Key-Independent Secret-Key Distinguisher for 7-Round AES based on the Joint Generalized Zero-Difference Property

Hanbeom Shin[1(✉)], Sunyeop Kim[1,2], Byoungjin Seok[3], Deukjo Hong[4], Jaechul Sung[5], Seokhie Hong[6], Sangjin Lee[1], and Dongjae Lee[7(✉)]

[1] Korea University, Seoul, Republic of Korea
newonetiger@korea.ac.kr
[2] Nanyang Technological University, Singapore
[3] Hansung University, Seoul, Republic of Korea
[4] Jeonbuk National University, Jeonju, Republic of Korea
[5] University of Seoul, Seoul, Republic of Korea
[6] SmartM2M, Seoul, Republic of Korea
[7] Kangwon National University, Chuncheon, Republic of Korea
dongjae.lee@kangwon.ac.kr

Abstract. A key-independent secret-key distinguisher identifies structural deviations from an ideal random permutation without discovering any information about the secret key. It is therefore of primary importance for understanding the inherent properties of a block cipher's round function. While numerous key-independent secret-key distinguishers have been proposed for 5- and 6-round AES, none has been proposed for 7-round AES to date. In this paper, we propose the first key-independent secret-key distinguisher for 7-round AES, which exploits solely the structural properties of the round function. We propose the *Joint Generalized Zero-Difference Property*, where a quartet constructed from related differences satisfies three distinct generalized zero-difference properties simultaneously. By leveraging this joint property, we construct a new 7-round differential characteristic that a right quartet follows with a probability of $2^{-250.4}$, whereas a random permutation satisfies the same conditions with a probability of $2^{-253.4}$. Based on this characteristic, we design a distinguishing attack requiring data, time, and memory complexities of $2^{126.2}$. Our analysis confirms that the proposed distinguisher achieves a success probability of at least 75.7%. We experimentally verify the joint property using small-scale AES, confirming that the theoretical predictions match the observed results. This work achieves the longest-round key-independent secret-key distinguisher for AES reported to date.

Keywords: AES · Secret-Key Distinguisher · Related Differences · Generalized Zero-Difference · Truncated Differential

1 Introduction

Evaluating the security of a block cipher requires analyzing both its resistance to key recovery and its indistinguishability from an ideal random permutation.

While a successful key recovery trivially implies distinguishability, the converse does not hold: a cipher may exhibit detectable structural deviations from a random permutation without revealing any information about the secret key. Specifically, distinguishing attacks are critical because they reveal intrinsic structural properties that should not exist in a secure cipher. Identifying such non-random behaviors proves that the cipher cannot be considered a secure instantiation of a random permutation. Furthermore, analyzing reduced-round variants is essential for quantifying the security margin of the full cipher. By determining the maximum number of rounds that exhibit distinguishable properties, we can establish a concrete boundary of security and validate the robustness of the full specification.

Among various attack models, *secret-key distinguishers* are of particular importance in evaluating the structural security of block ciphers. A secret-key distinguisher operates under the standard black-box model: the adversary is given oracle access to either the cipher (with an unknown random key) or a truly random permutation, and the goal is to determine which one is the cipher [16]. Unlike key recovery attacks, which aim to find the secret key, a secret-key distinguisher reveals structural deviations from an ideal random permutation. Grassi et al. [16] further classified secret-key distinguishers into two categories: (1) *key-independent* distinguishers, which verify a structural property that holds for every key without discovering any key information, and (2) *key-dependent* distinguishers, which are derived from key-recovery attacks and rely on the existence of a (partial) key satisfying a certain condition. A key-independent distinguisher is particularly significant because it exposes inherent weaknesses of the round function itself, independent of the key schedule, and remains valid even in the independent subkey model where key-recovery attacks may be invalidated.

In 2017, Grassi et al. proposed the multiple-of-8 property [16], establishing the first key-independent secret-key distinguisher for 5-round AES. This property has also been applied to key recovery attacks [25]. In the same year, Rønjom et al. proposed the Yoyo game [22], constructing key-independent distinguishers for both 5- and 6-round AES. However, Mondal et al. [20] later showed that the originally proposed complexities yield a success probability of approximately 0.5, making the distinguishers ineffective. They revised the 5-round distinguisher with corrected complexities, while the 6-round Yoyo distinguisher was invalidated as achieving a meaningful success probability requires complexity exceeding 2^{128}. Subsequently, Grassi et al. proposed mixture differential cryptanalysis [15], which generalizes the multiple-of-8 property. Boura et al. [6] further established a general proof framework that systematically explains these structural properties for a wide class of SPN ciphers, and this framework was later revisited for SKINNY [23]. In 2019, Bardeh et al. extended the distinguishing range to 6 rounds by proposing the exchange attack [4]. More recently, Bariant et al. proposed the truncated boomerang attack [5], and Yan et al. proposed the boomerang chain distinguisher [26], which currently holds the best complexity record for 6-round AES. For 5-round AES, Shin et al. proposed the multiple exchanged boomerang distinguisher [24], achieving the best-known complexity.

To date, most cryptanalytic results on 7-round AES have been key recovery attacks utilizing the properties of the key schedule, such as impossible differential attacks [18,7,17] and meet-in-the-middle attacks [13,11]. These attacks inherently depend on the AES key schedule; if the subkeys were chosen independently, most of them would be invalidated.

Although Bardeh and Rijmen proposed a 7-round key recovery attack independent of the key schedule [3], it necessitates an exhaustive search to verify the remaining key candidates and thus falls into the key-dependent category. Consequently, a key-independent secret-key distinguisher for 7-round AES that relies solely on the structural properties of the round function, without discovering any key information, has not been proposed previously.

In this paper, we propose the first key-independent secret-key distinguisher for 7-round AES. Our starting point is the generalized zero-difference property of Bardeh and Rijmen [3], which we extend from a single structural relation to three that hold simultaneously. Our main contributions are summarized as follows:

- **The Joint Generalized Zero-Difference Property:** We propose the *Joint Generalized Zero-Difference Property*, demonstrating that a quartet constructed from related differences satisfies three distinct generalized zero-difference properties simultaneously, as stated in Theorem 3. Among the three relations, two are exploited as filtering conditions in the proposed distinguisher. This property serves as the fundamental basis for constructing the first key-independent secret-key distinguisher for 7-round AES.
- **New 7-Round Differential Characteristic:** By leveraging this joint property, we construct a new 7-round differential characteristic. A right quartet follows this characteristic with a probability of $2^{-250.4}$, which is distinguishable from the probability of $2^{-253.4}$ for a random permutation. This distinct gap enables the construction of a valid distinguisher.
- **The First 7-Round Key-Independent Secret-Key Distinguisher:** Based on the proposed characteristic, we design a distinguishing attack requiring data, time, and memory complexities of $2^{126.2}$. Our analysis confirms a success probability of at least 75.7%. To the best of our knowledge, this is the longest-round key-independent secret-key distinguisher for AES reported to date.
- **Experimental Verification:** We experimentally verify the joint generalized zero-difference property using small-scale AES [8] with a 4-bit S-box and a 64-bit block size. Across 10^8 quartets over 100 independent keys, all three relations hold without exception, and the observed joint hit counts closely match the theoretical Poisson model, confirming the validity of the proposed property.

Table 1 compares our results with existing secret-key distinguishers for reduced-round AES. As shown in the table, while there are numerous distinguishers for 5 and 6 rounds, our work represents the first key-independent secret-key distinguisher reaching 7 rounds.

Table 1. Comparison of key-independent secret-key distinguishers for reduced-round AES.

Rounds	Method	Data	Time	Mem.	Succ.	Ref.
5	Multiple-of-8	2^{32} CP	$2^{35.6}$ M	2^{32}	99%	[16]
	Exchange	2^{30} CP	2^{30} E	2^{30}	63%	[4]
	Retracing Boom.	2^{31} ACC	2^{31} E	2^{31}	–	[12]
	Boomeyong	2^{47} ACC	2^{46} X	2^{23}	–	[21]
	Yoyo (revisited)	$2^{29.95}$ ACPC	$2^{29.95}$ M	–	55%	[20]
	Yoyo (revisited)	$2^{30.65}$ ACPC	$2^{29.95}$ M	–	81%	[20]
	Massive Exch. Boom.	2^{31} ACC	2^{31} M	2^{31}	70%	[24]†
	Multiple Exch. Boom.	$2^{27.1}$ ACC	$2^{27.1}$ M	$2^{27.1}$	80%	[24]†
6	Exchange	$2^{88.2}$ CP	$2^{88.2}$ E	$2^{88.2}$	73%	[4]
	Exchange	2^{84} ACC	2^{83} E	–	63%	[2]†
	Trunc. Diff.	$2^{89.4}$ CP	$2^{96.5}$ M	2^{32}	95%	[1]
	Trunc. Boom.	2^{87} ACC	2^{87} E	2^{68}	84%	[5]
	Re-boomerang	$2^{82.33}$ ACPC	$2^{82.33}$ E	–	64%	[26]
	Triple Boomerangs	$2^{77.82}$ ACPC	$2^{77.82}$ E	–	66%	[26]
	Boomerang Chain	$2^{76.57}$ ACPC	$2^{76.57}$ E	–	60%	[26]
7	Joint Zero-Diff (Ours)	$2^{126.2}$ CP	$2^{126.2}$ E	$2^{126.2}$	75.7%	Sect. 4

CP: Chosen Plaintexts, ACC: Adaptive Chosen Ciphertexts, ACPC: Adaptive Chosen Plaintexts and Ciphertexts, E: Encryptions, M: Memory accesses, X: XOR operations, Mem.: Memory, measured in 128-bit states, Succ.: Success probability. A dash (–) indicates that the quantity is not reported in the original work. †: ePrint.

The foundation of our approach is the observation that the generalized zero-difference property, which was previously used to derive a single structural relation from a quartet of related plaintexts, in fact yields three such relations simultaneously. When four plaintexts P_0, P_1, P_2, P_3 are constructed from two related differences Δx and $\Delta x'$, their intermediate states satisfy a zero-sum condition $A \oplus B \oplus C \oplus D = 0$ after the internal round layers. This single equation admits three distinct rearrangements, $(A \oplus B) = (C \oplus D)$, $(A \oplus C) = (B \oplus D)$, and $(A \oplus D) = (B \oplus C)$, each of which forces the corresponding pair of ciphertext differences to exhibit an identical zero-difference pattern through the subsequent non-linear layer. Each such pattern equality is called a μ-relation. Since all three μ-relations are deterministic consequences of one algebraic identity, they hold simultaneously at no additional probabilistic cost once the underlying zero-sum condition is met.

This joint property gives rise to a gap between 7-round AES and a random permutation. For AES, once a quartet conforms to the 4-round related-differential trail with probability 2^{-189}, we additionally require that two of the three ciphertext difference pairs each exhibit exactly two inactive inverse diagonals on disjoint index sets. Because these index sets are disjoint, the conditions on each set are independent. For AES, two of the three μ-relations each guarantee that the partner pair shares the same inactive pattern on its respective index set, yielding an overall characteristic probability of $2^{-250.4}$. For a random permutation, the absence of the zero-sum structure means that all four conditions must be satisfied independently, resulting in a lower probability of $2^{-253.4}$. To exploit this 2^3-fold gap, we construct $2^{126.2}$ chosen plaintexts whose structure implicitly defines approximately $2^{250.4}$ quartets. A sequential filtering procedure then identifies ciphertext pairs sharing inactive inverse diagonals, combines matching pairs into quartets, and verifies the joint condition on the complementary diagonal set.

