

Intelligent Fault Tolerance

Parham Ghayour

`monsieur.parham.ghayour@gmail.com`

Sorbonne University <https://orcid.org/0009-0001-5607-1532>

Research Article

Keywords: quantum error correction; adaptive decoding; Steane code; hidden Markov model; correlated Pauli noise; identifiability; Bayesian risk; uncertainty.

Posted Date: October 27th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11494585/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Intelligent Fault Tolerance

Parham Ghayour
Independent Researcher

24 October 2026

Research manuscript — computational and theoretical proof of concept

Abstract

An intelligent quantum error-correction system should improve recovery using information available during computation, but syndrome observations do not necessarily identify the physical mechanism producing them. We analyze this tension through a precisely specified, reproducible model of a genuine $[[7, 1, 3]]$ Steane CSS code subjected to hidden, temporally persistent, correlated Pauli noise. The simulation exhaustively enumerates all $4^7 = 16,384$ physical Pauli patterns and the 64 syndrome classes. A causal hidden-Markov filter uses only the observed syndrome history to update a posterior over two correlated-error environments; a logical-class maximum-a-posteriori decision then conditions on this posterior. A learned logistic temporal model provides a separate data-driven comparator. On 220,000 locked, sequentially correlated instances, the stationary noise-aware decoder fails in 2.7923% of cases versus 2.5218% for the hidden-Markov decoder, an observed relative reduction of 9.69%. A moving-block bootstrap yields a paired absolute difference of 0.2705 percentage points (approximate 95% interval $[0.2354, 0.3064]$ points). A state-oracle decoder achieves 1.9809%; deliberate model misspecification increases the hidden-Markov failure rate to 2.8809%. We derive an information-theoretic bound on the value of history, prove a decoder-regret bound controlled by posterior-state error, and give an explicit syndrome-aliasing construction on the same code: two physically different fault channels induce identical entire syndrome-history laws while requiring incompatible recoveries, forcing a strictly positive failure probability for every syndrome-only controller. These mathematical facts arise from established Bayesian decision and indistinguishability principles; we make no claim to a new fault-tolerance theorem. This is an original, small-code model study, *not* a circuit-level, real-time, scaling, or hardware result. It identifies both an opportunity and a non-removable information limit for intelligent QEC.

Keywords: quantum error correction; adaptive decoding; Steane code; hidden Markov model; correlated Pauli noise; identifiability; Bayesian risk; uncertainty.

1 Scientific question and scope

Fault-tolerant quantum computing requires reliable logical operations, fast syndrome processing, and a realistic accounting of faults and classical control overhead. The development of high-accuracy neural decoders [3, 7], adaptive QEC hardware control [4], decoder ensembling [6], real-time FPGA decoding [8, 9], and noise estimation from syndrome statistics [5, 10] establishes a substantial existing literature. It would be incorrect to claim that neural, adaptive, selective, or syndrome-only decoding is itself novel. A particularly close 2026 paper develops analytical sliding-window methods for learning drifting Pauli noise from QEC syndromes [5]. Our study does not supersede that work.

We investigate a more specific pair of questions: (i) how much can knowledge of a *hidden persistent noise regime* improve exact logical-class recovery in a fully specified stabilizer code; and (ii) when

is such a hidden regime *unidentifiable* from the syndromes, rendering further machine learning powerless without another observable? We build an exactly enumerable $[[7, 1, 3]]$ stabilizer instance, a causal model-based controller, a trained syndrome-only comparator, explicit posterior-risk formulas, and stress tests.

Scope boundary. Each indexed observation below is a newly sampled seven-qubit encoded-memory *instance* under a shared slowly varying simulated noise environment. It is not a further syndrome-extraction round on one continuously evolving logical qubit. Stabilizer measurements are perfect, there is no gate-level fault model or decoder latency measurement, and code distance does not scale. This distinction is crucial: the numbers must not be interpreted as logical error *per hardware QEC cycle*.

