta, b$ -bit  $\gamma, b$ -bit  $y, (b-1)$ -bit  $\alpha, 1$ -bit carry  $c$ ), the values are  $(1$ -bit  $\alpha_{\text{next}}, \text{non-zero } pr)$ . //  $\mathbf{A}_{\text{next}}^c$  will be used for the intermediate bits when LSB carry  $c$  is known.
  4. Output  $\mathbf{A}_0, \mathbf{A}_{\text{next}}, \mathbf{A}_{\text{next}}^c$
-

**Table 9.** Performance of the improved DDT-based distinguishers ( $\mathcal{YD}$ s and  $\mathcal{AD}_{\mathcal{YD}}$ s) on SPECK32/64 and comparisons with pure DDT-based distinguishers ( $\mathcal{DD}$ s), neural distinguishers ( $\mathcal{ND}$ s), and DDT-based key-averaging distinguishers ( $\mathcal{AD}_{\mathcal{KD}}$ s)

#R	Name	ACC	TPR	TNR	Mem (GBytes)	Time (Secs per $2^{20}$ )
4	$\mathcal{DD}^{\text{SPECK}_{4R}}$	0.9869	0.9869	0.9870	32.5	$2^{-4.98}$
4	$\mathcal{YD}^{\text{SPECK}_{4R}}$	0.9907	0.9887	0.9928	32.5	$2^{-2.37}$
5	$\mathcal{DD}^{\text{SPECK}_{5R}}$	0.9107	0.8775	0.9440	32.5	$2^{-4.94}$
5	$\mathcal{YD}^{\text{SPECK}_{5R}}$	0.9215	0.8947	0.9484	32.5	$2^{-1.87}$
5	$\mathcal{ND}^{\text{SPECK}_{5R}}$	0.9273	0.9011	0.9536	0.0277	$2^{+3.56}$
5	$\mathcal{AD}_{\mathcal{YD}}^{\text{SPECK}_{5R}}$	0.9362	0.9173	0.9552	32.5	$2^{+5.46}$
5	$\mathcal{AD}_{\mathcal{KD}}^{\text{SPECK}_{5R}}$	0.9364	0.9171	0.9557	32.5	$2^{+7.03}$
6	$\mathcal{DD}^{\text{SPECK}_{6R}}$	0.7584	0.6795	0.8371	32.5	$2^{-4.53}$
6	$\mathcal{YD}^{\text{SPECK}_{6R}}$	0.7663	0.7118	0.8207	32.5	$2^{-2.05}$
6	$\mathcal{ND}^{\text{SPECK}_{6R}}$	0.7876	0.7197	0.8554	0.0277	$2^{+3.54}$
6	$\mathcal{AD}_{\mathcal{YD}}^{\text{SPECK}_{6R}}$	0.7949	0.7309	0.8587	32.5	$2^{+5.12}$
6	$\mathcal{AD}_{\mathcal{KD}}^{\text{SPECK}_{6R}}$	0.7946	0.7309	0.8583	32.5	$2^{+7.03}$
7	$\mathcal{DD}^{\text{SPECK}_{7R}}$	0.5913	0.5430	0.6397	32.5	$2^{-4.49}$
7	$\mathcal{YD}^{\text{SPECK}_{7R}}$	0.5962	0.5582	0.6343	32.5	$2^{-2.18}$
7	$\mathcal{ND}^{\text{SPECK}_{7R}}$	0.6155	0.5325	0.6985	0.0277	$2^{+3.57}$
7	$\mathcal{AD}_{\mathcal{YD}}^{\text{SPECK}_{7R}}$	0.6237	0.5428	0.7048	32.5	$2^{+5.33}$
7	$\mathcal{AD}_{\mathcal{KD}}^{\text{SPECK}_{7R}}$	0.6240	0.5435	0.7046	32.5	$2^{+7.04}$
8	$\mathcal{DD}^{\text{SPECK}_{8R}}$	0.5116	0.4963	0.5268	32.5	$2^{-4.64}$
8	$\mathcal{YD}^{\text{SPECK}_{8R}}$	0.5117	0.4967	0.5268	32.5	$2^{-2.99}$
8	$\mathcal{ND}^{\text{SPECK}_{8R}}$	0.5135	0.5184	0.5085	0.0277	$2^{+3.55}$
8	$\mathcal{AD}_{\mathcal{YD}}^{\text{SPECK}_{8R}}$	0.5187	0.4914	0.5460	32.5	$2^{+5.51}$
8	$\mathcal{AD}_{\mathcal{KD}}^{\text{SPECK}_{8R}}$	0.5194	0.4919	0.5469	32.5	$2^{+7.04}$

– ACC: Accuracy, TPR: True Positive Rate, TNR: True Negative Rate

– For  $\mathcal{YD}$ s, the thresholds  $\tau$ 's for  $\sigma_{\alpha}(i)$ 's in building  $\mathcal{YD}^{\text{SPECK}_{4R}}$ ,  $\mathcal{YD}^{\text{SPECK}_{5R}}$ ,  $\mathcal{YD}^{\text{SPECK}_{6R}}$ ,  $\mathcal{YD}^{\text{SPECK}_{7R}}$  are 0.50, 0.30, 0.20, and 0.02, respectively. The number of samples for the accuracy testing is  $2^{24}$ .

– The number of samples for benchmark is  $2^{20}$ . Thus, the times are seconds taken by making predictions on  $2^{20}$  samples.

– All DDT-based distinguishers ( $\mathcal{DD}$ s,  $\mathcal{YD}$ s,  $\mathcal{AD}_{\mathcal{YD}}$ s, and  $\mathcal{AD}_{\mathcal{KD}}$ s) are implemented in C++ (compiled using g++ 9.4.0 with optimization option '-O3'), whereas  $\mathcal{ND}$ -based distinguishers ( $\mathcal{ND}$ s) are implemented in Python.

– The benchmark environment is as follows: OS: Ubuntu 20.04; Processor: Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz; Memory: 256 GB DDR4 memory; all timings were restricted to run using a single CPU thread.

