
SuperEncoder: Towards Iteration-Free Approximate Quantum State Preparation

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Numerous quantum algorithms operate under the assumption that classical data
2 has already been converted into quantum states, a process termed Quantum State
3 Preparation (QSP). However, achieving precise QSP requires a circuit depth that
4 scales exponentially with the number of qubits, making it a substantial obstacle in
5 harnessing quantum advantage. Recent research suggests using a Parameterized
6 Quantum Circuit (PQC) to approximate a target state, offering a more scalable
7 solution with reduced circuit depth compared to precise QSP. Despite this, the need
8 for iterative updates of circuit parameters results in a lengthy runtime, limiting its
9 practical application. To overcome this challenge, we introduce SuperEncoder,
10 a pre-trained classical neural network model designed to directly estimate the
11 parameters of a PQC for any given quantum state. By eliminating the need for
12 iterative parameter tuning, SuperEncoder represents a pioneering step towards
13 iteration-free approximate QSP.

14 1 Introduction

15 Quantum Computing (QC) leverages quantum mechanics principles to address classically intractable
16 problems [47, 36]. Various quantum algorithms have been developed, encompassing quantum-
17 enhanced linear algebra [15, 48, 45], Quantum Machine Learning (QML) [26, 19, 1, 33, 50, 3],
18 quantum-enhanced partial differential equation solvers [31, 13], etc. A notable caveat is that those
19 algorithms assume that classical data has been efficiently loaded into a specific quantum state, a
20 process known as Quantum State Preparation (QSP).

21 However, the realization of QSP presents significant challenges. Ideally, we expect each element of
22 the classical data to be precisely transformed into an amplitude of the corresponding quantum state.
23 This precise QSP is also known as Amplitude Encoding (AE). However, a critical yet unresolved
24 problem of AE is that the required circuit depth grows exponentially with respect to the number of
25 qubits [34, 41, 29, 46, 49]. Extensive efforts have been made to alleviate this issue, but they fail to
26 address it fundamentally. For example, while some methods introduce ancillary qubits for shallower
27 circuit [57, 56, 2], they may encounter an exponential number of ancillary qubits. Other methods aim
28 at preparing *special* quantum states with lower circuit depth, being only effective for either sparse
29 states [12, 32] or states with some special distributions [14, 17]. To summarize, realizing AE for
30 *arbitrary* quantum states still remains *non-scalable* due to its exponential resource requirement with
31 respect to the number of qubits. Moreover, in the Noisy Intermediate-Scale Quantum (NISQ) era [42],
32 hardware has limited qubit lifetimes and confronts a high risk of decoherence errors when executing
33 deep circuits, further exacerbating the problem of AE.

34 In fact, precise QSP is unrealistic in the present NISQ era due to the inherent errors of quantum
35 devices. Hence, iteration-based Approximate Amplitude Encoding (AAE) emerges as a promising
36 technique [59, 35, 52]. Specifically, AAE constructs a quantum circuit with tunable parameters, then

37 it iteratively updates the parameters to approximate a target quantum state. Since the updating of
 38 parameters can be guided by states obtained from noisy devices, AAE is robust to noises, becoming
 39 especially suitable for NISQ applications. More importantly, AAE has been shown to have shallow
 40 circuit depth [35, 52], making it more scalable than AE.

41 Unfortunately, AAE possesses a drawback that significantly undermines its potential advantages — the lengthy
 42 runtime stemming from iterative optimizations of parameters. For example, when a Quantum Neural Network
 43 (QNN) [3] is trained and deployed, the runtime of AAE dominates the inference time as we demonstrated in Fig. 1.
 44 Since loading classical data into quantum states becomes the bottleneck, the potential advantage of QNN diminishes
 45 no matter how efficient the computations are done on quantum devices.

51 Compared to AAE, AE employs a pre-defined arithmetic decomposition procedure to construct a circuit, thereby
 52 becoming much *faster* than AAE at runtime. Therefore, it is natural to ask: can we realize both *fast* and *scalable*
 53 methods for *arbitrary* QSP? This is precisely the question we tackle in this paper. Overall, we present three major
 54 contributions.

- 58 • Given a Parameterized Quantum Circuit (PQC) $U(\theta)$ that approximates a target quantum state,
 59 with θ the parameter vector. We show that there exists a *deterministic* transformation f that could
 60 map an arbitrary state $|\mathbf{d}\rangle$ to its corresponding parameters θ . Consequently, the parameters can be
 61 designated by f without time-intensive iterations.
- 62 • We show that the mapping f is *learnable* by utilizing a classical neural network model, which
 63 we term as SuperEncoder. With SuperEncoder, you can have your cake and eat it too, i.e.,
 64 simultaneously realizing *fast* and *scalable* QSP. We develop a prototype model and shed light on
 65 insights into its training methodology.
- 66 • We verify the effectiveness of SuperEncoder on both synthetic dataset and representative down-
 67 stream tasks, paving the way toward iteration-free approximate quantum state preparation.

68 2 Preliminaries

69 In this section, we commence with some basic concepts about quantum computing [36], and then
 70 proceed to a brief retrospect of existing QSP methods.

71 2.1 Quantum Computation

72 We use Dirac notation throughout this paper. A *pure quantum state* is defined by a vector $|\cdot\rangle$ named
 73 ‘ket’, with the unit length. A state can be written as $|\psi\rangle = \sum_{j=1}^N \alpha_j |j\rangle$ with $\sum_j |\alpha_j|^2 = 1$, where
 74 $|j\rangle$ denotes a computational basis state and N represents the dimension of the complex vector
 75 space. *Density operators* describe more general quantum states. Given a mixture of m pure states
 76 $\{|\psi_i\rangle\}_{i=1}^m$ with probabilities p_i and $\sum_i p_i = 1$, the density operator ρ denotes the *mixed state* as
 77 $\rho = \sum_{i=1}^m p_i |\psi_i\rangle\langle\psi_i|$ with $\text{Tr}(\rho) = 1$, where $\langle\cdot|$ refers to the conjugate transpose of $|\cdot\rangle$. Generally,
 78 we use the term *fidelity* to describe the similarity between an erroneous quantum state and its
 79 corresponding correct state.

80 The fundamental unit of quantum computation is the quantum bit, or *qubit*. A qubit’s state can be
 81 expressed as $\psi = \alpha|0\rangle + \beta|1\rangle$. Given n qubits, the state is generalized to $|\psi\rangle = \sum_j |j\rangle$, where
 82 $|j\rangle = |j_1 j_2 \dots j_n\rangle$ with j_k the state of k th qubit in computational basis, and $j = \sum_{k=1}^n 2^{n-k} j_k$.
 83 Applying *quantum operations* evolves a system from one state to another. Generally, these operations
 84 can be categorized into quantum gates and measurements. Typical single-qubit gates include the
 85 Pauli gates $X \equiv \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, $Y \equiv \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$, $Z \equiv \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$. These gates have associated rotation operations
 86 $R_P(\theta) \equiv e^{-i\theta P/2}$, where θ is the rotation angle and $P \in \{X, Y, Z\}$ ¹. Muti-qubit operations create

¹In this paper, R_z, R_y are equivalent to R_Z, R_Y .




