A SLR of Modern AI-Driven SCA on NIST's PQC Standards

Edgard Nicéas Arcoverde Gusmão Lima[1,2][0000−0002−5241−8092], Fábio Wlademir Monteiro Maia[1,2][0000−0001−9733−1591], Tiago Alessandro Espínola Ferreira[3][0000−0002−2131−9825], and Danilo Monteiro Ribeiro[2][0000−0001−7393−729X]

[1] CISSA, CESAR, Recife, Brazil
enagl@cesar.org.br, fwmm@cesar.org.br
[2] CESAR School, Recife, Brazil
enagl@cesar.edu.br, fwmm@cesar.edu.br, dmr@cesar.school
[3] PPGIA, Universidade Federal Rural de Pernambuco (UFRPE), Recife, Brazil
tiago.espinola@ufrpe.br

Abstract. The imminent threat of quantum computing has catalyzed a global migration to Post-Quantum Cryptography (PQC), guided by the U.S. National Institute of Standards and Technology (NIST) standardization process. While these new cryptographic algorithms are designed for mathematical resilience against quantum adversaries, their physical implementations expose a critical vulnerability to Side-Channel Attacks (SCAs). This threat is fundamentally amplified by Artificial Intelligence (AI), which has transformed the adversary model from one requiring statistical analysis of numerous device interactions to one capable of potent, sometimes single-trace, key extraction. This paper presents a systematic literature review (SLR) mapping the collision of AI-driven SCAs with the first cohort of NIST-selected PQC standards: CRYSTALS-Kyber, CRYSTALS-Dilithium, SPHINCS+, Falcon,and HQC. From the synthesized literature, we provide a structured synthesis of attacks by linking specific AI methodologies to the unique vulnerabilities in each cryptographic family, such as the repetitive arithmetic in lattice-based schemes. Our analysis reveals three prevailing trends: the rise of practical single-trace and low-trace attacks; the growing ability of deep-learning models to bypass low-strength masking, shuffling, noise, and jitter; and the continued absence of standardized benchmarks, which forces studies to rely on ad-hoc labeling and heterogeneous evaluation setups.

Keywords: Post-Quantum Cryptography · PQC · Side-Channel Analysis · SCA · Artificial Intelligence · AI · Machine Learning · Deep Learning · NIST PQC Standardization · ML-KEM · CRYSTALS-Kyber · ML-DSA · CRYSTALS-Dilithium · FN-DSA · Falcon · SLH-DSA · SPHINCS+ · HQC.

1 Introduction

Modern digital infrastructure relies on public-key cryptography (PKC) schemes such as RSA and Elliptic Curve Cryptography (ECC), whose security stems from problems that are classically hard to solve [73]. Large-scale quantum computers running Shor's algorithm threaten this foundation by solving those problems in polynomial time [74]. To anticipate that "quantum day," the U.S. National Institute of Standards and Technology (NIST) launched an open competition to standardize quantum-resistant algorithms [51]. In 2022, NIST announced Kyber for public-key encryption and Dilithium for signatures, adding Falcon and SPHINCS+ as alternatives, while the code-based HQC remains under evaluation in the fourth round [3,22,4]. A summary of the algorithms considered in this review is provided in Table 7 of the online supplementary material [44].

Mathematical soundness alone does not guarantee security. PQC schemes must run on physical devices that leak information via side channels such as timing, power consumption, and electromagnetic emanations [41,14]. These leakages let adversaries bypass theoretical guarantees by recovering keys from real deployments [67]. NIST therefore required side-channel resilience throughout the PQC evaluation process [1], effectively crowdsourcing a global red-team exercise. Discoveries such as the Falcon-targeting attacks [35] were treated as valuable feedback that shaped countermeasures, while several candidates, summarized in Table 8 of the online supplementary material [44], were eliminated after practical implementation-level breaks. PQC's novel arithmetic widens the attack surface, making practical defenses a primary concern.

Side-channel analysis itself is being transformed by Artificial Intelligence (AI). Profiled attacks can be cast as supervised learning where traces are features and secret-dependent intermediates act as labels [15]. Deep learning—especially MLP- and CNN-based classifiers—learns leakage patterns directly from raw traces, eliminating much of the manual feature engineering required by classical statistics [45]. These models regularly defeat masking, shuffling, and other countermeasures by exploiting residual higher-order dependencies [60]. The integration of AI therefore lowers the barrier to executing high-efficacy attacks and must be assumed in any realistic threat model [67,30].

This systematic literature review (SLR) maps AI-driven side-channel attacks against Kyber, Dilithium, Falcon, SPHINCS+, and HQC. We focus on peer-reviewed or IACR ePrint-style works describing practical attacks that leverage Machine Learning, Deep Learning, or related modern AI techniques. Our research questions are:

- **RQ1:** What are the modern AI techniques most commonly used in side-channel attacks against PQC selected algorithms?
- **RQ2:** What are the targeted cryptographic operations and leakage sources in these attacks?
- **RQ3:** What are the attack scenarios, evaluation metrics, and experimental setups employed?
- **RQ4:** Which attack trends, research gaps, and challenges emerge from the synthesized data?

The remainder of this SLR is organized as follows. Section 2 provides background. Sections 3-4 follow the Kitchenham protocol [38,39] to detail the methodology and study selection. Sections 5 review related surveys. Section 6 report and synthesize the results with our taxonomy and analysis. Section 7 concludes and outlines future work.

2 Theoretical Background

Large-scale quantum computers pose a direct threat to modern public-key cryptography because Shor's algorithm can factor RSA moduli and solve elliptic-curve discrete logarithms in polynomial time [73,74,63,24], while Grover's algorithm quadratically accelerates brute-force search on symmetric primitives [25]. In contrast to symmetric schemes—which can be repaired by doubling key sizes—public-key systems such as RSA, Diffie-Hellman and ECC offer no viable parameter scale-up once a cryptanalytically relevant quantum computer becomes available. This creates an urgent need for post-quantum cryptography (PQC) [51], particularly for protecting long-term confidentiality against "harvest now, decrypt later" adversaries who can store encrypted data today and decrypt it once quantum capabilities mature [16,80,47]. More broadly, quantum algorithms fundamentally challenge assumptions underlying classical complexity theory [19,61,12,49], reinforcing the necessity of transitioning to cryptographic schemes whose security does not collapse under quantum computation.

Post-quantum cryptography (PQC) aims to provide cryptographic primitives secure against both classical and quantum adversaries [12,88]. Unlike quantum cryptography, PQC algorithms are classical and designed to run efficiently on existing processors [49,19]. Their security is based on mathematical problems for which no quantum algorithm is known to provide exponential speedups, in contrast to the integer factorization and discrete logarithm problems broken by Shor's algorithm. The main families of PQC include lattice-based schemes grounded in problems such as SVP, LWE and Module-LWE [8,9,3,54,67]; code-based ([59,58]) cryptography originating with McEliece's 1978 system [48]; hash-based signatures built from minimal and well-understood assumptions [7,12]; as well as multivariate and isogeny-based constructions [31,71]. These families offer complementary trade-offs in security reductions, key sizes and implementation characteristics, and collectively form the foundation of quantum-resistant cryptography.

To orchestrate a global and transparent transition to PQC, NIST launched its public standardization process in 2016 through an open call for proposals [51,50,76]. From 69 complete submissions entering Round 1 [2], candidates were evaluated under criteria encompassing security (including resistance to classical, quantum and side-channel attacks) [53,1], cost and performance across platforms [52], and practical considerations such as design simplicity, IP status and implementability. After multiple competitive rounds, NIST selected a diverse portfolio for standardization: ML-KEM/CRYSTALS-Kyber for key establishment [55,56,8]; ML-DSA/CRYSTALS-Dilithium as the primary signature

scheme [54,56,9]; FN-DSA/Falcon for compact signatures based on NTRU lattices [3,59,21,35]; and SLH-DSA/SPHINCS+ as a conservative, stateless hash-based alternative [57,56,7]. In March 2025, NIST further expanded the portfolio by selecting HQC as a code-based KEM [59,58,62,22],adding a non-lattice standard whose security relies on the hardness of quasi-cyclic syndrome decoding and whose design offers well-analyzed decryption-failure behavior and long-studied code-based security. These five algorithms—Kyber, Dilithium, Falcon, SPHINCS+, and HQC—constitute the standardized suite that will secure next-generation communication systems in the presence of quantum-capable adversaries.

While the theoretical security of a cryptographic primitive rests on its mathematical hardness, its practical security depends on the behavior of its physical implementation. Executing an algorithm on real hardware inevitably produces side effects—power consumption, electromagnetic (EM) radiation, and timing variations—that correlate with the secret-dependent internal state [89,41]. Side-channel attacks (SCA) exploit these leakages to recover long-term keys, often completely bypassing the underlying cryptographic assumptions.