– We profiled the memory requirements of  $\mathcal{ND}$ s using the Tracemalloc module in Python. We specifically measured the peak allocated memory, excluding the memory allocated for storing the testing dataset. This was calculated by determining the memory usage when loading the  $\mathcal{ND}$  and making predictions, and then subtracting the memory usage when these operations were excluded (for example, 0.246898 GB – 0.219219 GB). For other distinguishers, we assessed memory requirements by referencing the 'RES' column associated with the process in the 'htop' command.

**Table 10.** The accuracy of differential-neural distinguishers using distinct differences obtained by (0040, 0000) after  $i$  rounds of propagation. Prob. represents the probability of the highest probability differential (0040, 0000)  $\rightarrow$  “Diff.”.

$i$	Diff.	Prob.	Acc.	$i$	Diff.	Prob.	Acc.
0	(0040, 0000)	1	0.6137	3	(8000, 840a)	$2^{-3}$	0.7394
1	(8000, 8000)	1	0.6137	4	(850a, 9520)	$2^{-7}$	0.9166
2	(8100, 8102)	$2^{-1}$	0.6705				

ciphertext values and intermediate state differences. Note that the neural distinguishers are not aware of the specific details of the ciphers, including their non-linear components and structure. Therefore, these neural distinguishers can be used for ciphers that have unknown components.

*On the Performance of Various Distinguishers.* Experiments showed that  $\mathcal{ND}$ s can be more efficient while achieving comparable accuracy to sophisticated manual methods (Alg. 2). Please refer to Table 9 for detailed benchmarks. Note that in benchmarks listed in Table 9, all DDT-based distinguishers are implemented in C++, whereas  $\mathcal{ND}$ -based distinguishers are implemented in Python Tensorflow. Although implementations in C++ might be inherently faster than its Python counterpart,  $\mathcal{ND}^{\text{SPECK}^*R_S}$  in Python are still more efficient than  $\mathcal{AD}_{\text{YD}}^{\text{SPECK}^*R}$  and  $\mathcal{AD}_{\text{KD}}^{\text{SPECK}^*R}$  in C++ (all restricted to run in a single CPU thread). Therefore, we can conclude that the neural network-based distinguishers provide a good trade-off between efficiency and accuracy.

## 4 Insights and Improvements on Training Differential-Neural Distinguisher

### 4.1 Relations Between Distinguisher Accuracy and Differential Distribution

Traditional differential cryptanalysis predominantly utilizes high-probability differentials as distinguishers. However, differential-neural cryptanalysis exploits all output differences for distinguishing while fixing input differences for plaintext pairs. In EUROCRYPT 2021, Benamira *et al.* [7] argued that differential-neural distinguisher is inherently building a very good approximation of the DDT during the learning phase.

Our study delves into the relation between the accuracy of the differential-neural distinguisher and the differential distribution of ciphertext pairs. We modify the input difference of plaintext pairs, inspired by Gohr’s staged training method [15]. In [15], while the basic training method can produce a valid 7-round distinguisher, an 8-round distinguisher must be trained using the staged training approach. The core of the staged training method is training a pre-trained 7-round distinguisher to learn 5-round SPECK32/64’s output pairs with

**Table 11.** The accuracy of the differential-neural distinguisher using distinct differences obtained by (0040, 0000) after 2 rounds of propagation. Prob. represents the probability of differential (0040,0000)  $\rightarrow$  “Diff.”. Round 2 +  $i$  represents the positive sample of the training set is the ciphertext pair obtained by encrypting the plaintext pair that satisfies this difference for  $i$  rounds

Diff.	Prob.	Acc. (2+4)	Acc. (2+5)	Acc. (2+6)	Diff.	Prob.	Acc. (2+4)	Acc. (2+5)	Acc. (2+6)
(8100,8102)	$2^{-1}$	0.8720	0.6811	0.5270	(8f00,8f02)	$2^{-4}$	0.6746	Fail	Fail
(8300,8302)	$2^{-2}$	0.8191	0.6218	Fail	(9f00,9f02)	$2^{-5}$	Fail	Fail	Fail
(8700,8702)	$2^{-3}$	0.7492	Fail	Fail	(bf00,bf02)	$2^{-6}$	Fail	Fail	Fail

the input difference (8000,804a) (the most likely difference to appear three rounds after the input difference (0040,0000)). Employing such plaintext pairs aims to concentrate the difference distribution of ciphertext pairs, escalating the output difference’s likelihood and simplifying the distinguisher’s learning task.

In our work, we first introduce a 4-round highest probability differential trail starting from (0040,0000).

$$(0040,0000) \rightarrow (8000,8000) \rightarrow (8100,8102) \rightarrow (8000,840a) \rightarrow (850a,9520)$$

Our experiments (see Table 10) initially employ a 4-round high-probability differential trail starting from (0040,0000), leading to (850a,9520).

By default, we use (0040,0000) as the input difference of the plaintext pair to generate the ciphertext pair. Here, in Table 10, we use the difference of the highest probability of (0040,0000) after  $i$  ( $1 \leq i \leq 4$ ) rounds of propagation as the input difference of the plaintext pair, respectively.

From Table 8, we can observe that the larger  $i$  is, the higher the accuracy of the differential-neural distinguisher. As  $i$  increases, the difference distribution in the ciphertext becomes more concentrated, and the probability of each difference increases. Therefore, the more significant the difference between the ciphertext and the random number, the accuracy of the differential-neural distinguisher is continuously improved.

To more comprehensively demonstrate the relation between the accuracy of the differential-neural distinguisher and the differential distribution of the ciphertext pairs, we conducted some experiments from another perspective. We fixed the number of rounds of differential but chose multiple 2-round differences with gradually decreasing probabilities. In Table 11, we notice that the higher the fixed probability of the differential, the higher the accuracy of the differential-neural distinguisher obtained. In other words, a lower probability means that after  $i$  rounds of encryption, the differential distribution of the ciphertext is more dispersed, and the neural network is more difficult to learn, resulting in a continuous decrease in the number of rounds and accuracy of the differential-neural distinguisher.