Figure 1: Breakdown of normalized runtime for QNN inference. Original data are listed in Table 1.

87 *entanglement* between qubits, allowing one qubit to interfere with others. In this work, we focus on
 88 the controlled-NOT (CNOT) gate, with the mathematical form of CNOT $\equiv |0\rangle\langle 0| \otimes \mathbf{I}_2 + |1\rangle\langle 1| \otimes X$.
 89 Quantum measurements extract classical information from quantum states, which is described by
 90 a collection $\{M_m\}$ with $\sum_m M_m^\dagger M_m = \mathbf{I}$. Here, m refers to the measurement outcomes that may
 91 occur in the experiment, with a probability of $p(m) = \langle \psi | M_m^\dagger M_m | \psi \rangle$. The post-measurement state
 92 of the system becomes $M_m |\psi\rangle / p(m)$.

93 A *quantum circuit* is the graphical representation of a series of quantum operations, which can be
 94 mathematically represented by a unitary matrix U . In the NISQ era, PQC plays an important role
 95 as it underpins variational quantum algorithms [11, 39]. Typical PQC has the form of $U(\theta) =$
 96 $\prod_i U_i(\theta_i) V_i$, where θ is its parameter vector, $U_i(\theta_i) = e^{-i\theta_i P_i/2}$ with P_i denoting a Pauli gate, and
 97 V_i denotes a fixed gate such as CNOT. For example, a PQC composed of R_y gates and CNOT gates
 98 is depicted in Fig. 2.

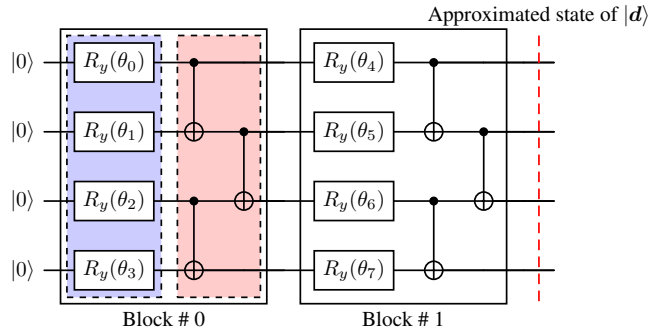


Figure 2: An example PQC with two blocks, with each block consisting of a rotation layer (filled blue) plus an entangler layer (filled red).

99 2.2 Quantum State Preparation

100 Successful execution of many quantum algorithms requires an initial step of loading classical data
 101 into a quantum state [5, 15], a process known as *quantum state preparation*. This procedure involves
 102 implementing a quantum circuit to evolve a system to a designated state. Here, we focus on *amplitude*
 103 *encoding* and formalize its procedure as follows. Let \mathbf{d} be a real-valued N -dimensional classical
 104 vector, AE encodes \mathbf{d} into the amplitudes of an n -qubit quantum state $|\mathbf{d}\rangle$, where $N = 2^n$. More
 105 specifically, the data quantum state is represented by $|\mathbf{d}\rangle = \sum_{j=0}^{N-1} d_j |j\rangle$, where d_j denotes the j th
 106 element of the vector \mathbf{d} , and $|j\rangle$ refers to a computational basis state. The main objective is to generate
 107 a quantum circuit U that initializes an n -qubit system by $U|0\rangle^{\otimes n} = \sum_{j=0}^{N-1} \alpha_j |j\rangle$, whose amplitudes
 108 $\{\alpha_j\}$ are equal to $\{d_j\}$. It is widely recognized that constructing such a circuit generally necessitates
 109 a circuit depth that scales exponentially with n [34, 41]. This property makes AE impractical in
 110 current NISQ era, as decoherence errors [23] can severely dampen the effectiveness of AE as the
 111 number of qubits increases [52].

112 In response to the inherent noisy nature of current devices, *approximate amplitude encoding* has
 113 emerged as a promising technique [59, 35, 52]. Specifically, AAE utilizes a PQC (a.k.a. ansatz) to
 114 approximate the target quantum state by iteratively updating the parameters of circuit, following
 115 a similar procedure of other variational quantum algorithms [39, 11]. AAE has been shown to be
 116 more advantageous for NISQ devices due to its ability to mitigate coherent errors through flexible
 117 adjustment of circuit parameters, coupled with its lower circuit depth [52]. We denote an ansatz as
 118 $U(\theta)$, where θ refers to a vector of tunable parameters for optimizations. A typical ansatz consists
 119 of several blocks of operations with the same structure. For example, a two-block ansatz with 4
 120 qubits is shown in Fig. 2, where the rotation layer is composed of single-qubit rotational gates
 121 $R_y(\theta_r) = e^{-i\theta_r Y/2}$, and the entangler layer comprises CNOT gates. Note that the entangler layer is
 122 configurable and hardware-native, which means that we can apply CNOT gates to physically adjacent
 123 qubits, thereby eliminating the necessity of additional SWAP gates to overcome the topological
 124 constraints [27]. This type of PQC is also known as *hardware-efficient ansatz* [20], being widely
 125 adopted in previous studies of AAE [59, 35, 52].

126 **3 SuperEncoder**

127 **3.1 Motivation**

128 Although AAE can potentially realize high fidelity QSP with $O(\text{poly}(n))$ circuit depth [35] with n
 129 the number of qubits, it requires repetitive *online* tuning of parameters to approximate the target
 130 state, which may result in an excessively long runtime that undermines its feasibility. Specifically, we
 131 could consider a simple application scenario in QML. The workflow with AAE is depicted in Fig. 3a.
 132 During the inference stage, we must iteratively update the parameters of the AAE ansatz for each
 133 input classical data vector, which may greatly dampen the performance. To quantify this impact, we
 134 measure the runtime of AAE-based data loading and the total runtime of model inference. As one can
 135 observe from Table 1, AAE dominates the runtime, thereby becoming the performance bottleneck.

n	T_{AAE} (s)	$T_{\text{total}} - T_{\text{AAE}}$ (s)
4	5.0086	0.0397
6	20.1810	0.0573
8	59.4193	0.0978

Table 1: **Performance overhead of AAE.** We break down the averaged inference runtime per sample from the MNIST dataset. T_{AAE} denotes time spent on loading classical data into quantum state using AAE, and T_{total} refers to total runtime.

136 The necessity of time-intensive iterations is grounded in the following assumption — Given an
 137 arbitrary quantum state $|\psi\rangle$, there *does not* exist a deterministic transformation $f : |\psi\rangle \rightarrow \theta$, where
 138 θ refers to the vector of parameters enabling a PQC to prepare an approximated state of $|\psi\rangle$. This
 139 assumption seems intuitively correct given the randomness of target states. However, we argue that a
 140 universal mapping f exists for any arbitrary data state $|\psi\rangle$. Taking a little thought of AE, we see that
 141 it implies the following conclusion: given an arbitrary state $|\psi\rangle$, there exists an universal arithmetic
 142 decomposition procedure $g : |\psi\rangle \rightarrow U$ satisfying $U|0\rangle = |\psi\rangle$. Inspired by this deterministic
 143 transformation, it is natural to ask: is there an universal transformation $g' : |\psi\rangle \rightarrow U'$ satisfying
 144 $E(U'|0, |\psi\rangle) \leq \epsilon$? Here E denotes the deviation between the prepared state by a circuit U' and the
 145 target state, and ϵ refers to certain acceptable error threshold. Since the structure of PQC in AAE
 146 is the same for any target state, U' is determined by θ . Then, the problem is reduced to exploring
 147 the existence of $f : |\psi\rangle \rightarrow \theta$. Should f exist, the overhead of online iterations could be eliminated,
 148 resulting in a novel QSP method being both fast and scalable.

