


AI-assisted cryptanalysis | documented findings, 2024-2026

Actual disclosure / presentation dates • Each numbered bubble links to a published finding below • Through 13 October 2026

Bubble height and area reflect a disclosed outcome category (rubric below), NOT a published CVSS score, a probability, or comparable bits-of-security loss.



Documented findings and source references

Dates are the verified public availability date (online publication, preprint posting, or scheduled conference session).

2024

01 02-02 Salsa Fresca: sparse LWE

ML secret recovery at $n=1024$; preprocessing 25x faster and $\sim 10\times$ more sample-efficient.

Ref: Stevens et al., IACR 2024/150

02 04-01 Dilithium2: CNN-assisted key recovery

On an STM32F4, secret s_1 recovered from two signatures in about five minutes.

Ref: Qiao et al., IACR 2024/512

03 04-29 Neural cryptanalysis: improved input selection

PCA/K-means data selection improved SIMON/SPECK neural distinguisher performance; $\sim 30\%$ faster...

Ref: Seok & Lee, IEEE TDSC (online 2024)

04 07-05 Fusion-channel side-channel learning

Learning-based fusion of several leakage channels improved side-channel attack classification.

Ref: Li et al., IACR 2024/1092

05 08-01 Neural differential distinguishers

Reported $\sim 2\text{--}8$ percentage-point accuracy gains for reduced-round SPECK/SIMON.

Ref: Wang et al., IEICE E107-D

06 08-09 Sparse Kyber-parameter LWE secret recovery

Recovered secrets of Hamming weight 9–11 in 28–36h using Salsa/Cool&Cruel.

Ref: Wenger et al., Meta AI / IEEE S&P

2025

07 03-04 Kyber side-channel beneath protective cover

Neural distinguisher accuracy $>99\%$ through 18-mm cover and $>89\%$ through 24-mm cover.

Ref: Chen et al., IEEE Trans. Comput.

08 04-01 Neural linear attacks against DES

Neural-aided partial key attacks extended to 8-round DES and improved attack success.

Ref: Hou et al., Cybersecurity 8, 22

09 05-16 SKINNY: neural-inspired integral attack

Neural-derived integral properties enable a 15-round SKINNY key-recovery attack (+1 round).

Ref: Zhang et al., arXiv:2505.10790

10 07-13 ML improves LWE attacks on special secrets

Custom loss/data distributions enabled recovery of LWE secrets twice as hard in tests.

Ref: Saxena et al., ICML 2025

11 08-15 AES-128: one key byte via ASCAD CNN

CNN achieved Rank 0 for one AES key byte after ~ 65 traces on the ASCAD dataset.

Ref: Poudel & Rahimi, arXiv:2508.11817

12 09-16 Kyber: neural recovery of secret coefficients

Non-profiled deep learning recovered key coefficients on unmasked and masked ARM Cortex-M4...

Ref: Li et al., IEEE TCAD

13 09-24 Neural differential-linear DES / SPECK

Demonstrated better ML-aided key-recovery success for 8-round DES and reduced SPECK.

Ref: Hou et al., JKSUCIS 37, 228

2026

14 04-05 LWE: repetition + stepwise regression

Improved ML recovery of denser, preprocessed sparse LWE secrets.

Ref: Alfarano et al., arXiv:2604.03903

15 04-21 PRESENT: neural round-key recovery

Recovered 64-bit round keys in 402/433 trials (92.8%) of reduced-round PRESENT.

Ref: Gauthier-Umana et al., Neural Comput. Appl.

16 05-19 SIMON / SIMECK: neural distinguishers

100% distinguishing accuracy up to 7 rounds; key-hypothesis experiment reported.

Ref: Mirzaali et al., Cybersecurity 9, 70

17 05-25 ML-DSA-44: deep-learning hardware key recovery

Reported full secret key from 7,400 power traces; 50.67% fewer than prior attacks.

Ref: Zhao et al., IEEE ISCAS (presented 25 May)

18 06-11 ML-KEM: 400-trace neural side-channel

Recovered secret coefficient pairs in unprotected CRYSTALS reference code using ~ 400 EM traces.

Ref: Ledger Donjon, 11 June 2026

19 07-28 HAWK: AI-assisted mathematical weakness

HAWK-256 attack estimate reduced from 2^{64} to 2^{38} ; candidate subsequently withdrawn.

Ref: Anthropic (28 Jul 2026); corroboration: NIST

20 07-28 AES: AI-found 7-round attack

New meet-in-the-middle method accelerates attack by 200–800x versus prior work.

Ref: Anthropic, 28 Jul 2026

21 09-13 SHA-3 / Ascon: neural distinguishers

3-round SHA-3-256 neural distinguisher reached 99.99% with 16 sample pairs.

Ref: Electronics 15(18), 4142

Impact rubric: 1 = distinguisher/method only; 2 = reduced rounds or specially chosen keys/parameters; 3 = partial key via leakage; 4 = full key via physical side-channel; 5 = algorithmic weakness caused candidate withdrawal.

Curated documented examples—not an exhaustive survey, attack prevalence estimate, or evidence of a continuous increase in severity. Physical side-channel attacks do not break the underlying mathematical primitive.