LEARNING TO REGULARIZE VARIATIONAL QUANTUM DYNAMICS UNDER DISTRIBUTION SHIFT

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantum-geometric variational dynamics repeatedly solve a local linear inverse problem whose spectrum can change substantially across Hamiltonians, circuit geometries, trajectory states, and estimator conditions. Regularization is therefore not merely a numerical safeguard: it determines which tangent-space directions are trusted and can become a deployment bottleneck when each new target requires retuning. We formulate *regularization transfer*: learn from source quantum geometric tensor (QGT)/force systems a task-agnostic map from local inverse geometry to a constrained regularization action, then reuse that map without target-side model selection. ScalarGuard predicts a global damping scale, SpectralGuard predicts mode-wise spectral gates, and a deterministic resolved-subspace verifier rejects inadmissible learned proposals in favor of a frozen quantum-natural-gradient Levenberg–Marquardt reference. Source-validation diagnostics show that preferred absolute Tikhonov scales span more than six orders of magnitude across families, establishing substantial heterogeneity in the quantity to be transferred. In frozen confirmatory evaluation, source-selected Tikhonov remains strongest on clean in-distribution data, whereas SpectralGuard improves over the reference on pooled clean structural shifts and attains the smallest descriptive worst-case relative regret, 20.9%, across seven non-floor regimes. Under strong synthetic estimator corruption, ScalarGuard is more stable and SpectralGuard mostly falls back to the reference, revealing an expressivity–robustness boundary rather than uniform spectral superiority. These results position learned regularization as an amortized numerical adaptation layer between fixed inverse solvers and fully learned optimizers. The QEC branch evaluates transfer across stabilizer-derived Hamiltonian families, not decoding or fault-tolerance performance.

1 INTRODUCTION

Variational quantum algorithms replace an exponentially large state-space evolution by dynamics on a parameterized manifold. In variational quantum imaginary-time evolution (VarQITE), McLachlan projection turns each integration step into the local inverse problem $G_t v_t = f_t$, where G_t is the real quantum geometric tensor (QGT), f_t is the imaginary-time force, and v_t is the parameter velocity (McArdle et al., 2019; Gomes et al., 2021). The same metric defines quantum natural gradient (QNG) descent in state space (Stokes et al., 2020; Koczor & Benjamin, 2022). This geometry is useful precisely because it distinguishes physically redundant and anisotropic directions, but it also creates the numerical bottleneck: weak or poorly estimated QGT modes receive large inverse weight. Consequently, regularization changes more than solver conditioning; by deciding which tangent directions are trusted, it changes the realized variational trajectory itself (van Straaten & Koczor, 2021).

The difficulty is amplified when the inverse problem is solved repeatedly across related but non-identical workloads. Hamiltonian family, circuit width and depth, trajectory location, and metric-estimation quality all change the QGT spectrum, so a damping scale that is appropriate on one source problem need not remain appropriate after deployment. This issue is especially visible in structured Hamiltonians derived from quantum error-correction (QEC) constraints: repetition, surface, and low-density parity-check constructions induce different parity and connectivity structure before the resulting detector constraints are reduced to diagonal-Ising VarQITE instances (Google Quantum

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

AI, 2021; Fowler et al., 2012; Panteleev & Kalachev, 2022). We use these stabilizer-derived Hamiltonians as controlled geometry shifts, not as a decoding benchmark: logical error rate, syndrome reconstruction, and fault-tolerance thresholds are outside the scope of the study.

This setting exposes a gap between two familiar strategies. Classical Tikhonov, truncated spectral, or Levenberg–Marquardt regularization preserves the analytic inverse structure but typically relies on problem-specific tuning; fully learned optimizers amortize update rules across tasks but can obscure the geometry and failure modes of the underlying inverse. Our source validation makes the portability problem concrete: the family-wise optimum of the absolute Tikhonov scale spans more than six orders of magnitude. A single fixed value can therefore be excellent when the target matches validation and systematically mismatched after a geometry shift, while per-target retuning defeats zero-shot deployment. The scientific question is whether *the regularization rule itself* can be transferred: can local QGT/force geometry provide enough information to adapt the inverse without target labels or target-side validation?

We address this question with *guarded learned regularization*. ScalarGuard preserves the low-dimensional inductive bias of Tikhonov damping and predicts one trajectory-conditioned scale. SpectralGuard instead predicts bounded mode-wise gates from normalized spectral scale, force–mode alignment, and global trajectory context. Both controllers act only on task-local numerical descriptors; Hamiltonian-family and shift identity are excluded. A deterministic resolved-subspace verifier then separates adaptation from admissibility by comparing the learned proposal with a trusted QNG–Levenberg–Marquardt (QNG–LM) reference before the common final step control. The method therefore learns how to regularize an existing geometric inverse rather than replacing the inverse with an unconstrained learned optimizer.

