

TABLE OF CONTENTS

S. NO	PAPER TITLE WITH AUTHOR	PAGE NO
01	Equivariant Graph Neural Networks for Scale-Invariant Quantum Error Decoding in Fault-Tolerant Surface Codes <i>V. Udhayakumar, and Dr.A. Karunamurthy</i>	01
02	Distributed Online Convex Optimization for Adaptive Real-Time Decoding in Fault-Tolerant Quantum Error Correction <i>Dr. T. Coumaressin, and Dr.A. Karunamurthy</i>	27
03	Differentiable Graph Rewiring for Real-Time Fault-Tolerant Quantum Error Decoding <i>Dr. N. Moganarangan, and Dr.A. Karunamurthy</i>	60
04	Variational Tensor Network Decoding with Reinforcement Learning for Scalable Fault-Tolerant Quantum Error Correction Under Non-Markovian Noise <i>P. Sudharsan, and Dr.A. Karunamurthy</i>	95
05	Stochastic Approximate Bayesian Decoding for Real-Time Fault-Tolerant Quantum Error Correction <i>Dr J. VinothKumar, and Dr.A. Karunamurthy</i>	125
06	Symmetry-Driven Block Diagonalization of Quantum Fisher Information for Efficient Optimization in NISQ Devices <i>S. Padmapriya, and Dr.A. Karunamurthy</i>	155
07	Differentiable Surrogate Noise Modeling for Error-Mitigated Ansatz Architecture Search in Near-Term Variational Eigensolvers <i>D. Sudha and Dr.A. Karunamurthy</i>	185

S. NO	PAPER TITLE WITH AUTHOR	PAGE NO
08	Graph-Embedded Transfer Learning for Noise-Aware Ansatz Architecture Search in Near-Term Variational Eigensolvers <i>K. Nirmal, V. Hemachandran, and Dr.A. Karunamurthy</i>	214
09	Cost-Aware Multi-Objective Bayesian Optimization for Error-Mitigated Ansatz Architecture Search in Near-Term Variational Eigensolvers <i>Dr. G. Ramya, and Dr.A. Karunamurthy</i>	247
10	Topology-Guided Noise Prediction and Probabilistic Error Cancellation for Adaptive Variational Quantum Eigensolvers <i>A. Anandh Babu, and Dr.A. Karunamurthy</i>	279

LIST OF TABLES

TABLE NO.	TABLES	PAGE NO.
1.1	Logical error rates at physical error rate $p = 0.005$ for varying code distances. The proposed equivariant GNN decoder is trained on $d \in \{3,5,7,9\}$ and evaluated zero-shot on $d = 11$. Standard deviations are computed over three independent runs.	14
1.2	Zero-shot generalization performance at physical error rate $p = 0.005$ for code distances not seen during training. The proposed equivariant GNN and generic GNN are trained on $d \in \{3,5,7,9\}$. MWPM serves as the optimal reference.	16
1.3	Average inference time per syndrome measurement (in milliseconds) for varying code distances. Neural network decoders are evaluated on GPU, while classical decoders run on CPU.	18
1.4	Ablation study results at physical error rate $p = 0.005$ and code distance $d = 9$. The full model corresponds to our proposed equivariant GNN decoder with all components.	19
2.1	Logical Error Rates at Code Distance $d = 7$ Across Three Noise Scenarios	46
2.2	Ablation Study: Logical Error Rates for Varying Coupling Strength β and Patch Size L at $d = 7$ Under Spatial Gradient Noise	50
3.1	Logical error rates for the rotated surface code with distance $d = 7$ under circuit-level depolarizing noise at three physical error rates.	81
3.2	Ablation study results for the DGR-BP decoder on the surface code with distance $d = 7$ under depolarizing noise at $p = 5 \times 10^{-3}$.	85