In conclusion, controlling differential propagation is imperative to enhance the differential-neural distinguisher’s accuracy and the number of rounds. We thus propose a method to control the differential propagation and reduce the

diffusion of features, thereby increasing the number of rounds of the differential-neural distinguisher. However, before the formal introduction, we introduce one method that can simplify the training process of high round distinguisher.

## 4.2 Freezing Layer Method

In existing experiments on SPECK32/64, especially with an input difference of (0040,0000), there has been a notable limitation. Researchers have been able to directly train a differential-neural distinguisher for up to only 7 rounds. Direct training for higher rounds from scratch has been challenging. A potential avenue that has garnered attention is the utilization of various network fine-tuning strategies. Specifically, continuing the training phase from pre-trained models has been proposed to potentially overcome these limitations and expand the distinguisher’s round capability. Examples include the staged training method in [15] and the staged pipeline method in [6].

The inability to directly train the 8-round distinguisher likely stems from feature diffusion associated with the input difference (0040,0000) over increasing rounds. This makes the 8-round features considerably challenging for the distinguisher to learn directly from limited data, as compared to lower rounds. One approach is to either mitigate feature diffusion or narrow the distinguisher’s solution space. While a technique to constrain feature diffusion is discussed in the subsequent chapter, in this context, we employ the classic network fine-tuning strategy, the freezing layer method, to limit the solution space.

Our distinguishers consist of two parts: the convolutional layers and fully connected layers. In the field of artificial intelligence, all convolutional layers are viewed as feature extractors, while all fully connected layers are viewed as a classifier. We argue that the feature extractor can be reused, and the classifiers are relatively similar in adjacent rounds. Therefore, to train an 8-round distinguisher for SPECK32/64, we can simply load a well-trained 7-round model and freeze all its convolutional layers, meaning that only parameters in fully connected layers can be updated. Then, we can obtain an 8-round distinguisher with accuracy identical to the ones in [6, 15], remaining all hyperparameters in the training process unchanged.

Relative to the staged training method [15], our approach maintains the same hyperparameters and does not require more samples in the final stage. In comparison with the method in [6], we only need two training rounds instead of multiple rounds in a row as required by the simple training pipeline in [6]. Besides, the simple training pipeline [6] did not produce  $\mathcal{ND}$ s with the same accuracy as Gohr’s on 8-round SPECK32/64; it needs a further polishing step to achieve similar accuracy, demanding more time and data. Our freezing layer method also speeds up the training process due to the reduction of trainable parameters. Therefore, we recommend trying the freezing layer method once the number of the distinguisher is too high to train directly.

## 5 Related-Key Differential-Neural Cryptanalysis

The  $\mathcal{ND}$  explainability concept serves as a fundamental theoretical underpinning when aiming to enhance and leverage its capabilities. With the outcome being that  $\mathcal{ND}$ s can effectively capture additional features and provide a better trade-off between efficiency and accuracy, there is substantial motivation for us to continue refining and exploiting their potential.

In this section, we introduce the related-key into differential-neural cryptanalysis, enabling control over differential propagation and facilitating the training of high-round  $\mathcal{ND}$ s. Furthermore, we enhance the DDT-based distinguisher under the  $\mathcal{RK}$  setting by employing the analytical methods and conclusions outlined in Sect. 3. As a result of these advancements, we successfully implement a 14-round key recovery attack for SPECK32/64 using the proposed  $\mathcal{RK}\text{-}\mathcal{ND}$ s.

### 5.1 Related-Key Differential-Neural Distinguisher for SPECK32/64

Here we present the related-key differential-neural distinguishers on SPECK32/64 obtained in this work.

*The Choice of the Input Difference.* The input difference is a crucial and central component of differential-neural cryptanalysis, and numerous papers delve into the study of the input difference, such as [3, 6, 15, 16, 21]. To maximize the number of rounds for both  $\mathcal{ND}$  and  $\mathcal{CD}$ , as well as the weak key space as large as possible to perform the longest key recovery attack, we use the SMT-based method to search for appropriate  $\mathcal{RK}$  differential or differential trails. It is important to note that the largest weak key space does not necessarily equate to the largest  $\mathcal{ND}$  or  $\mathcal{CD}$ , thus requiring a compromise between the three factors. In this paper, the choice of the best input difference is given under different compromises. Table 12 lists the  $\mathcal{RK}$  differential trails used to constrain the key space in SPECK32/64, where we label each distinguisher with an ID. Specifically, ID<sub>1</sub> is used to restrict the weak key space for the 13-round, ID<sub>2</sub> and ID<sub>3</sub> are used to restrict the 14-round. Note that part of the ID<sub>2</sub>/ID<sub>3</sub> (2-round to 11-round)  $\mathcal{RK}$  difference are same as the 10-round optimal  $\mathcal{RK}$  differential trail for speck32/64 given in Table 9 of [23]. In addition, the round-reduced of the trails are used to restrict the weak key space for shorter rounds, *e.g.*, ID<sub>2</sub> and ID<sub>3</sub> are used to restrict the weak key space for 13-round starting from the second round.

*Network Architecture.* Given the success of the neural network consisting of the Inception block and residual network in SPECK, SIMON and SIMECK [25, 26], as well as its superior performance in differential-neural distinguisher, we use this neural network proposed in [26] to train  $\mathcal{RK}$  differential-neural distinguisher. However, we also made some modifications to the network architecture. In deep learning, odd numbers such as 3, 5, and 7 are often used as the size of the convolution kernel. However, according to the cyclic shift of the round function of SPECK32/64, we choose 2 and 7 as the size of the convolution kernel. Furthermore, using 2 as the convolution kernel size can make the model's accuracy

**Table 12.** Related-key differential trails used to constrain the key space in SPECK32/64 where we label each distinguisher with an ID. For example, ID<sub>1</sub> represents the 13-round  $\mathcal{RK}$  differential trail for the key schedule algorithm with  $(\Delta l^2, \Delta l^1, \Delta l^0, \Delta k^0) = (0044, 0011, 4000, 0080)$