The remainder of this paper is organized as follows. Section 2 describes the specification of AES and introduces the theoretical background of the attack, including related differences and the generalized zero-difference property. Section 3 proposes the joint generalized zero-difference property, which is the core contribution of this paper. Furthermore, it presents the new 7-round differential characteristic constructed using this property and provides experimental verification using small-scale AES. Section 4 details the specific procedure of the first key-independent secret-key distinguisher for 7-round AES, designed based on the proposed 7-round characteristic, and analyzes its complexity and success probability. Finally, Section 5 concludes the paper.

2 Preliminaries

2.1 Description of AES

The Advanced Encryption Standard (AES) [9] is an SPN block cipher with a 128-bit block size and key lengths of 128, 192, or 256 bits, corresponding to $N_r = 10, 12, 14$ rounds, respectively. The internal state is a 4×4 matrix of bytes, transformed by the round function $R = AK \circ MC \circ SR \circ SB$, where SB (SubBytes) applies a non-linear 8-bit S-box to each byte, SR (ShiftRows) rotates the i-th row left by i bytes, MC (MixColumns) multiplies each column by a fixed matrix over $\mathbb{F}_{2^8}$, and AK (AddRoundKey) XORs a 128-bit round key. The encryption begins with an initial AK, followed by $N_r - 1$ standard rounds and a final round that omits MC. Since our distinguisher is independent of the key schedule, we omit its description. The round function is illustrated in Figure 1.

The bytes of the AES state are indexed $0, 1, \ldots, 15$. We denote the position of a byte at row r and column c as (r, c), which corresponds to the linear index $r + 4c$ $(0 \leq r, c \leq 3)$. To describe the proposed distinguisher, we define the notions of *column*, *diagonal*, and *inverse diagonal* using this coordinate notation as follows:

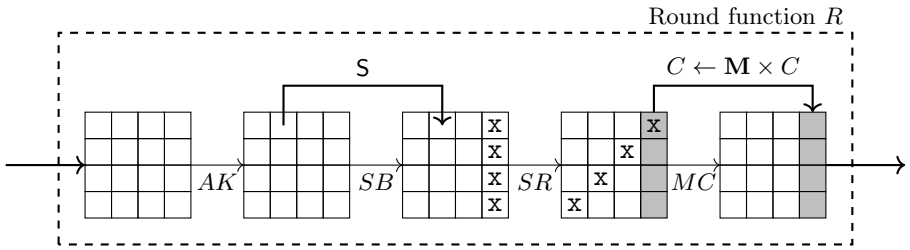

Fig. 1. The round function of AES.

- A *column* is the set of coordinates $\{(0,j),(1,j),(2,j),(3,j)\}$ for $0 \leq j \leq 3$.
- A *diagonal* is defined as the set of coordinates $\{(0,j),(1,j+1),(2,j+2),(3,j+3)\}$ (column indices mod 4) for $0 \leq j \leq 3$. This set maps to the j-th column after the SR operation.
- An *inverse diagonal* is defined as the set of coordinates $\{(0,j),(1,j-1),(2,j-2),(3,j-3)\}$ (column indices mod 4) for $0 \leq j \leq 3$. This set maps to the j-th column after the SR^{-1} operation.

For a state X, we denote the i-th byte as $X[i]$. Furthermore, we denote the i-th column, diagonal, and inverse diagonal as $X[col(i)]$, $X[dia(i)]$, and $X[inv(i)]$, respectively. For a set of indices A, we define $X[A] = \{X[i] \mid i \in A\}$, and similarly, $X[col(A)]$, $X[dia(A)]$, and $X[inv(A)]$ denote the union of columns, diagonals, and inverse diagonals determined by A.

2.2 Related Differences and Differentials

Daemen and Rijmen introduced the concept of related differences to study structural properties of the superbox in AES-like ciphers [10]. When two differences are related, any quartet formed by them contains at most two distinct values at each word position.

Definition 1 (Related Differences [10]). *Let $\Delta x, \Delta x' \in \mathbb{F}_{2^w}^n$ be two differences. They are called related differences if and only if the following condition holds for all $0 \leq i < n$:*

$$\Delta x_i \cdot \Delta x'_i \cdot (\Delta x_i \oplus \Delta x'_i) = 0.$$

This condition implies that for every word index i, at least one of Δx_i, $\Delta x'_i$, or $\Delta x_i \oplus \Delta x'_i$ must be zero. Consequently, the set of values at the i-th word position in any quartet formed by these differences is composed of two identical pairs, i.e., of the form $\{v_1, v_1, v_2, v_2\}$. Based on this, related differentials are defined as follows.

Definition 2 (Related Differentials [10]). *Two differentials $(\Delta x, \Delta y)$ and $(\Delta x', \Delta y')$ for a linear map L are called related differentials if and only if $\Delta y = L(\Delta x)$, $\Delta y' = L(\Delta x')$, and both the input pair $(\Delta x, \Delta x')$ and the output pair $(\Delta y, \Delta y')$ are pairs of related differences.*

Table 2. The sets of byte-related differentials over AES MC [10].

Δx	Δy	$\Delta x'$	$\Delta y'$	$\Delta x \oplus \Delta x'$	$\Delta y \oplus \Delta y'$
$[0, 1, 4, 7]$	$[0, 9, 0, B]$	$[5, 1, 0, 7]$	$[E, 0, D, 0]$	$[5, 0, 4, 0]$	$[E, 9, D, B]$
$[0, 1, 0, 3]$	$[0, 1, 4, 7]$	$[2, 0, 1, 0]$	$[5, 1, 0, 7]$	$[2, 1, 1, 3]$	$[5, 0, 4, 0]$
$[7, 0, 7, 7]$	$[9, E, 0, 0]$	$[7, 7, 7, 0]$	$[0, 0, 9, E]$	$[0, 7, 0, 7]$	$[9, E, 9, E]$
$[0, 3, 2, 0]$	$[7, 0, 7, 1]$	$[2, 0, 0, 3]$	$[7, 1, 7, 0]$	$[2, 3, 2, 3]$	$[0, 1, 0, 1]$

In the context of AES, Daemen and Rijmen analyzed the properties of the MC operation and identified the specific sets of differences that form related differentials. According to [10], the related differentials over the AES MC are limited to the four classes listed in Table 2 and their variations derived by rotation and scalar multiplication.

2.3 Generalized Zero-Difference Property

The zero-difference property, originally introduced by Rønjom et al. [22], identifies structural properties in SPN ciphers. To formalize this, we first define the zero-difference pattern for a general state in $\mathbb{F}_{2^w}^n$.

Definition 3 (Zero-Difference Pattern [22]). *Let $\alpha \in \mathbb{F}_{2^w}^n$ be a state. The zero-difference pattern, denoted by $\nu(\alpha) \in \mathbb{F}_2^n$, is defined as a vector where the i-th component is given by:*

$$\nu(\alpha)_i = \begin{cases} 1 & \text{if } \alpha_i = 0, \\ 0 & \text{otherwise.} \end{cases}$$

In an SPN structure, the zero-difference pattern at the output of the S-box layer is diffused by the subsequent linear layer L, making it no longer directly observable in the ciphertext. However, since the difference is defined by the XOR operation and L is linear, applying L^{-1} to the output difference exactly recovers the difference before L, and hence the original zero-difference pattern. Thus, we define the function μ using ν as follows.

Definition 4 (Zero-Difference Pattern of Inverse State [22]). *Let L be the linear layer of the SPN. For a state $\alpha \in \mathbb{F}_{2^w}^n$, we define the function $\mu(\alpha) \in \mathbb{F}_2^n$ as:*

$$\mu(\alpha) = \nu(L^{-1}(\alpha)).$$

This function effectively reveals the zero-difference pattern of the state before the final linear transformation L.

To effectively utilize the function μ for detecting structural properties, it is essential to understand how specific input structures propagate through the nonlinear layer. When inputs are constructed from related differences, they inherently possess a pairwise equality. The following lemma formalizes the behavior of the S-box with respect to such inputs, demonstrating that the zero-sum property is strictly preserved.

Lemma 1 (Preservation of Zero-Sum Property). *Let S be the non-linear layer consisting of bijective w-bit S-boxes applied independently to each word. For any quartet $\mathcal{Q} = \{\alpha, \beta, \gamma, \delta\}$ constructed from related differences, the zero-sum property is preserved after the S-layer:*

$$S(\alpha) \oplus S(\beta) \oplus S(\gamma) \oplus S(\delta) = 0.$$

Proof. By the definition of related differences, the multiset of inputs at each word position i consists of two identical pairs $\{x_i, x_i, y_i, y_i\}$. Since S is bijective, the outputs are $\{S(x_i), S(x_i), S(y_i), S(y_i)\}$, whose XOR sum is zero at every position. □

Since S is a permutation, it maps non-zero input differences to non-zero output differences, preserving the zero-difference pattern. We formalize this as follows.

Lemma 2 (Preservation of Zero-Difference Pattern). *Let S be the non-linear layer applying a bijective w-bit S-box to n words independently. Let Δ_{in}, $\Delta'_{in} \in \mathbb{F}_{2^w}^n$ be two input differences, and let $\Delta_{out}, \Delta'_{out} \in \mathbb{F}_{2^w}^n$ denote the corresponding output differences through the S-layer. If $\Delta_{in} = \Delta'_{in}$, then the zero-difference patterns of the output differences are identical:*

$$\nu(\Delta_{out}) = \nu(\Delta'_{out}).$$

Proof. Since S applies a permutation to each word independently, the output difference at index i is zero if and only if the input difference is zero. Thus $\Delta_{in} = \Delta'_{in}$ implies $\nu(\Delta_{out}) = \nu(\Delta'_{out})$. □

Bardeh and Rijmen [3] generalized the zero-difference property by combining it with related differences, building upon Lemma 1 and Lemma 2.

Theorem 1 (Generalized Zero-Difference Property [3]). *Let E denote a 2-round SPN encryption function defined as $E = L \circ S \circ L \circ S$, where S is the non-linear layer and L is the linear layer. If the input differences $\Delta x, \Delta x' \in \mathbb{F}_{2^w}^n$ form related differences, then for any plaintext $\alpha \in \mathbb{F}_{2^w}^n$, the following relation holds:*

$$\mu(E(\alpha) \oplus E(\alpha \oplus \Delta x)) = \mu(E(\alpha \oplus \Delta x') \oplus E(\alpha \oplus \Delta x \oplus \Delta x')).$$

Proof. Let $\beta = \alpha \oplus \Delta x$, $\gamma = \alpha \oplus \Delta x'$, $\delta = \alpha \oplus \Delta x \oplus \Delta x'$. By Lemma 1, the zero-sum property is preserved after the first S-layer. Since L is linear, $L(S(\alpha)) \oplus L(S(\beta)) \oplus L(S(\gamma)) \oplus L(S(\delta)) = 0$, which rearranges to $L(S(\alpha)) \oplus L(S(\beta)) = L(S(\gamma)) \oplus L(S(\delta))$. This means the input differences to the second S-layer are identical, so by Lemma 2, $\nu(S(L(S(\alpha))) \oplus S(L(S(\beta)))) = \nu(S(L(S(\gamma))) \oplus S(L(S(\delta))))$. By the definition of μ, the stated equality follows. □

This generalized property is directly applicable to AES. It is well known that 4-round AES can be represented as a 2-round SPN by utilizing a superbox representation [10,14,22]. Concretely, $S' = SR \circ SB \circ AK \circ MC \circ SR \circ SB$ and

$L' = AK \circ MC$ satisfy $L' \circ S' = R^2$, hence $R^4 = L' \circ S' \circ L' \circ S'$; here S' acts as four parallel 32-bit super-boxes, mapping the i-th diagonal onto the i-th inverse diagonal. Consequently, the result of Theorem 1 holds for 4-round AES, which leads to the following corollary.

