

Vectorized Attention with Learnable Encoding for Quantum Transformer

Ziqing Guo^{1,2}, Ziwen Pan¹, Alex Khan³, Jan Balewski²

¹Texas Tech University, TX, USA

²National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory, CA, USA

³National Quantum Laboratory, MD, USA

ziqguo@ttu.edu, ziwen.pan@ttu.edu, askhan@umd.edu, balewski@lbl.gov

Abstract

Vectorized quantum block encoding provides a way to embed classical data into Hilbert space, offering a pathway for quantum models, such as Quantum Transformers (QT), that replace classical self-attention with quantum circuit simulations to operate more efficiently. Current QTs rely on deep-parameterized quantum circuits (PQCs), rendering them vulnerable to QPU noise, and thus hindering their practical performance. In this paper, we propose the Vectorized Quantum Transformer (VQT), a model that supports ideal masked-attention matrix computation through quantum approximation simulation and efficient training via vectorized nonlinear quantum encoder, yielding shot-efficient and gradient-free quantum circuit simulation (QCS) and reduced classical sampling overhead. In addition, we demonstrate an accuracy comparison for IBM and IonQ in quantum circuit simulation and competitive results in benchmarking natural language processing tasks on the state-of-the-art Kingston QPU of IBM. Our noise intermediate-scale quantum (NISQ)-friendly VQT approach unlocks a novel architecture for end-to-end machine learning in quantum computing.

Introduction

Quantum computing (QC) holds the promise of solving database searching and RSA decrypting problems faster than classical methods, as demonstrated by Grover’s and Shor’s algorithms (Grover 1996; Shor 1994). This advantage stems from the unique quantum mechanical features, such as superposition, entanglement, and interference, which enable rich parallelism and non-classical correlations. Recent advances in quantum data encoding and quantum arithmetic operations (Amankwah et al. 2022; Balewski et al. 2024, 2025) further extend this capability by efficiently embedding classical data into high-dimensional Hilbert spaces and enhancing representational expressiveness. Vectorized quantum computation dovetails with classical self-attention (Vaswani et al. 2017), allowing high-capacity models to capture intricate nonlinear dependencies, thereby improving generalization, lowering perplexity, and increasing training efficiency.

Among the techniques in quantum computing, quantum circuit learning (QCL) (Mitarai et al. 2018) has emerged

as a leading approach for developing data-driven quantum models because of its capacity to PQC for learning complex patterns. By leveraging controlled unitary operations, QCL enables an efficient representation of structured data, facilitating expressive and compact quantum models that support downstream QT models. These benefits are amplified by recent advances in fault-tolerant quantum architectures (Acharya et al. 2024) and the integration of high-performance computing with QPU (Cacheiro et al. 2025), which collectively provide the infrastructure necessary to scale QCL and QT beyond the limitations of the NISQ era.

The core limitation is that the hybrid QCL paradigm currently exploits the restricted usage of unique quantum phenomena, notably the non-cloning theorem and quantum coherence. In practice, PQCs are trained by adjusting the gate rotations to minimize the classical loss computed from repeated projective measurements that collapse the state and discard full-wavefunction information. This is because neural network feedback is drawn only from classical statistics, entanglement and interference are largely unexploited, and intermediate quantum states remain unused. In addition, crosstalk and gate noise further erode coherence, particularly in deep or non-local circuits that exceed NISQ capabilities (Huang and et al 2024). Although studies of quantum kernel expressivity imply possible benefits (Sim, Johnson, and Aspuru-Guzik 2019), empirical evidence is inconclusive, and current PQC models typically resemble shallow classical networks and have yet to demonstrate a scalable quantum advantage.

In this study, we introduce a Vectorized Quantum Transformer (VQT) that integrates observable-based quantum arithmetic approximation with nonlinear encoding for the vectorized quantum dot product (VQDP) and a vectorized nonlinear quantum encoder (VNQE). Specifically, our quantum self-attention approach is implemented using uniformly controlled entangling gates with single-qubit rotations and CNOT gates, while an expressive feed-forward network supports nonlinear embedded quantum encoding. The resulting circuit batches key–query pairs, requires no trainable quantum parameters, and remains fully compatible with the current NISQ hardware. Experiments on IBM state-of-the-art superconducting devices demonstrated that multiple quantum heads yield accurate attention scores and efficient model training by leveraging VQDP and VNQE.