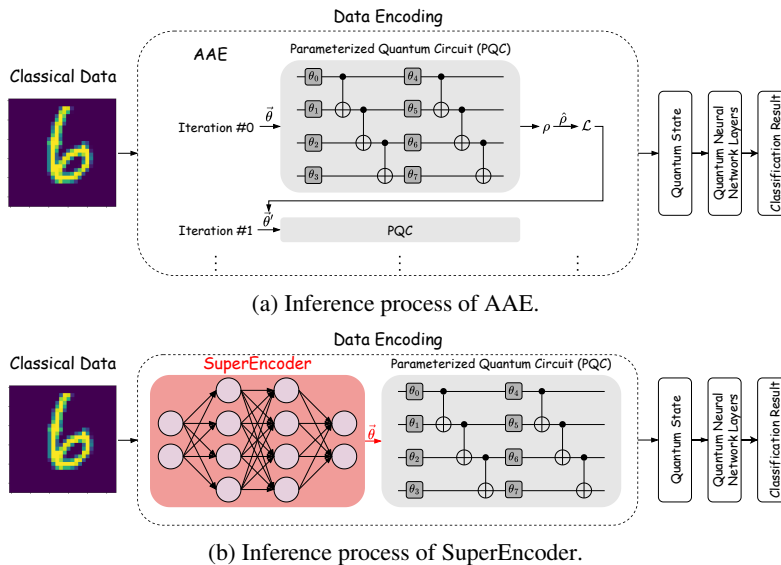


Figure 3: Comparison between AAE and SuperEncoder.

149 3.2 Design Methodology

150 Let $|\psi\rangle$ be the target state, and $U(\theta)$ be the PQC used in AAE with θ the optimized parameters.
 151 Our goal is to develop a model, termed SuperEncoder, to approximate the mapping $f : |\psi\rangle \rightarrow \theta$.
 152 Referring back to the scenario in QML, the workflow with SuperEncoder becomes iteration-free, as
 153 depicted in Fig. 3b.

154 Since neural networks could be used to approximate any continuous function [6], a natural solution is
 155 to use a neural network to approximate f . Specifically, we adopt a Multi-Layer Perceptron (MLP) as
 156 the backbone model for approximating f . However, training this model is nontrivial. Particularly, we
 157 find it challenging to design a proper loss function. In the remainder of this section, we explore three
 158 different designs and analyze their performance.




Figure 4: Virtualization of states generated by SuperEncoder trained with different loss functions. \mathcal{L}_2 is omitted as it produces very similar results to \mathcal{L}_3 .

159 The first and most straightforward method is *parameter-oriented* training — setting the loss function
 160 \mathcal{L}_1 as the MSE between the target parameters θ from AAE and the output parameters $\hat{\theta}$ from
 161 SuperEncoder. To evaluate the performance of \mathcal{L}_1 , we train a SuperEncoder using MNIST dataset,
 162 and test if it could load a test digit image into a quantum state with high fidelity. All images are
 163 downsampled and normalized into 4-qubit states for quick evaluation.

164 Unfortunately, results in Table 2 show that \mathcal{L}_1 achieves poor
 165 performance. The average fidelity of prepared quantum states
 166 is only 0.6208. As demonstrated in Fig. 4, \mathcal{L}_1 generates a state
 167 that loses the patterns of the original state. Additionally, utiliz-
 168 ing \mathcal{L}_1 implies that we need to first generate target parameters
 169 using AAE, of which the long runtime hinders pre-training on
 170 larger datasets. Consequently, required is a more effective loss
 171 function design without involving AAE.

\mathcal{L}_1	\mathcal{L}_2	\mathcal{L}_3
0.6208	0.9873	0.9908

Table 2: Fidelity comparison between SuperEncoders trained with different loss functions.

172 To address this challenge, we propose a *state-oriented* training
 173 methodology, which employs quantum states as targets to guide
 174 optimizations. Specifically, we may apply $\hat{\theta}$ to the circuit and execute
 175 it to obtain the prepared state $\hat{\psi}$. Then it is possible to calculate
 176 the difference between $\hat{\psi}$ and ψ as the loss to optimize SuperEncoder.
 177 In contrast to parameter-oriented training, this approach applies to
 178 larger datasets as it decouples the training procedure from AAE. We
 179 utilize two different state-oriented metrics, the first being the MSE
 180 between $\hat{\psi}$ and ψ , denoted as \mathcal{L}_2 , and the second is the *fidelity* of
 181 $\hat{\psi}$ relative to ψ , expressed as $\mathcal{L}_3 = 1 - |\langle \hat{\psi} | \psi \rangle|^2$ [25]. Results in
 182 Table 2 show that \mathcal{L}_2 and \mathcal{L}_3 achieve remarkably higher fidelity than
 183 \mathcal{L}_1 . Besides, we observe that \mathcal{L}_3 prepares a state very similar to the
 184 target one (Fig. 4), verifying that state-oriented training is more effective than parameter-oriented
 185 training.

186 **Landscape Analysis.** To understand the efficacy of these loss functions, we further analyze their
 187 landscapes following previous studies [28, 40, 18]. To gain insight from the landscape, we plot Fig. 6
 188 using the same scale and color gradients [18]. Compared to state-oriented losses (\mathcal{L}_2 and \mathcal{L}_3), \mathcal{L}_1 has
 189 a largely flat landscape with non-decreasing minima, thus the model struggles to explore a viable
 190 path towards a lower loss value, a similar pattern can also be observed in Fig. 5. In contrast, \mathcal{L}_2




Figure 5: Convergence of different loss functions.




Figure 6: Landscape virtualization of different loss functions.

191 and \mathcal{L}_3 have much lower minima and successfully converge to smaller loss values. Furthermore, we
 192 observe from Fig. 6 that \mathcal{L}_3 has a wider minima than \mathcal{L}_2 , which may indicate a better generalization
 193 capability [40].

194 **Gradient Analysis.** Based on the landscape analysis, we adopt \mathcal{L}_3 as the loss function to train
 195 SuperEncoder. We note that \mathcal{L}_3 can be written as $1 - \langle \psi | \hat{\psi} \rangle \langle \hat{\psi} | \psi \rangle$. If $\hat{\rho}$ is a pure state, it is equivalent
 196 to $|\hat{\psi}\rangle\langle\hat{\psi}|$. Then \mathcal{L}_3 is given by $\mathcal{L}_3 = 1 - \langle \psi | \hat{\rho} | \psi \rangle$.