2 A correct quantum code and exact logical classes

Let $H \in \mathbb{F}_2^{3 \times 7}$ be the binary matrix

$$H = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}. \quad (1)$$

The CSS stabilizer generator matrices are $H_X = H_Z = H$. Direct multiplication gives $HH^\top = 0$ over $\mathbb{F}_2$. The rank of H is three, so the code encodes $k = 7 - 3 - 3 = 1$ logical qubit. Its unrestricted Pauli distance is three [1, 2]; this was also checked by exhaustive enumeration over every seven-qubit Pauli operator. This corrects a substantive weakness in the earlier one-sided repetition-code pilot, where a bit-flip distance of seven did not imply a quantum distance of seven.

Represent a Pauli error modulo phase by $(x, z) \in \mathbb{F}_2^7 \times \mathbb{F}_2^7$, with Y corresponding to $x_i = z_i = 1$. The six-bit measured syndrome is

$$s(x, z) = (Hx, Hz) \in \mathbb{F}_2^6. \quad (2)$$

For any three-bit component u , let $r(u)$ be zero for $u = 0$ and the unique single-site binary vector with $Hr(u) = u$ otherwise. The joint reference recovery is $(r(Hx), r(Hz))$. Because each row of H has even weight, the stabilizer span has even binary parity; its two cosets inside $\ker H$ are distinguished by parity. Define the two logical residual bits

$$\ell_X = \text{wt}(x \oplus r(Hx)) \bmod 2, \quad \ell_Z = \text{wt}(z \oplus r(Hz)) \bmod 2. \quad (3)$$

Thus every Pauli fault maps unambiguously into one of 64×4 syndrome-logical-class pairs (s, ℓ) . The four values of $\ell = (\ell_X, \ell_Z)$ represent the four logical Pauli cosets modulo stabilizers. Syndrome-compatible decoding chooses a predicted class; it does not require recovering the exact physical error. Exhaustive verification finds all 21 single-qubit X, Y, Z faults belong to the identity residual class under the reference recovery and the minimum-weight undetectable logical Pauli has weight three.

3 A hidden, correlated Pauli-noise model

At each *fresh encoded instance* t let the hidden environment $Z_t \in \{0, 1\}$ follow a stationary symmetric two-state Markov chain:

$$\mathbb{P}(Z_t = Z_{t-1}) = \rho, \quad \rho = 0.993, \quad \mathbb{P}(Z_1 = 0) = \frac{1}{2}. \quad (4)$$

The base channel applies independent depolarizing noise of strength $p = 0.027$ to seven sites. In state 0, with probability $q = 0.075$, an additional perfectly correlated X_1X_2 burst is composed with the base channel; in state 1, it is an X_4X_5 burst. Here sites are numbered starting at one. Writing B_z for the binary X -mask of that state and E_0 for the independent depolarizing Pauli error,

$$E_t = \begin{cases} E_{0,t}, & 1 - q, \\ X(B_{Z_t})E_{0,t}, & q. \end{cases} \quad (5)$$

This is a deliberately simplified *Pauli memory-channel ensemble*. The switching process is inter-instance noise drift, not a microscopic device model. A correlated two-qubit burst can have the same syndrome as a single-qubit fault while differing by a nontrivial logical operator; this makes inference consequential. We compute the exact joint tables

$$Q_z(s, \ell) = \sum_{E: s(E)=s, \ell(E)=\ell} \mathbb{P}(E \mid Z = z) \quad (6)$$

by enumerating $4^7 = 16,384$ Paulis, with no sampling error in Q_z . Let $e_z(s) = \sum_{\ell} Q_z(s, \ell)$ denote the observable syndrome emission distribution. The total variation between the two *syndrome* distributions in the primary setting is 0.06820, and $D_{\text{KL}}(e_0 \| e_1) = 0.14080$ nats. The regimes are distinguishable, but only weakly from a single observation.