The experiments are designed around the motivation above. Hyperparameters of the conventional baselines are selected once on source validation and frozen on all confirmatory targets, so transfer is not confounded with target-side oracle tuning. Family out-of-distribution evaluation removes the target family from source training and validation before fitting the learned checkpoint. The resulting evidence separates three regimes. On clean in-distribution data, a well-tuned fixed Tikhonov solver remains strongest. Across clean structural shifts, SpectralGuard benefits from geometry-conditioned mode selection and improves over QNG–LM in the pooled confirmatory comparison. Under strong estimator corruption, however, the fine-grained spectral state itself becomes unreliable: ScalarGuard is more stable and the verifier rejects most SpectralGuard proposals. The central conclusion is therefore not universal learned superiority, but a more specific one: adaptive regularization is most useful when the inverse geometry changes yet remains informative, and deterministic verification is essential when that information degrades.

2 Related Work and Positioning

From quantum geometry to the regularized inverse. VarQITE and QNG establish why the QGT is the appropriate local geometry for variational dynamics and optimization (McArdle et al., 2019; Stokes et al., 2020). Subsequent work improves the variational manifold, extends geometric optimization to noisy or non-unitary settings, or reduces the cost of acquiring QGT information (Gomes et al., 2021; Koczor & Benjamin, 2022; Gacon et al., 2024; Anuar et al., 2024). These approaches determine *which* geometry is represented or *how* it is estimated. Our work addresses the next numerical layer: once a QGT/force estimate exists, which spectral directions should be trusted in its inverse, and can that decision transfer across tasks? This distinction isolates regularization as an object of learning rather than treating it as a fixed implementation detail of VarQITE.

Between classical regularization and learned optimization. Tikhonov damping and truncated spectral inversion encode strong, interpretable priors for ill-posed problems (Tikhonov, 1963; Hansen, 1998); their limitation in the present setting is not lack of stability but lack of portability when one hyperparameter must serve changing spectra. Learning-to-optimize methods attack the opposite end of the spectrum by training update rules over task distributions (Andrychowicz et al., 2016; Wichrowska et al., 2017), with later analyses emphasizing that inductive bias and stability determine whether such rules generalize out of distribution (Harrison et al., 2022). Guarded learned regularization occupies the narrower middle ground. It retains the analytic QGT inverse, learns only a constrained scalar or spectral filter from dimensionless local descriptors, and evaluates

transfer without task identity or target retraining. The intended benefit is therefore amortized solver adaptation with a substantially smaller hypothesis class than a fully learned optimizer.

Safeguards as a deployment principle. Trust-region and Levenberg–Marquardt methods distinguish proposing a step from accepting it when the local numerical model is unreliable (Levenberg, 1944; Marquardt, 1963; Nocedal & Wright, 2006). We use the same principle to make learning compatible with a geometric solver: the controller supplies an adaptive regularization proposal, while a deterministic test on numerically resolved directions decides whether that proposal is admissible relative to QNG–LM. The safeguard is thus not an auxiliary heuristic added after learning; it is the mechanism that preserves a conservative deployment path when the learned spectral state leaves its reliable regime.

QEC-derived structure as a transfer test. Repetition and surface codes provide sparse but distinct stabilizer geometries, while quantum LDPC constructions introduce different connectivity and rank structure (Google Quantum AI, 2021; Fowler et al., 2012; Panteleev & Kalachev, 2022). Related numerical work in QEC often targets decoding or noise-model inference itself; for example, differentiable maximum-likelihood estimation fits circuit-level noise parameters from syndrome statistics (Cao et al., 2026). Our use of QEC is deliberately different. Detector constraints are converted to diagonal-Ising Hamiltonians and serve as structured generators of QGT geometry; the transferred object is the inverse regularization policy, not a decoder, a noise model, or a fault-tolerance primitive. This role places the QEC branch alongside random k-local and Sherrington–Kirkpatrick families as a controlled test of numerical transfer while retaining physically structured source and target instances.

3 Guarded Regularization and Source-to-Target Transfer

Guarded learned regularization acts on the local inverse problem rather than predicting unconstrained optimizer steps. Figure 1 summarizes the two separable parts of the construction: local numerical adaptation and source-to-target deployment. The controller never receives task, family, or shift identity.

3.1 Local inverse geometry and spectral state

At step t, VarQITE solves

$$G_t v_t = f_t, \tag{1}$$

with real QGT G_t, imaginary-time force f_t, and parameter velocity v_t. The implementation first symmetrizes the estimated metric, $\widetilde{G}_t = (\widehat{G}_t + \widehat{G}_t^\top)/2$, and diagonalizes

$$\widetilde{G}_t = U_t \,\mathrm{diag}(\sigma_{t,1}, \ldots, \sigma_{t,d}) U_t^\top. \tag{2}$$

Small or negative estimated modes signal directions in which direct inversion is unreliable. Global context z_t summarizes trajectory progress, conditioning, effective rank, force scale, and recent update state; mode descriptors add relative spectral scale and force projection $|u_{t,i}^\top f_t|$. Training-set normalization is frozen before evaluation, and all task/family identifiers are excluded.