	ID <sub>1</sub>	ID <sub>1</sub>	ID <sub>2</sub> /ID <sub>3</sub>	ID <sub>2</sub> /ID <sub>3</sub>
$r$	Differential in Key	$\log_2 \text{Pr}$	Differential in Key	$\log_2 \text{Pr}$
0	(0044, 0011, 4000, 0080)		(0200, 0080, 0011, 4a00)	
1	(0000, 0044, 0011, 0200)	-1	(2800, 0200, 0080, 0001)	-4
2	(2000, 0000, 0044, 2800)	-2	(0000, 2800, 0200, 0004)	-1
3	(a000, 2000, 0000, 0000)	--2	(0000, 0000, 2800, 0010)	-1
4	(0000, a000, 2000, 0000)	-0	(0040, 0000, 0000, 0000)	-2
5	(0040, 0000, a000, 0040)	-1	(0000, 0040, 0000, 0000)	0
6	(0100, 0040, 0000, 0000)	-2	(0000, 0000, 0040, 0000)	0
7	(0000, 0100, 0040, 0000)	0	(8000, 0000, 0000, 8000)	0
8	(8000, 0000, 0100, 8000)	0	(8000, 8000, 0000, 8002)	0
9	(8002, 8000, 0000, 8000)	-1	(8002, 8000, 8000, 8008)	-1
10	(8000, 8002, 8000, 8002)	0	(8108, 8002, 8000, 812a)	-2
11	(8102, 8000, 8002, 8108)	-2	(802a, 8108, 8002, 8480)	-4
12	(8408, 8102, 8000, 802a)	-3	(8180, 802a, 8108, 9382)/(8280, 802a, 8108, 9082)	-3/-4
13			(8180, 8180, 802a, cf8a)/(8080, 8280, 802a, c28a)	-4/-4
	$\log_2 (P_r(Q_K))$ : -14		$\log_2 (P_r(Q_K))$ : -22/-23	

converge faster than 3. In [26], the size of the convolution kernel continues to increase as the depth of the residual network increases. We think it is reasonable to increase the convolution kernel’s size to improve the network’s receptive field, but it cannot always be increased. Therefore, we will limit the size of the convolution kernel to less than or equal to 7.

*The Training of Related-Key Differential-Neural Distinguisher.* This work still uses the basic training method to train short-round distinguishers. When the basic training method fails, we train the  $r$ -round distinguisher with the  $(r - 1)$ -round distinguisher by using the freezing layer method. Please refer to Appendix F in [4] for the detailed training method.


*Performance Evaluation of the Distinguisher.* In artificial intelligence, the model’s accuracy is the most critical evaluation indicator. In differential-neural cryptanalysis, it is judged whether the guessed key is correct based on the score of the distinguisher. Therefore, we evaluate the performance of the differential-neural distinguisher regarding both the accuracy and the score.

- *Test accuracy.* We summarize the accuracy of the differential-neural distinguisher in Table 13. The 8, and 9-round distinguishers were trained using the basic training method, while the 10-round distinguishers were trained using the freezing layer method. For more insight on related-key differential-neural distinguishers, please refer to Appendix F.2 in [4].



**Table 13.** The summary of related-key differential-neural distinguishers on SPECK32/64, where the plaintext difference is (0000,0000).

Diff.	#R	Name	Accuracy	True Positive Rate	True Negative Rate
ID <sub>1</sub>	8	$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{8R}}$	0.7584	0.6836	0.8332
ID <sub>1</sub>	9	$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}$	0.5620	0.5212	0.6028
ID <sub>2</sub> /ID <sub>3</sub>	8	$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{8R}}$	0.9259	0.9063	0.9455
ID <sub>2</sub>	9	$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}$	0.7535	0.7035	0.8036
ID <sub>2</sub>	10	$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{10R}}$	0.5643	0.5382	0.5893
ID <sub>3</sub>	9	$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}$	0.7726	0.7247	0.8206
ID <sub>3</sub>	10	$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{10R}}$	0.5562	0.5361	0.5765

**Fig. 3.** Wrong key response profile of ID<sub>1</sub>

- *Wrong key response profile (WKRP)*. In [15], the key search policy depends on the observation that a distinguisher’s response to wrong-key decryption varies with the bitwise difference between the guessed and real key. Instead of exhaustive trial decryption, it suggests specific subkeys and scores them. Figure 3 shows the mean response for varying Hamming distances between guessed and actual keys in ID<sub>1</sub>. Notably, high scores emerge when differences in keys are small, especially if the difference relates to {16384, 32768, 49152}. This indicates that errors in the 14th and 15th bits of the subkey minimally impact scores, allowing for a reduced key guessing space. This accelerated key recovery in [15]. For WKRPs of ID<sub>2</sub> and ID<sub>3</sub>, see Appendix B.2 in [4].

**On  $\mathcal{RK}\text{-}\mathcal{ND}$ ’s Explainability.** Beyond constructing and comparing various  $\mathcal{RK}$  distinguishers (see Appendix F.2 in [4]), we further undertook experiments analogous to Gohr’s *aaaa*-blinding experiment. Some  $\mathcal{RK}\text{-}\mathcal{ND}$ s behaved similarly to single-key setting  $\mathcal{ND}$ s, while others varied. Refer to  $\mathcal{RK}\text{-}\mathcal{ND}_{\text{ID}(2,9182)}^{\text{SPECK}_{9R}}$  and  $\mathcal{RK}\text{-}\mathcal{ND}_{\text{ID}(2,9382)}^{\text{SPECK}_{9R}}$  in Table 14 for example for the former and latter case, where the differential trail ID<sub>(2,9182)</sub> differs from ID<sub>(2,9382)</sub> only at the last round key, and ID<sub>(2,9382)</sub> is ID<sub>2</sub> from round 4 to 12. Notably, the behavior of  $\mathcal{RK}\text{-}\mathcal{ND}_{\text{ID}(3,9082)}^{\text{SPECK}_{9R}}$  presented intriguing phenomena (ID<sub>(3,9082)</sub> is ID<sub>3</sub> from round 4 to 12):

**Table 14.** Experiments detailing the information harnessed by  $\mathcal{RK}\text{-}\mathcal{ND}$ s using 9 round  $\text{ID}_{(2,9182)}$ ,  $\text{ID}_{(2,9382)}$ , and  $\text{ID}_{(3,9082)}$ , with similar settings in Table 2.