4 Causal intelligence: filtering and logical-class inference

Define $b_t^- = \mathbb{P}(Z_t = 0 \mid S_{1:t-1})$ and $b_t^+ = \mathbb{P}(Z_t = 0 \mid S_{1:t})$, where S_t is the observed syndrome of the current fresh code block. Given the correct Markov model, Bayesian filtering yields

$$b_t^+ = \frac{b_t^- e_0(S_t)}{b_t^- e_0(S_t) + (1 - b_t^-) e_1(S_t)}, \quad (7)$$

$$b_{t+1}^- = \rho b_t^+ + (1 - \rho)(1 - b_t^+). \quad (8)$$

The controller never receives Z_t , the underlying physical error E_t , or the true logical class. Those are retained only by the simulator to score its decisions. The exact history-conditioned logical-class MAP decision is

$$\hat{\ell}_t = \arg \max_{\ell \in \{0,1\}^2} [b_t^- Q_0(S_t, \ell) + (1 - b_t^-) Q_1(S_t, \ell)]. \quad (9)$$

This is not an invented neural quantum decoder. It is a rigorous Bayesian reference system against which an intelligent learned policy can be assessed. We compare five actions: (i) standard single-error CSS recovery (residual class zero); (ii) stationary marginal MAP with $b^- = 1/2$; (iii) Eq. (9); (iv) a *trained* logistic classifier on four windows of syndrome-frequency features (3, 10, 30, 90 observations) whose predicted state belief feeds the same class posterior; and (v) an unattainable state oracle using Z_t .

The logistic classifier is trained on 55,000 independent synthetic training instances. Latent regime labels are available *only during simulated offline training*; its test-time inputs are observed syndromes alone. This is a small interpretable supervised-learning comparator, not an architectural claim about deep neural networks. For all policy decisions, only current/past syndromes are used.

5 Two analytical results and a decisive information limit

Let $J_b(s, \ell) = bQ_0(s, \ell) + (1 - b)Q_1(s, \ell)$ for a prior belief b *before* observing the current syndrome, and let $d_b(s)$ be its Bayes-optimal logical-class decision. Let $R_b(d) = 1 - \sum_s J_b(s, d(s))$ be the one-shot logical failure risk of decision rule d .

Proposition 1 (Posterior-error-to-decoder-regret bound). *For any true belief b , approximate belief $\tilde{b}$, and logical-class MAP rule $d_{\tilde{b}}$, the excess risk satisfies*

$$0 \leq R_b(d_{\tilde{b}}) - R_b(d_b) \leq 2 \text{TV}(J_b, J_{\tilde{b}}) = 2|b - \tilde{b}| \text{TV}(Q_0, Q_1). \quad (10)$$

Here total variation uses the full joint syndrome–logical-class distributions, not merely the observed syndrome marginal.

Proof. For each syndrome s , let $a_s = d_b(s)$ and $c_s = d_{\tilde{b}}(s)$. By optimality of c_s for $J_{\tilde{b}}$,

$$\begin{aligned} J_b(s, a_s) - J_b(s, c_s) &\leq [J_b - J_{\tilde{b}}](s, a_s) \\ &\quad + [J_{\tilde{b}} - J_b](s, c_s). \end{aligned}$$

Summing and taking absolute values bounds the regret by $\sum_{s,\ell} |J_b - J_{\tilde{b}}|(s, \ell) = 2 \text{TV}(J_b, J_{\tilde{b}})$. Finally $J_b - J_{\tilde{b}} = (b - \tilde{b})(Q_0 - Q_1)$, proving equality in the last expression. $\square$

This is an application of standard Bayesian classification stability, *not* a claimed new quantum threshold theorem. It quantifies why accurate state estimation can matter: the maximum possible benefit of belief refinement is controlled by how differently the regimes distribute *logical classes jointly with syndromes*. Moreover, it does not guarantee that any particular approximate filter is better than the fixed decoder.