Corollary 1. *Let R^4 denote the 4-round AES encryption. If the input differences $\Delta x, \Delta x' \in \mathbb{F}_{2^{32}}^4$ form related differences, then for any plaintext $\alpha \in \mathbb{F}_{2^{32}}^4$, the following relation holds for $\mu : \mathbb{F}_{2^{32}}^4 \to \mathbb{F}_2^4$:*

$$\mu(R^4(\alpha) \oplus R^4(\alpha \oplus \Delta x)) = \mu(R^4(\alpha \oplus \Delta x') \oplus R^4(\alpha \oplus \Delta x \oplus \Delta x')).$$

2.4 7-Round Characteristic with Related Differentials

Bardeh and Rijmen [3] proposed a 7-round differential characteristic[8] for AES, constructed by embedding a 4-round related-differential trail into the first four rounds. The mathematical foundation of this characteristic relies on the 8-round property established in Theorem 2 (Theorem 5 in [3]), which is derived by concatenating the 4-round related-differential trail with the 4-round generalized zero-difference property from Corollary 1.

Theorem 2 ([3]). *Let $\alpha \in \mathbb{F}_{2^{32}}^4$, and let $\Delta x, \Delta x' \in \mathbb{F}_{2^{32}}^4$ be two related differences. Assume that all diagonals in Δx are non-zero, and two non-consecutive diagonals in $\Delta x'$ are zero. Let $\mathcal{Q} = (\alpha,\ \alpha \oplus \Delta x,\ \alpha \oplus \Delta x',\ \alpha \oplus \Delta x \oplus \Delta x')$ be the input quartet. Then, the following two relations hold with a probability of approximately 2^{-189}:*

1. *For all rounds t such that $0 < t \leq 6$, the XOR sum of the intermediate states of the quartet is zero:*
$$\bigoplus_{P \in \mathcal{Q}} R^t(P) = 0.$$
2. *The generalized zero-difference pattern in $\mathbb{F}_2^4$ for the 8-round output differences satisfies:*
$$\mu(R^8(\alpha) \oplus R^8(\alpha \oplus \Delta x)) = \mu(R^8(\alpha \oplus \Delta x') \oplus R^8(\alpha \oplus \Delta x \oplus \Delta x')).$$

The probability 2^{-189} follows from [3]: each of the two superbox layers spanned by the 4-round trail (Figure 2) contributes probability $(2^8 - 1)^4 \cdot (2^{32} - 1)^{-4} \approx 2^{-96}$, and there are $4 \times 2 = 8$ valid trail configurations (4 byte-related differential classes in Table 2, each admitting 2 non-consecutive zero-diagonal positions in $\Delta x'$: $\{0, 2\}$ or $\{1, 3\}$). Multiplying yields $8 \times (2^{-96})^2 = 2^{-189}$. We refer to [3] (Theorem 5 and its proof) for the complete derivation.

This theorem relies on the specific 4-round related-differential trail depicted in Figure 2. Specifically, if an input quartet defined as

$$\mathcal{Q} = (\alpha, \alpha \oplus \Delta x, \alpha \oplus \Delta x', \alpha \oplus \Delta x \oplus \Delta x')$$

[8] Throughout this paper, by "differential characteristic" we refer to a truncated differential trail that specifies the zero/non-zero pattern of state differences at each round, rather than the exact difference values.

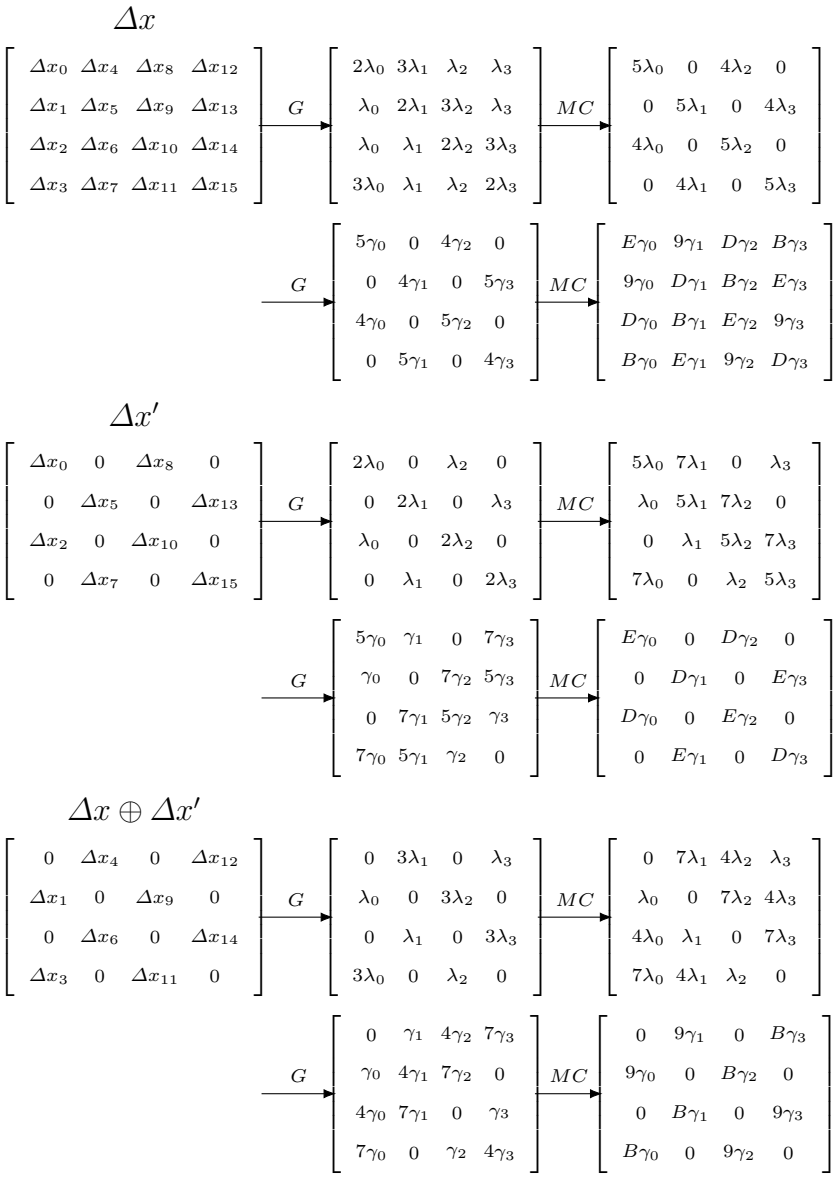

Fig. 2. 4-round related differential trails starting with Δx, $\Delta x'$, and $\Delta x \oplus \Delta x'$ respectively, with a probability of 2^{-189} [3], where $G = SR \circ SB \circ AK \circ MC \circ SR \circ SB \circ AK$ and $\lambda_i, \gamma_i \in \mathbb{F}_{2^8}$ for $0 \leq i \leq 3$.

conforms to this trail, the output differences after 4 rounds preserve the related-difference structure. This ensures that the conditions of Corollary 1 are met for the subsequent rounds.

The zero-sum property established in Theorem 2 directly implies a structural property for 7-round AES, which we formalize in the following corollary.

Corollary 2 (7-Round Generalized Zero-Difference Property). *Under the assumptions of Theorem 2, the following relation regarding the generalized zero-difference pattern holds for the 7th round output with a probability of* 2^{-189}*:*

$$\mu(R^7(\alpha) \oplus R^7(\alpha \oplus \Delta x')) = \mu(R^7(\alpha \oplus \Delta x) \oplus R^7(\alpha \oplus \Delta x \oplus \Delta x')).$$

Proof. By Theorem 2, the zero-sum property is preserved up to 6 rounds (i.e., $\bigoplus_{P \in \mathcal{Q}} R^6(P) = 0$) with a probability of 2^{-189}. This implies:

$$R^6(\alpha) \oplus R^6(\alpha \oplus \Delta x') = R^6(\alpha \oplus \Delta x) \oplus R^6(\alpha \oplus \Delta x \oplus \Delta x').$$

This equation indicates that the input differences to the 7th round are identical for the two pairs. By Lemma 2, the equality of input differences guarantees that the zero-difference patterns of the output differences of SB are identical. Consequently, the stated relation for μ is satisfied. □

To apply this property to a key-recovery attack on 7-round AES, [3] constructed the full 7-round characteristic depicted in Figure 3. The total probability that a random quartet conforms to this 7-round characteristic is approximately $2^{-189} \times 2^{-29.4} = 2^{-218.4}$. A quartet satisfying this characteristic preserves the zero-sum property up to 6 rounds by Theorem 2, which serves as the basis for the key recovery attack.

It is worth noting that the probability that a random permutation satisfies the input and output conditions of the 7-round characteristic is higher than the probability of the characteristic itself. For a random permutation, the probability of satisfying the conditions is $2^{-61.4} \times 2^{-64} = 2^{-125.4}$. Let C_i denote the ciphertext. This probability corresponds to the event where $MC^{-1}(C_0 \oplus C_2)$ is inactive in two inverse diagonals, and $MC^{-1}(C_1 \oplus C_3)$ is also inactive in the same two inverse diagonals. Since the probability of the characteristic is $2^{-218.4}$, which is significantly lower than the probability of $2^{-125.4}$ for a random permutation, it is disadvantageous to use this property as a standalone distinguisher.

However, [3] leveraged the fact that identifying even a single right quartet that satisfies the characteristic allows for the recovery of the secret key. By filtering data to find a right quartet, the attacker can exploit the structural zero-sum property to efficiently recover the key.

3 Joint Generalized Zero-Difference Property and 7-Round Characteristic for AES

In this section, we present the structural property utilized to distinguish 7-round AES. We demonstrate that a quartet constructed using related differences satisfies generalized zero-difference properties jointly. Furthermore, we describe the 7-round differential trail that enables this simultaneous occurrence.

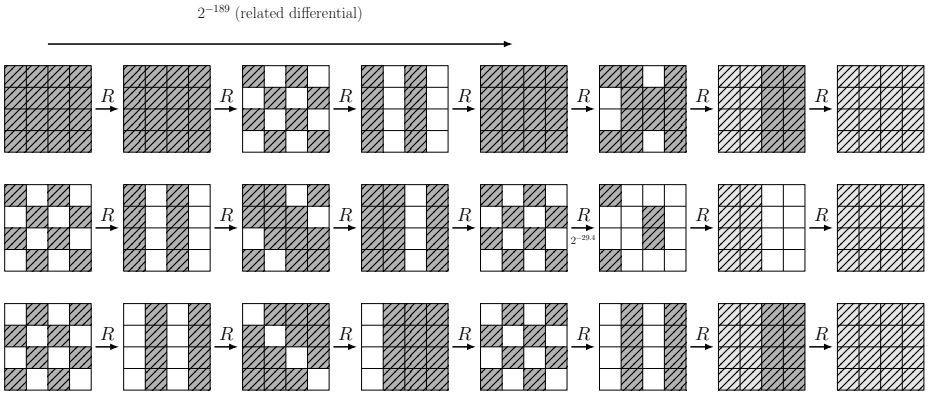

Fig. 3. The 7-round differential characteristics for the quartet, where the rows from top to bottom represent Δx, $\Delta x'$, and $\Delta x \oplus \Delta x'$, respectively [3]. The probability that a quartet follows these characteristics is $2^{-218.4}$. A white cell indicates that the state difference is zero for those bytes. Both darker and lighter line-patterned cells indicate that the state difference is non-zero; specifically, a darker cell indicates that the difference value in that cell is equal to the value of one of the other two differences in the same cell location.