Side-channel attacks fall into non-profiled and profiled categories. **Non-profiled attacks** operate directly on the target and typically require many traces to exploit weak leakage: Simple Power Analysis (SPA) inspects individual traces for coarse patterns such as square-and-multiply in RSA [40,41], while Differential Power Analysis (DPA) and Correlation Power Analysis (CPA) apply statistical tests by partitioning traces or correlating hypothetical leakage models with measurements [41,14,42]. **Profiled attacks** assume access to an identical device and proceed in a profiling and attack phase [15]; Template Attacks model leakage with multivariate Gaussians and can succeed with a single well-profiled trace [15,94]. **Profiled AI-based attacks** extend this paradigm by using ML and DL models to learn complex, non-linear leakage distributions, improving robustness to noise and enabling efficient key recovery against increasingly protected implementations. See Table 1.

Building on these attack capabilities, **countermeasures** against SCA operate at software and hardware levels and attempt either to remove secret-dependent behavior or to randomize it so that leakage becomes statistically ineffective. **Software protections** dominate PQC libraries due to portability: masking splits sensitive variables into random shares [60,70], shuffling randomizes operation order to induce desynchronization [60,82], constant-time coding removes timing and cache dependencies [40,11,5], and blinding decorrelates intermediate values [10,36]. **Hardware techniques** offer stronger leakage suppression, including balanced or equalizing logic and netlist-level augmentation [78,6], on-chip noise injection via CPNGs [79], and desynchronization through randomized clocking and jitter [72]. For PQC specifically, hardware-efficient shuffling schemes such as optimized Fisher–Yates for Kyber reduce alignment leakage at low cost [92], while adaptive FPGA mitigation using Dynamic Partial Reconfiguration (DPR) combined with DL-based monitoring disrupts attacker profiling and enables proactive protection [13]. Despite their effectiveness, these counter-

Table 1. Modern AI-based models applied in SCAs against PQC algorithms.

Classical	**Random Forest** [68,64] Ensemble of decision trees that models complex non-linear decision boundaries.
Reinforcement / Decision	**Deep Q-Network (DQN)** [87] Neural agent that learns optimal actions through value-based reinforcement learning.
Deep / Neural	**Convolutional Neural Network (CNN)** [69,81,28,77,17,65,66] Learns hierarchical spatial features from grid-structured data.
	Recurrent Neural Network (RNN) [23] Captures temporal dependencies in sequential data.
	Multi-Layer Perceptron (MLP) [33,20,23,29,32,34,83,93,69,81,26,46,85,84,37,18] Feedforward network that learns discriminative representations via dense layers.
	Graph Neural Network (GNN) [75] Learns from data represented as graphs by propagating information across nodes.
	Large Language Model (LLM) [96,75] Transformer model trained on large corpora for language understanding and generation.
	Deep Neural Network (DNN) [86] General multi-layer architecture that extracts increasingly abstract features.

measures are not free—many incur notable performance, area, randomness, or implementation overheads—and recent research shows that some classical protections can be partially or fully defeated by modern AI-based SCAs, which learn to bypass masking, undo shuffling, or exploit residual microarchitectural leakage [60,20,17,32].

3 Review Methodology

This chapter delineates the systematic methodology employed to identify, select, appraise, and synthesize the existing body of literature concerning the application of Artificial Intelligence (AI) techniques to conduct Side-Channel Attacks (SCAs) against the cryptographic algorithms standardized or finalized by the NIST Post-Quantum Cryptography (PQC) project. The protocol is designed to be transparent, replicable, and rigorous, following the guidelines for systematic literature reviews proposed by Kitchenham and Charters [38,39], ensuring the resulting review is comprehensive and unbiased. The review process is structured into three primary phases: planning, conducting, and reporting [38]. This section details the first two phases.

3.1 Search Strategy

A multi-stage search strategy was executed to ensure a comprehensive and unbiased collection of relevant primary studies. The search was concluded on the cutoff date of August 16, 2025.

Database Selection The selection of information sources is critical for capturing high-quality, peer-reviewed research in this fast-moving field. The following digital libraries and archives were chosen based on their prevalence as publication

venues for top-tier cryptography, hardware security, and computer security research: **IEEE Xplore:**, **ACM Digital Library:**, **Engineering Village:**, **ScienceDirect & Scopus:**, **SpringerLink (including LNCS):**, **IACR ePrint Archive:**. Further details are provided in Subsection *Databases Used* of the online supplementary material [44].

Search String The Population, Intervention and Context keywords from the PICOC [91] framework (see Table 1 in the online supplementary material [44]) were combined using Boolean operators (AND, OR) to form the final search query depicted in Figure 1.

Fig. 1. Base search string.

```
("post-quantum cryptography" OR "PQC" OR "quantum-resistant"
    ↪ OR "Kyber" OR "Dilithium" OR "SPHINCS+" OR
    ↪ "F\textsc{alcon}" OR "Hamming-Quasi-Cyclic" OR
    ↪ "ML-KEM" OR "ML-DSA" OR "SLH-DSA" OR "FN-DSA" OR
    ↪ "HQC")
AND
("artificial intelligence" OR "AI" OR "machine learning" OR
    ↪ "ML" OR "deep learning" OR "DL" OR "neural network"
    ↪ OR "CNN" OR "reinforcement learning" OR "generative
    ↪ AI" OR "transformer" OR "generative adversarial
    ↪ network" OR "GAN" OR "recurrent neural network" OR
    ↪ "RNN" OR "support vector machine" OR "SVM" OR
    ↪ "gradient boosting" OR "random forest" OR "graph
    ↪ neural network" OR "GNN" OR "computational
    ↪ intelligence" OR "data mining" OR "predictive
    ↪ modeling")
AND
("side-channel" OR "side channel" OR "SCA" OR "power
    ↪ analysis" OR "timing attack" OR "electromagnetic
    ↪ analysis" )
```

To enhance the precision of the search and ensure the relevance of the retrieved studies, the query was executed against the title, abstract, and keyword fields of the articles within the selected databases. The final string was designed to be adaptable to the specific syntax of each database.

3.2 Gold Set Validation

To validate the effectiveness of the search strategy, a "Gold Set" [95] of known relevant studies was used. The complete list of Gold Set papers is provided in Table 2 of the online supplementary material [44].

3.3 Study Selection Criteria and Process

The screening process was conducted in two main phases by two independent reviewers: a title and abstract screening, followed by a full-text screening. Any disagreements were resolved through discussion to reach a consensus. The criteria is detailed in Table 2.

Table 2. Inclusion and Exclusion Criteria

Criteria	Inclusion (Must Have)	Exclusion (Must Not Have)
Subject Matter	The study must describe a side-channel attack that uses an AI/ML/DL model as the primary distinguisher or analysis tool.	The study describes a conventional SCA (e.g., CPA, TA) without an AI/ML component. The study describes a different type of attack (e.g., fault injection, cryptanalysis). The study is about AI security in general, not SCA.
Target Algorithm	The target of the attack must be an implementation of a NIST PQC Round 4 finalist or a standardized algorithm (Kyber, Dilithium, SPHINCS+, Falcon, HQC, etc.).	The target is exclusively a pre-quantum algorithm (e.g., AES, RSA) UNLESS the technique is explicitly presented as a foundational method later applied to PQC in other works.
Publication Type	The study must be a full-text, peer-reviewed research paper (journal, conference, or workshop) or a relevant technical pre-print (e.g., from IACR ePrint).	The publication is a poster, abstract, presentation slide deck, editorial, or summary without sufficient technical detail.
Language	The study must be written in English.	The study is written in any language other than English.
Availability	The full text of the study must be accessible.	The full text cannot be obtained.

3.4 Quality Assessment Checklist

Each included study was evaluated for rigor, validity, and risk of bias following established SLR guidelines [38]. The criteria used appear in Table 3. A three-point scale was applied to each item—**1.0 (Yes)**, **0.5 (Partially)**, and **0.0 (No)**—to provide a quantitative indication of methodological completeness [39]. No study was excluded based on its score; the assessment is used solely to contextualize and interpret the strength of the evidence across the reviewed literature.

3.5 Data Extraction Form

A data extraction form was designed to systematically collect relevant information from each included study. This structured approach, detailed in Table 4, ensures consistency and facilitates the aggregation and comparison of data across studies to answer the research questions. Systematically documenting study characteristics and findings in this manner is a crucial step that supports clarity, precise reporting, and the ability to replicate the examination, as mandated by SLR guidelines [38].