197 This re-formalization is important as only the mixed state $\hat{\rho}$ could be obtained in noisy environments.
 198 Suppose an n -qubit circuit is parameterized by m parameters $\hat{\theta} = [\hat{\theta}_1, \dots, \hat{\theta}_k, \dots, \hat{\theta}_m]$. Let \mathbf{W} be
 199 the weight matrix of MLP, with k, l the element indices. We analyze the gradient of \mathcal{L}_3 w.r.t. $W_{k,l}$ to
 200 showcase its feasibility in different quantum computing environments.

$$\begin{aligned}
 \nabla_{W_{k,l}} \mathcal{L}_3 &= \frac{\partial \mathcal{L}_3}{\partial W_{k,l}} = -\langle \psi | \frac{\partial \hat{\rho}}{\partial W_{k,l}} | \psi \rangle \\
 &= -\langle \psi | \begin{bmatrix} \sum_{j=1}^m \frac{\partial \hat{\rho}_{1,1}}{\partial \theta_j} \frac{\partial \theta_j}{\partial W_{k,l}} & \cdots & \sum_{j=1}^m \frac{\partial \hat{\rho}_{1,N}}{\partial \theta_j} \frac{\partial \theta_j}{\partial W_{k,l}} \\ \vdots & \ddots & \vdots \\ \sum_{j=1}^m \frac{\partial \hat{\rho}_{N,1}}{\partial \theta_j} \frac{\partial \theta_j}{\partial W_{k,l}} & \cdots & \sum_{j=1}^m \frac{\partial \hat{\rho}_{N,N}}{\partial \theta_j} \frac{\partial \theta_j}{\partial W_{k,l}} \end{bmatrix} | \psi \rangle, \quad (1)
 \end{aligned}$$

201 The calculation of $\frac{\partial \theta_j}{\partial W_{k,l}}$ can be easily done on classical devices using backpropagation supported by
 202 automatic differentiation frameworks. Therefore, we only focus on $\frac{\partial \hat{\rho}_{i,j}}{\partial \theta_k}$. In a simulation environ-
 203 ment, the calculation of $\hat{\rho}$ is conducted via noisy quantum circuit simulation, which is essentially a
 204 series of tensor operations on state vectors. Therefore, the calculation of $\frac{\partial \hat{\rho}_{i,j}}{\partial \theta_k}$ is compatible with
 205 backpropagation. The situation on real devices becomes more complicated. On real devices, the
 206 mixed state $\hat{\rho}$ is reconstructed through *quantum tomography* [7] based on classical shadow [55, 16].
 207 Here, for notion simplicity, we denote the process of classical shadow as a transformation \mathcal{S} , and
 208 denote the measurement expectations of the ansatz as $U(\hat{\theta})$. Thus the reconstructed density ma-
 209 trix is given by $\hat{\rho} = \mathcal{S}(U(\hat{\theta}))$. Then the gradient of $\hat{\rho}_{i,j}$ with respect to $\hat{\theta}_k$ is $\sum_u \frac{\partial \hat{\rho}_{i,j}}{\partial U(\hat{\theta})} \frac{\partial U(\hat{\theta})}{\partial \theta_k}$.
 210 Here $\frac{\partial \hat{\rho}_{i,j}}{\partial U(\hat{\theta})}$ can be efficiently calculated on classical devices using backpropagation, as \mathcal{S} operates
 211 on expectation values on classical devices. However, $U(\hat{\theta})$ involves state evolution on quantum
 212 devices, where back-propagation is impossible due to the No-Cloning theorem [36]. Fortunately,
 213 it is possible to utilize the *parameter shift* rule [8, 4, 53] to calculate $\frac{\partial U(\hat{\theta})}{\partial \theta_k}$. In this way, the
 214 gradients of the circuit function U with respect to θ_j are $\frac{\partial U(\hat{\theta})}{\partial \theta_k} = \frac{1}{2} (U(\theta_+) - U(\theta_-))$, where
 215 $\theta_+ = [\theta_1, \dots, \theta_k + \frac{\pi}{2}, \dots, \theta_m]$, $\theta_- = [\theta_1, \dots, \theta_k - \frac{\pi}{2}, \dots, \theta_m]$. To summarize, training SuperEn-
 216 coder with \mathcal{L}_3 is theoretically feasible on both simulators and real devices.

217 4 Numerical Results

218 4.1 Experiment Setup

219 **Datasets.** To train a SuperEncoder for arbitrary quantum states, we need a dataset comprising a wide
220 range of quantum states with different distributions. To our knowledge, there is no dataset dedicated
221 for this special purpose. A natural solution is to use readily available datasets from classical machine
222 learning domains (e.g., ImageNet [9], Places [58], SQuAD [44]) by normalizing them to quantum
223 states. However, QSP is essential in various application scenarios besides QML. The classical data to
224 be loaded may not only contain natural images or languages but also contain arbitrary data (e.g., in
225 HHL algorithm [15]). Therefore, we construct a training dataset adapted from FractalDB-60 [21] with
226 60k samples, a formula-driven dataset originally designed for computer vision without any natural
227 images. We also construct a separate dataset to test the performance of QSP, which consists of data
228 sampled from different statistical distributions, including uniform, normal, log-normal, exponential,
229 and Dirichlet distributions, with 3000 samples per distribution. Hereafter we refer this dataset as the
230 *synthetic dataset*.

231 **Platforms.** We implement SuperEncoder using PennyLane [34], PyTorch [37] and Qiskit [43].
232 Simulations are done on a Ubuntu server with 768 GB memory, two 32-core Intel(R) Xeon(R) Silver
233 4216 CPU with 2.10 GHz, and 2 NVIDIA A-100 GPUs. IBM quantum cloud platform² is adopted to
234 evaluate the performance on real quantum devices.

235 **Metrics.** We evaluate SuperEncoder and compare it to AE and AAE in terms of runtime, scalability,
236 and fidelity. *Runtime* refers to how long it takes to prepare a quantum state. *Scalability* refers to how
237 the circuit depth grows with the number of qubits. *Fidelity* evaluates the similarity between prepared
238 quantum states and target quantum states. Specifically, the fidelity for two mixed states given by
239 density matrices ρ and $\hat{\rho}$ is defined as $F(\rho, \hat{\rho}) = \text{Tr}(\sqrt{\sqrt{\rho}\hat{\rho}\sqrt{\rho}})^2 \in [0, 1]$. A larger F indicates a
240 better fidelity.

241 **Implementation.** We implement SuperEncoder using an MLP consisting of two hidden layers.
242 The dimensions of input and output layers are respectively set to 2^n and m , where n refers to the
243 number of qubits and m refers to the number of parameters. We adopt \mathcal{L}_3 as the loss function.
244 Training data are down-sampled, flattened, and normalized to 2^n -dimensional state vectors. We
245 adopt the hardware efficient ansatz [20] (Fig. 2) as the backbone of quantum circuits and use the
246 same structure for AAE. Given a target state, a pre-trained SuperEncoder model is invoked to
247 generate parameters and thus the circuit for QSP. While for AAE, we employ online iterations for
248 each state. For AE, the arithmetic decomposition method in PennyLane [34, 4] is adopted. We
249 defer more details about implementation to Appendix A. Our framework is open-source at <https://anonymous.4open.science/r/SuperEncoder-A733> with detailed instructions to reproduce
250 our results.
251

252 4.2 Evaluation on Synthetic Dataset


253 For simplicity and without loss of generality, we focus our discussion on the results of 4-qubit QSP
254 tasks. The outcomes for larger quantum states are detailed in Appendix B.1. The parameters of both
255 AAE and SuperEncoder are optimized based on ideal quantum circuit simulation.