3.1 Joint Generalized Zero-Difference Property

The core contribution of this work is the observation that a quartet structure formed by related differences satisfies three distinct generalized zero-difference properties jointly, rather than a single relationship.

Theorem 3 (Joint Generalized Zero-Difference Property). *Let E denote a 2-round SPN encryption function defined as $E = L \circ S \circ L \circ S$, where S is the non-linear layer and L is the linear layer. If the input differences $\Delta x, \Delta x' \in \mathbb{F}_{2^w}^n$ form related differences, then for any plaintext $\alpha \in \mathbb{F}_{2^w}^n$, the following three relations hold simultaneously:*

$$\begin{aligned}
\mu(E(\alpha) \oplus E(\alpha \oplus \Delta x)) &= \mu(E(\alpha \oplus \Delta x') \oplus E(\alpha \oplus \Delta x \oplus \Delta x')),\\
\mu(E(\alpha) \oplus E(\alpha \oplus \Delta x')) &= \mu(E(\alpha \oplus \Delta x) \oplus E(\alpha \oplus \Delta x \oplus \Delta x')),\\
\mu(E(\alpha) \oplus E(\alpha \oplus \Delta x \oplus \Delta x')) &= \mu(E(\alpha \oplus \Delta x) \oplus E(\alpha \oplus \Delta x')).
\end{aligned}$$

Proof. Let $\beta = \alpha \oplus \Delta x$, $\gamma = \alpha \oplus \Delta x'$, $\delta = \alpha \oplus \Delta x \oplus \Delta x'$. By Lemma 1, the zero-sum property is preserved after the first S-box layer, and since L is linear, we obtain:

$$L(S(\alpha)) \oplus L(S(\beta)) \oplus L(S(\gamma)) \oplus L(S(\delta)) = 0.$$

By rearranging the terms of this equation, we can derive the equalities for three different cases.

1. **Case 1 (Relation regarding Δx):** Grouping the terms as (α, β) and (γ, δ) yields:
$$L(S(\alpha)) \oplus L(S(\beta)) = L(S(\gamma)) \oplus L(S(\delta)).$$
By Lemma 2 and the definition of μ, exactly as in the proof of Theorem 1, this yields
$$\mu(E(\alpha) \oplus E(\beta)) = \mu(E(\gamma) \oplus E(\delta)).$$
2. **Case 2 (Relation regarding $\Delta x'$):** Grouping the terms as (α, γ) and (β, δ) gives $L(S(\alpha)) \oplus L(S(\gamma)) = L(S(\beta)) \oplus L(S(\delta))$, and the identical argument yields $\mu(E(\alpha) \oplus E(\gamma)) = \mu(E(\beta) \oplus E(\delta))$.
3. **Case 3 (Relation regarding $\Delta x \oplus \Delta x'$):** Grouping the terms as (α, δ) and (β, γ) gives $L(S(\alpha)) \oplus L(S(\delta)) = L(S(\beta)) \oplus L(S(\gamma))$, and the same argument yields $\mu(E(\alpha) \oplus E(\delta)) = \mu(E(\beta) \oplus E(\gamma))$.

Therefore, a quartet constructed from related differences satisfies all three relations simultaneously. □

The observation of this *joint* property significantly enhances the performance of the distinguishing attack. For a random permutation, the probability that an arbitrary quartet satisfies multiple relations simultaneously is substantially lower compared to the probability of satisfying a single relation. We emphasize that the three μ-relations in Theorem 3 are not probabilistically independent events but rather three deterministic rearrangements of a single zero-sum equation. Consequently, once the underlying zero-sum property holds, all three relations are satisfied simultaneously with no additional probabilistic cost.

3.2 New 7-Round Characteristic for AES

In this section, we propose a new 7-round differential characteristic for AES by utilizing the joint generalized zero-difference property presented in Theorem 3.

First, Theorem 3 naturally extends to 8-round AES by combining the 4-round related-differential trail with the fact that 4-round AES can be represented as a 2-round SPN. We formally state this extension as follows.

Theorem 4 (Joint Generalized Zero-Difference Property for 8-Round AES). *Let R^t denote the t-round AES encryption function. Assume that the input differences $\Delta x, \Delta x' \in \mathbb{F}_{2^{32}}^4$ are chosen such that all diagonals in Δx are non-zero, and two non-consecutive diagonals in $\Delta x'$ are zero. Then, for any plaintext $\alpha \in \mathbb{F}_{2^{32}}^4$, the following two relations hold with a probability of approximately 2^{-189}:*

1. *For all rounds t such that $0 < t \leq 6$, the XOR sum of the intermediate states of the quartet is zero:*
$$\bigoplus_{P \in \mathcal{Q}} R^t(P) = 0.$$

2. *The following three relations regarding the generalized zero-difference pattern* $\mu : \mathbb{F}_{2^{32}}^4 \to \mathbb{F}_2^4$ *hold simultaneously at the 8th round:*

$$\mu(R^8(\alpha) \oplus R^8(\alpha \oplus \Delta x)) = \mu(R^8(\alpha \oplus \Delta x') \oplus R^8(\alpha \oplus \Delta x \oplus \Delta x')),$$
$$\mu(R^8(\alpha) \oplus R^8(\alpha \oplus \Delta x')) = \mu(R^8(\alpha \oplus \Delta x) \oplus R^8(\alpha \oplus \Delta x \oplus \Delta x')),$$
$$\mu(R^8(\alpha) \oplus R^8(\alpha \oplus \Delta x \oplus \Delta x')) = \mu(R^8(\alpha \oplus \Delta x) \oplus R^8(\alpha \oplus \Delta x')).$$

Proof. **Zero-sum preservation up to 6 rounds.** Let $\mathcal{Q} = (\alpha,\ \alpha \oplus \Delta x,\ \alpha \oplus \Delta x',\ \alpha \oplus \Delta x \oplus \Delta x')$ be the input quartet. By assumption, Δx and $\Delta x'$ satisfy the conditions of Theorem 2: all diagonals in Δx are non-zero, and two non-consecutive diagonals in $\Delta x'$ are zero. By Theorem 2, the quartet conforms to the 4-round related-differential trail in Figure 2 with probability 2^{-189}, and conditioned on this event, the XOR sum of the intermediate states is zero for all rounds $0 < t \leq 6$:

$$\bigoplus_{P \in \mathcal{Q}} R^t(P) = 0, \quad \forall\, 0 < t \leq 6.$$

From zero-sum at round 6 to three joint μ-relations at round 8. Denote the four intermediate states after 4 rounds as:

$$A = R^4(\alpha), \quad B = R^4(\alpha \oplus \Delta x), \quad C = R^4(\alpha \oplus \Delta x'), \quad D = R^4(\alpha \oplus \Delta x \oplus \Delta x').$$

Step 1: The output differences after 4 rounds form related differences. Conditioned on the quartet following the 4-round trail (Figure 2), the output differences $\Delta y = A \oplus B$, $\Delta y' = A \oplus C$, and $\Delta y \oplus \Delta y' = A \oplus D$ preserve the related difference structure. Specifically, the trail guarantees that at every 32-bit word position i (in the superbox representation), the multiset $\{A[i], B[i], C[i], D[i]\}$ contains at most two distinct values. This is precisely the condition required for $(\Delta y, \Delta y')$ to be related differences at the superbox level.
Step 2: The encryption from the 5th round to the 8th round forms a 2-round SPN in the superbox representation. It is well established [10,14,22] that 4 consecutive rounds of AES can be represented as a 2-round SPN $E' = L' \circ S' \circ L' \circ S'$, where S' is the 32-bit superbox non-linear layer and L' is the corresponding linear layer. The states (A, B, C, D) serve as inputs to this 2-round SPN.
Step 3: Applying the Joint Generalized Zero-Difference Property. Since the quartet (A, B, C, D) is formed from related differences (Step 1), it satisfies the assumptions of Theorem 3. Applying Theorem 3 to the 2-round SPN E', we obtain that the following three relations hold deterministically (i.e., with probability 1):

$$\mu(E'(A) \oplus E'(B)) = \mu(E'(C) \oplus E'(D)), \tag{1}$$
$$\mu(E'(A) \oplus E'(C)) = \mu(E'(B) \oplus E'(D)), \tag{2}$$
$$\mu(E'(A) \oplus E'(D)) = \mu(E'(B) \oplus E'(C)). \tag{3}$$

Since $E'(A) = R^8(\alpha)$, $E'(B) = R^8(\alpha \oplus \Delta x)$, $E'(C) = R^8(\alpha \oplus \Delta x')$, and $E'(D) = R^8(\alpha \oplus \Delta x \oplus \Delta x')$, (1)–(3) are precisely the three μ-relations stated in the theorem. Since the 4-round trail is the only probabilistic event, both properties hold simultaneously with probability 2^{-189}. □

The zero-sum property up to 6 rounds presented in Theorem 4 implies a significant structural property for 7-round analysis. We formalize this in the following corollary.

Corollary 3 (7-Round Joint Generalized Zero-Difference Property). *Under the assumptions of Theorem 4, the following three relations regarding the generalized zero-difference pattern hold simultaneously for the 7th round output with a probability of* 2^{-189}*:*

$$\begin{aligned} \mu(R^7(\alpha) \oplus R^7(\alpha \oplus \Delta x)) &= \mu(R^7(\alpha \oplus \Delta x') \oplus R^7(\alpha \oplus \Delta x \oplus \Delta x')), \\ \mu(R^7(\alpha) \oplus R^7(\alpha \oplus \Delta x')) &= \mu(R^7(\alpha \oplus \Delta x) \oplus R^7(\alpha \oplus \Delta x \oplus \Delta x')), \\ \mu(R^7(\alpha) \oplus R^7(\alpha \oplus \Delta x \oplus \Delta x')) &= \mu(R^7(\alpha \oplus \Delta x) \oplus R^7(\alpha \oplus \Delta x')). \end{aligned} \tag{4}$$

Proof. According to Theorem 4, the XOR sum of the intermediate states is zero for all rounds $0 < t \leq 6$ (i.e., $\bigoplus_{P \in \mathcal{Q}} R^6(P) = 0$) with a probability of 2^{-189}. By applying the same logic as in the proof of Theorem 3, we can rearrange the terms of this zero-sum equation for each pair combination, resulting in identical input differences for all three cases. By Lemma 2, the equality of input differences guarantees that the generalized zero-difference patterns of the output differences of SB are identical. Consequently, the three stated relations hold simultaneously. □

We propose a new 7-round differential characteristic that exploits this property, as illustrated in Figure 4. We specifically target 7 rounds instead of 8 because the probability of the property holding for 8 rounds is lower than that of a random permutation, rendering it unsuitable as a distinguisher.