Background

The original idea of the end-to-end QT workflow stems from the proposed AQT (Cha et al. 2021) inspired by the original transformer work (Vaswani et al. 2017) proved by the work demonstrating the 60 qubits Greenberger-Horne-Zeilinger (GHZ) state (Carrasquilla et al. 2021). This was followed by a fast simulation for an attention mechanism study (Gao et al. 2023) using a Grover search. Such QSAN enables quantum data encoding for natural language processing (NLP) tasks (Zheng, Gao, and Miao 2023) using a trainable parameterized quantum circuit (QPC) improved by variational QT using a quantum fourier transform (QFT) kernel (Evans et al. 2024). Such the application of quantum singular value transform (QSVT) (Khatri et al. 2024) supported quantum signal processing information theory (Eldar and Oppenheim 2002) provides a quantizing dot-product attention mechanism. In addition, the adaptive attention method (Chen and Kuo 2025) optimizes the inner dot-product relationship during the PQC training. Although the theoretical protocols here are too finicky, the success of recent NLP applications on quantum computers is likely to make AI algorithms plausible in the future production-scaling quantum computers (Widdows et al. 2024) with the tool (Guo, Pan, and Balewski 2025). These approaches typically require a learnable long-distance QPC, which is not plausible for near-term noisy quantum computers.

Nonlinear Quantum Data Encoding

Basic quantum data encoding techniques can be categorized into basis encoding, angle encoding, and amplitude encoding (Weigold et al. 2021). Note that the encoding qubit and repetitive data feeding overhead are listed in Table 1. The encoding techniques stem from the transformation of the mapping of classical data into a quantum state $x \mapsto |\psi(x)\rangle$. More generally, let us define a N-dimensional state vector $x = (x_0, \dots, x_{N-1})^\top$. Here, x as the classical input which maps to the multi-party quantum state $\frac{1}{\|x\|_2} \sum_{i=0}^{N-1} x_i |i\rangle$, where $\|x\|_2 = \left(\sum_{i=0}^{N-1} |x_i|^2\right)^{1/2}$ as the coefficients of each party defined by Born’s rule. Specifically, given N features, the unitary operation $U(x)$ acts on the initial fiducial state $|0\rangle^{\otimes n}$. Here, n is the number of qubits derived by $n = \lceil \log_2 N \rceil$. In our model, we inherit the spirit of the expectation-value encoding (EVEN) scheme (Balewski et al. 2025) because of the natural value constraint correlation between the machine learning (ML) activation function tanh and the real value encoding input $x_i \in [-1, 1]$ (see details in Section). The EVEN encoding fits the category of angle encoding because the constructed quantum state can be subjected to single qubit $R_y(\theta_i)$ rotation, where $\theta_i = \arccos(x_i)$.

Quantum State Evolution

To build the quantum circuit (Nielsen and Chuang 2010) model for computation, in which computation is a sequence of quantum gates including single-qubit operation, two-qubit entanglement gate, and multi-qubit gates. In general, we use unitary evolution to preserve the probability of the quantum state, namely Hamiltonian $\mathcal{H} |\psi(t)\rangle =$

Encoding	Number of Qubits	Number of Passes
Basic	N	1
Amplitude	$\log_2(MN)$	1
Angle	N	M

Table 1: Resource scaling for three common data encoding strategies when processing M samples with N features each.

$e^{-iHt/\hbar}|\psi(0)\rangle$. Specifically, the quantum state evolution can be denoted by the observable $O = \langle \psi | H(t) | \psi \rangle$. Here, \mathcal{H} is a sequence of Hermitian matrices. Note that R_y gate specifically introduce a real, continuous rotation of the qubit about y-axis of the Bloch sphere by the angle θ

$$R_y(\theta) = \begin{bmatrix} \cos(\frac{\theta}{2}) & -\sin(\frac{\theta}{2}) \\ \sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2}) \end{bmatrix}, \quad R_y(\frac{\pi}{2}) = \begin{bmatrix} \cos(\frac{\pi}{4}) & -\sin(\frac{\pi}{4}) \\ \sin(\frac{\pi}{4}) & \cos(\frac{\pi}{4}) \end{bmatrix}. \quad (1)$$

By following the EVEN encoding, we consider two classical input values $\{x_0, x_1\}$ in the range $[-1, 1]$. The encoding process involves feeding each value x_i to the parameterized rotation gates $R_y(\theta_i)$, which are then applied to a 2-qubit quantum state initialized in $|00\rangle$. Here, we denote the quantum state as $|\phi_i\rangle$ input composed of tensor products of unitary operations acting on the initial state.

$$|\phi_i\rangle = |00\rangle [R_y\theta_0 \otimes R_y\theta_1] = \frac{1}{2} \begin{bmatrix} \sqrt{1+x_0}\sqrt{1+x_1} \\ \sqrt{1+x_0}\sqrt{1-x_1} \\ \sqrt{1-x_0}\sqrt{1+x_1} \\ \sqrt{1-x_0}\sqrt{1-x_1} \end{bmatrix}, \quad (2)$$

thereby achieving classical information embedded into the two-qubit quantum state. The encoding scheme can be represented as a quantum circuit.