3.2 Learned regularization controllers

ScalarGuard and SpectralGuard share the same local numerical state but differ in action dimension. ScalarGuard adapts a single source-selected anchor,

$$\lambda_t = \lambda_{\mathrm{ref}} e^{\phi_g(z_t)}, \qquad v_t^{\mathrm{sc}} = (\widetilde{G}_t + \lambda_t I)^{-1} f_t, \tag{3}$$

where $\lambda_{\mathrm{ref}} = 10^{-3}$ is the single absolute Tikhonov value selected on the source validation split. The learned log residual is bounded to $\pm \log 10^6$; neither the anchor nor the bound is retuned on a target regime.

SpectralGuard predicts $g_{t,i} \in [0, 1]$ and applies

$$v_t^{\mathrm{sp}} = U_t \,\mathrm{diag}\left(\frac{g_{t,i}}{[\sigma_{t,i}]_+ + \epsilon_t} \right) U_t^\top f_t, \qquad \epsilon_t = 10^{-10} \max_i [\sigma_{t,i}]_+ + 10^{-12}. \tag{4}$$

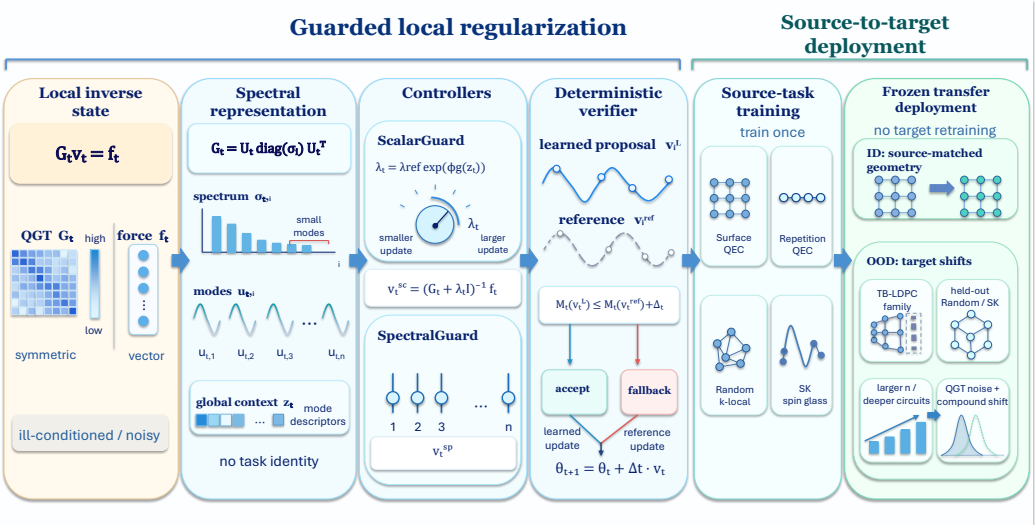

Figure 1: Guarded regularization and source-to-target deployment. The observed QGT/force system is converted to a task-agnostic spectral state, which drives either ScalarGuard or SpectralGuard. SpectralGuard is a filtered generalized inverse: negative estimated eigenvalues are not inverted, and a scale-aware floor stabilizes weak positive modes. The deterministic verifier compares a learned proposal with a trusted QNG–LM reference and either accepts the learned update or falls back before the common final step control. The source-task panel depicts the full source pool used by the principal ID/size/depth/noise models. Family-OOD evaluation instead follows a leave-one-family protocol in which the held-out family is omitted from source training and validation before fitting the frozen checkpoint. ID denotes source-matched geometry; OOD denotes changes in family, instance cohort, size/depth, or estimator quality.

For an indefinite estimate this is a learned spectral filter, not the exact inverse of $\widetilde{G}_t$: negative modes contribute no inverse scale. Gates are shared within numerically degenerate eigenspaces to remove dependence on the arbitrary basis chosen by the eigensolver. We do not claim that the filtered proposal alone guarantees monotone imaginary-time descent under an indefinite noisy metric; deployment admissibility is handled separately by the verifier below. Supplementary Section A gives the complete feature and invariance definitions.

3.3 Deterministic verification and safe update

The learned proposal v_t^L is compared with a condition-adaptive QNG–LM reference v_t^{ref}. Let $\mathcal{R}_t$ contain positive modes above

$$\kappa_t = \max\Big(\tau_{\mathrm{abs}}, \tau_{\mathrm{rel}} \max_i [\sigma_{t,i}]_+\Big). \tag{5}$$

On this resolved subspace the verifier scores

$$M_t(v) = q_{\mathcal{R}_t}(v) + \rho r_{\mathcal{R}_t}(v)^2 + \eta n_{\mathcal{R}_t^\perp}(v)^2, \tag{6}$$

where q is the local quadratic McLachlan model, r is normalized residual, and n is the fraction of velocity in unresolved directions. A finite, norm-feasible learned proposal is accepted only when $M_t(v_t^L) \le M_t(v_t^{\mathrm{ref}}) + \Delta_t$; otherwise the QNG–LM reference is used. The tolerance increases conservatively as resolved rank decreases or the observed metric becomes more indefinite. All verifier constants were selected on BO development data and locked before confirmatory evaluation; their exact values and the tolerance formula are given in Supplementary Section A.