Proposition 2 (Information bound on the value of syndrome history). *Let $U = S_{1:t-1}$ be past syndrome history, $S = S_t$ the current syndrome, L the actual logical error class, and $Z = Z_t$ the contemporaneous hidden noise regime. Suppose $L \perp U \mid (Z, S)$, as in the fresh-instance hidden-Markov model above. Let $R^*(S)$ denote optimal classification risk with the current syndrome alone and $R^*(S, U)$ optimal classification risk with both observations. Using natural logarithms,*

$$0 \leq R^*(S) - R^*(S, U) \leq \sqrt{\frac{1}{2}I(L; U \mid S)} \leq \sqrt{\frac{1}{2}I(Z; U \mid S)}. \quad (11)$$

Proof. For every observation pair (s, u) compare the posterior distribution $p(\ell \mid s, u)$ to $p(\ell \mid s)$. The difference between their maximum coordinates is no greater than their total variation distance. Taking expectations gives an upper bound on the Bayes improvement by $\mathbb{E}_{s,u} \text{TV}(p(L \mid s, u), p(L \mid s))$. Applying Pinsker’s inequality pointwise and Jensen’s inequality yields the middle term of (11), because the expected posterior divergence equals the conditional mutual information $I(L; U \mid S)$. The conditional Markov relation $U \rightarrow Z \rightarrow L$ given S and data processing give the last inequality. $\square$

This is a standard information-theoretic argument, not a novel information inequality. It supplies a necessary *informational budget* for adaptation: if the past reveals nothing about Z once the current syndrome is known, it cannot improve the logical-class Bayes decision in this model. Conversely the bound may be extremely loose and positive information alone does not guarantee a large improvement.

Theorem 3 (Syndrome-aliasing impossibility on the Steane code). *There exist two distinct Pauli channels acting on the Steane $[[7, 1, 3]]$ code whose distributions over any finite or infinite syndrome history are identical, while their optimal recoveries disagree. If the hidden regime is equally likely, no decoder observing syndromes alone can attain logical failure probability smaller than $q/2$ in the construction below, while a decoder given the hidden regime attains zero.*

Table 1: Locked simulation on 220,000 correlated instances. Bootstrap intervals account approximately for serial dependence; the oracle is unavailable in practice.

Decoder	Failures	Risk (%)	Block-bootstrap 95% CI (%)
Single-error CSS	17,244	7.8382	[7.7277, 7.9514]
Stationary MAP	6,143	2.7923	[2.7227, 2.8632]
Causal HMM-MAP	5,548	2.5218	[2.4586, 2.5905]
Learned window classifier	5,673	2.5786	[2.5127, 2.6455]
State oracle	4,358	1.9809	[1.9223, 2.0414]
Misspecified HMM	6,338	2.8809	[2.8114, 2.9517]

Proof. Fix $q \in (0, 1)$. In regime 0, apply identity with probability $1 - q$ and X_3 with probability q . In regime 1, apply identity with probability $1 - q$ and X_1X_2 with probability q . From the actual matrix (1), its third column is the binary XOR of its first two columns. Therefore the errors X_3 and X_1X_2 have *identical* nonzero Z -check syndrome, and neither creates an X -check syndrome. But their product $X_1X_2X_3$ has zero syndrome and odd weight, hence implements a nontrivial logical X coset; their desired logical-class decisions are opposite. The two regimes both emit the same syndrome with probability q and the zero syndrome with probability $1 - q$. Thus any sequence of syndromes has the same probability under either regime; the observations convey zero information about its identity, even if the hidden state has an arbitrary stationary Markov persistence. Under equal stationary regime priors and syndrome-only inference, a nonzero syndrome occurs with probability q and must be assigned one of the two incompatible logical classes. At least half these events are decoded incorrectly, giving risk at least $q/2$. A regime oracle chooses the matching class and has zero risk. $\square$

Remark 4. *This is a constructive information-theoretic limitation, not a new threshold theorem. It rules out unconditional claims that an arbitrarily powerful neural or generative model can infer all useful error information from raw syndromes. Additional trusted calibration, an independently informative hardware signal, a different physical measurement protocol, or changed experimental assumptions can break this indistinguishability.*

6 Experimental protocol and numerical evidence

We use deterministic NumPy seeds and exact noise/syndrome tables. The simulator draws a training trajectory of 55,000 and a *separate locked evaluation* trajectory of 220,000 code-block instances. The latter contains 1,571 environment switches. The primary comparison uses a fixed, model-known HMM and a trained syndrome-only logistic model. No true error or hidden regime is passed to either at evaluation. No parameter tuning uses the locked test sequence.