This new characteristic builds upon the existing 7-round characteristic described in Figure 3. While the original characteristic induces zero differences in two diagonals for $\Delta x'$ after the 5th round, we impose an additional condition requiring that $\Delta x \oplus \Delta x'$ also exhibits zero differences in two diagonals after the 5th round. Crucially, the set of zero diagonals for $\Delta x'$ and that for $\Delta x \oplus \Delta x'$ must be disjoint.

For $\Delta x'$, two inactive inverse diagonals at the 7th round output correspond to two inactive diagonals after the 5th round encryption (since SR maps diagonals to columns). The probability that exactly two out of four diagonals become inactive after the 5th round is $\binom{4}{2} \times 2^{-32} = 6 \times 2^{-32} \approx 2^{-29.4}$. The per-pair probability 2^{-32} arises because $\Delta x'$ has only 2 active columns entering the 5th round SB, so each target pair of inverse diagonals contains only $2 \times 2 = 4$ non-zero bytes that must collide independently through SB (each with probability 2^{-8}). To ensure that *exactly* two diagonals are inactive (excluding cases where three or more are simultaneously inactive), we subtract the overcounting term $\binom{4}{3} \times 2^{-48} = 2^{-46}$, yielding a precise probability of $2^{-29.4} - 2^{-46}$. Since the subtraction term is negligible, we approximate this as $2^{-29.4}$.

For $\Delta x \oplus \Delta x'$, once the two inactive inverse diagonals of the 7th round ciphertexts for $\Delta x'$ are determined (i.e., the set $\mathcal{I}$ is fixed), the disjointness condition forces $\mathcal{J} = \{0, 1, 2, 3\} \backslash \mathcal{I}$. The two inactive inverse diagonals at the 7th round output correspond to two inactive diagonals after the 5th round encryption. Since

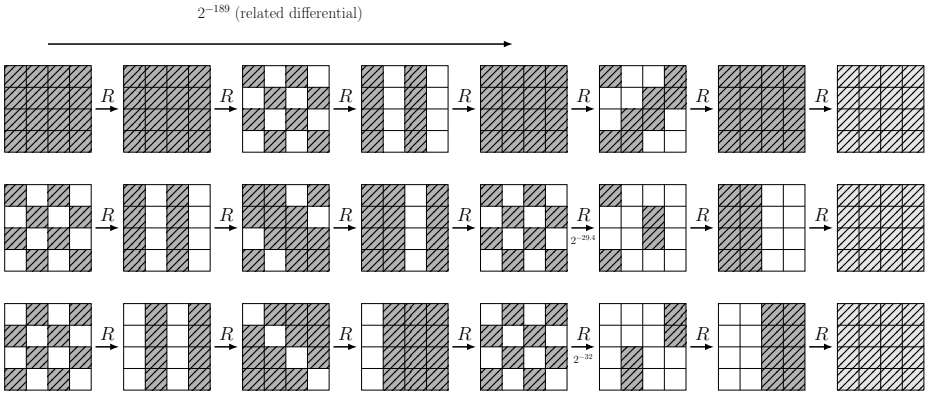

Fig. 4. The proposed 7-round differential characteristics for the quartet, where the rows from top to bottom represent Δx, $\Delta x'$, and $\Delta x \oplus \Delta x'$, respectively. Unlike the existing characteristic in Figure 3, this new characteristic additionally constrains $\Delta x \oplus \Delta x'$ to exhibit two inactive inverse diagonals disjoint from those of $\Delta x'$ at the 7th round output. The probability that a quartet follows these characteristics is $2^{-250.4}$. A white cell indicates that the state difference is zero for those bytes. Both darker and lighter line-patterned cells indicate that the state difference is non-zero; specifically, a darker cell indicates that the difference value in that cell is equal to the value of one of the other two differences in the same cell location.

the output of the 4th round for $\Delta x \oplus \Delta x'$ also has two active diagonals, the probability that $\Delta x \oplus \Delta x'$ has exactly the two inactive diagonals corresponding to $\mathcal{J}$ after the 5th round is 2^{-32}.

Consequently, the total probability that a quartet conforms to this new 7-round characteristic is:

$$2^{-189} \times 2^{-29.4} \times 2^{-32} = 2^{-250.4}.$$

To use this characteristic as a distinguisher, we formally define the property that the distinguisher verifies on the ciphertexts and compare its probability under 7-round AES versus a random permutation.

Definition 5 (7-Round Distinguishing Property). *Let (C_0, C_1, C_2, C_3) be a quartet of ciphertexts. Let $\mathcal{I}, \mathcal{J} \subset \{0, 1, 2, 3\}$ be two subsets satisfying $|\mathcal{I}| = |\mathcal{J}| = 2$ and $\mathcal{I} \cap \mathcal{J} = \emptyset$. We say that the quartet satisfies the 7-round distinguishing property with respect to $(\mathcal{I}, \mathcal{J})$ if the following four conditions hold simultaneously:*

- ***C1.*** $MC^{-1}(C_0 \oplus C_2)[inv(k)] = 0$ *for all* $k \in \mathcal{I}$, *and* $MC^{-1}(C_0 \oplus C_2)[inv(k)] \neq 0$ *for all* $k \notin \mathcal{I}$.
- ***C2.*** $MC^{-1}(C_1 \oplus C_3)[inv(k)] = 0$ *for all* $k \in \mathcal{I}$, *and* $MC^{-1}(C_1 \oplus C_3)[inv(k)] \neq 0$ *for all* $k \notin \mathcal{I}$.

C3. $MC^{-1}(C_0 \oplus C_3)[inv(k)] = 0$ *for all* $k \in \mathcal{J}$, *and* $MC^{-1}(C_0 \oplus C_3)[inv(k)] \neq 0$ *for all* $k \notin \mathcal{J}$.

C4. $MC^{-1}(C_1 \oplus C_2)[inv(k)] = 0$ *for all* $k \in \mathcal{J}$, *and* $MC^{-1}(C_1 \oplus C_2)[inv(k)] \neq 0$ *for all* $k \notin \mathcal{J}$.

Conditions C1 and C2 correspond to the μ-relation for $\Delta x'$ (the first equation in (4)), while Conditions C3 and C4 correspond to the μ-relation for $\Delta x \oplus \Delta x'$ (the third equation in (4)).

The inactive (zero) conditions in Definition 5 directly reflect the generalized zero-difference property from Corollary 3. The active (non-zero) conditions arise from the structure of the 7-round characteristic (Figure 4): the characteristic prescribes that the output differences have exactly two zero inverse diagonals for $\Delta x'$ and exactly two zero inverse diagonals for $\Delta x \oplus \Delta x'$, with the remaining inverse diagonals being strictly non-zero.

We note that Definition 5 exploits only the first and third μ-relations from Equation (4) (corresponding to $\Delta x'$ and $\Delta x \oplus \Delta x'$, respectively). The second μ-relation (corresponding to Δx) is not used because Δx has all four diagonals non-zero after the 4-round trail; consequently, it does not produce inactive inverse diagonals at the 7th round output, and thus provides no additional filtering power in the distinguisher.

We now establish the probability that this property is satisfied under 7-round AES and under a random permutation.

Proposition 1 (Probability under 7-Round AES). *Let (C_0, C_1, C_2, C_3) be the ciphertexts of a quartet $\mathcal{Q} = (\alpha, \alpha \oplus \Delta x, \alpha \oplus \Delta x', \alpha \oplus \Delta x \oplus \Delta x')$ encrypted under 7-round AES with a fixed unknown key. Under the conditions on Δx and $\Delta x'$ specified in Theorem 4, for any disjoint subsets $\mathcal{I}, \mathcal{J} \subset \{0, 1, 2, 3\}$ with $|\mathcal{I}| = |\mathcal{J}| = 2$, the probability that the quartet satisfies the 7-round distinguishing property (Definition 5) with respect to $(\mathcal{I}, \mathcal{J})$ is approximately $2^{-250.4}$.*

Proof. The probability is determined by three events occurring in sequence.

First, by Theorem 4, the quartet conforms to the 4-round related-differential trail with probability 2^{-189}, which guarantees the zero-sum property up to 6 rounds and consequently the three μ-relations in (4).

Second, conditioned on the 4-round trail, we analyze the probability that the 7th round ciphertext difference for $\Delta x'$ has exactly two inactive inverse diagonals. After 4 rounds, $\Delta x'$ has two active diagonals. After SR in the 5th round, these two active diagonals correspond to two active columns. For a specific pair of diagonals to become inactive after the 5th round encryption, only 4 bytes (the bytes in those two diagonals originating from the active columns) need to have their differences collide through SB, each with probability 2^{-8}. This gives a probability of $(2^{-8})^4 = 2^{-32}$ for a specific pair of diagonals. Since there are $\binom{4}{2} = 6$ choices for which two diagonals become inactive, the total probability is $6 \times 2^{-32} \approx 2^{-29.4}$. Let $\mathcal{I}$ denote the indices of the resulting two inactive inverse diagonals at the 7th round output.

Third, by the same reasoning, the 7th round ciphertext difference for $\Delta x \oplus \Delta x'$ must also have exactly two inactive inverse diagonals. After 4 rounds, $\Delta x \oplus \Delta x'$ also has two active diagonals. Since $\mathcal{I}$ is already determined, the disjointness condition forces $\mathcal{J} = \{0,1,2,3\} \setminus \mathcal{I}$, leaving only one admissible pair of diagonals. By the same argument as above, the probability that this specific pair becomes inactive after the 5th round is 2^{-32}.

Once these three events occur, the quartet follows the 7-round characteristic in Figure 4, and by Corollary 3 all four conditions of Definition 5 are satisfied.

The three events are independent: the 5th-round S-box transitions in the second and third events depend only on the input differences at the byte positions involved, not on the intermediate values fixed by the 4-round trail, and the second and third events act on the disjoint byte positions indexed by $\mathcal{I}$ and $\mathcal{J}$. Therefore, all four conditions are satisfied with a combined probability of $2^{-189} \times 2^{-29.4} \times 2^{-32} = 2^{-250.4}$. □

Proposition 2 (Probability under a Random Permutation). *Let π be a random permutation, and let (C_0, C_1, C_2, C_3) be the images of a quartet under π. For any disjoint subsets $\mathcal{I}, \mathcal{J} \subset \{0,1,2,3\}$ with $|\mathcal{I}| = |\mathcal{J}| = 2$, the probability that the quartet satisfies the 7-round distinguishing property (Definition 5) is approximately $2^{-253.4}$.*

Proof. We analyze the four conditions and establish their independence.

Individual probabilities. For a random permutation, the outputs on four distinct inputs are modeled as independent uniformly random 128-bit values. Each inverse diagonal consists of 4 bytes (32 bits).

For Condition C1, the set $\mathcal{I}$ is not fixed *a priori*; the distinguisher accepts any two inverse diagonals of $MC^{-1}(C_0 \oplus C_2)$ being zero, provided the other two are non-zero. There are $\binom{4}{2} = 6$ choices, and for each specific pair the probability is 2^{-64}. Thus the probability of Condition C1 is $6 \times 2^{-64} \approx 2^{-61.4}$.

For Conditions C2, C3, and C4, the index sets $\mathcal{I}$ and $\mathcal{J}$ are already determined by Condition C1. Each condition requires that two specific inverse diagonals of the respective ciphertext difference are zero and the other two are non-zero, yielding a probability of approximately 2^{-64} for each.