3.4 Source-task training and frozen transfer

The full source pool comprises 100 dependency groups: 10 surface-code-derived, 10 repetition-code-derived, 40 random k-local, and 40 Sherrington–Kirkpatrick (SK) groups. QEC memory circuits are converted with Stim (Gidney, 2021) to detector models and then through the

Rank/Fourier construction to diagonal-Ising Hamiltonians. The QEC branch therefore supplies structured stabilizer-derived Hamiltonians; decoder accuracy is not an evaluation target. Controllers are learned offline by task-balanced imitation followed, for SpectralGuard, by short-horizon refinement. Non-Tikhonov spectral teacher actions occur in roughly 12% of selected snapshots, so the learned spectral capacity is used selectively rather than universally.

Deployment uses no target-side model selection. For ID, size, depth, and noise tests, the checkpoint trained on the full source pool is frozen. Family-OOD tests use a preregistered leave-one-family protocol: a new source checkpoint is trained after removing the target family from both source training and validation, and that checkpoint is then frozen on the held-out family. In particular, the TB-LDPC target is evaluated by a leave-QEC-out checkpoint trained on random k-local and SK source groups, not by a checkpoint that has seen surface/repetition QEC groups. The post-confirmatory compound panel combines size/depth shift with estimator corruption and remains exploratory. Supplementary Section C gives exact task counts, source/target maps, and frozen configuration provenance.

4 Experimental Results

The results test the portability question posed in the Introduction: whether a source-selected regularization rule remains useful when the local inverse geometry changes. All conventional hyperparameters are selected before confirmatory evaluation and frozen across targets. Absolute Tikhonov uses $\lambda = 10^{-3}$, spectrum-scaled Tikhonov uses $\lambda/\sigma_{\mathrm{ref}} = 0.03$, and QNG–LM uses the fixed Phase-1 configuration; no target label or target validation statistic enters baseline selection. Training and evaluation dependency groups are disjoint, with the dependency group—not an integration step—as the statistical unit. This design makes clean ID a test of source-matched efficiency, structural OOD a test of geometric portability, and estimator-noise regimes a test of how each controller behaves when the spectral information on which adaptation relies becomes unreliable. Supplementary Section C gives the exact provenance; among the planned inferential tests, only pooled structural SpectralGuard-versus-QNG-LM superiority is established.

4.1 Clean dynamics expose the source of transfer error

Figure 2a establishes the source-matched reference point. Validation-selected Tikhonov has the lowest late-trajectory error and the lowest clean-ID aggregate mean in Table 1; ScalarGuard and SpectralGuard remain close, whereas QNG–LM is higher. This ordering is expected when validation and deployment share the same geometry: a fixed scalar chosen directly for that regime already supplies the appropriate bias, so additional learned degrees of freedom have little room to help and may only add approximation variance. The transfer problem becomes substantive only when the spectrum moves away from the source regime.

That movement is visible in Figure 2d. The family-wise diagnostic optimum of the absolute Tikhonov scale changes from 3×10^{-10} on surface-code-derived tasks to 10^{-3} on Sherrington–Kirkpatrick tasks, a span of approximately 6.5 orders of magnitude. These optima are not used to tune the frozen baselines; instead, they reveal the numerical heterogeneity that a single source-selected ridge must absorb. The result gives a concrete mechanism for fixed-scale transfer error: after a family or geometry shift, the balance between suppressing weak modes and preserving informative directions changes even when the underlying VarQITE equation is unchanged.

Figures 2b–c test whether geometry-conditioned adaptation can absorb part of this mismatch. The held-out TB-LDPC target uses the leave-QEC-out checkpoint, so neither surface nor repetition QEC groups appear in learned source training or validation. Scaled Tikhonov and ScalarGuard attain 2.73×10^{-6} and 2.88×10^{-6} mean infidelity, SpectralGuard 5.88×10^{-6}, and QNG–LM 1.51×10^{-5}. Because all principal regularized solvers lie near the numerical floor, this cell supports compatibility with a distinct stabilizer-derived Hamiltonian family rather than a multiplicative QEC advantage. In the non-floor structural shifts, the paired SpectralGuard-minus-QNG–LM differences are -0.00421 for held-out random instances, -0.00225 for held-out SK instances, -0.01254 for larger systems, and -0.00613 for deeper circuits, with all four confidence intervals below zero; the pooled difference is -0.00606 with 95% CI $[-0.00716, -0.00504]$. The larger gain on the size shift is consistent with the motivating mechanism: changes in parameter-space geometry alter the spectral scale and