256 **Runtime.** The runtime and fidelity results, evaluated on the synthetic dataset, are presented in Table 3.
257 We observe that SuperEncoder runs faster than AAE by orders of magnitudes and has a similar
258 runtime to AE, affirming that SuperEncoder effectively overcomes the main drawback of AAE.

	AE		AAE		SuperEncoder	
	Fidelity	Runtime	Fidelity	Runtime	Fidelity	Runtime
Uniform			0.9996		0.9731	
Normal			0.9992		0.8201	
Log-normal			0.9993		0.9421	
Exponential			0.9996		0.9464	
Dirichlet			0.9995		0.9737	
Average	1.0000	0.0162 s	0.9994	5.0201 s	0.9310	0.0397 s

Table 3: Comparison between AE, AAE and SuperEncoder in terms of runtime and fidelity.

²<https://quantum-computing.ibm.com/>



(a) Scaling of circuit depth w.r.t. # qubits. (b) Fidelity of different QSP methods on `ibm_osaka`.

Figure 7: Comparison between AE, AAE, and SuperEncoder in terms of circuit depth and fidelity on real devices.

259 **Scalability.** Although AE runs fast, it exhibits poor scalability since the circuit depth grows exponentially with the number of qubits. The depth of AAE is empirically determined by increasing depth until the final fidelity does not increase, same depth is adopted for SuperEncoder. We deter the details of determining the depth of AAE/SuperEncoder to Appendix A. As shown in Fig. 7a, the depth of AE grows fast and becomes much larger than AAE/SuperEncoder, e.g., the depth of AE for a 8-qubit state is 984, whereas the depth of AAE/SuperEncoder is only 120.

265 **Fidelity.** From Table 3, it is evident that SuperEncoder experiences notable fidelity degradation when compared with AAE and AE. Specifically, the average fidelity of SuperEncoder is 0.9307, whereas AAE and AE achieve higher average fidelities of 0.9994 and 1.0, respectively. Note that, although AE demonstrates the highest fidelity under ideal simulation, its performance deteriorates significantly in noisy environments. Fig. 7b presents the performance of these three QSP methods on quantum states with 4, 6, and 8 qubits on the `ibm_osaka` machine. While the fidelity of AE is higher than AAE/SuperEncoder on the 4-qubit and 6-qubit states, its fidelity on the 8-qubit state is only 0.0049, becoming much lower than AAE/SuperEncoder. This decline is primarily attributed to its large circuit depth as shown in Fig. 7a.




Figure 8: Schematic of a QNN (above) and test accuracies of QSP methods on the QML task (below).

279 4.3 Application to Downstream Tasks

280 **Quantum Machine Learning.** We first apply SuperEncoder to a QML task. MNIST dataset is adopted for demonstration, we extract a sub-dataset composed on digits 3 and 6 for evaluation. The quantum circuit that implements a QNN is depicted in Fig. 8, which consists of an encoder block and m entangler layers. Here the encoder block is implemented via QSP circuits, either AE, AAE, or SuperEncoder, of which the parameters are frozen during the training of QNN. The test results are shown in Fig. 8, we observe that SuperEncoder achieves similar performance with AAE and AE. The reason lies in the fact that classification tasks can be robust to noises. Consequently, approximate QSP (AAE and SuperEncoder) with a certain degree of fidelity loss is tolerable.

292 Figure 9: Schematic of HHL.

293 **HHL Algorithm.** Besides QML, quantum-enhanced linear algebra algorithms are another important set of applications that heavily rely on QSP. The most famous algorithm is the HHL algorithm [15]. The problem can be defined as, given a matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$, and a vector $\mathbf{b} \in \mathbb{C}^N$, find $\mathbf{x} \in \mathbb{C}^N$ satisfying $\mathbf{A}\mathbf{x} = \mathbf{b}$. A typical implementation of HHL utilizes the circuit depicted in Fig. 9. The outline of HHL is as follows. (i) Apply a QSP circuit to prepare the quantum state $|\mathbf{b}\rangle$. (ii) Apply Quantum Phase Estimation [10] (QPE) to estimate the eigenvalue of \mathbf{A} (iii) Apply conditioned rotation gates on ancillary qubits based on the eigenvalues (R). (iv) Apply an inverse QPE (QPE_inv) and measure the ancillary qubits to reconstruct the solution vector \mathbf{x} . Note that, HHL does not return the solution \mathbf{x} itself, but rather an approximation of the expectation value of some operator \mathbf{M} associated with \mathbf{x} , e.g.,

302 $\mathbf{x}^\dagger \mathbf{M} \mathbf{x}$. Here, we adopt an optimized version of HHL proposed by Vazquez et al. [51] for evaluation.
 303 To compare the performance between different QSP methods, we construct linear equations with
 304 fixed matrix \mathbf{A} and operator \mathbf{M} , while we sample different vectors from our synthetic dataset as \mathbf{b} .
 305 Results are concluded in Table 4. Unlike QML, HHL expects precise QSP, thus we take the results
 306 from AE as the ground truth values and compare the relative error between AAE/SuperEncoder and
 307 AE. The relative error of SuperEncoder is 2.4094%, while the error of AAE is only 0.3326%.

308 4.4 Discussion and Future Work

309 The results of our evaluation can be concluded in two folds.
 310 (i) SuperEncoder effectively eliminates the iteration over-
 311 head of AAE, thereby becoming both fast and scalable.
 312 However, it has a notable degradation in fidelity. (ii) The
 313 impact of fidelity degradation varies across different down-
 314 stream applications. For QML, the fidelity degradation is
 315 affordable as long as the prepared states are distinguish-
 316 able across different classes. However, algorithms like
 317 HHL rely on precise QSP to produce the best result. In
 318 these algorithms, SuperEncoder suffers from higher error
 319 ratio than AAE.

	AE	AAE	SuperEncoder
\mathbf{b}_0	0.7391	0.7404	0.7355
\mathbf{b}_1	0.7449	0.7445	0.7544
\mathbf{b}_2	0.7492	0.7469	0.8134
\mathbf{b}_3	0.7164	0.7099	0.7223
\mathbf{b}_4	0.7092	0.7076	0.7155
Avg err		0.3326%	2.4094%

Table 4: Performance of different QSP methods in HHL algorithm. ‘Avg err’ denotes the average relative errors between AAE/SuperEncoder and AE.

320 Note that, the current evaluation results may not reflect the
 321 actual performance of SuperEncoder on real NISQ devices.
 322 Recent work has shown that AAE achieves significantly better fidelity than AE does [52]. This is due
 323 to the intrinsic noise awareness of AAE, as it could obtain states from noisy devices to guide updating
 324 parameters with better robustness. In essence, the proposed SuperEncoder possesses the same nature
 325 as AAE. Unfortunately, although the noise-robustness of AAE can be evaluated on a small set of test
 326 samples, it is difficult to perform noise-aware training for SuperEncoder as it requires a large dataset
 327 for pre-training. Consequently, SuperEncoder relies on huge amounts of interactions with noisy
 328 devices, thereby becoming extremely time-consuming. As a result, the effectiveness of SuperEncoder
 329 in noisy environments remains largely unexplored, which we leave for future exploration. More
 330 discussion about this perspective is in Appendix C.