Table 3. Quality Assessment Checklist

QA#	Quality Assessment Question
QA1	Are the research objectives and attack goals clearly stated?
QA2	Is the target PQC algorithm, implementation (software/hardware), and platform described in sufficient detail to allow for replication?
QA3	Is the side-channel data acquisition setup (e.g., equipment, sampling rate, number of traces) clearly documented?
QA4	Are the AI/ML model architecture, hyperparameters, and training process adequately specified?
QA5	Are the results evaluated using standard, appropriate metrics for both SCA (e.g., Guessing Entropy, Success Rate) and ML (e.g., accuracy, loss)?
QA6	Does the study discuss its limitations and potential threats to the validity of its findings?
QA7	Is the source code or dataset made publicly available to support reproducibility?

4 Study Selection

This section presents the findings obtained by executing the review protocol detailed in the previous section. We begin by reporting the results of the literature search and selection process, followed by a summary of the quality assessment. Then we synthesize the key findings from the selected primary studies, focusing on the most relevant trends and insights regarding AI-based side-channel attacks against post-quantum cryptographic algorithms. The section concludes with a bibliometric analysis derived from the data systematically collected from the final set of primary studies.

4.1 Literature Search and Initial Results

The initial phase of the systematic literature review consisted of querying seven electronic databases: *IEEE Xplore*, *ACM Digital Library*, *Engineering Village*, *ScienceDirect*, *Springer Link*, the *IACR ePrint Archive*, and *Scopus*. This comprehensive search strategy yielded a total of **1,495 articles**. A subsequent de-duplication process eliminated 157 duplicate records, resulting in a set of **1,338 unique primary studies**.

These studies were then assessed according to the predefined inclusion and exclusion criteria. As shown in Table 5, the majority of records originated from *Springer Link* (1,000), followed by *ScienceDirect* (175) and *Engineering Village* (89). After de-duplication, *Springer Link* contributed 990 unique entries, while *IEEE Xplore* and *ScienceDirect* accounted for 60 and 175 unique entries, respectively.

During the screening process, **nine articles were excluded due to lack of access**, leaving **1,329 accessible papers** for further evaluation. Of these, **27 were identified as potentially relevant** and subsequently confirmed as the final set of **27 selected studies** for detailed analysis.

4.2 Quality Assessment Results

The Quality Assessment Checklist defined in the methodology (Table 3) have been applied to all 27 selected primary studies. The purpose of this step was not

Table 4. Data Extraction Form

Category	Data Field	Description
Bibliographic Data	Study ID	Unique identifier for the study within this SLR.
	Authors & Year	Authors and publication year.
	Title & Venue	Full title and publication venue (e.g., TCHES 2025).
PQC Target Details	PQC Algorithm	Kyber, Dilithium, SPHINCS+, Falcon, HQC.
	Security Level	e.g., NIST Level 1 (Kyber-512), Level 3, Level 5.
	Attack Vector	e.g. Re-encrption of Decriptation, Fujisaki–Okamoto (FO) transformation, etc...
	Implementation Type	FPGA, ASIC, RTL, Software.
	Target Platform	8-bit AVR, ARM Cortex-M4, Artix-7 FPGA.
	Target Platform Details	The specific processor, board or setup.
	Countermeasures	None, Masking, Shuffling, Hiding, Constant-Time, Anti-tampering or a combination.
	Countermeasures Details	e.g. Masking order.
Side-Channel Context	Side-Channel Type	Power, Electromagnetic (EM), Timing.
	Fault Injection	Yes or No.
	Tool	Setup. e.g. ChipWhisperer, Ociloscope, etc...
AI/ML Intervention	Model Type	e.g., CNN, MLP, SVM, RNN.
	Model Details	Key details (e.g., number of layers, filter sizes for CNNs).
	Profiling Traces	Number of traces used for training the model.
	Attack Traces	Number of traces used to perform the key recovery attack.
Attack Outcome	Primary Metric	Guessing Entropy (GE), Success Rate (SR), Key Rank.
	Result	The reported value of the primary metric (e.g., "GE drops to 0 after 50 traces").
	Key Findings	A qualitative summary of the main conclusions drawn by the authors.

to exclude low-quality papers, but to provide a systematic appraisal of the rigor of the included research and to inform the data synthesis phase.

The complete quality assessment results (Table 9) are available in the online supplementary material [44]. The results show that a large majority of papers (23 out of 27) scored 5.0 or higher. The most common reason for a score below 7 was the missing of source code or dataset to allow replication (QA7).

5 Related Works

The paper by Hernández-Álvarez et al. [27] surveys recent AI-based side-channel attacks targeting the NIST-standardized PQC schemes CRYSTALS-Kyber and CRYSTALS-Dilithium. It reviews the cryptographic structures of both algorithms, identifies specific implementation components susceptible to leakage (e.g., Kyber's Encode and re-encryption steps, Dilithium's NTT and bit-unpacking), and summarizes how machine-learning models—primarily MLPs and clustering techniques—are used to recover messages or secret keys from power or EM traces. The work provides algorithmic details, attack configurations, and comparative performance results across several profiling attacks.

Table 5. Summary of the Study Selection Process by Source

Source	Found	Unique*	No Access	Potentially Relevant	Selected
IEEE Xplore	61	60	0	10	10
ACM Digital Library	5	3	0	1	1
Engineering Village	89	54	4	3	3
ScienceDirect	175	175	0	1	1
Springer Link	1000	990	5	2	2
IACR ePrint Archive	56	36	0	7	7
Scopus	109	20	0	3	3
Total	**1495**	**1338**	**9**	**27**	**27**

(*) After de-duplication process Cutoff date of August 2025

The work by Li [43] provides an extensive review of attacks targeting CRYSTALS-Kyber, covering common cryptanalytic attacks, side-channel attacks, SCA-assisted chosen-ciphertext attacks, and fault-injection techniques. It summarizes the feasibility of timing, SASCA, message-encoding attacks, deep-learning-based attacks, LDPC-based frameworks, EM-based plaintext-checking oracles, CPA on NTT operations, and fault attacks such as Roulette and error-tolerant key recovery. The survey highlights that several attacks—including DL-based single-trace methods and fault-injection approaches—remain effective even against protected or masked implementations, underscoring the evolving threat landscape around Kyber.

Ravi et al. [67] present a comprehensive survey of side-channel and fault-injection attacks on lattice-based post-quantum schemes, focusing primarily on CRYSTALS-Kyber and CRYSTALS-Dilithium. The work categorizes timing attacks, power and electromagnetic analysis, chosen-ciphertext side-channel attacks, and a range of fault-based key-recovery methods, while also reporting new experimental results illustrating the feasibility of practical attacks on constrained hardware. The survey highlights how both classical and modern techniques—spanning DPA/CPA, template attacks, and fault-based perturbations—can compromise implementations of leading PQC algorithms when protections are insufficient.

Despite their breadth, these surveys do not answer our research questions because they lack a systematic, AI-focused synthesis: they do not categorize modern AI techniques used in PQC SCAs (RQ1), generalize targeted operations and leakage sources (RQ2), or extract comparable attack scenarios, metrics, and setups (RQ3), nor do they identify cross-paper trends or research gaps (RQ4). Moreover, all three works date from 2023-2024, and research in AI-based PQC side-channel analysis has been growing at a pace that effectively doubles each year, making their coverage increasingly outdated.

6 Synthesis of Key Findings

6.1 RQ1: What are the modern AI techniques most commonly used in side-channel attacks against PQC selected algorithms?

Figure 2 show that the overwhelming majority of AI-based PQC side-channel attacks rely on Multi-Layer Perceptrons (MLPs) and Convolutional Neural Networks (CNNs), which together constitute the core of practical deep-learning SCAs. **MLP**s appear most frequently and remain the default architecture for message and key recovery across Kyber and Dilithium due to their simplicity, efficiency, and strong performance on segmented or weakly structured leakage. **CNN**s form the second major class and achieve some of the best distinguishers in scenarios where spatially structured leakage is present. Notably, Rezaeezade et al. [69] report that a CNN outperforms an MLP under identical blind-SCA conditions; however, such isolated results cannot be generalized because meaningful model-to-model comparison is hindered by heterogeneous leakage conditions, countermeasure settings, labeling strategies, and evaluation methodologies across the literature. Beyond these two dominant families, other AI models—including **Random Forest**s, **DNN**s, **RNN**s, **GNN**s, and **DQN**-based reinforcement learners—appear far less frequently and typically serve specialized roles such as segmentation, oracle construction, or leakage assessment rather than full key recovery. Overall, current evidence indicates a clear consolidation around deep feed-forward and convolutional models as the most commonly used AI techniques in PQC-focused SCAs. Finally, Zhou et al. [96] demonstrate that, with carefully engineered expert prompts, even a general-purpose **LLM** (ChatGPT) can perform operation-type classification well enough to enable full private-key recovery. Table 1 complements this synthesis by detailing each model category alongside its associated references.

Fig. 2. Distribution of AI-based attacks.