ID	Set	Positive Samples	Negative Samples	Acc.
	1-1	$(A, B, C, D)$	<i>Random</i>	0.7531
$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(2,9182)}}$	1-2	$(\mathcal{AR}_1, \mathcal{BR}_1, \mathcal{CR}_1, \mathcal{DR}_1)$	<i>Random</i>	0.7534
	1-1	$(A, B, C, D)$	<i>Random</i>	0.7574
$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(2,9382)}}$	1-2	$(\mathcal{AR}_1, \mathcal{BR}_1, \mathcal{CR}_1, \mathcal{DR}_1)$	<i>Random</i>	0.7529
	1-1	$(A, B, C, D)$	<i>Random</i>	0.7746
$\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(3,9082)}}$	1-2	$(\mathcal{AR}_1, \mathcal{BR}_1, \mathcal{CR}_1, \mathcal{DR}_1)$	<i>Random</i>	0.7539

1.  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(3,9082)}}$  performed differently on  $\text{Set-1-1} := \{\Gamma_A, \Gamma_B, \Gamma_C, \Gamma_D\}$  and  $\text{Set-1-2} := \{\Gamma_{\mathcal{AR}_1}, \Gamma_{\mathcal{BR}_1}, \Gamma_{\mathcal{CR}_1}, \Gamma_{\mathcal{DR}_1}\}$ , which, under the assumption of a random last-round key  $K$ , defines the same information set per Sect. 3.1 (please refer to Table 14).
2.  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(3,9082)}}$  showed superior performance over  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(2,9382)}}$  (0.7726 vs. 0.7535, refer to Table 22 in [4]), while theoretically, if there is no information on the key being revealed beyond the key difference,  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(3,9082)}}$  should perform exactly the same as  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(2,9382)}}$ , since the two differential trails differ only at the last round key difference thus the two output difference distributions are affine-equivalent.
3. Surprisingly,  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(3,9082)}}$  even outperformed our manually enhanced distinguisher  $\mathcal{RK}\text{-}\mathcal{AD}^{\text{SPECK}_{9R}}_{\text{YD}}$  (0.7726 vs. 0.7574, refer to Table 22 in [4]).

Upon closer examination of the differential trail of  $\text{ID}_{(3,9082)}$ , we identified the causative factor. Let’s denote input/output differences and values around the last  $\boxplus$  in the key schedule producing the 8-round (counting start from 0) key  $k^8$  as  $\alpha, \beta, \gamma, x, y, z$ . Then from the differential trail  $\text{ID}_{(3,9082)}$ , specifically focus on

$$\text{the 7- and 8-round, we have } \begin{cases} \alpha = 0x8002 \lll 7 & = 0b\ 0000\ 0101\ 0000\ 0000, \\ \beta = 0x8480 & = 0b\ 1000\ 0100\ 1000\ 0000, \\ \gamma = 0x8280 & = 0b\ 1000\ 0010\ 1000\ 0000, \end{cases}$$

According to Tables 4 and 5, we have follows.

1. The (8, 7)-th bit position is in case  $\text{Cxc}1_{(8,7)}$ , we have  $c_8 = y_7$ .
2. The (9, 8)-th bit position is in case  $\text{Cxc}0_{(9,8)}$ , we have  $c_9 = c_8$ .
3. The (10, 9)-th bit position is in case  $\text{Cxy}1_{(10,9)}$ , we have  $x_9 \oplus y_9 = z_9 \oplus c_9 = 1$ .

Consequently, we have  $z_9 \oplus y_7 = 1$ . Note that  $z_9 \oplus y_7 = 1$  implies that the 9th bit of the last round key is constantly 1. This does not obscure 1-bit information of the output of the last  $\boxplus$  in the encryption path, allowing for better accuracy of the resulting distinguisher. This explains all the odds on  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(3,9082)}}$ .

Additionally, for  $\mathcal{RK}\text{-}\mathcal{ND}^{\text{SPECK}_{9R}}_{\text{ID}_{(2,9382)}}$ , the 10th bit of the last-round key conforming to the round difference has a bias towards 0 (equals 0 with a probability of

3/4), which could explain its slightly differed accuracy between **Set.1–1** and **Set.1–2** (refer to Table 14). After fixing the 10th bit to be 0 and re-training the distinguisher, it achieves almost the same accuracy as  $\mathcal{RK}\text{-}\mathcal{ND}_{\text{ID}(3,9082)}^{\text{SPECK9R}}$ . When analyzing the probability of related-key pairs under these conditions, we deduced that restricting the 10th bit for  $\text{ID}_{(2,9382)}$  still results in a larger weak-key space compared with  $\text{ID}_{(3,9082)}$  while achieving the same high  $\mathcal{RK}\text{-}\mathcal{ND}$  accuracy.

## 5.2 Key Recovery Attack on Round-Reduced SPECK32/64

This subsection describes the implementation of  $\mathcal{RK}$  differential-neural crypt-analysis using the trained distinguisher. The key recovery framework is similar to [3, 15, 26]. Since the whole attack is in the  $\mathcal{RK}$  setting, we need to specify the difference between each round of subkeys. Specifically, it is unclear how to perform a key recovery attack if only applying a difference to the master key without specifying the difference in the round-key state. In such cases, the guessed one last-round key cannot directly infer the other last-round key in the related pair, as the difference in the last-round key is not specified.

We first introduce some preparatory work before officially implementing the key recovery attack.

*Generalized Neutral Bits.* We incorporate  $\mathcal{CD}$  before  $\mathcal{ND}$  to increase the number of rounds for the key recovery attack. Furthermore, to enhance predictive performance, we employ the distinguisher to estimate the scores of multiple ciphertexts with the same distribution (ciphertext structure) and combine them to obtain the scores for the guessed subkey. However, the  $\mathcal{CD}$  is probabilistic, and the randomly generated plaintext structure does not retain the same distribution after encryption. Hence, we require neutral bits to generate the plaintext structure, which we encrypt to obtain the ciphertext structure, achieving a successful key recovery attack. Therefore, the  $\mathcal{CD}$  should have a high probability and a sufficient number of neutral bits. Appendix B.3 in [4] lists the NBs/SNBs we used to perform the key recovery attack.