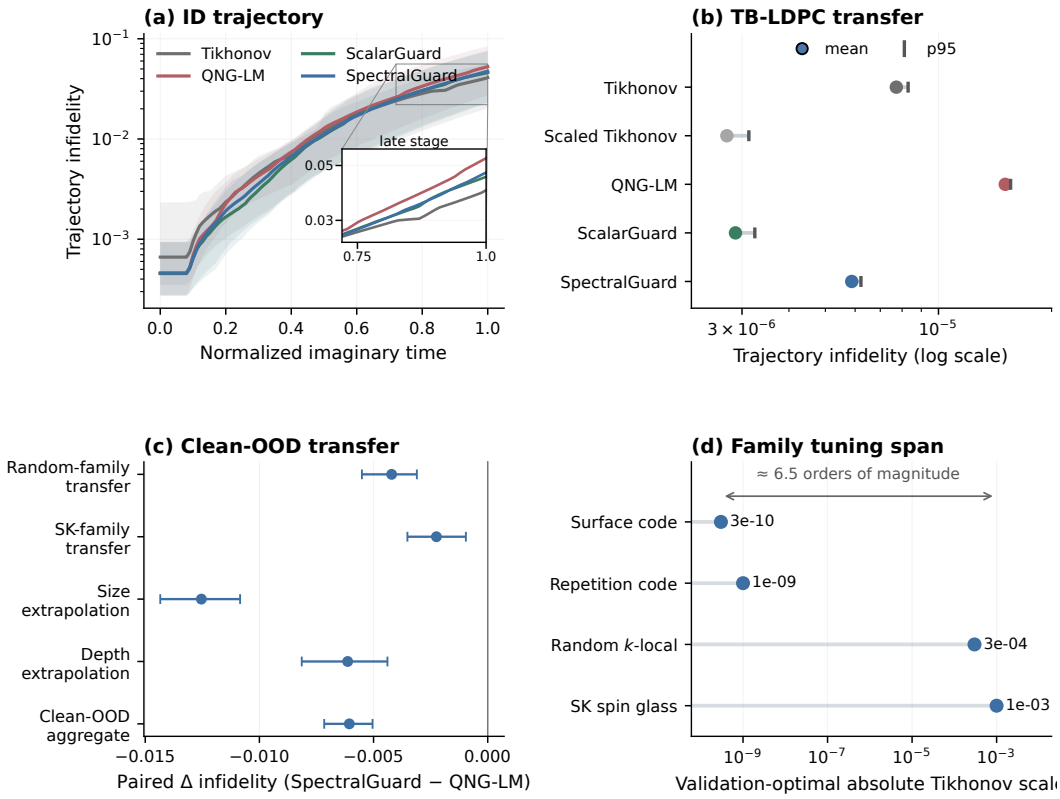

Figure 2: Clean dynamics and transfer evidence. (a) Clean-ID median trajectory infidelity with 25th–75th percentile bands; the inset enlarges the late-trajectory region where the methods separate most clearly. (b) Zero-shot transfer to the held-out TB-LDPC QEC family; circles are mean trajectory infidelity and vertical ticks are the 95th percentile. (c) Paired clean structural-transfer effects for SpectralGuard relative to QNG-LM; each horizontal interval is a 95% confidence interval and values left of zero favor SpectralGuard. The final row pools the four structural OOD splits. (d) Validation-optimal absolute Tikhonov damping for the four source families, showing the approximately 6.5-order-of-magnitude scale mismatch that motivates adaptive regularization.

resolved-rank structure, making a local mode-conditioned filter more useful than a reference rule whose damping response is not learned from cross-task variation.

Table 1 makes the resulting specialization explicit. "Clean ID" denotes held-out source-family tasks at source geometry; "QEC family" is held-out TB-LDPC under the leave-QEC-out learned checkpoint; "Random inst." and "SK inst." are unseen source-scale instances; size and depth alter random/SK geometry; and the final columns perturb QGT/force estimates. Tikhonov is strongest where its source-selected bias remains well matched, scaled Tikhonov is strongest on the near-floor QEC target and size shift, and ScalarGuard is strongest at both frozen noise strengths. SpectralGuard is not the single-column winner, but it avoids the large regime-specific degradation exhibited by methods whose inductive bias is better suited to only one part of the benchmark. This is the pattern anticipated by the transfer motivation: learning is valuable here not because it dominates a correctly tuned fixed solver on its home distribution, but because it reduces sensitivity to which inverse geometry is encountered after deployment.

Supplementary Section E expands the clean-transfer analysis with additional comparator-level intervals and the per-seed compound diagnostics. Its split-level results preserve the same direction as Figure 2c, while the compound table shows that retraining heterogeneity is substantial and must be interpreted separately from the frozen confirmatory transfer results.

Table 1: Mean trajectory infidelity on the frozen confirmatory benchmarks (lower is better). Clean ID denotes held-out tasks from the source families at the source size/depth range with no added estimation noise. QEC family denotes the held-out two-block LDPC target; learned-controller entries use the predeclared leave-QEC-out checkpoint, while conventional baselines retain their single source-selected frozen hyperparameters. Random and SK instance columns use unseen instances at source geometry; size and depth shifts change random/SK system size or circuit depth; the final two columns are mixed-source-family QGT-estimation-noise benchmarks. Bold entries mark the best principal regularized solver in each column.