$$|0^i\rangle \left\{ \begin{array}{l} \boxed{R_y\theta_0} \\ \boxed{R_y\theta_1} \end{array} \right\} |\phi_e\rangle, \quad i \in \{0, 1\}, \quad \theta_i = \arccos(x_i), \quad (3)$$

Furthermore, the entanglement gate enables state modification between multiple qubits. For instance, one of the best techniques to showcase the state changes is to add a CNOT gate after the encoded state $|\phi\rangle$ with an R_z gate. By doing this, the second qubit acquires a phase that is conditioned on the logical value of the first qubit. Consequently, the two qubits are no longer in a simple product state. Here, the $|\phi\rangle$ state from (2) can be evolved as

$$\begin{aligned} |\phi_\Pi\rangle &:= |\phi\rangle_e \cdot \text{CNOT}_{0,1} \cdot [I \otimes R_z\pi/2] = \frac{e^{-i\frac{\pi}{4}}}{2} \begin{bmatrix} \sqrt{1+x_0}\sqrt{1+x_1} \\ i\sqrt{1+x_0}\sqrt{1-x_1} \\ i\sqrt{1-x_0}\sqrt{1-x_1} \\ \sqrt{1-x_0}\sqrt{1+x_1} \end{bmatrix} \\ &= e^{-i\pi/4} \left(\cos\frac{\theta_0}{2} \cos\frac{\theta_1}{2} |00\rangle + i \cos\frac{\theta_0}{2} \sin\frac{\theta_1}{2} |01\rangle \right. \\ &\quad \left. + i \sin\frac{\theta_0}{2} \sin\frac{\theta_1}{2} |10\rangle + \sin\frac{\theta_0}{2} \cos\frac{\theta_1}{2} |11\rangle \right). \end{aligned} \quad (4)$$

To expand the encoding into a broader data permutation scenario, the quantum circuit can be represented by a sequence of multiple controlled qubit gates, such as the

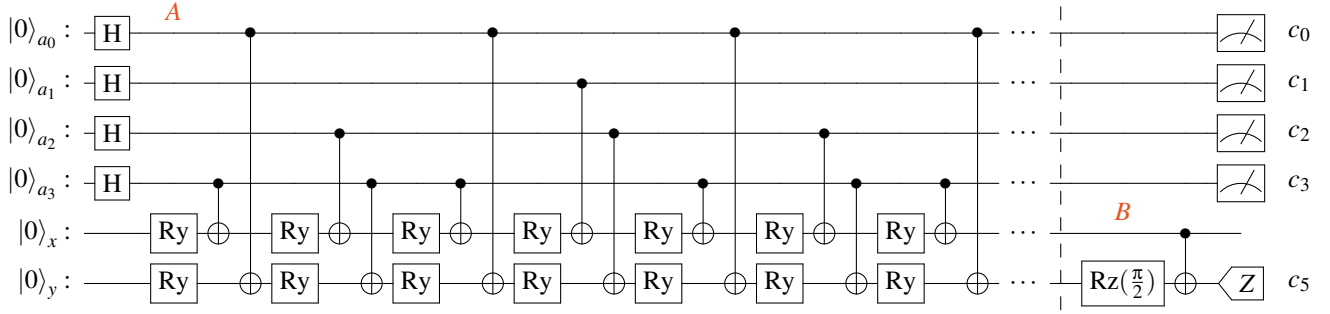


Figure 2: Example quantum circuit computing vectorized product (**VQDP**) $x_i \cdot y_i$ for the batch size of 32 (x_i, y_i) pairs. It uses four qubits for addresses and two qubits for input values of x_i, y_i . **QCrank** (A) encoding is before barrier and **EHands** (B) multiplier after it. The expectation value of register c_4 yields $x_i \cdot y_i$ for index i given by classical registers c_0, \dots, c_3 . Here a_0, \dots, a_3 are address qubits (nq_{addr}) and x, y are data qubits (nq_{data}).

loaded with all-zero vectors; therefore, any state beyond the BT^2 valid pairs contributes nothing to the computed dot products. We recall Eq. (2), the second qubit final measured result $|\phi_{output}\rangle$ is measured by O_z (Z-basis measurement).

$$O_z := I \otimes \sigma_z = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \quad (7)$$