331 5 Related Work

332 Besides QSP, there are other methods for loading classical data into quantum states. These methods
 333 can be roughly regarded as *quantum feature embedding* primarily used in QML, which maps classical
 334 data to a completely different distribution encoded in quantum states. A widely used embedding
 335 method is known as angle embedding. Li et al. have proven that this method has a concentration issue,
 336 which means that the encoded states may become indistinguishable as the circuit depth increases [26].
 337 Lei et al. proposed an automatic design framework for efficient quantum feature embedding, resolving
 338 the issue of concentration [24]. The central idea of this framework is to search for the most efficient
 339 circuit architecture for a given classical input, which is also known as Quantum Architecture Search
 340 (QAS) [38, 30, 54]. While the application scenario of quantum feature embedding is largely limited
 341 to QML, QSP has broader usage in general quantum applications, distinguishing SuperEncoder from
 342 all aforementioned work.

343 6 Conclusion

344 In this work, we propose SuperEncoder, a neural network-based QSP framework. Instead of iteratively
 345 tuning the circuit parameters to approximate each quantum state, as is done in AAE, we adopt a
 346 different approach by directly learning the relationship between target quantum states and the required
 347 circuit parameters. SuperEncoder combines the scalable circuit architecture of AAE with the fast
 348 runtime of AE, as verified by a comprehensive evaluation on both synthetic dataset and downstream
 349 applications.

References

- 350
- 351 [1] Amira Abbas, David Sutter, Christa Zoufal, Aurélien Lucchi, Alessio Figalli, and Stefan
352 Woerner. The power of quantum neural networks. *Nature Computational Science*, 1(6):403–409,
353 2021.
- 354 [2] Israel F Araujo, Daniel K Park, Teresa B Ludermir, Wilson R Oliveira, Francesco Petruccione,
355 and Adenilton J Da Silva. Configurable sublinear circuits for quantum state preparation.
356 *Quantum Information Processing*, 22(2):123, 2023.
- 357 [3] Johannes Bausch. Recurrent quantum neural networks. *Advances in neural information*
358 *processing systems*, 33:1368–1379, 2020.
- 359 [4] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, Shahnawaz Ahmed, Vishnu
360 Ajith, M Sohaib Alam, Guillermo Alonso-Linaje, B AkashNarayanan, Ali Asadi, et al. Pen-
361 nylane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint*
362 *arXiv:1811.04968*, 2018.
- 363 [5] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth
364 Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.
- 365 [6] Tianping Chen and Hong Chen. Universal approximation to nonlinear operators by neural
366 networks with arbitrary activation functions and its application to dynamical systems. *IEEE*
367 *Transactions on Neural Networks*, 6(4):911–917, 1995.
- 368 [7] Marcus Cramer, Martin B Plenio, Steven T Flammia, Rolando Somma, David Gross, Stephen D
369 Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. Efficient quantum state
370 tomography. *Nature communications*, 1(1):149, 2010.
- 371 [8] Gavin E Crooks. Gradients of parameterized quantum gates using the parameter-shift rule and
372 gate decomposition. *arXiv preprint arXiv:1905.13311*, 2019.
- 373 [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-
374 scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern*
375 *recognition*, pages 248–255. Ieee, 2009.
- 376 [10] Uwe Dorner, Rafal Demkowicz-Dobrzanski, Brian J Smith, Jeff S Lundeen, Wojciech
377 Wasilewski, Konrad Banaszek, and Ian A Walmsley. Optimal quantum phase estimation.
378 *Physical review letters*, 102(4):040403, 2009.
- 379 [11] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization
380 algorithm. *arXiv preprint arXiv:1411.4028*, 2014. [https://doi.org/10.48550/arXiv.
381 1411.4028](https://doi.org/10.48550/arXiv.1411.4028).
- 382 [12] Niels Gleinig and Torsten Hoefler. An efficient algorithm for sparse quantum state preparation.
383 In *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pages 433–438. IEEE, 2021.
- 384 [13] Javier Gonzalez-Conde, Ángel Rodríguez-Rozas, Enrique Solano, and Mikel Sanz. Simulating
385 option price dynamics with exponential quantum speedup. *arXiv preprint arXiv:2101.04023*,
386 2021.
- 387 [14] Javier Gonzalez-Conde, Thomas W Watts, Pablo Rodriguez-Grasa, and Mikel Sanz. Efficient
388 quantum amplitude encoding of polynomial functions. *Quantum*, 8:1297, 2024.
- 389 [15] Aram W Harrow, Avinandan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems
390 of equations. *Physical Review Letters*, 103(15):150502, 2009. [https://doi.org/10.1103/
391 PhysRevLett.103.150502](https://doi.org/10.1103/PhysRevLett.103.150502).
- 392 [16] Hsin-Yuan Huang. Learning quantum states from their classical shadows. *Nature Reviews*
393 *Physics*, 4(2):81–81, 2022.
- 394 [17] Jason Iaconis, Sonika Johri, and Elton Yechao Zhu. Quantum state preparation of normal
395 distributions using matrix product states. *npj Quantum Information*, 10(1):15, 2024.