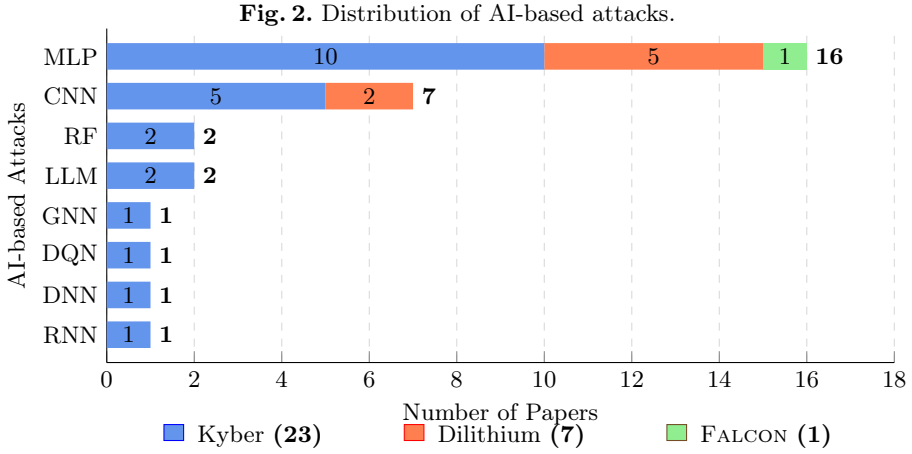

6.2 RQ2: What are the targeted cryptographic operations and leakage sources in these attacks?

Across the surveyed literature, AI-based side-channel attacks concentrate overwhelmingly on decapsulation for Kyber and signing for Dilithium, reflecting the operations where secret-dependent arithmetic is most exposed.

For **Kyber**, attacks commonly target the message or coefficient encoding pipeline—including polytomsg, masked message encoding, and FO-based re-encryption—as well as NTT/INTT-related arithmetic, coefficient multiplications, and PRF evaluations used for re-encryption or oracle construction [83,23,29,32,34,28,81,17,77]. The labels used for learning typically encode message bits (e.g., HD(m[i-1],m[i]) or HW(msg byte)), secret-key coefficients, PRF input classes, or oracle outcomes distinguishing valid vs. random ciphertexts [33,29,34,69,77,93]. These operations frequently leak through structured patterns in modular reductions, byte-wise encodings, or reference-versus-random comparisons, enabling both message recovery and full secret-key extraction [93,32,28,23,83,34].

For **Dilithium**, the primary targets include NTT, INTT, Montgomery reduction, sampling, addition/rounding, and unpacking of t_0 and secret polynomials during signing [26,46,85,84,37,65,66]. Correspondingly, labels often represent coefficient values (e.g., $-2..2$), LSBs of t_0, s_1/s_2 coefficients, or sparsity/zero-nonzero patterns of sampled vectors [85,84,26,46]. These operations leak coefficient magnitudes, polynomial structure, or sign information that AI models can learn from segmented traces [26,85,65,66]. Leakage sources span power, EM, FPGA-level switching activity, RTL simulations, and oracle outputs, covering both unmasked and first-order masked implementations [37,75,77].

For **HQC**, the literature contains only analytical discussion rather than a realized AI-based attack. Works studying FO-based re-encryption leakage (e.g., PRF execution during equality checking) note that HQC uses SHAKE during re-encryption, and that leakage in this component could—in principle—enable a plaintext-checking oracle [81].

For **Falcon**, AI models target leakage from CDT sampling, using labels that represent sampling outcomes or comparison results, enabling high-accuracy leakage classification; however, no key recovery has been achieved [18].

6.3 RQ3: What are the attack scenarios, evaluation metrics, and experimental setups employed?

Across the surveyed literature, AI-based side-channel attacks on NIST-selected PQC schemes fall into a small set of recurring scenarios, with clear patterns in how traces, labels, and auxiliary oracles are used.

Attack Scenarios

– **Profiling SCA (dominant setting).** The attacker collects labeled traces on a profiling device and applies the trained model to a target device. This setting underlies most Kyber and Dilithium attacks,

including message recovery, coefficient recovery, and full key extraction [26,46,83,85,84,34,33,28,23].
- **Blind or device-agnostic SCA.** Attacks where labels are absent (e.g., MC-labeling via GMMs) or where training and target devices differ. Evaluation is performed using guessing entropy (GE) or rank metrics, with CNNs outperforming MLPs in blind Kyber SCA [69].
- **Oracle-assisted scenarios.** Leakage from PRF-based re-encryption enables binary or multi-valued plaintext-checking oracles that drive adaptive decryption [81,77]. Belief-propagation analysis [64] formalizes how such oracles propagate constraints in CCA-secure KEMs.
- **Fault-assisted SCA.** Voltage glitches and desynchronization faults bypass protections such as shuffling, enabling single-trace attacks on Kyber decapsulation [32].

Evaluation Metrics

- **Bit-level accuracy.** Used when predicting labels such as $\mathrm{HD}(m[i-1], m[i])$, $\mathrm{HW}(\mathrm{msg})$, or sampled bits from arithmetic stages [33,32].
- **Coefficient-recovery rates.** Accuracy over secret polynomials (s_0, s_1, s_2) or over unpacked coefficients (e.g., t_0 LSBs) during Dilithium and Kyber signing/decapsulation [26,85,84].
- **Message-recovery probability.** Particularly relevant in Kyber decapsulation, sometimes boosted by enumeration or majority voting (e.g., 0.887→0.969 under 32/64-bit enumeration) [32].
- **Full key-recovery success.** Reported rates range from single-trace partial recovery (9% for Dilithium-2 under restrictive assumptions) to 100% recovery with enough traces or signatures [46,83,37].
- **Guessing Entropy (GE).** Used in blind-SCA; CNNs achieve GE=2.01 vs. GE=9.2 for MLP under identical labeling noise [69].
- **Oracle reliability and query complexity.** Evaluated via PRF-classifier accuracy, ciphertext-validity prediction accuracy, and the number of required oracle queries, e.g., multi-valued PC oracles in Tanaka et al. [77] and analysis of contraction propagation in [64].

Experimental Setups

- **Target Platforms.** Common setups include ARM Cortex-M4 microcontrollers [26,46,83,85,84,28,69], FPGA implementations [34,33], RTL-level simulations and power models [75], and full EM benches [17,81]. ChipWhisperer is widely used for synchronized trace capture [32,23].
- **Trace Acquisition.** Profiling involves from 1k to over 500k traces, with segmentation yielding 255k–2.56M labeled samples for training [33,32,23]. Some works rely on **oscilloscopes or EM probes** (hundreds of MS/s to multi-GS/s) to capture high-frequency leakage that ChipWhisperer cannot resolve [18,29,81,17].

– **Trace Windowing and Sampling.** MLPs typically use 400–800-sample windows, whereas CNNs operate on 600-sample windows or full-length traces (e.g., 100k EM samples) [28,17].
– **Training Methodology.** Most works train neural distinguishers using standard deep-learning optimizers such as Adam [29,37,66,18,83,65], Nadam [23,34,20,32,85,33], or RMSprop [23,28], depending on the architecture and task.

Countermeasure Context

– **Masking.** First- to fifth-order masking is evaluated across Kyber and Dilithium, with attacks demonstrating partial or full key recovery even in masked settings [20,33,29,23,37].
– **Hiding techniques.** Shuffling, random delays, artificial noise, and clock jitter are commonly assessed; accuracy may drop from 88% to below 1% under strong jitter [29].
– **Structural protections.** These include constant-time implementations, anti-tamper covers used in EM settings [17], TI masking in hardware [81], and RTL-level leakage mitigation [75].

Overall, AI-based SCAs are conducted under profiling-heavy, high-trace-volume experimental setups, evaluated through accuracy-, rank-, and oracle-centric metrics, and deployed on diverse hardware platforms ranging from embedded microcontrollers to EM benches and RTL simulations. Despite heterogeneous conditions, a consistent methodology emerges: fine-grained segmentation, supervised learning pipelines, and multi-metric reporting tied to message recovery, coefficient correctness, and end-to-end key extraction.