Dependence-aware uncertainty. A hidden Markov regime induces temporal dependence, so ordinary binomial Clopper–Pearson intervals would not be valid as exact sequence-level confidence bounds. We therefore report *approximate* stationary-process uncertainty with a non-overlapping block bootstrap: 440 contiguous blocks of length 500, resampled 4,000 times, percentile 95% intervals. The block length exceeds the order of the model’s characteristic mixing time. This procedure is not a finite-sample theorem and does not address misspecified stationarity.

The exact one-shot class tables give state-conditional risks of 7.7913% for minimum-weight recovery, 2.7424% for stationary MAP, and 1.9424% for the state oracle; the two states have equal risks by symmetry. These provide enumeration-based checks on the corresponding Monte Carlo figures.




Figure 1: Locked logical failure probabilities for the exact Steane code with hidden correlated Pauli bursts. Each point reports one fresh encoded-memory instance, not a QEC cycle.




Figure 2: Illustrative causal belief trajectory. The red state sequence is plotted *only for retrospective analysis*; no hidden-state labels are supplied to test-time decoding.

The HMM decreases observed failure by 0.27045 percentage points relative to stationary MAP, or 9.69% relative, with paired block-bootstrap difference interval $[0.2354, 0.3064]$ percentage points. These are conditional results for the *specified synthetic Markov channel*. The oracle gap remains 0.5409 percentage points. The trained window classifier improves over stationary MAP but underperforms the known-model HMM by 0.0568 percentage points (paired block interval $[0.0350, 0.0786]$). HMM state-belief Brier score is 0.10097 versus 0.11475 for the learned classifier, demonstrating imperfect but informative context inference. The empirical conditional information $I(Z_t; S_{<t} | S_t)$ estimated from the exact-filter posterior is approximately 0.3300 nats. The bound in Proposition 2 gives $\sqrt{I/2} \approx 0.4062$, far above the observed absolute risk gain 0.00270: this bound is conceptually useful but numerically very loose in this regime. These numbers do not show that an HMM is always preferable to neural sequence models; the model is exactly specified here, making the HMM unusually strong.




Figure 3: Sensitivity to correlated-burst probability, each point from a separate 90,000-instance synthetic trajectory. The HMM parameters are given the correct generating model in this experiment.

7 Sensitivity, stress tests, and the value of additional information

Under a deliberately incorrect HMM emission model (burst probability assumed 0.018 rather than 0.075, with persistence assumed 0.96 rather than 0.993), failure rises to 2.8809%, slightly *worse* than stationary MAP. Consequently, a learned or inferred noise model is not a universal safeguard. Without reliable model checking and fallback, intelligence can increase logical risk.

We additionally vary the burst probability in $\{0.02, 0.04, 0.06, 0.075, 0.10, 0.14\}$ with 90,000 simulated instances each, retracing the correct HMM posterior for each generating channel (Fig. 3). This is exploratory sensitivity, not a statistically powered multi-configuration claim. For state persistence $\rho \in \{0.90, 0.97, 0.993, 0.999\}$, the respective pairs of static-versus-HMM empirical failure probabilities are approximately (2.773, 2.769)%, (2.703, 2.648)%, (2.770, 2.540)%, and (2.771, 2.101)%. Short-lived regimes yield almost no history advantage; high persistence allows the controller to accumulate useful evidence. These experiments use separate draws, so small variations should not be overinterpreted.