- 396 [18] Christian Cmeheil-Warn Jacob Hansen. Loss landscapes. In *ICLR Blog Track*, 2022. [https://loss-](https://loss-landscapes.github.io/Loss-Landscapes-Blog/2022/12/01/loss-landscapes/)
397 [landscapes.github.io/Loss-Landscapes-Blog/2022/12/01/loss-landscapes/](https://loss-landscapes.github.io/Loss-Landscapes-Blog/2022/12/01/loss-landscapes/).
- 398 [19] Weiwen Jiang, Jinjun Xiong, and Yiyu Shi. A co-design framework of neural networks and
399 quantum circuits towards quantum advantage. *Nature Communications*, 12(1):579, 2021.
400 <https://doi.org/10.1038/s41467-020-20729-5>.
- 401 [20] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M.
402 Chow, and Jay M. Gambetta. Hardware-efficient variational quantum eigensolver for small
403 molecules and quantum magnets. *Nature*, 549(7671):242–246, September 2017.
- 404 [21] Hirokatsu Kataoka, Kazushige Okayasu, Asato Matsumoto, Eisuke Yamagata, Ryosuke Yamada,
405 Nakamasa Inoue, Akio Nakamura, and Yutaka Satoh. Pre-training without natural images. In
406 *Proceedings of the Asian Conference on Computer Vision*, 2020.
- 407 [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
408 *arXiv:1412.6980*, 2014.
- 409 [23] Philip Krantz, Morten Kjaergaard, Fei Yan, Terry P Orlando, Simon Gustavsson, and William D
410 Oliver. A quantum engineer’s guide to superconducting qubits. *Applied physics reviews*, 6(2),
411 2019.
- 412 [24] Cong Lei, Yuxuan Du, Peng Mi, Jun Yu, and Tongliang Liu. Neural auto-designer for enhanced
413 quantum kernels. In *The Twelfth International Conference on Learning Representations*, 2023.
- 414 [25] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster. Speedup for quantum
415 optimal control from automatic differentiation based on graphics processing units. *Physical*
416 *Review A*, 95(4):042318, 2017. <https://doi.org/10.1103/PhysRevA.95.042318>.
- 417 [26] Guangxi Li, Ruilin Ye, Xuanqiang Zhao, and Xin Wang. Concentration of data encoding in
418 parameterized quantum circuits. *Advances in Neural Information Processing Systems*, 35:19456–
419 19469, 2022.
- 420 [27] Gushu Li, Yufei Ding, and Yuan Xie. Tackling the qubit mapping problem for nisq-era quantum
421 devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural*
422 *Support for Programming Languages and Operating Systems*, pages 1001–1014, 2019. <https://doi.org/10.1145/3297858.3304023>.
- 423
- 424 [28] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss
425 landscape of neural nets. *Advances in neural information processing systems*, 31, 2018.
- 426 [29] Gui-Lu Long and Yang Sun. Efficient scheme for initializing a quantum register with an
427 arbitrary superposed state. *Physical Review A*, 64(1):014303, 2001.
- 428 [30] Xudong Lu, Kaisen Pan, Ge Yan, Jiaming Shan, Wenjie Wu, and Junchi Yan. Qas-bench:
429 rethinking quantum architecture search and a benchmark. In *International Conference on*
430 *Machine Learning*, pages 22880–22898. PMLR, 2023.
- 431 [31] Michael Lubasch, Jaewoo Joo, Pierre Moinier, Martin Kiffner, and Dieter Jaksch. Variational
432 quantum algorithms for nonlinear problems. *Physical Review A*, 101(1):010301, 2020.
- 433 [32] Rui Mao, Guojing Tian, and Xiaoming Sun. Towards optimal circuit size for sparse quantum
434 state preparation. *arXiv e-prints*, pages arXiv–2404, 2024.
- 435 [33] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. Quantum circuit
436 learning. *Physical Review A*, 98(3):032309, 2018.
- 437 [34] Mikko Möttönen, JJ Vartiainen, Ville Bergholm, and Martti M Salomaa. Transformation of
438 quantum states using uniformly controlled rotations. *Quantum Information and Computation*, 5,
439 2005.
- 440 [35] Kouhei Nakaji, Shumpei Uno, Yohichi Suzuki, Rudy Raymond, Tamiya Onodera, Tomoki
441 Tanaka, Hiroyuki Tezuka, Naoki Mitsuda, and Naoki Yamamoto. Approximate amplitude
442 encoding in shallow parameterized quantum circuits and its application to financial market
443 indicators. *Physical Review Research*, 4(2):023136, 2022.

- 444 [36] Michael A Nielsen and Isaac L Chuang. Quantum computation and quantum information. 2010.
- 445 [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
446 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative
447 style, high-performance deep learning library. *Advances in neural information processing*
448 *systems*, 32, 2019.
- 449 [38] Yash J. Patel, Akash Kundu, Mateusz Ostaszewski, Xavier Bonet-Monroig, Vedran Dunjko,
450 and Onur Danaci. Curriculum reinforcement learning for quantum architecture search under
451 hardware errors. In *The Twelfth International Conference on Learning Representations*, 2024.
- 452 [39] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J
453 Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic
454 quantum processor. *Nature communications*, 5(1):4213, 2014. [https://doi.org/10.1038/
455 ncomms5213](https://doi.org/10.1038/ncomms5213).
- 456 [40] Henning Petzka, Michael Kamp, Linara Adilova, Cristian Sminchisescu, and Mario Boley.
457 Relative flatness and generalization. *Advances in neural information processing systems*,
458 34:18420–18432, 2021.
- 459 [41] Martin Plesch and Āslav Brukner. Quantum-state preparation with universal gate decomposi-
460 tions. *Physical Review A*, 83(3):032302, 2011.
- 461 [42] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- 462 [43] Qiskit contributors. Qiskit: An open-source framework for quantum computing, 2023.
- 463 [44] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions
464 for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- 465 [45] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. Prediction by linear regression on a
466 quantum computer. *Physical Review A*, 94(2):022342, 2016.
- 467 [46] Vivek V Shende, Stephen S Bullock, and Igor L Markov. Synthesis of quantum logic circuits. In
468 *Proceedings of the 2005 Asia and South Pacific Design Automation Conference*, pages 272–275,
469 2005.
- 470 [47] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms
471 on a quantum computer. *SIAM review*, 41(2):303–332, 1999. [https://doi.org/10.1137/
472 S0036144598347011](https://doi.org/10.1137/S0036144598347011).
- 473 [48] Siddarth Srinivasan, Carlton Downey, and Byron Boots. Learning and inference in hilbert space
474 with quantum graphical models. *Advances in Neural Information Processing Systems*, 31, 2018.
- 475 [49] Xiaoming Sun, Guojing Tian, Shuai Yang, Pei Yuan, and Shengyu Zhang. Asymptotically
476 optimal circuit depth for quantum state preparation and general unitary synthesis. *IEEE*
477 *Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2023.
- 478 [50] Jinkai Tian, Xiaoyu Sun, Yuxuan Du, Shanshan Zhao, Qing Liu, Kaining Zhang, Wei Yi, Wan-
479 rong Huang, Chaoyue Wang, Xingyao Wu, et al. Recent advances for quantum neural networks
480 in generative learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- 481 [51] Almudena Carrera Vazquez, Ralf Hiptmair, and Stefan Woerner. Enhancing the quantum linear
482 systems algorithm using richardson extrapolation. *ACM Transactions on Quantum Computing*,
483 3(1):1–37, 2022.
- 484 [52] Hanrui Wang, Yilian Liu, Pengyu Liu, Jiaqi Gu, Zirui Li, Zhiding Liang, Jinglei Cheng,
485 Yongshan Ding, Xuehai Qian, Yiyu Shi, et al. Robuststate: Boosting fidelity of quantum state
486 preparation via noise-aware variational training. *arXiv preprint arXiv:2311.16035*, 2023.
- 487 [53] David Wierichs, Josh Izaac, Cody Wang, and Cedric Yen-Yu Lin. General parameter-shift rules
488 for quantum gradients. *Quantum*, 6:677, 2022.

- 489 [54] Wenjie Wu, Ge Yan, Xudong Lu, Kaisen Pan, and Junchi Yan. Quantumdarts: differentiable
490 quantum architecture search for variational quantum algorithms. In *International Conference*
491 *on Machine Learning*, pages 37745–37764. PMLR, 2023.
- 492 [55] Ting Zhang, Jinzhao Sun, Xiao-Xu Fang, Xiao-Ming Zhang, Xiao Yuan, and He Lu. Ex-
493 perimental quantum state measurement with classical shadows. *Physical Review Letters*,
494 127(20):200501, 2021.
- 495 [56] Xiao-Ming Zhang, Man-Hong Yung, and Xiao Yuan. Low-depth quantum state preparation.
496 *Physical Review Research*, 3(4):043200, 2021.
- 497 [57] Jian Zhao, Yu-Chun Wu, Guang-Can Guo, and Guo-Ping Guo. State preparation based on
498 quantum phase estimation. *arXiv preprint arXiv:1912.05335*, 2019.
- 499 [58] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A
500 10 million image database for scene recognition. *IEEE transactions on pattern analysis and*
501 *machine intelligence*, 40(6):1452–1464, 2017.
- 502 [59] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. Quantum generative adversarial networks
503 for learning and loading random distributions. *npj Quantum Information*, 5(1):103, 2019.