6.4 RQ4: Which attack trends, research gaps, and challenges emerge from the synthesized data?

The surveyed literature reveals a clear upward trend in AI-based side-channel attacks on post-quantum cryptographic schemes. As illustrated in Figure 3, publication activity remained low in 2020-2022 but increased sharply from 2023 onward, reaching a peak of ten works in 2024. Despite 2025 being incomplete, the current publication rate of approximately 0.75 works per month already exceeds earlier years by a wide margin. This growth reflects both increasing maturity in attack methodology and expanding interest in evaluating the practical resilience of lattice-based schemes beyond their algorithmic security guarantees

A second trend is the consolidation around deep-learning profiling attacks—particularly MLPs and CNNs—which continue to dominate successful key-recovery demonstrations. These models support single-trace or low-trace adversaries, remain effective even under masking, shuffling, noise, jitter, and anti-tamper covers, and exploit structured leakage in encoding pipelines, PRF-based re-encryption, NTT/INTT arithmetic, and sampling. Simultaneously, blind-SCA and device-agnostic settings are becoming more prominent, with noisy-label

Fig. 3. Number of Publications per Year with Monthly Average

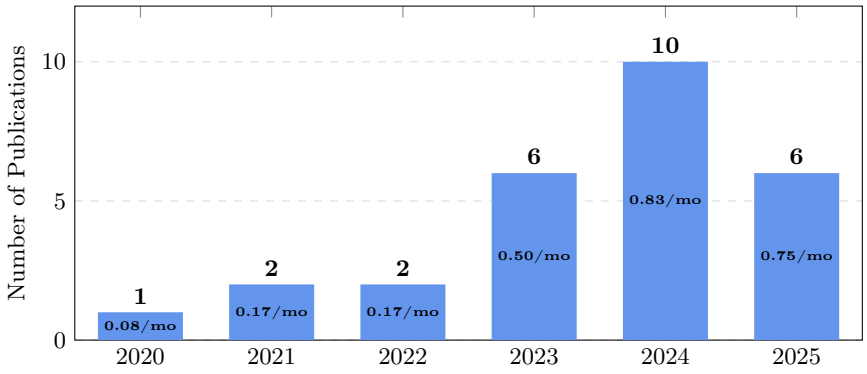

*includes publications up to August 2025.

pipelines, GMM-based MC-labeling, ciphertext-knowledge augmentation, and reinforcement-learning-based segmentation reducing the need for precise labels or manual trace engineering.

A recurring limitation across the surveyed literature is that, although several works do evaluate attacks against *combinations* of countermeasures—most commonly first-order masking together with shuffling, artificial noise, or clock jitter—these combinations remain relatively lightweight. No study analyzes higher-order masking combined with stronger hiding strategies, nor do they consider more entropy-rich variants of these protections (e.g., polynomial-level coefficient shuffling, high-variance randomized delays, or diversified execution paths). Existing attacks therefore demonstrate that AI models can bypass *specific* low-strength combinations, but they do not provide evidence regarding the robustness of such attacks when countermeasures are jointly deployed at higher complexity or randomness.

Only a single work to date demonstrates a combined AI-assisted side-channel and fault-injection attack. Jendral et al. [32] show that a voltage-glitch can bypass Fisher-Yates shuffling in masked Kyber decapsulation, simplifying the leakage sufficiently for an MLP to achieve near single-trace message recovery. However, no study performs a systematic evaluation of how neural models behave in joint SCA+FI settings—such as robustness to fault noise, parameter sensitivity, or generalization across devices. As a result, the broader design space of fault-assisted deep-learning attacks remains largely unexplored, despite clear potential to reduce trace complexity or defeat stronger countermeasure combinations.

Another noticeable gap is the near-complete absence of hybrid *DSP–AI* pipelines. Across all surveyed works, preprocessing remains minimal—typically limited to segmentation, normalization, and occasional alignment—while more structured DSP techniques such as Fourier or wavelet transforms, band-pass filtering, jitter compensation, or spectral-domain denoising are virtually never

employed. Although this reflects the prevailing assumption that CNNs and MLPs can learn temporal and spectral features directly from raw traces, it also leaves open whether DSP-enhanced features could improve robustness under high-noise EM settings, cross-device variability, or higher-order masked implementations. At present, the effectiveness of DSP-driven preprocessing for PQC-focused SCAs remains largely unexplored.

A further challenge concerns the *lack of meaningful, controlled comparisons between AI models*. Although some papers report local superiority (e.g., CNN outperforming MLP under identical conditions in Rezaeezade et al. [69]), such findings cannot be generalized because studies differ widely in leakage sources, countermeasure configurations, trace segmentation, labeling strategies, preprocessing pipelines, and evaluation metrics. Even within a single scheme, variations in architecture depth, optimizer selection, and profiling budgets make it impossible to derive reliable cross-paper rankings of model performance. The absence of standardized benchmarks and unified evaluation protocols therefore represents a major limitation, preventing principled assessments of which AI techniques are genuinely more robust, sample-efficient, or countermeasure-resistant.

Despite these advances, several research gaps persist. Cross-device generalization remains fragile, with most attacks degrading significantly under hardware, environmental, or process variations. Evaluations of robust countermeasures—including higher-order masking for Dilithium, threshold implementations, domain-oriented masking, and hardened EM designs—are still sparse. Falcon and HQC receive limited practical analysis, with existing work focusing mostly on CDT leakage characterization or analytical PRF-based oracle reasoning rather than demonstrated key recovery. Heterogeneous setups, inconsistent labeling strategies, and differing accuracy or GE metrics further complicate meaningful comparisons across studies.

A notable cross-cutting gap in the surveyed literature is the absence of hybrid or ensemble AI pipelines. Despite the diversity of models appearing across individual works—MLPs, CNNs, RFs, GNNs, DQNs, and even LLMs—every practical attack adopts a single-model architecture, with no combination of complementary techniques. No paper integrates CNN feature extraction with MLP or RF classifiers, uses GNN-based PoI selection prior to deep learning, couples reinforcement-learning segmentation with downstream neural inference, or blends DSP-based filtering with learned feature hierarchies. This stands in sharp contrast to other ML-heavy domains (e.g., computer vision, speech, anomaly detection), where ensembles and multi-stage architectures significantly improve robustness and cross-domain generalization. Given the nonstationary, noisy, device-dependent nature of PQC leakage—and the challenges posed by masking, shuffling, jitter, and anti-tamper protections—such hybrid approaches are natural candidates for improved resilience, yet remain unexplored. This lack of model integration constitutes a clear research opportunity for more powerful, noise-tolerant, and countermeasure-resistant PQC side-channel analysis.

6.5 Threats to Validity and Limitations

As in any systematic mapping study, this work is inherently exposed to certain limitations and potential validity threats. To address these, we followed guidance from relevant methodological literature [90,97]

Search and Selection Bias: As described in Section 3, we mitigated this risk through a multi-database search, PICOC-based query construction (Figure 1), and explicit inclusion/exclusion criteria. To maximize recall, we deliberately incorporated synonyms and acronyms into the search string, accepting the trade-off of retrieving a larger number of potentially unrelated records, which were subsequently filtered during the screening phase. Nevertheless, some studies—particularly those in less accessible venues or using unconventional terminology—may not have been captured.

Data Extraction Bias: The classification of studies and interpretation of extracted evidence may be subject to researcher judgment. To mitigate this risk, data extraction was conducted independently by two reviewers, with disagreements resolved through discussion until reach a consensus as described in Section 3.3.

Publication Bias: Our study only included published research, which may favor positive results and failed attempts or negative results. This limitation is well recognized in systematic reviews [97] and should be considered when interpreting the conclusions. Future work could mitigate this bias by incorporating gray literature and explicitly searching for studies reporting effective countermeasures against AI-based SCAs. Although countermeasures are closely related to this mapping, they were not the primary focus of the present study nor directly addressed by the research questions.

Paper/database inaccessible: Some papers were not accessible due to paywalls or other restrictions. A total of 9 papers (Table 5), published in 2025, were inaccessible. While the abstracts showed potential relevance, they where excluded by our inclusion/exclusion criteria (Figure 2). This may have led to the omission of recent developments in the field, and thus represents a limitation of this study.

Restricted Time Span: The search process was concluded in August 2025, while the manuscript is being published in February 2026. Given the rapid evolution of AI-based side-channel attacks against PQC implementations, relevant studies published after the cutoff date may not have been captured. This temporal limitation may affect the external validity of the review by restricting the completeness and contemporaneity of the mapped evidence.

Scope and Generalizability Limitation: This mapping study focused exclusively on AI-based side-channel attacks targeting NIST PQC algorithms and did not include comparisons with classical attacks (e.g., DPA, CPA) or non-PQC cryptographic schemes. While this delimitation was intentional and aligned with the research questions, it may limit the external validity of the findings, as conclusions cannot be directly generalized to the broader landscape of traditional side-channel analysis techniques.

7 Conclusion

This systematic review synthesized the rapidly expanding body of research on AI-based side-channel attacks (SCAs) against NIST-selected post-quantum cryptographic schemes, with a focus on CRYSTALS-Kyber and CRYSTALS-Dilithium. Our analysis shows that the field is accelerating at an exceptional rate, with the number of publications roughly doubling each year since 2021. Despite differences in methodology, datasets, and countermeasure assumptions, several overarching patterns clearly emerge.