Independence of the four conditions. The crucial point is that $\mathcal{I}$ and $\mathcal{J}$ are disjoint. Since the MC operation acts on each column independently, and distinct inverse diagonals correspond to distinct columns after MC^{-1}, constraints on inverse diagonals indexed by $\mathcal{I}$ are statistically independent from those indexed by $\mathcal{J}$. Although the four pairwise differences $C_0 \oplus C_2$, $C_1 \oplus C_3$, $C_0 \oplus C_3$, and $C_1 \oplus C_2$ satisfy the componentwise linear relation $(C_0 \oplus C_2) \oplus (C_1 \oplus C_3) \oplus (C_0 \oplus C_3) \oplus (C_1 \oplus C_2) = 0$, this relation imposes no cross-constraint between Conditions C1–C2 (acting on $\mathcal{I}$) and Conditions C3–C4 (acting on $\mathcal{J}$) precisely because $\mathcal{I} \cap \mathcal{J} = \emptyset$.

For Conditions C1 and C2 (both involving $\mathcal{I}$): the ciphertext differences $C_0 \oplus C_2$ and $C_1 \oplus C_3$ involve disjoint pairs of ciphertexts under a random permutation, so Conditions C1 and C2 are independent.

For Conditions C3 and C4 (both involving $\mathcal{J}$): by the same reasoning, $C_0 \oplus C_3$ and $C_1 \oplus C_2$ are independent.

It remains to verify that there is no cross-dependency between Conditions C1, C2 and Conditions C3, C4. If there existed $k \in \mathcal{I} \cap \mathcal{J}$, then the linear relation

$$MC^{-1}(C_1 \oplus C_2)[inv(k)] = MC^{-1}(C_0 \oplus C_2)[inv(k)] \oplus MC^{-1}(C_1 \oplus C_3)[inv(k)] \\ \oplus MC^{-1}(C_0 \oplus C_3)[inv(k)]$$

would force Condition C4 to hold automatically whenever Conditions C1, C2, and C3 are satisfied, eliminating its filtering power. However, since $\mathcal{I} \cap \mathcal{J} = \emptyset$, the constraints of Conditions C1 and C2 act on inverse diagonals in $\mathcal{I}$, while Conditions C3 and C4 act on inverse diagonals in $\mathcal{J}$. These column sets are entirely disjoint under MC^{-1}, so all four conditions are independent.

Combined probability. Multiplying the individual probabilities yields:

$$2^{-61.4} \times 2^{-64} \times 2^{-64} \times 2^{-64} = 2^{-253.4}.$$

□

Consequently, the probability ratio between 7-round AES and a random permutation is $2^{-250.4}/2^{-253.4} = 2^3$, confirming that the proposed characteristic yields a valid distinguisher. The specific distinguishing attack algorithm is presented in Section 4.

3.3 Experimental Verification

To support the theoretical analysis, we experimentally verify the joint generalized zero-difference property using small-scale AES [8] with a 4-bit S-box and a 64-bit block size. In this variant, each inverse diagonal consists of 4 nibbles (16 bits), making the truncated probabilities computationally feasible to observe.

Setup. We consider 4-round small-scale AES, which corresponds to a 2-round SPN in the superbox representation and thus directly instantiates the setting of Theorem 3. The related differences Δx and $\Delta x'$ are chosen such that all nibbles of Δx are non-zero, and $\Delta x'$ has two zero diagonals (diagonals 0 and 2) with the remaining nibbles equal to the corresponding nibbles of Δx. For each experiment, we sample a fresh random key, fresh random related differences, and encrypt 2^{30} quartets constructed from these related differences. We repeat this procedure for 100 independent keys.

Experiment 1: Deterministic μ-relation (Theorem 3). For each key, we verify whether the three μ-relations hold simultaneously on the first 10^6 quartets. Across all 100 keys (10^8 quartets in total), every quartet satisfied all three μ-relations without exception, confirming the deterministic nature of Theorem 3.

Table 3. Experimental results over 100 independent keys with 2^{30} quartets per key. Left: summary statistics. Right: distribution of joint hit counts compared with Poisson($\lambda=3$).

	Small AES	**Random**
μ-relation	$10^8/10^8$	—
Hits (mean)	2.86	0.00
Hits (var.)	2.38	0.00
Poisson λ	3.0	≈ 0

Hits	**Observed**	**Poisson**
0	4	5.0
1	16	14.9
2	26	22.4
3	20	22.4
4	18	16.8
5	11	10.1
6	4	5.0
7	1	2.2

Experiment 2: Joint zero inverse-diagonal hits. We count the number of quartets whose ciphertext differences satisfy the following joint condition: $MC^{-1}(C_0 \oplus C_2)$ and $MC^{-1}(C_1 \oplus C_3)$ share a common zero inverse diagonal, and simultaneously $MC^{-1}(C_0 \oplus C_3)$ and $MC^{-1}(C_1 \oplus C_2)$ share a common zero inverse diagonal at a disjoint position. There are $\binom{4}{1} \times \binom{3}{1} = 12$ ordered choices for such a disjoint pair of inverse diagonal indices.

For a **random permutation**, all four ciphertext differences are independent. The probability that both $MC^{-1}(C_0 \oplus C_2)$ and $MC^{-1}(C_1 \oplus C_3)$ are zero on the same inverse diagonal is $2^{-16} \times 2^{-16} = 2^{-32}$. Similarly, the probability that both $MC^{-1}(C_0 \oplus C_3)$ and $MC^{-1}(C_1 \oplus C_2)$ are zero on the same inverse diagonal is 2^{-32}. Since the two pairs involve disjoint inverse diagonals, these events are independent, yielding a per-quartet probability of 12×2^{-64} and an expected count of $2^{30} \times 12 \times 2^{-64} \approx 0$ per run.

For **4-round small-scale AES**, the μ-relation from Theorem 3 guarantees that if $MC^{-1}(C_0 \oplus C_2)$ has a zero inverse diagonal, then $MC^{-1}(C_1 \oplus C_3)$ is automatically zero on the same inverse diagonal with probability 1. Therefore, the probability that this pair jointly has a zero inverse diagonal reduces to 2^{-16} instead of 2^{-32}. By the same reasoning, the probability for the pair ($MC^{-1}(C_0 \oplus C_3)$, $MC^{-1}(C_1 \oplus C_2)$) also reduces to 2^{-16}. This yields a per-quartet probability of 12×2^{-32} and an expected count of $2^{30} \times 12 \times 2^{-32} \approx 3.0$ per run. The results of 100 independent experiments are summarized in Table 3.

The observed mean of 2.86 and variance of 2.38 closely match the theoretical Poisson distribution with $\lambda = 12 \times 2^{-32} \times 2^{30} = 3.0$, where the Poisson model predicts equal mean and variance. The distribution of hit counts further confirms this agreement. In contrast, no joint hits were observed for random permutations across all 100 runs, consistent with the negligible expected value of $12 \times 2^{-64} \times 2^{30} \approx 2.0 \times 10^{-8}$. These results experimentally confirm that the μ-relation from Theorem 3 holds as predicted, and that the joint zero inverse-diagonal counts under small-scale AES match the theoretical expectations derived from the property. We note that Theorem 3 is a structural property of any 2-round SPN with related differences, independent of the word size or block size.

Since small-scale AES and full AES share the same SPN round structure (SubBytes, ShiftRows, MixColumns, AddRoundKey), the deterministic μ-relations validated in Experiment 1 hold identically for full AES. The quantitative agreement observed in Experiment 2 further confirms that the truncated probability model accurately captures the behavior of the joint property.

We remark that an end-to-end verification of the full 7-round distinguisher (Algorithm 1) on small-scale AES is computationally infeasible. For 7-round small-scale AES with 4-bit S-boxes, the 4-round trail probability scales to $4 \times 2 \times 2^{-48} \times 2^{-48} = 2^{-93}$, and the 7-round characteristic probability becomes $2^{-93} \times 2^{-13.4} \times 2^{-16} = 2^{-122.4}$. The resulting data complexity of $2^{62.2}$ chosen plaintexts is close to the full codebook size of 2^{64}, rendering such an experiment impractical. Nonetheless, the experiments above validate the core structural property that underlies the distinguisher.

The source code for all experiments is publicly available for reproducibility.[9]

4 New Key-Independent Secret-Key Distinguisher for 7-Round AES

In this section, we propose a new distinguisher for 7-round AES using the new 7-round differential characteristic proposed in Section 3.2. It is worth noting that the proposed distinguisher is the first known key-independent secret-key distinguisher for 7-round AES. Previous results on 7-round AES were mostly key recovery attacks or key-dependent distinguishers, and no key-independent distinguisher has been proposed for 7-round AES to date. The proposed attack has a complexity of $2^{126.2}$ in terms of data, time, and memory.

4.1 Data Collection

To perform the attack, the attacker must collect input quartets with specific structural features. The input plaintext quartet $\mathcal{Q} = (P_0, P_1, P_2, P_3)$ has the following form:

$$P_0, \quad P_1 = P_0 \oplus \Delta x, \quad P_2 = P_0 \oplus \Delta x', \quad P_3 = P_0 \oplus \Delta x \oplus \Delta x'.$$

Here, all diagonals in Δx are non-zero, and two non-consecutive diagonals in $\Delta x'$ are zero. In this context, the differences Δx, $\Delta x'$, and $\Delta x \oplus \Delta x'$ correspond to the input differences of the 7-round characteristic illustrated in Figure 4. This quartet structure can be viewed as two pairs of plaintexts (P_0, P_1) and (P_2, P_3), where the second pair (P_2, P_3) is generated by exchanging two diagonals between the plaintexts in the first pair (P_0, P_1).

We need to generate approximately $2^{250.4}$ quartets to find a right quartet satisfying the 7-round characteristic proposed in Figure 4. To do this efficiently, we employ the following structural data generation method in [3].

First, we pick two random subsets A_0 and A_1 of size m from $\mathbb{F}_{2^8}^8$. Then, we generate all m^2 possible plaintexts according to the following rule:

[9] `https://github.com/shb115/joint-zero-diff-aes`

– The diagonals with indices 0 and 2 take elements from the set A_0.
– The diagonals with indices 1 and 3 take elements from the set A_1.

Since we select m elements from each set, we can generate $\binom{m}{2}$ unique pairs for each diagonal group. According to [3], the total number of unique quartets generated by combining sets A_0 and A_1 is:

$$N_{quartet} = \binom{m}{2} \times \binom{m}{2} \approx \frac{m^2}{2} \cdot \frac{m^2}{2}.$$

Since our goal is to obtain a right quartet, the total number of quartets $N_{quartet}$ must be equal to the inverse of the characteristic's probability, $2^{250.4}$. To achieve this, we set the parameter m to $2^{63.1}$. Then, the number of generated quartets is calculated as:

$$\binom{2^{63.1}}{2} \times \binom{2^{63.1}}{2} \approx 2^{125.2} \times 2^{125.2} = 2^{250.4}.$$

Therefore, with $m = 2^{63.1}$, the data complexity (number of plaintexts) required for the attack is $m^2 = (2^{63.1})^2 = 2^{126.2}$. Consequently, the attacker queries $2^{126.2}$ chosen plaintexts to structurally construct $2^{250.4}$ quartets, and the expected number of quartets satisfying the 7-round characteristic in Figure 4 among them is one.

4.2 Distinguishing Attack

In this section, we present the specific procedure to distinguish 7-round AES from a random permutation by identifying quartets that satisfy the characteristic proposed in Figure 4 (Section 3.2). We employ the *Search for Double Collisions* technique proposed in Section 5.2 of [3] to efficiently perform the attack.