*The Parameters for Key Recovery Attack.* The attacks follow the framework of the improved key recovery attacks in [15]. An  $r$ -round main and an  $(r-1)$ -round helper  $\mathcal{ND}$ s are employed, and an  $s$ -round  $\mathcal{CD}$  is prepended. The key guessing procedure applies a simple reinforcement learning procedure. The last subkey and the second to last subkey are to be recovered without exhaustively using all candidate values to perform one-round decryption. Moreover, a Bayesian key search employing the *wrong key response profile* will be used. We count a key guess as successful if the last round key was guessed correctly and if the second round key is at the hamming distance at most two of the real keys. The parameters to recover the last two subkeys are indicated below.

Parameter	Definition
$n_{cts}$	The number of ciphertext structures
$n_b$	The number of ciphertext pairs in each ciphertext structure, that is, $2^{ \text{NB} }$
$n_{it}$	The total number of iterations in the ciphertext structures
$c_1, c_2$	The cutoffs with respect to the scores of the recommended last subkey and second to last subkey, respectively
$n_{byit1/2}$	The number of iterations, the default value is 5
$n_{cand1/2}$	The number of key candidates within each iteration, default value is 32

*Complexity Evaluation of Key Recovery Attack.* The experiment is conducted by Python 3.7.15 and Tensorflow 2.5.0 in Ubuntu 20.04. The device information is Intel Xeon E5-2680V4\*2 with 2.40 GHz, 256GB RAM, and NVIDIA RTX3080Ti 12 GB\*7. To reduce the experimental error, we perform 210 key recovery attacks for each parameter setting, take the average running time  $rt$  as the running time of an experiment, and divide the number of successful experiments by the total experimental number as the success rate  $sr$  of the key recovery attack.

1. *Data complexity.* The data complexity of the experiment is calculated using the formula  $n_b \times n_{ct} \times 2$ , which is a theoretical value. In the actual experiment, when the accuracy of the differential-neural distinguisher is high, the key can be recovered quickly and successfully. Not all data are used, so the actual data complexity is lower than theoretical.
2. *Time complexity.* We use  $2^{32}$  data to test the speed of encryption and decryption on our device, and each core can perform  $2^{26.814}$  rounds of decryption operations per second for SPECK32/64. The formula for calculating the time complexity in our experiments:  $2^{26.814} \times rt$ .

*The Result of Key Recovery Attacks.* We list the results of key recovery attacks in multiple differential modes in Table 15. We calculate the corresponding weak key space  $wks$  according to the probabilities of  $ID_1$ ,  $ID_2$ , and  $ID_{(3,9082)}$ . Adv. represents the advantage compared to the time complexity of brute forcing. The time and data complexity can be reduced by reducing  $n_{cts}$  and  $n_{it}$ , but the success rate  $sr$  also decreases accordingly. The first metric for our experiment is to reduce the time complexity.

*Remark 3 (The profiling information of the key-recovery attack).* To pinpoint the attack's bottleneck, we profiled a 14-round key-recovery attack using  $ID_3$ . The main result is detailed in Table 16. From the profiling result, the performance of our implementation of the attack is mostly limited by the speed of neural network evaluation (the proportion taken by  $\mathcal{ND}$  making the prediction is  $79.18\% + 5.17\% = 84.35\%$ ). The next limiting factor is the speed of computing the weighted Euclidean distance with the wrong key response profile.

**Table 15.** Summary of key recovery attacks on SPECK32/64

Diff.	#R	Configure	$wks$	$n_{cts}$	$n_{it}$	$n_b$	$c_1$	$c_2$	$sr$	Time	Data	Advantage
ID <sub>1</sub>	13	1+2+9+1	$2^{50}$	$2^7$	$2^8$	$2^8$	8	5	54.28%	$2^{34.57}$	$2^{16}$	$2^{15.43}$
ID <sub>2</sub>	13	1+2+9+1	$2^{46}$	$2^6$	$2^7$	$2^5$	5	5	93.33%	$2^{33.95}$	$2^{12}$	$2^{12.05}$
ID <sub>2</sub>	13	1+2+9+1	$2^{46}$	$2^5$	$2^6$	$2^5$	5	5	72.86%	$2^{33.01}$	$2^{11}$	$2^{12.99}$
ID <sub>2</sub>	13	1+2+9+1	$2^{46}$	$2^4$	$2^5$	$2^5$	10	1	44.28%	$2^{31.79}$	$2^{10}$	$2^{14.21}$
ID <sub>2</sub>	14	1+3+9+1	$2^{42}$	$2^9$	$2^{10}$	$2^6$	8	10	75.71%	$2^{35.59}$	$2^{16}$	$2^{6.41}$
ID <sub>2</sub>	14	1+3+9+1	$2^{42}$	$2^9$	$2^{10}$	$2^5$	8	5	55.71%	$2^{35.32}$	$2^{15}$	$2^{6.68}$
ID <sub>3</sub>	13	1+2+9+1	$2^{45}$	$2^6$	$2^7$	$2^5$	5	5	95.24%	$2^{34.26}$	$2^{12}$	$2^{10.75}$
ID <sub>3</sub>	13	1+2+9+1	$2^{45}$	$2^5$	$2^6$	$2^5$	5	5	77.62%	$2^{33.55}$	$2^{11}$	$2^{11.45}$
ID <sub>3</sub>	13	1+2+9+1	$2^{45}$	$2^4$	$2^5$	$2^5$	10	1	46.67%	$2^{32.20}$	$2^{10}$	$2^{12.80}$
ID <sub>3</sub>	14	1+3+9+1	$2^{41}$	$2^{10}$	$2^{11}$	$2^7$	10	25	90%	$2^{36.39}$	$2^{18}$	$2^{4.61}$
ID <sub>3</sub>	14	1+3+9+1	$2^{41}$	$2^9$	$2^{10}$	$2^7$	10	15	71.43%	$2^{35.78}$	$2^{17}$	$2^{5.22}$
ID <sub>3</sub>	14	1+3+9+1	$2^{41}$	$2^9$	$2^{10}$	$2^5$	5	5	68.57%	$2^{35.40}$	$2^{15}$	$2^{5.6}$