First, modern AI techniques used in PQC SCAs are dominated by deep-learning architectures—most notably MLPs and CNNs—which consistently achieve the strongest message- and key-recovery results in both profiling and single-trace attack settings. Although alternative models such as Random Forests, GNNs, DQNs, and LLMs have appeared, their roles remain specialized, and no study employs hybrid or ensemble pipelines that combine complementary strengths across model families.

Second, the targeted leakages are distributed across a diversity of algorithmic sub-parts. In Kyber and Dilithium, attacks exploit sensitive operations throughout decapsulation or signing, including message-encoding routines, NTT/INTT arithmetic, PRF-based re-encryption, sampling, and coefficient unpacking. Labels derived from these operations—such as message bits, coefficient values, t_0 LSBs, oracle decisions, and sparsity indicators—enable high-accuracy inference even under masking, shuffling, jitter, noise, and anti-tamper covers.

Third, evaluation methodologies show strong convergence: high-volume profiling, fine-grained segmentation, accuracy- and GE-based metrics, and ChipWhisperer or oscilloscope-based data acquisition form the backbone of experimental practice. However, the field lacks standardized benchmarks, unified metrics, and reproducible leakage datasets, making inter-paper comparison difficult and often unreliable.

Finally, significant research gaps remain. No existing study evaluates higher-order masking combined with strong hiding countermeasures; cross-device generalization is still fragile; Falcon and HQC are scarcely explored, and SPHINCS+ is completely unaddressed; and DSP-enhanced or ensemble-AI pipelines—common in other ML domains—are entirely absent. Fault injection has appeared only once in combination with deep learning and has not been systematically evaluated as part of multi-layer attack strategies, leaving the broader landscape of AI-augmented FI almost completely unexplored. These gaps highlight that while current AI-driven SCAs demonstrate substantial practical risk for PQC implementations, our understanding of their limits, robustness, and countermeasure resilience is still incomplete.

Overall, this review highlights the growing role of AI in PQC security evaluation. At the same time, the lack of standardized evaluation protocols, the absence of hybrid AI methodologies, and the limited exploration of strong countermeasure combinations point to clear directions for future research. Continued progress in these areas will be important to establish realistic, evidence-based assessments of PQC implementation security in the coming decade.

References

1. Alagic, G., Alperin-Sheriff, J., Apon, D., Cooper, D., Dang, Q., Kelsey, J., Liu, Y.K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Smith-Tone, D.: Status report on the second round of the NIST post-quantum cryptography standardization process. Tech. Rep. NIST IR 8309, National Institute of Standards and Technology (2020). https://doi.org/10.6028/NIST.IR.8309
2. Alagic, G., Alperin-Sheriff, J., Apon, D., Cooper, D., Dang, Q., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Smith-Tone, D., Liu, Y.K.: Status report on the first round of the nist post-quantum cryptography standardization process (01 2019). https://doi.org/10.6028/NIST.IR.8240
3. Alagic, G., Apon, D., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Liu, Y.K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Smith-Tone, D.: Status report on the third round of the nist post-quantum cryptography standardization process. Tech. Rep. NIST IR 8413, National Institute of Standards and Technology (9 2022). https://doi.org/10.6028/NIST.IR.8413-upd1
4. Alagic, G., Bros, M., Ciadoux, P., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Liu, Y.K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Silberg, H.: Status report on the fourth round of the nist post-quantum cryptography standardization process. Tech. Rep. NIST IR 8545, National Institute of Standards and Technology (3 2025). https://doi.org/10.6028/NIST.IR.8545
5. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key Exchange—A new hope. In: 25th USENIX Security Symposium (USENIX Security 16). pp. 327–343 (08 2016), https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim
6. Asghar, A., Becher, A., Ziener, D.: Backing the wrong horse: How bit-level netlist augmentation can counter power side channel attacks (2025). https://doi.org/10.48550/arXiv.2510.04640
7. Aumasson, J.P., Bernstein, D.J., Beullens, W., Dobraunig, C., Eichlseder, M., Fluhrer, S., Gazdag, S.L., Hülsing, A., Kampanakis, P., Kölbl, S., Kudinov, M., Lange, T., Lauridsen, M.M., Mendel, F., Niederhagen, R., Rechberger, C., Rijneveld, J., Schwabe, P., Westerbaan, B.: SPHINCS+. Submission to the NIST Post-Quantum Cryptography Standardization Process (10 2020), https://sphincs.org/data/sphincs+-r3.1-specification.pdf
8. Avanzi, R., Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS-Kyber (version 3.02). Submission to the NIST Post-Quantum Cryptography Standardization Process (08 2021), https://pq-crystals.org/kyber/data/kyber-specification-round3-20210804.pdf
9. Bai, S., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS-Dilithium (version 3.1). Submission to the NIST Post-Quantum Cryptography Standardization Process (08 2021), https://pq-crystals.org/dilithium/data/dilithium-specification-round3-20210208.pdf
10. Bauer, A., Jaulmes, E., Prouff, E., Wild, J.: Horizontal and vertical side-channel attacks against secure rsa implementations. In: Proceedings of the 13th International Conference on Topics in Cryptology. pp. 1–17 (2013). https://doi.org/10.1007/978-3-642-36095-4_1
11. Bernstein, D.J.: Cache-timing attacks on aes (2005), https://cr.yp.to/papers.html#cachetiming

12. Bernstein, D.J., Lange, T.: Post-quantum cryptography. Nature **549**(7671), 188–194 (09 2017). https://doi.org/10.1038/nature23461
13. Bommana, S.R., et al.: Mitigating side-channel attacks on fpga through deep learning and dynamic partial reconfiguration. Scientific Reports (2025). https://doi.org/10.1038/s41598-025-98473-3
14. Brier, É., Clavier, C., Olivier, F.: Correlation power analysis with a leakage model. In: Cryptographic Hardware and Embedded Systems - CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings. vol. 3156, pp. 16–29 (2004). https://doi.org/10.1007/978-3-540-28632-5_2
15. Chari, S., Rao, J.R., Rohatgi, P.: Template attacks. In: Cryptographic Hardware and Embedded Systems - CHES 2002. vol. 2523, pp. 13–28 (2002). https://doi.org/10.1007/3-540-36400-5_3
16. Chen, L., Jordan, S., Liu, Y.K., Moody, D., Peralta, R., Perlner, R., Smith-Tone, D.: Report on post-quantum cryptography. Tech. Rep. NISTIR 8105, National Institute of Standards and Technology (2016). https://doi.org/10.6028/NIST.IR.8105
17. Chen, P., Li, J., Cheng, W., Cheng, C.: Uncover secrets through the cover: A deep learning-based side-channel attack against kyber implementations with anti-tampering covers. IEEE Transactions on Computers **74**(6), 2159–2167 (06 2025). https://doi.org/10.1109/TC.2025.3547610
18. Choi, K.h., Han, J., Han, D.: Single trace analysis of visible vs. invisible leakage for comparison-operation-based CDT sampling. Electronics (Switzerland) **13**(23) (2024). https://doi.org/10.3390/electronics13234681
19. Deutsch, D.: Quantum theory, the church-turing principle and the universal quantum computer. Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences **400**(1818), 97–117 (1985). https://doi.org/10.1098/rspa.1985.0070
20. Dubrova, E., Ngo, K., Gärtner, J., Wang, R.: Breaking a fifth-order masked implementation of CRYSTALS-kyber by copy-paste. In: Proceedings of the 10th ACM Asia Public-Key Cryptography Workshop. pp. 10–20. APKC '23 (2023). https://doi.org/10.1145/3591866.3593072
21. Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Prest, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON: Fast-fourier lattice-based compact signatures over NTRU. Submission to the NIST Post-Quantum Cryptography Standardization Process (10 2020), https://falcon-sign.info/falcon.pdf
22. Gaborit, P., Aguilar-Melchor, C., Aragon, N., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Persichetti, E., Zémor, G., Bos, J., Dion, A., Lacan, J., Robert, J.M., Véron, P., Barreto, P.L., Ghosh, S., Gueron, S., Güneysu, T., Misoczki, R., Richter-Brokmann, J., Sendrier, N., Tillich, J.P., Vasseur, V.: Hamming quasi-cyclic (hqc) fourth round version. NIST PQC Standardization Process Submission (08 2025), https://pqc-hqc.org/doc/hqc_specifications_2025_08_22.pdf
23. Ganesh, B.S., Ahmed, M.M., Mady, A.: Higher order leakage assessment and neural network-based attack on CRYSTALS-kyber. In: Proceedings of the International Conference on Security and Cryptography. pp. 373 – 380 (2024). https://doi.org/10.5220/0012715700003767
24. Gidney, C., Ekerå, M.: How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits. Quantum **5**, 433 (2021). https://doi.org/10.22331/q-2021-04-15-433

25. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing. pp. 212–219 (1996). https://doi.org/10.1145/237814.237866
26. Han, J., Lee, T., Kwon, J., Lee, J., Kim, I.J., Cho, J., Han, D.G., Sim, B.Y.: Single-trace attack on NIST round 3 candidate dilithium using machine learning-based profiling. IEEE Access **9**, 166283–166292 (2021). https://doi.org/10.1109/ACCESS.2021.3135600
27. Hernandez-Alvarez, L., de la Torre, M.A.G., Hernandez, E.I., Encinas, L.H.: How to attack a galaxy: From star wars to star trek. In: 2023 Congress in Computer Science, Computer Engineering, Applied Computing (CSCE). pp. 2347–2354 (07 2023). https://doi.org/10.1109/CSCE60160.2023.00381
28. Hoang, A.T., Kennaway, M., Pham, D.T., Mai, T.S., Khalid, A., Rafferty, C., O'Neill, M.: Deep learning enhanced side channel analysis on CRYSTALS-kyber. In: 2024 25th International Symposium on Quality Electronic Design (ISQED). pp. 1–8 (04 2024). https://doi.org/10.1109/ISQED60706.2024.10528674
29. Huang, Z., Wang, H., Cao, B., He, D., Wang, J.: A comprehensive side-channel leakage assessment of CRYSTALS-kyber in IIoT. Internet of Things **27**, 101331 (2024). https://doi.org/10.1016/j.iot.2024.101331
30. Iavich, M., Gnatyuk, S., Mukasheva, A.: Decoding the CRYSTALS-kyber attack using artificial intelligence: Examination and strategies for resilience. In: Proceedings of the Workshop on Cybersecurity Providing in Information and Telecommunication Systems II (CPITS II 2024). vol. 3826, pp. 342–349 (2024), https://ceur-ws.org/Vol-3826/short26.pdf
31. Jao, D., De Feo, L.: Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In: Post-Quantum Cryptography. pp. 19–34 (2011)
32. Jendral, S., Ngo, K., Wang, R., Dubrova, E.: Breaking SCA-protected CRYSTALS-kyber with a single trace. In: 2024 IEEE International Symposium on Hardware Oriented Security and Trust (HOST). pp. 70–73 (05 2024). https://doi.org/10.1109/HOST55342.2024.10545390
33. Ji, Y., Dubrova, E.: A side-channel attack on a masked hardware implementation of crystals-kyber. Journal of Cryptographic Engineering **15**(1), 7 (04 2025). https://doi.org/10.1007/s13389-025-00375-7
34. Ji, Y., Wang, R., Ngo, K., Dubrova, E., Backlund, L.: A side-channel attack on a hardware implementation of CRYSTALS-kyber. In: 2023 IEEE European Test Symposium (ETS). pp. 1–5 (05 2023). https://doi.org/10.1109/ETS56758.2023.10174000
35. Karabulut, E., Aysu, A.: Falcon down: Breaking falcon post-quantum signature scheme through side-channel attacks. In: 2021 58th ACM/IEEE Design Automation Conference (DAC). pp. 691–696 (2021), https://ieeexplore.ieee.org/document/9586131
36. Kim, H., Kim, T.H., Yoon, J.C., Hong, S.: Practical second-order correlation power analysis on the message blinding method and its novel countermeasure for rsa. ETRI Journal **32**(1), 102–111 (2010). https://doi.org/10.4218/etrij.10.0109.0249
37. Kim, I.J., Lee, T.H., Han, J., Sim, B.Y., Han, D.G.: Novel single-trace ML profiling attacks on NIST 3 round candidate dilithium (2020), https://eprint.iacr.org/2020/1383
38. Kitchenham, B., Charters, S.: Guidelines for performing systematic literature reviews in software engineering. Tech. Rep. EBSE-2007-01, Keele University and University of Durham (2007), https://legacyfileshare.elsevier.com/promis_misc/525444systematicreviewsguide.pdf

39. Kitchenham, B., Pearl Brereton, O., Budgen, D., Turner, M., Bailey, J., Linkman, S.: Systematic literature reviews in software engineering - a systematic literature review. Information and Software Technology **51**(1), 7–15 (2009). https://doi.org/10.1016/j.infsof.2008.09.009
40. Kocher, P.C.: Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In: Advances in Cryptology — CRYPTO '96. pp. 104–113 (1996). https://doi.org/10.1007/3-540-68697-5_9
41. Kocher, P.C., Jaffe, J., Jun, B.: Differential power analysis. In: Advances in Cryptology — CRYPTO' 99. vol. 1666, pp. 388–397 (1999). https://doi.org/10.1007/3-540-48405-1_25
42. Koeune, F., Standaert, F.X.: A Tutorial on Physical Security and Side-Channel Attacks, pp. 78–108 (2005). https://doi.org/10.1007/11554578_3
43. Li, S.: Overview and discussion of attacks on CRYSTALS-kyber (2023), https://eprint.iacr.org/2023/1952
44. Lima, E., Maia, F., Ferreira, T.A.E., Ribeiro, D.: Supplementary material for: A slr of modern ai- driven sca on nist's pqc standards (Feb 2026). https://doi.org/10.5281/zenodo.18697380
45. Maghrebi, H., Portigliatti, T., Prouff, E.: Breaking cryptographic implementations using deep learning techniques. In: Security, Privacy, and Applied Cryptography Engineering - SPACE 2016. vol. 10076, pp. 3–26 (2016). https://doi.org/10.1007/978-3-319-49445-6_1
46. Marzougui, S., Ulitzsch, V., Tibouchi, M., Seifert, J.P.: Profiling side-channel attacks on dilithium: A small bit-fiddling leak breaks it all (2022), https://eprint.iacr.org/2022/106
47. Mascelli, J., Rodden, M.: "harvest now, decrypt later": Examining post-quantum cryptography and the data privacy risks for distributed ledger networks. Tech. Rep. FEDS 2025-093, Federal Reserve Board, Washington, D.C. (09 2025). https://doi.org/10.17016/FEDS.2025.093
48. McEliece, R.J.: A public-key cryptosystem based on algebraic coding theory. DSN Progress Report **42-44**, 114–116 (1978), https://tmo.jpl.nasa.gov/progress_report/42-44/44N.PDF
49. Mosca, M.: Cybersecurity in an era with quantum computers: Will we be ready? IEEE Security & Privacy **16**(5), 38–41 (2018). https://doi.org/10.1109/MSP.2018.3761723
50. National Institute of Standards and Technology: Post-quantum cryptography, https://csrc.nist.gov/projects/post-quantum-cryptography
51. National Institute of Standards and Technology: Post-quantum cryptography standardization, https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization
52. National Institute of Standards and Technology: Evaluation criteria. NIST PQC Standardization Website (2016), https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria
53. National Institute of Standards and Technology: Security (evaluation criteria). NIST PQC Standardization Website (2016), https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-(evaluation-criteria)
54. National Institute of Standards and Technology: Module-lattice-based digital signature standard. Tech. Rep. FIPS 204, National Institute of Standards and Technology (2024). https://doi.org/10.6028/NIST.FIPS.204

55. National Institute of Standards and Technology: Module-lattice-based key-encapsulation mechanism standard. Tech. Rep. FIPS 203, National Institute of Standards and Technology (2024). https://doi.org/10.6028/NIST.FIPS.203
56. National Institute of Standards and Technology: Nist releases first 3 finalized post-quantum encryption standards. NIST News (8 2024), https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards
57. National Institute of Standards and Technology: Stateless hash-based digital signature standard. Tech. Rep. FIPS 205, National Institute of Standards and Technology (2024). https://doi.org/10.6028/NIST.FIPS.205
58. National Institute of Standards and Technology: Nist pqc standardization process: Hqc announced as a 4th round selection. NIST News (3 2025), https://www.nist.gov/news-events/news/2025/03/nist-pqc-standardization-process-hqc-announced-4th-round-selection
59. National Institute of Standards and Technology: Nist selects hqc as fifth algorithm for post-quantum encryption. NIST News (3 2025), https://www.nist.gov/news-events/news/2025/03/nist-selects-hqc-fifth-algorithm-post-quantum-encryption
60. Ngo, K., Dubrova, E., Johansson, T.: A side-channel attack on a masked and shuffled software implementation of saber. Journal of Cryptographic Engineering **13**(4), 443–460 (11 2023). https://doi.org/10.1007/s13389-023-00315-3
61. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information: 10th Anniversary Edition (2010)
62. PQShield: Nist selects hqc for standardization. PQShield Blog (2025), https://pqshield.com/nist-selects-hqc-for-standardization/
63. Proos, J., Zalka, C.: Shor's discrete logarithm quantum algorithm for elliptic curves. Quantum Info. Comput. **3**(4), 317–344 (07 2003), https://www.rintonpress.com/xqic3/qic-3-4/317-344.pdf
64. Qiao, K., Wang, Z., Chang, H., Sun, S., Wu, Z., Cheng, J., Ou, C., Wang, A., Zhu, L.: A closer look at the belief propagation algorithm in side-channel attack on cca-secure pqc kem. Science China Information Sciences **67**(11), 212302 (10 2024). https://doi.org/10.1007/s11432-024-4150-3
65. Qiao, Z., Liu, Y., Zhou, Y., Zhao, Y., Chen, S.: Single trace is all it takes: Efficient side-channel attack on dilithium (2024), https://eprint.iacr.org/2024/512
66. Qiao, Z., Liu, Y., Zhou, Y., Zhao, Y., Yuan, H., Du, D.: Efficient CNN-based side-channel attacks on dilithium without device access. In: 2025 IEEE International Symposium on Circuits and Systems (ISCAS). pp. 1–5 (05 2025). https://doi.org/10.1109/ISCAS56072.2025.11043271
67. Ravi, P., Chattopadhyay, A., D'Anvers, J.P., Baksi, A.: Side-channel and fault-injection attacks over lattice-based post-quantum schemes (kyber, dilithium): Survey and new results. ACM Trans. Embed. Comput. Syst. **23**(2) (03 2024). https://doi.org/10.1145/3603170
68. Ravi, P., Jap, D., Bhasin, S., Chattopadhyay, A.: Invited paper: Machine learning based blind side-channel attacks on PQC-based KEMs - a case study of kyber KEM. In: 2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD). pp. 01–07 (10 2023). https://doi.org/10.1109/ICCAD57390.2023.10323721
69. Rezaeezade, A., Yap, T., Jap, D., Bhasin, S., Picek, S.: Breaking the blindfold: deep learning-based blind side-channel analysis (2025)