Hence, according to Eq. (7) and Eq. (2), we proceed the final observable result

$$\langle O \rangle = \langle \Psi_{\Pi} | O_z | \Psi_{\Pi} \rangle. \quad (8)$$

Note that we only measure the second data qubit because the CNOT gate does not affect the first qubit. Given Eq. (7) and Eq. (8), the output state analytical result is $|\phi_o\rangle = \vec{x}_i \cdot \vec{y}_i$ which is the simplification of $\langle \phi_o \rangle$. But in quantum circuit simulation (QCS), the numerical expectation output using Z basis which can be typically reconstructed by the quantum Monte Carlo shots denoted by

$$\langle Z \rangle = \frac{n_0 - n_1}{n_0 + n_1}. \quad (9)$$

We would like to emphasize n_0 and n_1 are the number of shots corresponding to the $|0\rangle, |1\rangle$ measurements, respectively, which vary based on the real quantum hardware for each run.

Tanh Projection Head

For the classical input to the quantum attention head, we chose tanh to narrow the input down to output within -1 to 1 because we utilize arccos to encode the data into angles for quantum circuits. The benefits become clearer when the model uses a vectorized quantum dot product to calculate the attention score because of the polarized distributed multiplication. Additionally, recall that it is possible to replace the tanh projection into the QAF (Parisi et al. 2020) but in the case of nonlinear quantum encoding (see Section), the entire sequence of the data serves as the negative value contributing to the probability of the classical output after absolute probability calculation using Born's rule, where the

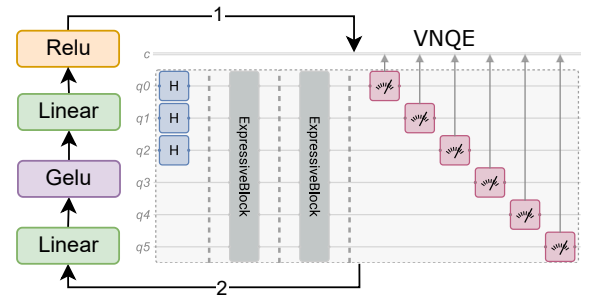


Figure 3: Default three-by-three expressive quantum head. The left deep learning network is the simplified representation of AngleMLP and right quantum circuit is vectorized nonlinear quantum encoder (expressive block shown in Figure 1 of **QCrank**).

QAF has suboptimal behavior. Simply, by constructing a tanh projection head, one can typically confront the tanh function encoding the query and key matrices through the **VQDP**. The remaining architecture follows with a value and self-attention matrix connected with another **VQDP** layer to compare the similarity between the attention and value matrices. Therefore, the entire quantum overhead is twice the number of attention matrices.

Expressive Quantum Head

For the expressive layer, we first model a simple multi-layer perceptron named AngleMLP, as shown in Figure 1 detailed in Figure 3. Arrow 1 represents the untrained sequence of three-dimensional tensors, which are parameterized by the batch size and sequence length, as output from AngleMLP and encoded into **VNQE**. Arrow 2 indicates an approximate observable, depicted as a superpositioned data output, which serves as the input for training AngleMLP. The quantum circuit is composed of the address and data qubits nq_{addr}, nq_{data} representing the number in which the matrix corresponding to the column index is followed by the number of expressive blocks consisting of permuted controlled R_y gates, where the gate enables complex phase transition that is encoded by

the input data angles (notation details referred in Eq. (5)). Consequently, the tensor can be compressed into a vectorized quantum circuit with one pass nonlinearly, which saves the classical nonlinear conversion time from $O(n)$ passes to $O(1)$. Given the unfreezed AngleMLP, during the training phase, the MLP layer enables parameter adjustment before the quantum data encoding, thereby promising zero-gradient transformation because the angles of each individual gate are exactly mapped with the tuned MLP output. This leads to QT improvement, in which quantum computing fulfills the nonlinearity encoding combined with classical neural networks trained using backpropagation.

For **VNQE**, we use qubits to encode the quantum embedded hidden dimension that is calculated with

$$Q_{dim} = 2^{nq_{addr}} * nq_{data}. \quad (10)$$

To maximize the encoded dimension, Eq. (10) can be denoted by

$$Q_{dim}^{max} = 2^{N-1}, \quad (11)$$

where $nq_{addr} = N - 1$, $nq_{data} = 1$, and N is the total number of qubits. In the ideal scenario, for a 156 IBM Marrakesh QPU, the encoded quantum dimension is $4.56719262 \times 10^{46}$; however, we save this for future research because the long-distance qubit crosstalk error cannot be eliminated in today's QPUs. Meanwhile, we also provide the lower bound for qubits concerning the encoded quantum dimension

$$N_{min}(Q_{dim}) = \min_{1 \leq k \leq \log_2 Q_{dim}} \left[k + \frac{Q_{dim}}{2^k} \right], \quad (12)$$

where $Q_{dim} - 2^{nq_{addr}} \geq 1$ and is integer. The loose lower bound is defined as

$$N_{min} = \lceil \log_2(Q_{dim}) \rceil + 1 \quad (13)$$

because we assume the encoded quantum dimension is slightly larger than the required dimension. In our simplest scenario, requiring an encode quantum dimension of 32, we set the default number of qubits to 6.