Method	Clean ID	TB-LDPC	Random	SK	Size	Depth	Noise .003	Noise .01
Tikhonov	**0.0259**	7.72e-06	**0.0433**	**0.0207**	0.0790	**0.0383**	0.7394	0.7100
Scaled Tikhonov	0.0284	**2.73e-06**	0.0839	0.0636	**0.0766**	0.1705	0.7652	0.6693
QNG-LM	0.0324	1.51e-05	0.0505	0.0272	0.0955	0.0465	0.6514	0.6165
ScalarGuard	0.0303	2.88e-06	0.0508	0.0243	0.0983	0.0462	**0.4999**	**0.5574**
SpectralGuard	0.0294	5.88e-06	0.0463	0.0250	0.0829	0.0404	0.5808	0.5996

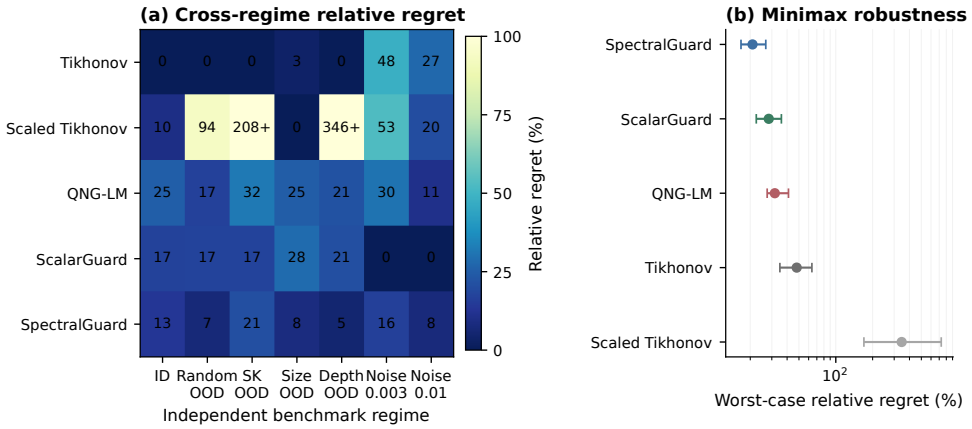

Figure 3: Frozen cross-regime robustness. (a) Relative regret from Eq. 7 across seven independent non-floor confirmatory regimes; darker cells are better, and the two noise columns are categorical settings rather than a fitted trend. (b) Worst-case relative regret with bootstrap 95% intervals; lower is more uniformly near the regime-specific best.

4.2 Cross-regime robustness favors consistently near-best behavior

To compare heterogeneous non-floor regimes on a common scale, let $\bar{I}_{m,r}$ be the mean trajectory infidelity of method m in regime r, and let $\bar{I}^*_r = \min_{m'} \bar{I}_{m',r}$ over the principal regularized methods. We define the *relative regret*

$$R_{m,r} = \frac{\bar{I}_{m,r}}{\bar{I}^*_r} - 1. \tag{7}$$

Here "regret" is descriptive normalized excess error on fixed benchmarks, not cumulative online-learning regret; the per-regime minimum is computed only after evaluation and never retunes a method. Near-floor TB-LDPC transfer and the post-confirmatory compound test are excluded, leaving seven non-floor confirmatory regimes. Figure 3 reveals why the aggregate matters. Fixed Tikhonov variants achieve very low regret in regimes aligned with their scalar bias but incur much larger mismatch elsewhere, while ScalarGuard trades some clean-structure accuracy for substantially better behavior under estimator perturbation. SpectralGuard occupies a different compromise: mode-wise adaptation is expressive enough to remain within roughly 5–21% of the regime-specific best whenever the observed spectrum is informative, without committing to one global damping scale. Its worst-case $R^{\max}_m = \max_r R_{m,r}$ is 20.9% (bootstrap 95% CI 16.8–26.8%), versus 28.3% for ScalarGuard, 31.7% for QNG–LM, and 47.9% for Tikhonov. Thus the cross-regime result answers the motivating portability question directly: SpectralGuard does not replace the best specialized solver in each regime, but it reduces the penalty for not knowing the target regime in advance.