70. Rivain, M., Prouff, E.: Provably secure higher-order masking of aes. In: Mangard, S., Standaert, F.X. (eds.) Cryptographic Hardware and Embedded Systems, CHES 2010. pp. 413–427 (2010). https://doi.org/10.1007/978-3-642-15031-9_28
71. Rostovtsev, A., Stolbunov, A.: Public-key cryptosystem based on isogenies. IACR Cryptol. ePrint Arch. p. 145 (2006), http://eprint.iacr.org/2006/145
72. Sajadi, A., Zidaric, N., Stefanov, T., Mentens, N.: A systematic comparison of side-channel countermeasures for risc-v-based socs. In: 2024 IEEE Nordic Circuits and Systems Conference (NorCAS). pp. 1–7 (2024). https://doi.org/10.1109/NorCAS64408.2024.10752477
73. Shor, P.W.: Algorithms for quantum computation: Discrete logarithms and factoring. In: Proceedings of the 35th Annual Symposium on Foundations of Computer Science. pp. 124–134 (1994)
74. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing **26**(5), 1484–1509 (1997). https://doi.org/10.1137/S0097539795293172
75. Srivastava, A., Das, S., Choudhury, N., Psiakis, R., Silva, P.H., Pal, D., Basu, K.: SCAR: Power side-channel analysis at RTL level. IEEE Transactions on Very Large Scale Integration (VLSI) Systems **32**(6), 1110–1123 (06 2024). https://doi.org/10.1109/TVLSI.2024.3390601
76. of Standards, N.I., Technology: Submission requirements and evaluation criteria for the post-quantum cryptography standardization process. Federal Register Notice (12 2016), https://csrc.nist.gov/csrc/media/projects/post-quantum-cryptography/documents/call-for-proposals-final-dec-2016.pdf
77. Tanaka, Y., Ueno, R., Xagawa, K., Ito, A., Takahashi, J., Homma, N.: Multiple-valued plaintext-checking side-channel attacks on post-quantum KEMs. IACR Transactions on Cryptographic Hardware and Embedded Systems **2023**(3), 473 – 503 (2023), https://eprint.iacr.org/2022/940
78. Tena-Sánchez, E., Potestad-Ordóñez, F.E., Jiménez-Fernández, C.J., Acosta, A.J., Chaves, R.: Gate-level hardware countermeasure comparison against power analysis attacks. Applied Sciences **12**(5) (2022). https://doi.org/10.3390/app12052390
79. Tena-Sánchez, E., Potestad-Ordóñez, F.E., Zúñiga-González, V., Acosta, A.J.: Low-cost full correlated-power-noise generator to counteract side-channel attacks. Applied Sciences **15**(6) (2025). https://doi.org/10.3390/app15063064
80. The White House: Report on post-quantum cryptography (7 2024), https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/07/REF_PQC-Report_FINAL_Send.pdf
81. Ueno, R., Xagawa, K., Tanaka, Y., Ito, A., Takahashi, J., Homma, N.: Curse of re-encryption: A generic power/EM analysis on post-quantum KEMs. IACR Transactions on Cryptographic Hardware and Embedded Systems **2022**(1), 296 – 322 (2021), https://eprint.iacr.org/2021/849
82. Veyrat-Charvillon, N., Medwed, M., Kerckhof, S., Standaert, F.X.: Shuffling against side-channel attacks: A comprehensive study with cautionary note. In: Advances in Cryptology – ASIACRYPT 2012. pp. 740–757 (2012)
83. Wang, J., Cao, W., Chen, H., Li, H.: Practical side-channel attack on masked message encoding in latticed-based KEM (2022), https://eprint.iacr.org/2022/859
84. Wang, R., Gärtner, J., Dubrova, E.: Decompressing dilithium's public key with fewer signatures using side channel analysis. In: 2025 IEEE 55th International Symposium on Multiple-Valued Logic (ISMVL). pp. 135–140 (06 2025). https://doi.org/10.1109/ISMVL64713.2025.00034

85. Wang, R., Ngo, K., Gärtner, J., Dubrova, E.: Single-trace side-channel attacks on CRYSTALS-dilithium: Myth or reality? (2023), https://eprint.iacr.org/2023/1931
86. Wang, Y., Huang, F., Duan, X., Hu, H.: Second-order side-channel attacks on kyber: Targeting the masked hash function. Journal of Cryptologic Research **11**(6), 1415 – 1436 (2024). https://doi.org/10.13868/j.cnki.jcr.000745
87. Wang, Z., Ding, Y., Wang, A., Zhang, Y., Wei, C., Sun, S., Zhu, L.: Spa-gpt: General pulse tailor for simple power analysis based on reinforcement learning. IACR Transactions on Cryptographic Hardware and Embedded Systems **2024**(4), 40–83 (09 2024). https://doi.org/10.46586/tches.v2024.i4.40-83
88. Wikipedia: Post-quantum cryptography. Online Encyclopedia (2025), https://en.wikipedia.org/wiki/Post-quantum_cryptography
89. Wikipedia: Side-channel attack. Online Encyclopedia (2025), https://en.wikipedia.org/wiki/Side-channel_attack
90. Wohlin, C.: Guidelines for snowballing in systematic literature studies and a replication in software engineering. In: Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering (2014). https://doi.org/10.1145/2601248.2601268
91. Wohlin, C., Runeson, P., Höst, M., Ohlsson, M.C., Regnell, B., Wesslén, A.: Experimentation in Software Engineering (2024). https://doi.org/10.1007/978-3-662-69306-3
92. Xu, D., Wang, K., Tian, J.: A hardware-friendly shuffling countermeasure against side-channel attacks for kyber. IEEE Transactions on Circuits and Systems II: Express Briefs **72**(3), 504–508 (2025). https://doi.org/10.1109/TCSII.2025.3528751
93. Yang, Y., Huang, J., Wang, Z., Ye, J., Sun, Z., Fan, J., Chen, S., Li, H., Li, X., Cao, Y.: A template attack on reduction without reference device on kyber. In: 2023 IEEE 32nd Asian Test Symposium (ATS). pp. 1–6 (10 2023). https://doi.org/10.1109/ATS59501.2023.10318019
94. You, S.C.: Single-trace template attacks on permutation-based cryptography. Ph.D. thesis, University of Cambridge (12 2022), https://www.cl.cam.ac.uk/~scy27/PhD_thesis.pdf
95. Zhang, H., Babar, M.A., Tell, P.: Identifying relevant studies in software engineering. Inf. Softw. Technol. **53**(6), 625–637 (06 2011). https://doi.org/10.1016/j.infsof.2010.12.010
96. Zhou, W., Wang, A., Ding, Y., Wei, C., Zhang, J., Zhu, L.: One solves all: Exploring ChatGPT's capabilities for fully automated simple power analysis on cryptosystems (2024), https://eprint.iacr.org/2024/2069
97. Zhou, X., Jin, Y., Zhang, H., Li, S., Huang, X.: A map of threats to validity of systematic literature reviews in software engineering. In: 2016 23rd Asia-Pacific Software Engineering Conference (APSEC). pp. 153–160 (2016). https://doi.org/10.1109/APSEC.2016.031