Metrics

Owing to the scarcity of peer-reviewed research articles, there is currently no standardized model evaluation for QT models. However, drawing from classical evaluation methods, we propose the concept of quantum perplexity (QPL) to assess the complexity of quantum models and evaluate the quality of QT prediction. Given by the perplexity definition $PP(p) = b^{H(p)}$, H is the entropy of the distribution, the QPL metric is defined as follows, given a sequence of tokens x_1, x_2, \dots, x_N

$$QPL = \exp \left(-\frac{1}{N} \sum_{t=1}^N \log p(x_t | x_{<t}) \right). \quad (14)$$

Note that $p(x_t | x_{<t})$ is informed by QCS because the final prediction is based on the measured bit string outcomes for reconstruction.

Experiment

We show that VQT is a NISQ-friendly quantum model that enables an end-to-end quantum gradient-free workflow on noisy quantum hardware. By leveraging Perlmutter's state-of-the-art GPU node, VQT supports classical optimization fully accelerated by running on GPU nodes at the scale (NERSC Documentation 2025). Below, we demonstrate the model performance for two representative cases: (1) quantum attention analysis and (2) model performance evaluation.

Quantum Attention Analysis

Multiplication statistical behaviour In **VQDP** ideally one obtains $\langle Z \rangle_{ideal} = \vec{x} \cdot \vec{y}$.




Figure 4: Accuracy of computation $x_i \cdot y_i$ using 6 qubits circuit shown in Fig. 2 from the ideal Qiskit simulator with $N_{shot}=80,000$. a) Correlation between true and measured values, averaged over 30 batches of 32 randomly sampled (x_i, y_i) pairs. b) distribution residuals have std dev 0.017, which scales with $1/\sqrt{N_{shot}}$.




Figure 5: Results of computation of product $x_i \cdot y_i$ on IBM Kingston using the same configuration as in Fig 4. a) Correlation between the truth and raw measured data is shown in orange. The blue data show the correlation after scaling of the measurement by a factor of 2.22. b) The residuals for the scaled measurement have std dev 0.059.

We observed that the noisy QPU is also able to produce low-error multiplication results with amplitude correction, as shown in Figure 5 compared with Figure 4. Note that halving the error requires approximately four times as many shots. The experimental settings and results with 30 iterated

batch size	num qubits	num shots	Ideal			IBM Kingston			IonQ Aria-1 [†]
			num CX*	CX depth	RMSE	num CZ	CZ depth	RMSE	RMSE
4	4	10K	9	5	0.017	18	14	0.022	0.057
8	5	20K	17	9	0.016	29	25	0.023	0.155
16	6	40K	33	17	0.017	78	69	0.037	0.215
32	7	80K	65	33	0.017	164	121	0.059	–
64	8	160K	129	65	0.016	348	249	0.274	–
128	9	320K	257	129	0.016	704	514	–	–

Table 2: The characterization of circuits employed in simulations and experiments, along with the achieved accuracy for IBM and IonQ quantum platforms, is presented. Note that CX* indicates that the number of CX (CNOT) gates for the corresponding ideal simulators and transpiled IonQ Aria-1 are configured identically. IonQ[†] has all-to-all connectivity with sequential gate implementation, rendering the CX gate count the primary determinant of execution efficiency, whereas IBM’s parallel architecture makes the native controlled-Z (CZ) gate depth more salient. The root mean squared error (RMSE) serves as a confidence interval for real hardware. Results denoted by – are considered to be meaningless.

batches are presented in Table 2, where IBM outperforms IonQ’s hardware in terms of RMSE.




Figure 6: The left panel depicts the distribution of absolute deviations between quantum and classical attention scores computed over all query–key pairs and batches, while the right panel visualizes the corresponding error matrix, in which lighter shades denote smaller discrepancies.

Quantum vs. classical attention We show that, with an input size of (10,10,10) corresponding to batches, sequence length, and features, the error between classical attention and VQDP is constantly below 1.2% with 3.0 million shots. Here, although the average two-qubit gate error for the QPU is approximately $2.5e^{-3}$ (note that the real-time error varies; see , our result indicates that the VQDP method allows for competitive results compared with the classical counterpart.