4.1	Logical error rate ($\times 10^{-4}$) at physical error rate $p_0 = 10^{-3}$ for varying noise correlation times τ_c . Lower values indicate better performance.	110
4.2	Normalized contraction cost and decoding latency (μs) for varying code distances at $\tau_c = 10$.	112
4.3	Action selection frequency, logical error rate ($\times 10^{-4}$), and normalized contraction cost for the RL agent at $p_0 = 10^{-3}$.	114
4.4	Ablation study results at $p_0 = 10^{-3}$ and $\tau_c = 20$. Checkmarks indicate the inclusion of each component.	116
4.5	Estimated error thresholds (%) for varying noise correlation times.	118
5.1	Decoder latency (in nanoseconds) for surface codes of varying distances at physical error rate $p = 10^{-3}$. The SABD latency is deterministic and measured on the Xilinx Versal ACAP platform. Baseline decoder latencies are measured on an Intel Xeon Gold 6248R CPU.	141
5.2	Logical error rates (per cycle) at physical error rate $p = 10^{-3}$ for surface codes of varying distances. The SABD parameters are $N_{\text{MCMC}} = 10d$ and $N_{\text{samp}} = 5d$.	143
5.3	Ablation study of SABD hyperparameters for a distance-13 surface code at physical error rate $p = 10^{-3}$. The baseline configuration is $N_{\text{MCMC}} = 10d = 130$ and $N_{\text{samp}} = 5d = 65$.	145
6.1	Average QFIM Inversion Time per Iteration (milliseconds)	170
6.2	Convergence Performance on VQE and QAOA Benchmarks	172
6.3	Final VQE Energy (Ha) under Varying Depolarizing Noise Rates	173
6.4	QAOA Iterations to Convergence under Varying Shot Counts	174
6.5	Ablation Study: Effect of Circuit Depth on Symmetry-Resolved QNG Performance	175

6.6	Ablation Study: Effect of Symmetry Breaking on Method Performance	176
7.1	Ground state energy errors (in Hartree) for the proposed method and baselines under noisy conditions. The FCI reference energies are -1.1373 Hartree for H ₂ , -7.8824 Hartree for LiH, and -75.0126 Hartree for H ₂ O.	202
7.2	Noise robustness scores (NRS) for the proposed method and baselines. Lower values indicate better noise robustness.	203
7.3	Circuit depths (number of two-qubit gate layers) of the final ansatz architectures.	205
7.4	Ablation study results on LiH, showing the impact of removing key components of the proposed framework.	207
8.1	Mean Absolute Error (MAE) of noise deviation prediction stratified by circuit depth and two-qubit gate count on the hardware test set.	233
8.2	Ground state energy estimation accuracy and success rate for molecular Hamiltonians on IBM Quantum hardware.	235
8.3	Ablation study of GENP components on noise deviation prediction accuracy	237
9.1	Performance comparison of optimal ansatz architectures identified by different search methods across three molecular systems. Energy error is reported in kcal/mol relative to FCI, and sampling overhead γ is computed from device-specific gate error rates. Results are averaged over 10 independent runs with standard deviations in parentheses.	266
9.3	Ablation study results for BeH ₂ comparing the full proposed method against variants without hardware-aware kernel modulation and with simplified noise metrics. Energy error is reported in kcal/mol, and overhead γ is the PEC sampling overhead. Results are averaged over 10 independent runs.	271

10.1	Absolute ground state energy errors ΔE (in kcal/mol) relative to FCI reference across molecular systems. Values represent means over 10 independent runs with standard deviations in parentheses.	299
------	---	-----

LIST OF FIGURES

FIGURE NO.	FIGURES	PAGE NO.
1.1	System Level Architecture of Scalable Logical Qubit Stabilization with GNN Decoder	08
1.2	Logical Error Rate vs. Code Distance for Quantum Error Correction Decoders with Zero-Shot Generalisation at $d=11$	15
1.3	Logical error rate as a function of physical noise strength and code distance for the equivariant GNN decoder, demonstrating the error suppression threshold where increasing code distance reduces logical errors below the threshold physical error rate.	16
1.4	Zero-shot logical error rate vs. code distance Comparison of MWPM, Generic GNN, and Equivariant GNN (Ours) at $d = 11, 13, 15$ (log scale)	17
1.5	Decoding time vs. code distance Average decoding time (ms) for MWPM, Union-Find, Generic GNN, and Equivariant GNN (Ours) at $d = 5, 9, 15, 21$	18
1.6	Ablation study — logical error rate by model variant Lower is better. Purple = full model (baseline); gray = ablated variants.	19
1.7	Learned message-passing vector fields for a representative syndrome pattern under original and 90-degree rotated configurations, demonstrating that the internal representations rotate consistently with the input lattice, verifying the geometric equivariance of the neural network updates.	20
2.1	System Level Architecture of Adaptive Quantum Decoding	35
2.2	Logical Error Rate Comparison of Decoding Methods Under Three Noise Conditions	46
2.3	: Spatial distribution of learned per-qubit error rates across the 2D lattice after convergence under a simulated magnetic field	47