**Table 16.** Profiling information of the key-recovery attack

Function	Time (Percentage)
<b>test_bayes</b>	242 s (100 %)
- bayesian_key_recovery	- 229.39 s (94.79 %)
-   - (GPU) net.predict	-   - 191.62 s (79.18 %)
-   - (CPU) bayesian_rank_kr	-   - 28.99 s (11.98 %)
- verifier_search	- 12.63 s (5.22 %)
-   - (GPU) net.predict	-   - 12.51 s (5.17 %)

- test\_bayes: a full run of the attack excluding the generation of the related-key, load models, and generation of ciphertext structures.
- bayesian\_key\_recovery: the run of the BAYESIANKEYSEARCH algorithm.
- verifier\_search: the run of the final improvement [3].
- net.predict: using  $\mathcal{N}\mathcal{D}$ s to score the ciphertext structures decrypted by one round.
- bayesian\_rank\_kr: computing the weighted Euclidean distance with WKRP.

*Remark 4 (Efficiency measures in symmetric-key cryptanalysis attacks).* Assessing the efficiency of distinguishers and key recovery attacks in symmetric-key cryptanalysis poses intricate challenges, particularly when pinpointing computational complexities based on real-time attack timings and then extrapolating these to equivalent primitive evaluations, as done in both  $\mathcal{N}\mathcal{D}$ -based and traditional attacks in [12, 14, 24] (listed in Table 1).

Factors influencing these complexities include architecture compatibility and algorithmic suitability, varied computation intensity and various operation costs across platforms, memory constraints and flexible trade-offs, and implementation factors. Given these complexities, it is a good idea to have secondary metrics for comparison, for instance, power consumption and cost efficiency (please refer to Appendix E in [4] for detailed discussions). While there’s a pressing need for uni-

versal metrics, formulating such benchmarks is challenging, warranting caution when interpreting the comparison results and warranting further exploration.

## 6 Conclusion

This paper provides explicit rules that a distinguisher can use beyond the full differential distribution table to achieve better distinguishing performance. These rules are based on high correlations between values of bits in right pairs of differential propagation through addition modular  $2^n$ . By leveraging the value-dependent differential probability, which is not typically applied in traditional differential distinguishers, we can equip additional knowledge to DDT-based distinguishers, enhancing their accuracy. These rules or their equivalent form are likely the additional features beyond full DDT that the neural distinguishers exploit. While these rules are not difficult to derive with careful analysis, they rely on non-trivial relations that traditional distinguishers often overlook. This indicates that neural networks help break the limitations of traditional cryptanalysis. Studying this unorthodox model can provide new opportunities to understand cryptographic primitives better.

Another investigation in this paper revealed that controlling differential propagation is crucial to enhance the accuracy of differential-neural distinguisher. It is typically believed that introducing differences into the keys provides chances to cancel differences in the encryption states, thus resulting in stronger differential propagations. However, unlike traditional differential attacks, differential-neural attacks do not specify the output difference and, thus, are not limited to a single differential trail. Therefore, it is unclear whether the difference in a key is helpful in differential-neural attacks. It is also unclear how resistant SPECK is against differential-neural attacks in the  $\mathcal{RK}$  setting. This work confirmed that differential-neural cryptanalysis in the  $\mathcal{RK}$  setting could be more powerful than in the single-key setting by conducting a 14-round key recovery attack on SPECK32/64.

**Acknowledgments.** The authors would like to thank anonymous reviewers for their insightful and helpful comments. Special thanks go to the anonymous shepherd, who provided invaluable constructive feedback and patient guidance, which helped us improve the manuscript significantly. This research is partially supported by Nanyang Technological University in Singapore under Start-up Grant 04INS000397C230, and Ministry of Education in Singapore under Grants RG91/20. Zhenzhen Bao was partially supported by the National Key R&D Program of China (Grant No. 2018YFA0704701), the Major Program of Guangdong Basic and Applied Research (Grant No. 2019B030302008), the High Performance Computing Center of Tsinghua University. Jinyu Lu was supported by the National Natural Science Foundation of China (Grant No. 62172427). Liu Zhang was supported by the National Natural Science Foundation of China (Grant No. 62172319), and the Fundamental Research Funds for the Central Universities (Grant No. QTZX23090).