In addition, replacing the classical attention kernel with a quantum sampler yields conceptual benefits. First, the variance in Eq. (17) acts as a data-dependent stochastic regularizer which stabilizes training in regimes where the classical model tends to overfit (see Section). Second, the quantum mechanism provides nonlinear kernels whose functional forms can be modified in hardware (choice of measurement basis or additional controlled phases) without changing the classical model architecture. Additionally, Table 3 compares the scaling rule of the classical matrix-multiplication (MatMul) algorithm with VQDP.

	Classical matmul	VQDP
Time	$O(BT^2d)$	$O(\log(BT^2)shots)$
Computation cost	1 fused-FLOP	$s \cdot (gt_{1Q} + 2t_{CX}) + t_{ro}$
Runtime	$\lesssim 0.2ms$	$\sim 0.1s$
Parallel scale	SIMD / GPU cores	feature per circuit

Table 3: Complexity analysis for classical and quantum attention (s =shots, g =circuit depth, t_{1Q} =one-qubit gate time, t_{CX} =CNOT time, t_{ro} =read-out time). The wall-clock cost of VQDP is ultimately limited by the circuit-layer-operations-per-second (CLOPS) rate of the QPU (Wack et al. 2021).

Model Evaluation

The Brown Corpus dataset (Ide and Macleod 2001) comprises 1.01 million words across 500 text samples spanning 15 genres of American English prose. We utilize Fasttext (Bojanowski et al. 2016) to compress the inputs as batch-processed matrices specified with three-dimensional tensors decoded for the quantum hidden dimension Q_{dim} . We refer to block A in Figure 2 as the quantum oracle for the expressive quantum head. Note that, the default VQT quantum hidden dimension is 32 (see Eq. (10)) with the qubits settings in Table 4 (we refer hyperparameter settings table for the VQT model evaluation in Appendix for more details).

The evidence demonstrates that VQT provides competitive results as a classical benchmark model denoted by NanoGPT (the smallest transformer). Note that we select the smallest GPT because of the limitation of current noisy qubits in the NISQ era. Interestingly, we observe that increasing the Q_{dim} does not help with the perplexity of the model but slightly improves accuracy. However, we configure no dropout rate for VQT, as indicated in Appendix , which results in a lower loss rate than prior quantum models. This is because VQDP allows nonlinear latent space transformation between classical and quantum layers, which enables the model to have a lower rate of overfitting problems, as discussed in (Kobayashi, Nakaji, and Yamamoto 2022).

Model	QPL	Loss	Param	Qubit
NanoGPT [†] 2023	92.5	0.94	125K	–
Q-LSTM 2022	125.3	1.18	*	6
Quixer 2024	117.1	1.61	*	•
Hybrid QT 2025	127.6	1.08	*	•
VQT (default)	105.4	1.12	*	•
VQT (large encoder)	108.2	1.08	*	12

Table 4: Performance comparison of models on the Brown corpus. Note that each column is averaged over ten runs after 50 epochs. More qubits indicates larger hidden quantum dimensions as described in Eq. (10).

- [†] Benchmark baseline for QT model evaluation.
 * Quantum models inherit classical layers and embedding.
 • Default qubit size of 6 was selected for noisy QPU compatibility.

Discussion

In conclusion, this study proposed a vectorized quantum model that allows data-encoded heuristic self-attention paradigm providing a end-to-end support for multi-head expansion in the future production quantum hardware. Unlike the prior quantum transformer method using learnable VQCs with parameter shift rule to calculate the gradients, our approach accurately and straightforwardly simulated the attention matrix using shot-based quantum circuit. By leveraging quantum data encoding techniques, the model largely ameliorate the overfitting problems, which enables a pathway for potential better reasoning model. We envision the quantum tokenizer might benefit our model at the scale.

Proof of Variance of VQDP

Denote by $b_i \in \{0, 1\}$ the ancilla bit of shot i and map it to $Z_i = +1 - 2b_i$. With $p = \Pr(b_i = 0) = \frac{1}{2}(1 + xy)$, the expectation and variance of each Z_i are

$$\mathbb{E}[Z_i] = (+1)(p) + (-1)(1 - p) = 2p - 1 = xy, \quad (15)$$

$$\text{Var}[Z_i] = \mathbb{E}[Z_i^2] - \mathbb{E}[Z_i]^2 = 1 - (xy)^2, \quad (16)$$

because $Z_i \in \{+1, -1\}$, so $Z_i^2 = 1$. Since the Z_i are independent, the variance of their empirical mean is

$$\text{Var}[\hat{Z}] = \text{Var}\left[\frac{1}{M} \sum_{i=1}^M Z_i\right] = \frac{1}{M^2} \sum_{i=1}^M \text{Var}[Z_i] = \frac{1 - (xy)^2}{M}. \quad (17)$$