	gradient, comparing the ground-truth noise profile with the parameters estimated by the distributed online convex optimization protocol.	
2.4	Temporal Evolution of The Cumulative Regret Normalized by The Square Root of Time Steps For Multiple Local Quantum Processing Units, Demonstrating The Convergence Bound Under Varying Degrees of Non-Stationary Noise Fluctuations.	48
3.1	System Level Architecture for Scalable Logical Qubit Stabilization	69
3.2	Logical error rate as a function of physical error rate and code distance, with overlaid iso-performance lines comparing the proposed DGR-BP decoder against standard BP and MWPM baselines.	80
3.3	Logical error rate vs. physical error rate (log scale) for Standard BP, NMS, NBP, MWPM, and DGR-BP decoders.	81
3.3a	Grouped bar comparison of logical error rates across decoders at $p = 10^{-3}$, 5×10^{-3} , and 10^{-2} .	82
3.4	Probability distribution of decoder convergence iterations as a function of noise correlation strength, illustrating the speed-up achieved by dynamic rewiring.	83
3.5	Relationship between local Lovász bound values of induced subgraphs and the bit-flip error correction success rate per check node, demonstrating the rewiring score's effectiveness in identifying problematic graph topologies.	84
3.6a	Ablation study: logical error rate per decoder variant at $p = 5 \times 10^{-3}$ (lower is better).	85
3.6b	Ablation study: BP convergence rate per decoder variant at $p = 5 \times 10^{-3}$ (higher is better).	86
4.1	System Level Architecture of Scalable Logical Qubit Stabilization with VTN-RL Decoder	103

4.2	Internal Architecture of the Variational Tensor Network Decoder	103
4.3	Decoding Latency vs. Syndrome Correlation Time for Quantum Error Correction Decoders	111
4.4	Logical error suppression landscape across varying noise correlation times and physical error rates, showing the suppression factor achieved by the VTN-RL decoder relative to the MWPM baseline.	112
4.5	Computational Cost–Latency Trade-off Across Surface Code Distances: (a) Normalised Computational Cost and (b) Decoding Latency	113
4.6	Adaptive Contraction Strategy Selection and Its Effect on Logical Error Rate and Computational Cost in VTN-RL Across Syndrome Correlation Times: (a) Strategy Usage, (b) LER, and (c) Normalised Cost	115
4.7	Operational Pareto frontier balancing decoding latency and logical fidelity against tensor network complexity, identifying the optimal region for real-time deployment.	115
4.8	Component-wise Ablation Study of VTN-RL: (a) Logical Error Rate and (b) Normalised Computational Cost	117
4.9	Convergence dynamics of variational parameters under RL-guided versus static contraction strategies, showing faster convergence paths selected by the RL agent.	117
4.10	Decoding Throughput Under Low and High Syndrome Correlation Time for Five Quantum Error Correction Decoders	118
5.1	System Level Integration of SABD in Fault Tolerant Quantum Error Correction	134
5.2	Decoding Latency Scaling vs. Code Distance for SABD and Baseline Decoders ($d = 5$ to 21)	141
5.3	Logical error rates as a function of physical error rate for the SABD and exact MWPM decoder across four code distances. The	142

	crossing point of the curves indicates the error threshold, which is preserved under the stochastic approximation.	
5.4	Logical Error Rate Scaling with Code Distance for SABD and Baseline Decoders ($d = 5$ to 17)	143
5.5	Runtime scaling of the SABD and baseline decoders as a function of code distance. The SABD exhibits quadratic scaling, while the exact MWPM decoder shows super-cubic scaling in the average case and quartic scaling in the worst case.	144
5.6	Trade-off between Logical Error Rate (LER) and decoding latency for different MCMC hyperparameter configurations.	146
6.1	Overall System Architecture of Symmetry-Driven VQA	163
6.2	QFIM Inversion Time Comparison (Lower Is Better)	170
6.3	Comparison of the full Quantum Fisher Information Matrix from a standard VQA run and the reconstructed block-diagonal matrix produced by the symmetry analyzer, showing the magnitude of matrix elements with zero-valued off-block regions that the method exploits to reduce inversion complexity.	171
6.4	Two-dimensional projection of the parameter space showing optimization trajectories of the proposed block-diagonal optimizer, conventional natural gradient descent, and standard gradient descent, with the cost function landscape represented by contour lines.	171
6.5	Convergence iterations comparison (Lower is Better)	172
6.6	VQE Final Energy vs. Noise Rate (Higher energy magnitude – better; lower noise sensitivity is better)	173
6.7	(a): Convergence Vs. Shots (line chart – lower is better). (b): Convergence Vs. Shots (Grouped bar – lower is better).	174
6.8	Scalability Analysis of the Proposed Symmetry-Resolved Quantum Natural Gradient (SR-QNC) Method	175