## References

1. Source codes in this work (2023). <https://www.dropbox.com/sh/yleufeu0wqwcjv/AADUpM15q86Uk1lM8z99fU2ia?dl=0>
2. Bao, Z., Guo, J., Liu, M., Ma, L., Tu, Y.: Enhancing differential-neural cryptanalysis. Cryptology ePrint Archive, Report 2021/719 (2021). <https://eprint.iacr.org/2021/719>
3. Bao, Z., Guo, J., Liu, M., Ma, L., Tu, Y.: Enhancing differential-neural cryptanalysis. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part I. LNCS, vol. 13791, pp. 318–347. Springer, Heidelberg (2022). [https://doi.org/10.1007/978-3-031-22963-3\\_11](https://doi.org/10.1007/978-3-031-22963-3_11)
4. Bao, Z., Lu, J., Yao, Y., Zhang, L.: More insight on deep learning-aided cryptanalysis. Cryptology ePrint Archive, Paper 2023/1391 (2023). <https://eprint.iacr.org/2023/1391>, full version of this paper
5. Beaulieu, R., Shors, D., Smith, J., Treatman-Clark, S., Weeks, B., Wingers, L.: The SIMON and SPECK families of lightweight block ciphers. Cryptology ePrint Archive, Report 2013/404 (2013). <https://eprint.iacr.org/2013/404>
6. Bellini, E., Gerault, D., Hambitzer, A., Rossi, M.: A cipher-agnostic neural training pipeline with automated finding of good input differences. Cryptology ePrint Archive, Report 2022/1467 (2022). <https://eprint.iacr.org/2022/1467>
7. Benamira, A., Gerault, D., Peyrin, T., Tan, Q.Q.: A deeper look at machine learning-based cryptanalysis. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 805–835. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-77870-5\\_28](https://doi.org/10.1007/978-3-030-77870-5_28)
8. Beyne, T., Rijmen, V.: Differential cryptanalysis in the fixed-key model. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part III. LNCS, vol. 13509, pp. 687–716. Springer, Heidelberg (2022). [https://doi.org/10.1007/978-3-031-15982-4\\_23](https://doi.org/10.1007/978-3-031-15982-4_23)
9. Biryukov, A., dos Santos, L.C., Teh, J.S., Udovenko, A., Velichkov, V.: Meet-in-the-filter and dynamic counting with applications to speck. IACR Cryptology ePrint Archive p. 673 (2022)
10. Chen, Y., Bao, Z., Shen, Y., Yu, H.: A deep learning aided key recovery framework for large-state block ciphers. Cryptology ePrint Archive, Report 2022/1659 (2022). <https://eprint.iacr.org/2022/1659>
11. Chen, Y., Yu, H.: Bridging machine learning and cryptanalysis via EDLCT. Cryptology ePrint Archive, Report 2021/705 (2021). <https://eprint.iacr.org/2021/705>
12. Dinur, I.: Improved differential cryptanalysis of round-reduced speck. In: Joux, A., Youssef, A. (eds.) SAC 2014. LNCS, vol. 8781, pp. 147–164. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-13051-4\\_9](https://doi.org/10.1007/978-3-319-13051-4_9)
13. Dinur, I., Dunkelman, O., Shamir, A.: Collision attacks on up to 5 rounds of SHA-3 using generalized internal differentials. In: Moriai, S. (ed.) FSE 2013. LNCS, vol. 8424, pp. 219–240. Springer, Heidelberg (2014). [https://doi.org/10.1007/978-3-662-43933-3\\_12](https://doi.org/10.1007/978-3-662-43933-3_12)
14. Feng, Z., Luo, Y., Wang, C., Yang, Q., Liu, Z., Song, L.: Improved differential cryptanalysis on speck using plaintext structures. In: Simpson, L., Rezazadeh Bae, M.A. (eds.) ACISP 2023. LNCS, vol. 13915, pp. 3–24. Springer, Cham (2023). [https://doi.org/10.1007/978-3-031-35486-1\\_1](https://doi.org/10.1007/978-3-031-35486-1_1)
15. Gohr, A.: Improving attacks on round-reduced Speck32/64 using deep learning. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 150–179. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-26951-7\\_6](https://doi.org/10.1007/978-3-030-26951-7_6)

16. Gohr, A., Leander, G., Neumann, P.: An assessment of differential-neural distinguishers. *Cryptology ePrint Archive, Report 2022/1521* (2022). <https://eprint.iacr.org/2022/1521>
17. Lee, H., Kim, S., Kang, H., Hong, D., Sung, J., Hong, S.: Calculating the approximate probability of differentials for ARX-based cipher using sat solver. *J. Korea Inst. Inf. Secur. Cryptol.* **28**(1), 15–24 (2018)
18. Leurent, G.: Analysis of differential attacks in ARX constructions. In: Wang, X., Sako, K. (eds.) *ASIACRYPT 2012*. LNCS, vol. 7658, pp. 226–243. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-34961-4\\_15](https://doi.org/10.1007/978-3-642-34961-4_15)
19. Leurent, G.: Construction of differential characteristics in ARX designs application to skein. In: Canetti, R., Garay, J.A. (eds.) *CRYPTO 2013, Part I*. LNCS, vol. 8042, pp. 241–258. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-40041-4\\_14](https://doi.org/10.1007/978-3-642-40041-4_14)
20. Lipmaa, H., Moriai, S.: Efficient algorithms for computing differential properties of addition. In: Matsui, M. (ed.) *FSE 2001*. LNCS, vol. 2355, pp. 336–350. Springer, Heidelberg (2002). [https://doi.org/10.1007/3-540-45473-X\\_28](https://doi.org/10.1007/3-540-45473-X_28)
21. Lu, J., Liu, G., Liu, Y., Sun, B., Li, C., Liu, L.: Improved neural distinguishers with (related-key) differentials: applications in SIMON and SIMECK. *Cryptology ePrint Archive, Report 2022/030* (2022). <https://eprint.iacr.org/2022/030>
22. Peyrin, T.: Improved differential attacks for ECHO and Grøstl. In: Rabin, T. (ed.) *CRYPTO 2010*. LNCS, vol. 6223, pp. 370–392. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-14623-7\\_20](https://doi.org/10.1007/978-3-642-14623-7_20)
23. Sadeghi, S., Rijmen, V., Bagheri, N.: Proposing an MILP-based method for the experimental verification of difference-based trails: application to speck, SIMECK. *Des. Codes Crypt.* **89**, 2113–2155 (2021)
24. Song, L., Huang, Z., Yang, Q.: Automatic differential analysis of ARX block ciphers with application to SPECK and LEA. In: Liu, J.K., Steinfeld, R. (eds.) *ACISP 2016*. LNCS, vol. 9723, pp. 379–394. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-40367-0\\_24](https://doi.org/10.1007/978-3-319-40367-0_24)
25. Zhang, L., Lu, J., Wang, Z., Li, C.: Improved differential-neural cryptanalysis for round-reduced simeck32/64. *arXiv preprint arXiv:2301.11601* (2023)
26. Zhang, L., Wang, Z., Wang, B.: Improving differential-neural cryptanalysis with inception blocks. *Cryptology ePrint Archive, Report 2022/183* (2022). <https://eprint.iacr.org/2022/183>
27. Zhang, Z., Hou, C., Liu, M.: Collision attacks on round-reduced SHA-3 using conditional internal differentials. In: Hazay, C., Stam, M. (eds.) *EUROCRYPT 2023, Part IV*. LNCS, vol. 14007, pp. 220–251. Springer, Heidelberg (2023). [https://doi.org/10.1007/978-3-031-30634-1\\_8](https://doi.org/10.1007/978-3-031-30634-1_8)