Error Mitigation and Error Correction

Distinguish protecting a logical state from mitigating an observable estimator.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Decode single-X syndromes of the three-qubit code.
- Calculate idealised majority-decoding failure.
- Quantify the variance cost of linear noise extrapolation.

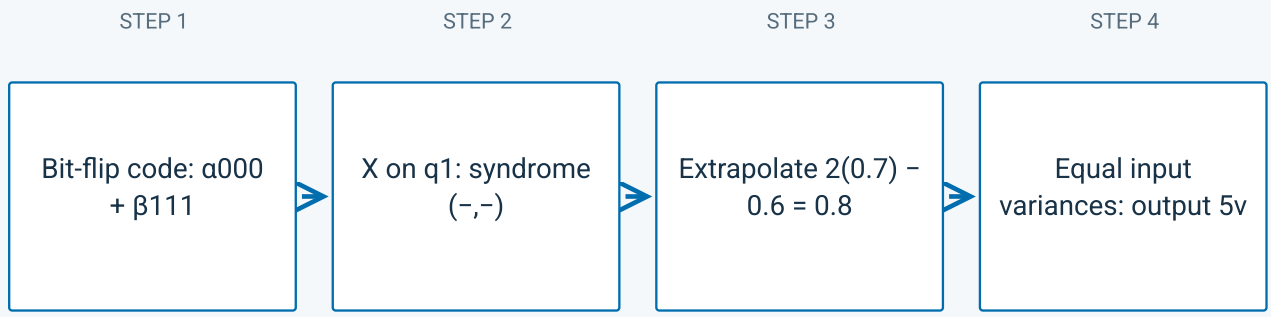
Two different responses to imperfect hardware

Error correction protects encoded quantum information during a computation. Error mitigation uses extra experiments and classical processing to improve estimates obtained from imperfect circuits. Both can be useful, but they offer different guarantees and consume different resources. A mitigated expectation value is not evidence that the underlying quantum state was fully protected.

Prerequisites: measurement, no-cloning and reversibility and noise models. The examples below are intentionally small enough to calculate without a general noise simulator.

WORKED DIAGRAM · A37

Correction and mitigation



The code corrects one X under its stated model but not arbitrary errors. Linear extrapolation removes bias only under its model and increases variance.

The correction setup

An error-correcting code specifies an encoded subspace, detectable error syndromes and a recovery procedure under an error model. Syndrome measurements reveal information about errors while preserving the encoded logical

information when the construction works. Fault tolerance further requires that imperfect correction and gate operations do not spread errors uncontrollably.

The output is a protected logical computation with a bounded logical error under stated assumptions. Code size alone is insufficient: syndrome extraction, decoding, measurement faults, correlations and logical-gate implementations matter. Stabiliser theory provides a primary formal framework for these ideas. Gottesman.

Worked example: the three-qubit bit-flip code

Encode a logical state by mapping alpha zero plus beta one to alpha 000 plus beta 111. This is an entangled encoding, not three independent copies of an unknown qubit. The no-cloning theorem is therefore not violated. The code is designed to correct a single X error, not an arbitrary quantum error.

Measure the two parity checks Z on q0 times Z on q1, and Z on q1 times Z on q2. With no X error, both signs are positive. An X on q0 gives negative, positive; an X on q1 gives negative, negative; an X on q2 gives positive, negative. Each syndrome identifies a different one-bit flip, allowing the matching X recovery.

The checks do not reveal whether the logical amplitudes were alpha or beta, because 000 and 111 have the same parity signs. That is the essential protection mechanism: learn an error label without directly measuring the logical basis label. A Z error, however, changes a relative logical phase without producing these bit-flip syndromes, so this small code does not protect against it.

Under independent bit flips with probability p, majority decoding fails when two or three physical bits flip. The failure probability is three p squared minus two p cubed. At p equal to 0.01, it is 0.000298, much smaller than 0.01 in this idealised model. Imperfect syndrome extraction and other error types would require a fuller calculation.

Encode one logical state into correlated physical amplitudes, then count the uncorrectable error patterns.

$$\alpha|0\rangle + \beta|1\rangle \longmapsto \alpha|000\rangle + \beta|111\rangle, \qquad P_{\text{logical}} = \binom{3}{2}p^2(1-p) + p^3 = 3p^2 - 2p^3.$$

The probability formula assumes independent physical X errors with probability p and ideal majority recovery. At $p = 0.01$, it gives 0.000298. Phase errors and noisy syndrome extraction lie outside this toy guarantee.