The core idea of the attack is to exploit the *Joint Generalized Zero-Difference Property* presented in Section 3.2 (specifically Equation (4)). Let (C_0, C_1, C_2, C_3) denote the quartet of 7-round ciphertexts corresponding to the plaintext quartet (P_0, P_1, P_2, P_3). According to this property, a right quartet simultaneously satisfies multiple relations regarding the zero-difference pattern of the inverse states (defined by μ). Specifically, the ciphertext pairs (C_0, C_2) and (C_1, C_3) share an identical zero-difference pattern on the inverse diagonals, and similarly, the pairs (C_0, C_3) and (C_1, C_2) share an identical zero-difference pattern on the inverse diagonals. Consequently, our strategy minimizes memory requirements by sequentially filtering for the $\binom{4}{2} = 6$ possible inverse diagonal combinations for Conditions C1 and C2 in Definition 5 (denoted by index set $\mathcal{I}$) and then verifying Conditions C3 and C4 (denoted by index set $\mathcal{J}$).

Step 1: Filtering Pairs for $\Delta x'$ (List N_1). First, we fix two indices out of the inverse diagonal index set $\{0, 1, 2, 3\}$ that are expected to be zero. Let $\mathcal{I}$ be this set of indices. We then search for pairs (C_a, C_b) that satisfy the following two conditions:

1. The input difference of the corresponding plaintext pair (P_a, P_b) is $\Delta x'$ (structurally satisfied by the data generation method).
2. $MC^{-1}(C_a \oplus C_b)$ is inactive in the inverse diagonals defined by $\mathcal{I}$, and active in the others. That is, $MC^{-1}(C_a \oplus C_b)[inv(k)] = 0$ for $k \in \mathcal{I}$ and $MC^{-1}(C_a \oplus C_b)[inv(k)] \neq 0$ for $k \notin \mathcal{I}$ (Probability: $2^{-64} - 2^{-96} \approx 2^{-64}$).

To perform this search efficiently, we utilize the data structure from Section 4.1. The collected $2^{126.2}$ plaintexts are divided into $2^{63.1}$ groups based on the values of the set A_1, which consists of the second and fourth diagonals (corresponding to the active diagonals of $\Delta x \oplus \Delta x'$). For each group, we construct a hash table using the 64-bit value of $MC^{-1}(C)$ at positions $\mathcal{I}$ as the key to detect collisions. The expected number of pairs stored in list N_1 is calculated as follows:

$$|N_1| \approx 2^{63.1} \times 2^{125.2} \times 2^{-64} = 2^{124.3}.$$

Here, $2^{63.1}$ represents the number of groups, and $2^{125.2}$ represents the number of possible pairs generatable within each group.

Step 2: Constructing Quartets (N_2). The list N_1 contains approximately $2^{124.3}$ pairs (P_a, P_b) with difference $\Delta x'$. These pairs are distributed across $2^{63.1}$ groups based on the values of the diagonals in A_1 (indices 1 and 3), with each group containing approximately $2^{61.2}$ pairs.

Our objective is to find two pairs, (P_a, P_b) and (P_c, P_d), both existing in N_1, that share the same values on the diagonals in A_0 (indices 0 and 2). To efficiently identify such pairs within N_1, we utilize a hash table indexed by the 128-bit value formed by concatenating the values of the diagonals in A_0 for both plaintexts. Specifically, for every pair in N_1, we compute the key $P_a[\text{dia}(A_0)] \parallel P_b[\text{dia}(A_0)]$ and insert the pair into the table.

We then look for collisions in the hash table. A collision occurs when two different pairs (P_a, P_b) and (P_c, P_d) from N_1 are mapped to the same key, satisfying:

$$P_a[\text{dia}(A_0)] = P_c[\text{dia}(A_0)] \quad \text{and} \quad P_b[\text{dia}(A_0)] = P_d[\text{dia}(A_0)].$$

This collision implies that the difference $P_a \oplus P_c$ is inactive (zero) on the diagonals in A_0. Furthermore, since the pairs originate from different groups defined by A_1, the following condition holds:

$$P_a[\text{dia}(A_1)] \neq P_c[\text{dia}(A_1)] \quad \text{and} \quad P_b[\text{dia}(A_1)] \neq P_d[\text{dia}(A_1)].$$

Consequently, the difference between these two plaintexts is active only on the diagonals in A_1, matching the required difference $\Delta x \oplus \Delta x'$. The identified pairs form a quartet (P_a, P_d, P_b, P_c) with the following difference relationships:

- $P_a \oplus P_b = \Delta x'$ and $P_c \oplus P_d = \Delta x'$ (since both pairs are from N_1).
- $P_a \oplus P_c = \Delta x \oplus \Delta x'$ (due to the collision on A_0).
- $P_b \oplus P_c = P_a \oplus \Delta x' \oplus P_c = (P_a \oplus P_c) \oplus \Delta x' = \Delta x$.
- $P_a \oplus P_d = P_a \oplus (P_c \oplus \Delta x') = \Delta x$.

The quartets satisfying these conditions are stored in list N_2. The total number of structurally possible quartets is $\binom{m}{2}^2 \approx 2^{250.4}$ (Section 4.1), and we estimate $|N_2|$ by considering the filtering probability over this space. A valid quartet requires that both pairs $(MC^{-1}(C_a), MC^{-1}(C_b))$ and $(MC^{-1}(C_c), MC^{-1}(C_d))$ are inactive on the two inverse diagonals in $\mathcal{I}$ (as defined in Step 1). Since the probability of this occurrence is 2^{-64} for each pair, the expected size of N_2 is:

$$|N_2| \approx 2^{250.4} \times 2^{-64} \times 2^{-64} = 2^{122.4}.$$

Remark (Consistency of $|N_2|$ via collision analysis). The estimate above can be independently verified by a bottom-up analysis based on the collision search performed on N_1. The list N_1 contains $|N_1| \approx 2^{124.3}$ pairs, and a quartet in N_2 is formed whenever two pairs from *different* A_1-groups collide on the hash key $P_a[\mathrm{dia}(A_0)] \,\|\, P_b[\mathrm{dia}(A_0)]$. Although this key occupies 128 bits, both components are drawn from the set A_0 of size $m = 2^{63.1}$. Consequently, the number of distinct keys is at most $\binom{m}{2} \approx 2^{125.2}$, not 2^{128}. Since there are approximately $|N_1|^2/2$ cross-group pair combinations, each colliding with probability $1/\binom{m}{2}$, the expected number of collisions is:

$$|N_2| \;\approx\; \frac{|N_1|^2}{2 \cdot \binom{m}{2}} \;=\; \frac{2^{248.6}}{2^{126.2}} \;=\; 2^{122.4},$$

which is consistent with the top-down estimate.

The procedure described so far is similar to the method proposed in [3].

Step 3: Verifying Remaining Conditions. Finally, for each quartet in N_2, we verify the remaining conditions (Conditions 3 and 4). We check whether the output differences of the pairs (C_a, C_c) and (C_b, C_d) are zero at the positions defined by the remaining inverse diagonal set $\mathcal{J}$, and non-zero at the other positions:

$$\begin{aligned} &MC^{-1}(C_a \oplus C_c)[\mathrm{inv}(k)] = 0, \quad MC^{-1}(C_b \oplus C_d)[\mathrm{inv}(k)] = 0, \quad \forall k \in \mathcal{J}, \\ &MC^{-1}(C_a \oplus C_c)[\mathrm{inv}(k)] \neq 0, \quad MC^{-1}(C_b \oplus C_d)[\mathrm{inv}(k)] \neq 0, \quad \forall k \notin \mathcal{J}. \end{aligned}$$

Strictly speaking, the probability of satisfying these conditions for a single pair is $2^{-64} - 2^{-96}$. However, since the subtraction term is negligible, we approximate this probability as 2^{-64}. Thus, the probability of satisfying these conditions for a specific combination is approximately $2^{-64} \times 2^{-64} = 2^{-128}$. Therefore, for a random permutation, considering all 6 column combinations, the total expected number of remaining quartets is:

$$6 \times 2^{122.4} \times 2^{-128} \approx 2^{-3},$$

which is negligible. In contrast, for 7-round AES, checking all 6 column combinations yields an expected value of one right quartet. The specific algorithm performing this process is Algorithm 1.

Algorithm 1 Distinguishing Attack on 7-Round AES

Require: Oracle access to the encryption function E.
Ensure: Decision whether E is 7-round AES or a Random Permutation.

Data Collection:
Ask for encryption of $2^{126.2}$ plaintexts constructed from sets $A_0, A_1 \subset \mathbb{F}_{2^8}^8$ ($|A_i| = 2^{63.1}$), where diagonals 0, 2 use A_0 and 1, 3 use A_1.
Group these plaintexts based on the values of the set A_1.
Attack Execution:
// Loop over all 6 inverse diagonal combinations
for all $\mathcal{I} \subset \{0, 1, 2, 3\}$ with $|\mathcal{I}| = 2$ **do**
    Let $\mathcal{J} = \{0, 1, 2, 3\} \setminus \mathcal{I}$.
    Initialize list $N_1 \leftarrow \emptyset$.
    // Step 1: Filter pairs for $\Delta x'$ (Conditions 1 & 2)
    for each group of plaintexts **do**
        Build hash table using key $K = MC^{-1}(C)[inv(k)]$ ($k \in \mathcal{I}$) for each ciphertext C.
        Store the plaintext pair (P_a, P_b) corresponding to the colliding ciphertexts (C_a, C_b) in N_1. ▷ Expected $|N_1| \approx 2^{124.3}$
    end for
    // Step 2: Construct Quartets
    Store in a list N_2 all quartets (C_a, C_d, C_b, C_c) from pairs in N_1 such that the corresponding plaintexts form (P_a, P_d, P_b, P_c) satisfying the differences. ▷ Expected $|N_2| \approx 2^{122.4}$
    // Step 3: Verify Remaining Conditions (Conditions 3 & 4)
    for each quartet $(C_a, C_d, C_b, C_c) \in N_2$ **do**
        if $MC^{-1}(C_a \oplus C_c)[inv(k)] = 0$ and $MC^{-1}(C_b \oplus C_d)[inv(k)] = 0$ for all $k \in \mathcal{J}$, and $MC^{-1}(C_a \oplus C_c)[inv(k)] \neq 0$ and $MC^{-1}(C_b \oplus C_d)[inv(k)] \neq 0$ for all $k \notin \mathcal{J}$ **then**
            return 7-round AES
        end if
    end for
end for
return Random Permutation

4.3 Complexity and Success Probability

In this section, we evaluate the complexity of the proposed 7-round distinguisher (Algorithm 1) and analyze its theoretical success probability using the binomial distribution model.

Complexity Analysis. The complexity of the attack is determined by the data collection and processing phases.

1. **Data Complexity:** As designed in Section 4.1, we require $2^{126.2}$ chosen plaintexts, constructed by combining two sets A_0 and A_1 where $|A_0| = |A_1| = 2^{63.1}$.
2. **Time Complexity:** The breakdown of the time complexity for the analysis phase is as follows:

- **Step 1:** Grouping $2^{126.2}$ ciphertexts based on A_1 values and detecting collisions via a hash table requires approximately $2^{126.2}$ memory accesses.
- **Step 2:** To construct quartets, we iterate through the list N_1 (size $2^{124.3}$) and perform a hash table lookup for each pair to find its counterpart. Thus, this step requires approximately $2^{124.3}$ memory accesses.
- **Step 3:** The final verification of Conditions 3 and 4 is performed on the quartets stored in N_2 (size $2^{122.4}$), adding approximately $2^{122.4}$ memory accesses.