8 Selective recovery and what *cannot* be claimed about cost

A fixed stationary action d_{stat} can be replaced with a history-aware action only when its *posterior* logical-success advantage exceeds a chosen threshold γ :

$$\Delta_t = \max_{\ell} \mathbb{P}(\ell \mid S_{1:t}) - \mathbb{P}(d_{\text{stat}}(S_t) \mid S_{1:t}), \quad d_t = \begin{cases} d_{\text{HMM}}, & \Delta_t > \gamma, \\ d_{\text{stat}}, & \text{otherwise.} \end{cases} \quad (12)$$

For $\gamma = 0$, the selected logical class differs from stationary MAP in approximately 0.758% of the held-out instances. The simulated failure rate is about 2.521%. This is an interesting *decision sparsity* observation: a small number of reassigned syndrome classes accounts for the available improvement.

However the Bayesian filter and posterior scores still have to be computed to know whether to reassign a decision. Therefore 0.758% is *not* a 99.2% saving in computation, energy, latency,




Figure 4: Logical error versus the fraction of logically changed decisions under threshold gating; no physical latency or compute saving is inferred.

or decoder invocations; no such saving is established. The threshold is a policy illustration with no measured hardware cost model. Harmonized decoder ensembles and confidence-conditioned escalation are important prior art [6]. Our selective gate is not a new decoder or a demonstrated system-level resource reduction.

9 Connection to 2026 research and novelty boundaries

The field has already demonstrated major components of adaptive fault tolerance. Google’s Willow experiment uses error-detection events as control-learning signals and reports a 3.5-fold improvement against injected drift [4]. Bhardwaj *et al.* analytically recover time-varying noise from syndrome histories and show noise-aware decoding improvements [5]; our use of syndrome-only drift information should not be represented as a first. Shutty *et al.* study confidence-aware selective decoder ensembling [6]. Attention-based neural decoding is evaluated for circuit-level algorithms [7]; FPGA systems demonstrate stringent feedback latency [8, 9]. Recent detector-model research also treats non-Pauli and correlated faults [10, 11], and an APS-accepted October 2026 study investigates syndrome-resampling improvements [12].

The specific *assembled computational exercise here* comprises exact four-class Steane likelihoods under switched correlated XX noise, a causal hidden-state reference decoder, a learned temporal baseline, approximate dependence-aware uncertainty, parameter sweeps, misspecification, and an explicit indistinguishable-syndrome Pauli-channel counterexample. The mathematical bounds are straightforward consequences of standard Bayesian decision and information theory; the explicit counterexample uses stabilizer syndrome equivalence. We do not assert priority for these general principles or claim that this is a competitive decoder result for surface codes or qLDPC codes. The scientific contribution is a transparent, falsifiable model analysis clarifying when extra temporal intelligence can and cannot improve decoding.

10 Limitations and implications for fault-tolerant computing

The code distance is fixed at three; the model knows p, q, ρ ; the base channel is depolarizing with hand-specified XX bursts; stabilizer extraction is perfect; memory blocks are independent given the noise regime; the logical objective is a single-shot Pauli-class decision. There are no faulty gates, syndrome-measurement faults, leakage, analog coherent errors, idles, entanglement consumption, magic-state production, real-time feedback requirements, or logical circuit benchmarks. Experimental intervals are approximate block-bootstrap estimates under stationarity. The learned baseline uses synthetic hidden-state labels during *training* (not inference), which might be unavailable for laboratory calibration. We have not demonstrated a physical-qubit saving, a lower decoding latency, a fault-tolerance threshold, or quantum computational advantage.

These limitations still yield a useful systems principle. *Model accuracy and information availability are independent requirements.* When the syndrome distributions retain information about slowly drifting regimes, a compact filter can exploit that information. When distinct logical-error mechanisms are syndrome-equivalent, no architecture, regardless of neural-network scale, can identify them using only those observations. One must introduce informative side measurements, active calibration, redesigned extraction, or carefully justified priors. Safety engineering must separately address model misspecification.

The next decisive experiment is a *circuit-level rotated surface code* with noisy stabilizer extraction and honest use of recorded detector events, followed by qLDPC tests against BP+OSD, BP+LSD, BP+OTF, and established neural or matching baselines. A successful follow-up should measure full logical operation failure and FPGA/ASIC-relevant $p99$ decoding latency under independent held-out drift distributions. Such experiments were *not* completed here and must not be implied by the title.