The mitigation setup

A mitigation method estimates an ideal observable from noisy data under additional assumptions, often at increased sampling cost. Examples include noise extrapolation, probabilistic cancellation and measurement-error correction. The method should state its noise model, calibration procedure, estimator bias and variance. It generally does not provide the same scalable protection theorem as fault-tolerant error correction.

Primary short-depth mitigation work introduces extrapolation and quasi-probability methods as ways to improve expectation estimates under defined conditions. Temme, Bravyi and Gambetta. The extra experiments and statistical overhead are part of the method, not a cost-free postprocessing detail.

Worked example: linear extrapolation

Suppose an observable follows the toy response E(lambda) equal to 0.8 minus 0.1 lambda, where lambda scales a known noise process. Measurements at noise scales one and two have ideal mean values 0.7 and 0.6. The linear extrapolator twice the first estimate minus the second has mean 0.8, the zero-noise intercept.

If the two independent estimates have equal variance v, the extrapolated estimate has variance five v. The bias is removed in this exact linear toy model, but statistical uncertainty increases. If the true response curves with noise scale,

the same formula can retain bias. More samples reduce variance, not a wrong extrapolation model's systematic error.

An extrapolated estimate may fall outside the physical range of an expectation in a finite sample. That does not mean a physical state has an impossible observable value; it means the estimator is a signed combination of noisy measurements. Report its uncertainty and assumptions rather than interpreting it as direct tomography of a protected state.

The extrapolator cancels the linear bias while increasing statistical variance.

$$E(\lambda) = 0.8 - 0.1\lambda, \qquad \widehat{E}_0 = 2\widehat{E}(1) - \widehat{E}(2), \qquad \mathbb{E}[\widehat{E}_0] = 0.8, \quad \mathrm{Var}(\widehat{E}_0) = 5v.$$

The variance equality assumes independent estimates at the two noise scales, each with variance v. The zero-bias statement depends on the exact linear model and does not automatically survive curvature.

Choose the method for the task

If a long coherent algorithm needs reliable intermediate states, correcting logical errors is a central requirement. If a short circuit only needs an expectation estimate, mitigation may improve accuracy under a manageable overhead. Hybrid strategies are possible, but the role of each layer should remain explicit.

Do not compare a corrected computation's logical error probability with a mitigated estimator's standard error as if they were the same quantity. One concerns the execution model and encoded information; the other concerns inference from experiments. An application-level error budget can include both, but it must explain how they affect the requested output.

A reproducible evaluation plan

For correction, test known error patterns, syndrome tables and recovery, then introduce faults in the extraction circuit. For mitigation, test the estimator on small circuits with independent ideal references and vary the noise assumptions. Report total shots, calibration data, model mismatch and uncertainty. A method that improves one selected circuit may not generalise to another noise regime.

Test the code against an excluded error

Apply Z to q0 of the encoded state. It leaves the 000 amplitude unchanged and negates the 111 amplitude, creating a logical phase change. Both measured bit-parity checks remain positive because Z commutes with them. The decoder therefore sees no bit-flip syndrome and cannot correct this error under the stated scheme.

This is a deliberate limitation of the three-qubit example. A general quantum code must address both bit and phase information with an appropriate syndrome structure. Presenting the small repetition code as protection against every physical error would teach the wrong lesson. A successful single-X test validates only the error family that the construction was designed to handle.

Exercises and worked answers

Exercise 1. Which bit-flip syndrome identifies q1 in the stated three-qubit code?

Check the worked answer

Both parity checks are negative. Flipping q1 changes its parity with each neighbour. Recovery applies X to q1, assuming the single-X-error model.

Exercise 2. Does linear extrapolation remain unbiased if the response includes an unknown quadratic term?

Check the worked answer

Generally no. Two points cancel the linear contribution but do not determine arbitrary curvature. Additional scales or a different justified model may help, at further sampling and modelling cost. Quantify residual bias rather than claiming that extrapolation universally removes noise.

Check your understanding

PAUSE & THINK

For independent estimates each with variance v, what is $\mathrm{Var}(2E_1 - E_2)$?

- [ ] v
- [ ] 3v
- [ ] 5v

Answer: 5v. The squared coefficients give 4v + v = 5v. Bias reduction can require more shots to control the larger statistical uncertainty.

Sources

- Gottesman
- Temme, Bravyi and Gambetta