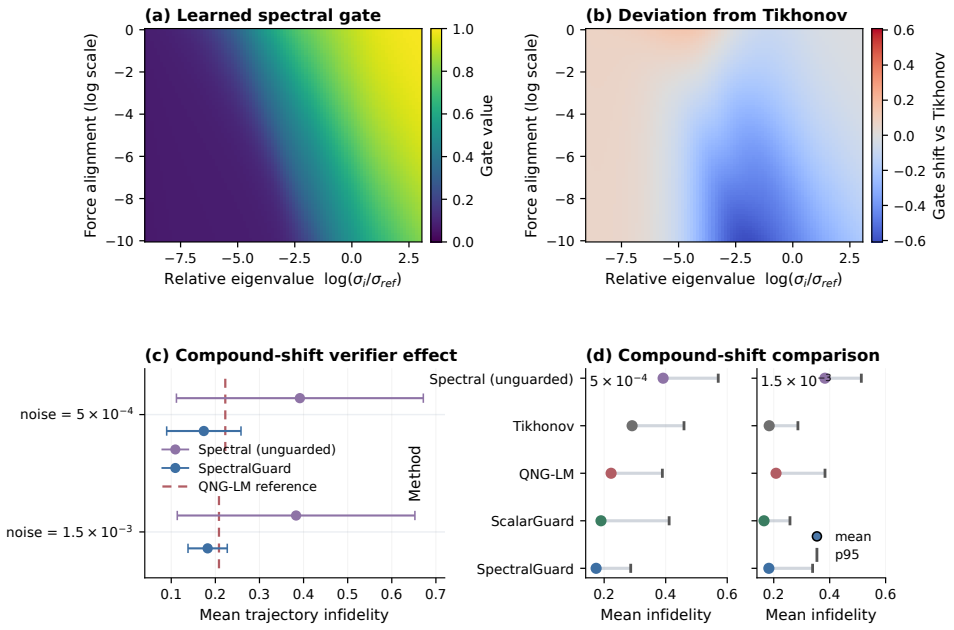

Figure 4: Learned spectral action and exploratory compound-shift behavior. (a) Mode-wise gate over relative QGT eigenvalue and force alignment. (b) Difference from matched scalar Tikhonov; sign changes indicate mode-dependent departure from global damping. (c) Mean ± 1 standard deviation across three training seeds for the unguarded spectral controller and SpectralGuard; the red dashed segment is the QNG-LM reference, not an uncertainty interval. (d) Method comparison in the same two compound-shift cells: circles are means and vertical ticks are 95th percentiles. The facets are categorical stress tests, not a fitted noise trend; (c–d) are post-confirmatory and descriptive.

Supplementary Section D contains the planned inferential tests: pooled structural SpectralGuard-versus-QNG-LM superiority is established, while the remaining planned hypothesis families are not; Figure 3 is therefore descriptive synthesis.

4.3 Spectral action explains both transfer and failure

Figure 4a–b provides a mechanistic view of the tradeoff above. The learned gate decreases toward weak QGT modes, reproducing the stabilizing role expected from classical spectral regularization, but its deviation from matched scalar Tikhonov changes sign with force–mode alignment. SpectralGuard therefore does not merely learn a different global ridge: it uses information about whether a mode is both geometrically weak and dynamically relevant. This additional degree of freedom explains why the method can accommodate structural shifts whose spectra move relative to the source distribution. It also identifies the likely failure mode: if the estimated spectrum or force projection is itself corrupted, the controller's most informative features become unreliable.

Figure 4c–d is consistent with that prediction. In the exploratory compound cells, verification reduces the seed-averaged unguarded spectral mean from 0.392 to 0.174 at noise 5×10^{-4} and from 0.383 to 0.183 at 1.5×10^{-3}; the large unguarded standard deviations show that fine-grained spectral control is also more sensitive to the learned policy when structural and estimator shifts coincide. SpectralGuard is best at the lower compound-noise cell, whereas ScalarGuard is best at the higher one. The two points are categorical stress tests rather than a noise trend, but their ordering matches the model-class distinction: when spectral information remains usable, mode-wise control can exploit it; when that information is degraded, the lower-dimensional scalar action is less exposed to mode-level estimation error.

The verifier changes the interpretation of this failure mode. On frozen clean structural OOD it lowers mean error by 0.0215 relative to the unverified spectral controller (95% CI $[-0.0528, -0.00018]$),

but under the strongest estimator perturbations it rejects nearly all spectral proposals: fallback is 94.7% at noise 0.003 and 97.7% at noise 0.01, versus 12.1% on clean ID and 14.8–38.3% across clean non-QEC structural shifts. Thus verification does not make the learned spectral proposal intrinsically noise-robust; it limits the consequence of leaving the regime in which spectral adaptation is trustworthy. This closes the loop with the motivating design choice: adaptation and admissibility are separated because the numerical features useful for zero-shot transfer are precisely the features that can become unreliable under estimator shift. Supplementary Sections E–G report the per-seed compound values, complete fallback table, strong-noise failure analysis, and simulator-side classical-control accounting.

5 Discussion and Conclusion

The study begins from a deployment problem rather than a clean-ID optimization objective: QGT regularization must be chosen repeatedly even though the inverse geometry changes from one Hamiltonian, circuit, or estimator condition to another. The experiments support that premise in two complementary ways. First, the source-family diagnostic shows that the damping scale preferred by a conventional solver is highly non-universal. Second, no single method is best in every frozen regime. A source-selected Tikhonov solver remains the strongest choice when deployment matches validation, whereas geometry-conditioned controllers become useful when the cost of carrying one fixed numerical bias across regimes dominates the benefit of specialization.

The difference between ScalarGuard and SpectralGuard sharpens this interpretation. SpectralGuard gains flexibility by assigning separate trust to QGT modes and is correspondingly more uniform across clean structural shifts. ScalarGuard discards that fine structure and is less adaptive, but its lower-dimensional action is more stable once estimator corruption makes individual modes unreliable. The learned gate visualization and compound stress tests therefore connect the empirical outcome to the method design: the relevant axis is not "learned versus fixed" in isolation, but how much spectral detail remains trustworthy after distribution shift.