6.9	Circuit Type Comparison U(1)- Symmetric Vs Symmetry-Broken Circuits	176
7.1	System Architecture of Error Mitigated Ansatz Search with Differentiable Surrogate Noise Model	193
7.2	Energy error comparison across VQE methods and molecular benchmarks	202
7.3	Circuit depth comparison across VQE methods and molecular benchmarks	203
7.4	Optimization trajectories on the noisy energy landscape comparing the proposed differentiable surrogate approach against conventional noise-agnostic optimization for a 2-qubit subsystem of LiH.	204
7.5	Metric scaling across VQE methods for H ₂ , LiH, and H ₂ O benchmarks	205
7.6	Gradient vector field in the projected architectural parameter space showing the influence of the differentiable noise model on ansatz selection direction for LiH.	206
7.7	Dependence of ground state energy error on circuit depth and noise intensity for the error-mitigated ansatz search on LiH.	206
7.8	Ablation study: contribution of each component to overall performance	208
8.1	System Architecture of Graph Embedded Transfer Learning for Noise Aware Ansatz Search	223
8.2	Correlation between predicted and actual noise-induced energy deviations across diverse ansatz topologies, with density contours highlighting regions of high and low prediction confidence.	232
8.3	Error Mitigation Performance Across Circuit Categories	233

8.4	Prediction accuracy of the GENP model as a function of the number of fine-tuning samples, demonstrating rapid convergence with limited hardware data.	234
8.5	Comparison of VQE Ground-State Estimation Methods	235
8.6	Pairwise similarity matrix of ansatz circuit embeddings in the latent space, revealing natural clusters that group circuits with comparable noise-resilience profiles.	236
8.7	Ablation study showing the effect of removing key architectural components from GENP	237
9.1	System Architecture of Cost Aware Bayesian Optimization Enhanced VQE	255
9.2	Performance Comparison of Quantum Circuit Optimization Methods Across Benchmark Molecules: (a) Energy Approximation Error and (b) Sampling Overhead	267
9.3	Distribution of evaluated ansatz architectures for LiH in the objective space, showing the trade-off between mitigated energy error and PEC sampling overhead. The empirical Pareto front identified by the proposed method dominates those of the baseline approaches, demonstrating superior navigation of the accuracy-cost trade-off.	268
9.4	Predicted energy landscape and sampling overhead contours in the latent architecture space for LiH, with the optimization trajectory of selected ansatz architectures overlaid. The path illustrates the algorithm's progression from initial exploration toward low-energy solutions within the feasible overhead region.	269
9.5	Convergence Efficiency of Quantum Circuit Optimization Methods: (a) Evaluations to Reach Chemical Accuracy Threshold and (b) Evaluation Budget Utilisation	270

9.6	Ablation study evaluating the contribution of individual components of the proposed method on BeH ₂ . (a) Energy approximation error (kcal/mol); (b) sampling overhead γ ; (c) Pareto hypervolume indicator. Values in parentheses denote standard deviations over 10 independent trials. Lower is better for (a) and (b); higher is better for (c).	272
10.1	Adaptive VQE Loop with Topology-Guided Noise Prediction	288
10.2	Comparison of mean absolute energy error (Hartree) for ADAPT-VQE with different error mitigation methods across LiH, BeH ₂ , and H ₄ molecules	299
10.3	Convergence of the energy error relative to the FCI ground state for BeH ₂ under moderate noise intensity, comparing the proposed TGNP-enhanced ADAPT-VQE against standard ADAPT-VQE and fixed-ansatz VQE.	300
10.4	Computational efficiency and accuracy comparison of ADAPT-VQE error mitigation methods on BeH ₂ (12 qubits) — final ansatz size, total circuit executions ($\times 10^3$), and energy error ΔE (kcal/mol).	301
10.5	GNN noise predictor vs. average baseline RMSE across single-qubit depolarizing rate, two-qubit depolarizing rate, and readout confusion (off-diagonal) — GNN achieves consistently lower prediction error across all noise types.	302
10.6	Ablation study of TGNP components — energy error ΔE (kcal/mol) and final ansatz size across four configurations on BeH ₂ , showing the contribution of GNN-based noise prediction and circuit-specific PEC to overall performance.	303