The bound for the deviation probability of the estimator $\hat{Z} = \frac{1}{M} \sum Z_i$ can be derived more tightly from the binomial Chernoff inequality by applying it to the sum $\sum b_i$, which is a binomial random variable with parameters M and $q = \Pr(b_i = 1) = \frac{1}{2}(1 - xy)$. For any $0 < \delta < 1$, the binomial Chernoff inequality provides

$$\Pr\left[\left|\frac{1}{M} \sum_{i=1}^M b_i - q\right| \geq \delta q\right] \leq 2 \exp\left(-\frac{\delta^2 q M}{3}\right). \quad (18)$$

Note that $\hat{Z} = \frac{1}{M} \sum (1 - 2b_i) = 1 - 2 \cdot \frac{1}{M} \sum b_i$, so the deviation in \hat{Z} from its mean $xy = 1 - 2q$ is directly tied to

the deviation in the average of the b_i from q . Therefore, by change of variables, for any $\epsilon > 0$,

$$\begin{aligned} \Pr\left[|\hat{Z} - xy| \geq \epsilon\right] &= \Pr\left[\left|\frac{1}{M} \sum b_i - q\right| \geq \frac{\epsilon}{2}\right] \\ &\leq 2 \exp\left(-\frac{\epsilon^2 M}{12q}\right), \end{aligned} \quad (19)$$

where the last step uses $\delta = \epsilon/(2q)$ and substitutes it into (18), providing an exponentially decaying bound on the probability that the empirical estimator \hat{Z} deviates from its expectation xy by more than ϵ , in terms of the number of shots M .

Model Settings

Hyperparameter	Default	Large
Vocabulary size (V)	100	100
Sequence length (T)	6	8
Embedding dimension (d)	32	32
Feed-forward hidden size (d_{ff})	128	128
VQT blocks (B)	1	1
Number of QSE (H)	2	2
Address, data qubits ($nq_{\text{addr}}, nq_{\text{data}}$)	3, 3	6, 6
Encoded dimension Q_{dim}	24	384
Shots per nq_{addr} (S)	1024	1024
AngleMLP hidden size (d_{mlp})	128	128
Dropout (ρ)	0	0
Optimizer	AdamW	AdamW
Learning rate (η)	1×10^{-3}	1×10^{-3}
Batch size (B_{train})	5	5

Table 5: Main hyperparameters for the vectorized quantum transformer (VQT) model configuration and training setup.

Acknowledgments

This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility, and UMD QLab. This study was supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC awards DDR-ERCAP0034486 and DDR-ERCAP0034385. All figures and numerical results reported in this study are publicly available at <https://zenodo.org/records/17178332> to facilitate its reproducibility.

References

- Acharya, R.; Abanin, D. A.; Aghababaie-Beni, L.; Aleiner, I.; Andersen, T. I.; Ansmann, M.; Arute, F.; Arya, K.; Asfaw, A.; Astrakhantsev, N.; et al. 2024. Quantum error correction below the surface code threshold. *Nature*.
 Amankwah, M. G.; Camps, D.; Bethel, E. W.; Van Beeumen, R.; and Perciano, T. 2022. Quantum pixel representations and compression for N-dimensional images. *Scientific reports*, 12(1): 7712.