The deterministic verifier is important for the same reason. It should not be interpreted as evidence that the learned spectral proposal itself remains reliable under severe noise; at the strongest perturbations the deployed method behaves predominantly as the QNG–LM fallback. Its contribution is instead to convert an uncertain learned proposal into a conservative numerical procedure with an explicit escape path. An automatic selector between scalar and spectral controllers, driven by resolved rank or estimator uncertainty, is a natural extension, but is not part of the evaluated method.

Several boundaries remain. The estimator-noise model is synthetic and norm-scaled rather than calibrated to a finite-shot hardware covariance model. The QEC branch evaluates transfer across stabilizer-derived diagonal-Ising Hamiltonians and does not measure decoder or fault-tolerance performance. The simulator study is also small scale ($n \leq 18$ in the size-shift panel); dense spectral control requires $O(d^3)$ classical eigendecomposition and $O(d^2)$ storage in parameter dimension d, in addition to the quantum cost of estimating G_t and f_t. Within these limits, the evidence supports a precise positioning: guarded learned regularization is an amortized numerical adaptation layer for repeated quantum-geometric inverse problems. It preserves the analytic solver structure, improves portability when local geometry changes, and reverts to a trusted reference when the information supporting that adaptation is no longer reliable.

AI use statement

Generative AI tools were used to improve manuscript wording and presentation, including figure-layout refinement, and to assist with code verification, debugging, and consistency checks between analysis/plotting scripts and the reported results. All AI-assisted edits and checks were manually reviewed by the authors. The authors verified the reported experiments and numerical results and take responsibility for the final content of the submission.

Reproducibility statement

Full protocol, source-only baseline selection, frozen statistical definitions, transfer diagnostics, noise ablations, resource accounting, and source/configuration fingerprints are provided in Supplementary Sections A–G and the accompanying reproducibility manifest.

References

Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

Aeishah Ameera Anuar, François Jamet, Fabio Gironella, Fedor IV Šimkovic, and Riccardo Rossi. Operator-projected variational quantum imaginary time evolution. *arXiv preprint arXiv:2409.12018*, 2024.

Hanyan Cao, Dongyang Feng, Cheng Ye, and Feng Pan. Differentiable maximum likelihood noise estimation for quantum error correction, 2026.

Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, 2012. doi: 10.1103/PhysRevA.86.032324.

Julien Gacon, Jannes Nys, Riccardo Rossi, Stefan Woerner, and Giuseppe Carleo. Variational quantum time evolution without the quantum geometric tensor. *Physical Review Research*, 6 (1):013143, 2024. doi: 10.1103/PhysRevResearch.6.013143.

Craig Gidney. Stim: a fast stabilizer circuit simulator. *Quantum*, 5:497, 2021. doi: 10.22331/q-2021-07-06-497.

Niladri Gomes, Anirban Mukherjee, Feng Zhang, Thomas Iadecola, Cai-Zhuang Wang, Kai-Ming Ho, Peter P. Orth, and Yong-Xin Yao. Adaptive variational quantum imaginary time evolution approach for ground state preparation. *Advanced Quantum Technologies*, 4(12):2100114, 2021. doi: 10.1002/qute.202100114.

Google Quantum AI. Exponential suppression of bit or phase errors with cyclic error correction. *Nature*, 595:383–387, 2021. doi: 10.1038/s41586-021-03588-y.

Per Christian Hansen. *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. SIAM, Philadelphia, PA, 1998. doi: 10.1137/1.9780898719697.

James Harrison, Luke Metz, and Jascha Sohl-Dickstein. A closer look at learned optimization: Stability, robustness, and inductive biases. In *Advances in Neural Information Processing Systems*, volume 35, pp. 3758–3773, 2022. doi: 10.52202/068431-0272.

Bálint Koczor and Simon C. Benjamin. Quantum natural gradient generalized to noisy and nonunitary circuits. *Physical Review A*, 106(6):062416, 2022. doi: 10.1103/PhysRevA.106.062416.

Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.

Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963. doi: 10.1137/0111030.

Sam McArdle, Tyson Jones, Suguru Endo, Ying Li, Simon C. Benjamin, and Xiao Yuan. Variational ansatz-based quantum simulation of imaginary time evolution. *npj Quantum Information*, 5(1): 75, 2019. doi: 10.1038/s41534-019-0187-2.

Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, NY, 2 edition, 2006. doi: 10.1007/978-0-387-40065-5.

Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical LDPC codes. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 375–388, 2022. doi: 10.1145/3519935.3520017.

James Stokes, Josh Izaac, Nathan Killoran, and Giuseppe Carleo. Quantum natural gradient. *Quantum*, 4:269, 2020. doi: 10.22331/q-2020-05-25-269.

Andrey N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics Doklady*, 4:1035–1038, 1963.

Barnaby van Straaten and Bálint Koczor. Measurement cost of metric-aware variational quantum algorithms. *PRX Quantum*, 2(3):030324, 2021. doi: 10.1103/PRXQuantum.2.030324.

Olga Wichrowska, Niru Maheswaranathan, Matthew W. Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Nando de Freitas, and Jascha Sohl-Dickstein. Learned optimizers that scale and generalize. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3751–3760. PMLR, 2017.