11 Conclusion

We tested an intelligent, causally valid decoder based on temporal inference of hidden correlated noise within a genuine small quantum code. With a correct noise model, syndrome history improves logical-class decisions over a stationary MAP reference on the selected synthetic channel. A learned classifier offers partial gains; deliberate model misspecification erases them. A constructive syndrome-aliasing theorem makes the key limitation precise: the quantum syndrome need not identify which recovery is correct, even with unlimited historical data. These observations motivate an *information-aware* rather than merely architecture-aware view of intelligent fault tolerance. They are a rigorous research prototype, not a demonstration of scalable fault-tolerant quantum computing or an established breakthrough.

Reproducibility and integrity statement

The complete source code `src/experiment.py`, numerical outputs `results/results.json`, and vector figures are provided in the accompanying archive. All 4^7 Paulis, the CSS commutation relation, code distance, syndrome representatives, and single-qubit correction cases are checked by assertions. Reproduction command: `OPENBLAS_NUM_THREADS=1 python src/experiment.py`. NumPy, SciPy, Matplotlib and scikit-learn are required; seeds and experimental sample counts are recorded in the script. Web research was conducted against original journal or primary-source records available as of 24 October 2026. Code/runtime review is internal; this is not external peer review. No physical experiment was performed.

References

- [1] A. M. Steane, “Error correcting codes in quantum theory,” *Phys. Rev. Lett.* **77**, 793 (1996). doi:10.1103/PhysRevLett.77.793.
- [2] V. V. Albert et al., “The Error Correction Zoo,” Steane $[[7,1,3]]$ code entry, <https://errorcorrectionzoo.org/c/steane> (accessed 2026-10-24).
- [3] J. Bausch et al., “Learning high-accuracy error decoding for quantum processors,” *Nature* **635**, 834 (2024). doi:10.1038/s41586-024-08148-8.
- [4] V. Sivak et al., “Reinforcement learning control of quantum error correction,” *Nature* **655**, 879 (2026). doi:10.1038/s41586-026-10759-2.
- [5] D. Bhardwaj, E. Takou, Y. Lin and K. R. Brown, “Adaptive estimation of drifting noise in quantum error correction,” *PRX Quantum* **7**, 033024 (2026). doi:10.1103/zlhc-nqw5.
- [6] N. Shutty, M. Newman and B. Villalonga, “Efficient near-optimal decoding of the surface code through ensembling,” *Phys. Rev. Lett.* **136**, 070603 (2026). doi:10.1103/77j6-32xx.
- [7] J. P. Bonilla Ataides, A. Gu, S. F. Yelin and M. D. Lukin, “Neural decoders for universal quantum algorithms,” *PRX Intelligence* **1**, 013008 (2026). doi:10.1103/vjnl-mbxi.
- [8] L. Caune et al., “Demonstrating real-time and low-latency quantum error correction with superconducting qubits,” *Nature Communications* **17**, 7383 (2026). doi:10.1038/s41467-026-73331-6.
- [9] X. Yang et al., “Real-time surface-code error correction using an FPGA-based neural-network decoder,” *arXiv:2605.04892* (2026). <https://arxiv.org/abs/2605.04892>.
- [10] E. Takou and K. R. Brown, “Correcting non-Pauli noise with detector error models estimated from syndrome measurements,” *Phys. Rev. A* **114**, 022449 (2026). doi:10.1103/s46l-pnlq.
- [11] A. S. Maan et al., “Decoding correlated errors in quantum LDPC codes,” *Nature Communications* **17**, 3965 (2026). doi:10.1038/s41467-026-70556-3.
- [12] L. Colmenarez, A. Márton and M. Müller, “Syndrome resampling enhances quantum error correction thresholds,” *Phys. Rev. Lett.*, accepted 20 October 2026. doi:10.1103/v3kk-rwk8.
- [13] A. deMarti iOlius et al., “An almost-linear time decoding algorithm for quantum LDPC codes under circuit-level noise,” *npj Quantum Information* **12**, 152 (2026). doi:10.1038/s41534-026-01292-1.