- Balewski, J.; Amankwah, M. G.; Bethel, E.; Perciano, T.; and Van Beeumen, R. 2025. EHands: Quantum Protocol for Polynomial Computation on Real-Valued Encoded States. *arXiv preprint arXiv:2502.15928*.
- Balewski, J.; Amankwah, M. G.; Van Beeumen, R.; Bethel, E. W.; Perciano, T.; and Camps, D. 2024. Quantum-parallel vectorized data encodings and computations on trapped-ion and transmon QPUs. *Scientific Reports*, 14(1): 3435.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2016. Enriching Word Vectors with Subword Information. *arXiv preprint arXiv:1607.04606*.
- Cacheiro, J.; Sánchez, Á. C.; Rundle, R.; Long, G. B.; Dold, G.; Friel, J.; and Gómez, A. 2025. QMIO: A tightly integrated hybrid HPCQC system. *arXiv preprint arXiv:2505.19267*.
- Camps, D.; Amankwah, M.; Van Beeumen, R.; Bethel, W.; and Perciano, T. 2022. Quantum pixel representations and compression for N-dimensional images. In *APS March Meeting Abstracts*, volume 2022, Y37–008.
- Carrasquilla, J.; Luo, D.; Pérez, F.; Milsted, A.; Clark, B. K.; Volkovs, M.; and Aolita, L. 2021. Probabilistic simulation of quantum circuits using a deep-learning architecture. *Physical Review A*, 104(3): 032610.
- Cha, P.; Ginsparg, P.; Wu, F.; Carrasquilla, J.; McMahon, P. L.; and Kim, E.-A. 2021. Attention-based quantum tomography. *Machine Learning: Science and Technology*, 3(1): 01LT01.
- Chen, C.-S.; and Kuo, E.-J. 2025. Quantum adaptive self-attention for quantum transformer models. *arXiv preprint arXiv:2504.05336*.
- Chen, S. Y.-C.; Yoo, S.; and Fang, Y.-L. L. 2022. Quantum long short-term memory. In *Icassp 2022-2022 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 8622–8626. IEEE.
- Eldar, Y. C.; and Oppenheim, A. V. 2002. Quantum signal processing. *IEEE Signal Processing Magazine*, 19(6): 12–32.
- Evans, E. N.; Cook, M.; Bradshaw, Z. P.; and LaBorde, M. L. 2024. Learning with sasquatch: a novel variational quantum transformer architecture with kernel-based self-attention. *arXiv preprint arXiv:2403.14753*.
- Gao, Y.; Song, Z.; Yang, X.; and Zhang, R. 2023. Fast quantum algorithm for attention computation. *arXiv preprint arXiv:2307.08045*.
- Grover, L. K. 1996. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 212–219.
- Guo, Z.; Pan, Z.; and Balewski, J. 2025. Q-GEAR: Improving quantum simulation framework. *arXiv preprint arXiv:2504.03967*.
- Huang, H.-Y.; and et al. 2024. Learning shallow quantum circuits. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, 1343–1351.
- Ide, N.; and Macleod, C. 2001. The american national corpus: A standardized resource of american english. In *Proceedings of corpus linguistics*, volume 3, 1–7. Lancaster University Centre for Computer Corpus Research on Language.
- Karpathy, A. 2023. nanoGPT. <https://github.com/karpathy/nanoGPT>. Accessed: 2025-06-26.
- Khatiri, N.; Matos, G.; Coopmans, L.; and Clark, S. 2024. Quixer: A quantum transformer model. *arXiv preprint arXiv:2406.04305*.
- Kobayashi, M.; Nakaji, K.; and Yamamoto, N. 2022. Overfitting in quantum machine learning and entangling dropout. *Quantum Machine Intelligence*, 4(2): 30.
- Mitarai, K.; Negoro, M.; Kitagawa, M.; and Fujii, K. 2018. Quantum circuit learning. *Physical Review A*, 98(3): 032309.
- NERSC Documentation. 2025. Perlmutter Architecture. <https://docs.nersc.gov/systems/perlmutter/architecture/>.
- NERSC offers GPU compute cabinet that contains 8 chassis, each containing 8 compute blades and 4 switch blades.
- Nielsen, M. A.; and Chuang, I. L. 2010. *Quantum computation and quantum information*. Cambridge university press.
- Parisi, L.; Neagu, D.; Ma, R.; and Campean, F. 2020. QReLU and m-QReLU: Two novel quantum activation functions to aid medical diagnostics. *arXiv preprint arXiv:2010.08031*.
- Shor, P. W. 1994. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th annual symposium on foundations of computer science*, 124–134. Ieee.
- Sim, S.; Johnson, P. D.; and Aspuru-Guzik, A. 2019. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12): 1900070.
- Smaldone, A. M.; Shee, Y.; Kyro, G. W.; Farag, M. H.; Chandani, Z.; Kyoseva, E.; and Batista, V. S. 2025. A hybrid Transformer architecture with a quantized self-attention mechanism applied to molecular generation. *Journal of Chemical Theory and Computation*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wack, A.; Paik, H.; Javadi-Abhari, A.; Jurcevic, P.; Faro, I.; Gambetta, J. M.; and Johnson, B. R. 2021. Quality, Speed, and Scale: three key attributes to measure the performance of near-term quantum computers. *arXiv:2110.14108*.
- Weigold, M.; Barzen, J.; Leymann, F.; and Salm, M. 2021. Encoding patterns for quantum algorithms. *IET Quantum Communication*, 2(4): 141–152.
- Widdows, D.; Aboumradi, W.; Kim, D.; Ray, S.; and Mei, J. 2024. Quantum natural language processing. *KI-Künstliche Intelligenz*, 1–18.
- Zheng, J.; Gao, Q.; and Miao, Z. 2023. Design of a Quantum Self-Attention Neural Network on Quantum Circuits. In *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 1058–1063. IEEE.