The analysis phase (Steps 1 to 3) is repeated for each of the 6 inverse diagonal combinations (the $\binom{4}{2} = 6$ possible choices of disjoint index pairs $(\mathcal{I}, \mathcal{J})$ with $|\mathcal{I}| = |\mathcal{J}| = 2$). Since Step 1 dominates Steps 2 and 3, the analysis phase costs approximately $6 \times 2^{126.2}$ memory accesses. To express this cost in units of encryptions, we use the standard T-table implementation of AES, in which each round is evaluated by 16 table lookups (four lookups for each of the four output columns). One 7-round encryption therefore costs about $7 \times 16 = 112$ memory accesses, and the analysis phase amounts to $6 \times 2^{126.2}/112 \approx 2^{122.0}$ encryptions. This is negligible compared with the $2^{126.2}$ encryption queries required for data collection, and hence the total time complexity is $2^{126.2}$ encryptions.

3. **Memory Complexity:** While the algorithm requires storage for the intermediate lists N_1 ($2^{124.3}$) and N_2 ($2^{122.4}$), the total memory requirement is dominated by the initial phase. Storing the collected $2^{126.2}$ data blocks and maintaining the hash table requires approximately $2^{126.2}$ 128-bit blocks.

Success Probability Analysis. To derive the success probability, let X be the random variable representing the number of right quartets found in the entire search space. The total number of trials (total quartets) N is:

$$N \approx 2^{250.4}.$$

We aim to distinguish whether the oracle is 7-round AES (H_{AES}) or a random permutation (H_{Rand}). The probability p that a single random quartet satisfies all conditions under each hypothesis is given by:

- **Hypothesis H_{AES}:** By Corollary 3, a quartet that follows the characteristic illustrated in Figure 4 satisfies Conditions C1–C4 deterministically, and this occurs with probability $2^{-250.4}$. A quartet that does not follow the characteristic may still satisfy the conditions by chance, which would only increase the probability. Estimating that contribution, however, requires assuming that AES behaves like a random permutation on such quartets, which we cannot guarantee. We therefore omit it and use the lower bound

$$p_{AES} \geq 2^{-250.4}.$$

- **Hypothesis H_{Rand}:** The probability corresponds to satisfying the conditions purely by chance:

$$p_{Rand} \approx 2^{-253.4} = 0.125 \times 2^{-250.4}.$$

Assuming each trial is independent, the random variable X follows a binomial distribution $B(N, p)$. The decision rule for the distinguisher is defined as follows: if at least one quartet satisfying the conditions is found (i.e., $X \geq 1$), the algorithm returns 7-round AES. Otherwise, if no such quartet is found (i.e., $X = 0$), it returns random permutation. According to the definition of the binomial distribution, the probability of finding zero quartets ($X = 0$) is $(1-p)^N$. Based on this rule, we calculate the success probabilities for each hypothesis:

1. **Probability of correctly identifying AES:** This is the probability of finding at least one quartet under H_{AES} (which triggers the AES decision).

$$\begin{aligned}\Pr[X \geq 1 \mid H_{AES}] &= 1 - \Pr[X = 0 \mid H_{AES}] \\ &= 1 - (1 - p_{AES})^N \\ &\approx 1 - e^{-N \cdot p_{AES}} \\ &\geq 1 - e^{-1} \approx 0.632.\end{aligned}$$

2. **Probability of correctly identifying Random Permutation:** This is the probability of finding no quartets under H_{Rand} (which triggers the Random Permutation decision).

$$\begin{aligned}\Pr[X = 0 \mid H_{Rand}] &= (1 - p_{Rand})^N \\ &\approx e^{-N \cdot p_{Rand}} \\ &= e^{-0.125} \approx 0.882.\end{aligned}$$

Finally, assuming equal prior probabilities for both hypotheses, the average success probability $P_{success}$ is:

$$P_{success} = \frac{\Pr[X \geq 1 \mid H_{AES}] + \Pr[X = 0 \mid H_{Rand}]}{2} \geq \frac{0.632 + 0.882}{2} \approx 0.757.$$

Therefore, the proposed distinguisher distinguishes 7-round AES from a random permutation with a success probability of at least 75.7%. We remark that if the quartets that do not follow the characteristic were assumed to satisfy the conditions under AES with the same probability as in the random case, then $p_{AES} \approx 2^{-250.4} + 2^{-253.4} = 1.125 \times 2^{-250.4}$, which would give $\Pr[X \geq 1 \mid H_{AES}] \approx 1 - e^{-1.125} \approx 0.675$ and $P_{success} \approx 77.8\%$. Our claim does not rely on this assumption.

4.4 Discussion

Relationship to key recovery attacks. While the proposed characteristic can be extended to a 7-round key recovery attack similar to [3], the complexity of our distinguisher ($2^{126.2}$) is higher than existing key recovery results, offering no advantage. Extending to 8 rounds is also impractical, as verifying the inverse diagonal pattern at the 7th round output requires guessing the full last round key. Therefore, the primary contribution lies in establishing the first key-independent distinguisher, rather than improving key recovery efficiency.

Comparison with exhaustive search. A key-independent distinguisher and exhaustive key search operate in fundamentally different models. Exhaustive search is a key recovery attack that requires evaluating the key schedule for each candidate, and therefore does not fall into the key-independent model. In contrast, a key-independent distinguisher detects structural non-randomness inherent in the round function without discovering any key information. Consequently, comparing the complexity of these two approaches directly is not meaningful.

Theoretical significance. The security of a block cipher is evaluated based on the analysis of all known attacks, and the discovery of new techniques is significant even when the complexity is close to exhaustive search, as it establishes concrete security boundaries. Our result presents the first key-independent secret-key distinguisher reaching 7 rounds of AES, demonstrating that 7-round AES possesses a detectable structural deviation from a random permutation in the key-independent model. This directly relates to the fundamental notion of pseudorandomness formalized by Luby and Rackoff [19], which requires a secure block cipher to be indistinguishable from a random permutation.

5 Conclusion

In this paper, we proposed the first key-independent secret-key distinguisher for 7-round AES. We introduced the joint generalized zero-difference property, demonstrating that a quartet constructed from related differences satisfies three distinct generalized zero-difference properties simultaneously. Based on this property, we constructed a new 7-round differential characteristic and designed a distinguishing attack with data, time, and memory complexities of $2^{126.2}$ and a success probability of at least 75.7%. The joint property was further validated through experiments on small-scale AES.

Since the complexity of the proposed distinguisher is below 2^{128}, 7-round AES is distinguishable from a random permutation in the key-independent model. The core observation that a zero-sum structure yields multiple simultaneous relations is not specific to AES and may be applicable to other structural properties of SPN ciphers. Reducing the complexity, extending key-independent distinguishers to more rounds, and developing key recovery attacks that exploit the joint property remain important directions for future research.

Acknowledgments. This work was supported in part by the Institute of Information Communications Technology Planning and Evaluation (IITP) Grant funded by the Korea Government (MSIT) under Grant RS-2023-00229028, and in part by a Korea University Grant.

References

1. Bao, Z., Guo, J., List, E.: Extended truncated-differential distinguishers on round-reduced aes. IACR Transactions on Symmetric Cryptology pp. 197–261 (2020)

2. Bardeh, N.G.: A key-independent distinguisher for 6-round aes in an adaptive setting. Cryptology ePrint Archive (2019)
3. Bardeh, N.G., Rijmen, V.: New key-recovery attack on reduced-round aes. IACR Transactions on Symmetric Cryptology pp. 43–62 (2022)
4. Bardeh, N.G., Rønjom, S.: The exchange attack: How to distinguish six rounds of aes with chosen plaintexts. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 347–370. Springer (2019)
5. Bariant, A., Leurent, G.: Truncated boomerang attacks and application to aes-based ciphers. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 3–35. Springer (2023)
6. Boura, C., Canteaut, A., Coggia, D.: A general proof framework for recent aes distinguishers. IACR Transactions on Symmetric Cryptology **2019**(1), 170–191 (2019)
7. Boura, C., Lallemand, V., Naya-Plasencia, M., Suder, V.: Making the impossible possible. Journal of Cryptology **31**(1), 101–133 (2018)
8. Cid, C., Murphy, S., Robshaw, M.J.: Small scale variants of the aes. In: International Workshop on Fast Software Encryption. pp. 145–162. Springer (2005)
9. Daemen, J., Rijmen, V.: The design of Rijndael, vol. 2. Springer (2002)
10. Daemen, J., Rijmen, V.: New criteria for linear maps in aes-like ciphers. Cryptography and Communications **1**(1), 47–69 (2009)
11. Derbez, P., Fouque, P.A., Jean, J.: Improved key recovery attacks on reduced-round aes in the single-key setting. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 371–387. Springer (2013)
12. Dunkelman, O., Keller, N., Ronen, E., Shamir, A.: The retracing boomerang attack. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 280–309. Springer (2020)
13. Dunkelman, O., Keller, N., Shamir, A.: Improved single-key attacks on 8-round aes-192 and aes-256. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 158–176. Springer (2010)
14. Gilbert, H.: A simplified representation of aes. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 200–222. Springer (2014)
15. Grassi, L.: Mixture differential cryptanalysis: a new approach to distinguishers and attacks on round-reduced aes. IACR Transactions on Symmetric Cryptology pp. 133–160 (2018)
16. Grassi, L., Rechberger, C., Rønjom, S.: A new structural-differential property of 5-round aes. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 289–317. Springer (2017)
17. Leurent, G., Pernot, C.: New representations of the aes key schedule. Journal of Cryptology **38**(1), 2 (2025)
18. Lu, J., Dunkelman, O., Keller, N., Kim, J.: New impossible differential attacks on aes. In: International Conference on Cryptology in India. pp. 279–293. Springer (2008)
19. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. SIAM journal on computing **17**(2), 373–386 (1988)
20. Mondal, S.K., Rahman, M., Sarkar, S., Adhikari, A.: Revisiting yoyo tricks on aes. IACR Transactions on Symmetric Cryptology **2023**(4), 28–57 (2023)
21. Rahman, M., Saha, D., Paul, G.: Boomeyong: Embedding yoyo within boomerang and its applications to key recovery attacks on aes and pholkos. IACR Transactions on Symmetric Cryptology **2021**(3), 137–169 (2021)

22. Rønjom, S., Bardeh, N.G., Helleseth, T.: Yoyo tricks with aes. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 217–243. Springer (2017)
23. Shin, H., Kim, I., Kim, S., Kim, S., Hong, D., Sung, J., Hong, S.: Revisiting the multiple-of property for skinny: The exact computation of the number of right pairs. IEEE Access **12**, 35502–35511 (2024)
24. Shin, H., Kim, S., Seok, B., Lee, D., Hong, D., Sung, J., Hong, S.: New exchanged boomerang distinguishers for 5-round aes. Cryptology ePrint Archive (2025)
25. Shin, H., Kim, S., Seok, B., Lee, D., Hong, D., Sung, J., Hong, S.: Key-recovery attack on 5-round aes with multiple-of-8 property. IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences (2025)
26. Yan, X., Tan, L., Xu, H., Qi, W.: The boomerang chain distinguishers: New record for 6-round aes. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 301–329. Springer (2024)