504 The structure of our Appendix is as follows. Appendix A provides more details of implementing
 505 SuperEncoder. Appendix B provides additional numerical results to illustrate the impact of state
 506 sizes, model architectures, and training datasets. Appendix C analyzes the estimated runtime of
 507 training SuperEncoder on real devices.

508 A Implementation Details

509 In this section, we elaborate the missing details of SuperEncoder in the main text.

510 The overarching workflow of SuperEncoder is illustrated in Fig. 10. The target quantum states are
 511 input to the MLP model. Then, the MLP model generates predicted parameters based on the target
 512 states. Afterwards, the parameters are applied to the PQC to obtain the prepared quantum states.
 513 Finally, we calculate the loss based on the prepared states and target states and optimize the weights
 514 of MLP through backpropagation.

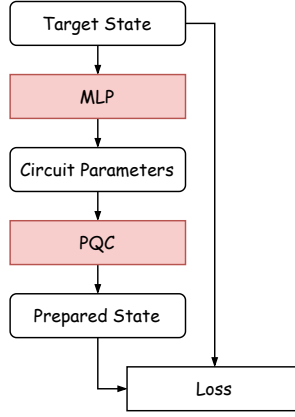


Figure 10: Detailed workflow of SuperEncoder.

515 The settings of MLP and PQC are as follows.

516 **MLP.** As listed in Table 5, we implement a two-layer MLP. Each layer consists of 512 neurons. We
 517 employ Tanh as the activation functions since θ represents the *angles* of rotation gates, ranging from
 518 $-\pi$ to π .

Linear	Input	$(\text{batch_size}, 2^n)$
	Output	$(\text{batch_size}, 512)$
Tanh	Input	$(\text{batch_size}, 512)$
	Output	$(\text{batch_size}, 512)$
Linear	Input	$(\text{batch_size}, 512)$
	Output	$(\text{batch_size}, \dim(\theta))$
Tanh	Input	$(\text{batch_size}, \dim(\theta))$
	Output	$(\text{batch_size}, \dim(\theta))$

Table 5: MLP based SuperEncoder. n refers to the number of qubits. θ denotes the parameter vector.

519 **PQC.** The circuit structure is the same with the one depicted in Fig. 2, except that the number of
 520 blocks is determined dynamically through empirical examinations. Specifically, we utilize AAE to
 521 approximate a target state while increasing the number of blocks. The number of blocks is designated
 522 when the resulting state fidelity no longer increases. For example, Fig. 11 demonstrates how fidelity
 523 changes while increasing the number of blocks. As one can observe, the fidelity converges when the
 524 number of layers is larger than 8. Hence, the number of layers is set to be 8 for 4-qubit quantum
 525 states. We follow the same procedure to set the number of blocks for other state sizes. Each block
 526 has the same structure, consisting of a rotation layer and an entangler layer. Given an n -qubit system,
 527 a rotation layer comprises $n R_y$ gates, each operating on a distinct qubit. The entangler layer is
 528 composed of two CNOT layers. The first CNOT layer applies CNOT gates to $\{(q_0, q_1), (q_2, q_3), \dots\}$,
 529 and the second CNOT layer applies CNOT gates to $\{(q_1, q_2), (q_3, q_4), \dots\}$. Hence, the depth of

530 a block is 3. Let l be the number of blocks; then the dimension of the parameter vector is given
 531 by $\dim(\theta) = n \times l$, and the depth of AAE/SuperEncoder is $3 \times l$. We conclude the settings of
 532 AAE/SuperEncoder used throughout this study in Table 6.




Figure 11: Fidelity vs. # blocks for 4-qubit states using AAE.

Number of Qubits	4	6	8
Number of Blocks	8	20	40
Depth	24	60	120

Table 6: Number of blocks and corresponding depth of AAE/SuperEncoder.

533 The hyperparameters for training SuperEncoder and optimizing AAE are as follows.

534 **Training Hyperparameters for SuperEncoder.** Throughout our experiments, the number of epochs
 535 are consistently set to be 10. For 4-qubit states, we set `bath_size` to 32, while we set it 64 for
 536 6-qubit and 8-qubit states. We adopt Adam optimizer [22] with a learning rate of 3e-3 and a weight
 537 decay of 1e-5.

538 **Hyperparameters for AAE.** To optimize the parameters of AAE, we also use the Adam optimizer,
 539 with a learning rate of 1e-2 and zero weight decay. For all quantum states, we train the AAE for 100
 540 steps.

541 B More Numerical Results

542 B.1 Results on Larger Quantum States

543 In line with the main text, we train the SuperEncoder for 6-qubit and 8-qubit quantum states using
 544 FractalDB-60 as the training dataset. Then we evaluate the performance of SuperEncoder on the
 545 synthetic test datasets. As shown in Table 7, the average fidelity on 6-qubit and 8-qubit states are
 546 0.8655 and 0.7624 respectively. In Appendix B.2, B.3, we discuss potential optimizations to alleviate
 547 this performance degradation.

Dataset	$n = 4$	$n = 6$	$n = 8$
Uniform	0.9731	0.9254	0.8648
Normal	0.8201	0.7457	0.6075
Log-normal	0.9421	0.8575	0.7122
Exponential	0.9464	0.8757	0.7613
Dirichlet	0.9737	0.9232	0.8663
Avg	0.9310	0.8655	0.7624
Avg-AAE	0.9994	0.9964	0.9910

Table 7: Performance evaluation on larger quantum states (6-qubit and 8-qubit). The last separate row shows the results of AAE for comparison.

548 **B.2 Impact of Model Architecture**

549 As a preliminary investigation, the optimal model architecture for SuperEncoder still requires further
 550 exploration. Currently, we have set the size of the hidden units at a constant 512 (Table 5). However,
 551 as the number of qubits, n , increases, a wider network architecture may become necessary. To
 552 showcase the impact of model width, we adjust the size to 4×2^n for 6-qubit states and 16×2^n for
 553 8-qubit states, and compare their performance with the original settings, as shown in Table 8. As
 554 evident from the results, this simple adjustment significantly enhances the fidelity of SuperEncoder,
 555 suggesting that there is substantial potential to boost SuperEncoder’s performance by developing a
 556 more tailored network architecture.

Dataset	$n = 6$		$n = 8$	
	$h = 512$	$h = 4 \times 2^6$	$h = 512$	$h = 16 \times 2^8$
Uniform	0.9254	0.9267	0.8648	0.8821
Normal	0.7457	0.7580	0.6075	0.6401
Log-normal	0.8575	0.8608	0.7122	0.7294
Exponential	0.8757	0.8732	0.7613	0.7781
Dirichlet	0.9232	0.9261	0.8663	0.8805
Avg	0.8655	0.8690	0.7624	0.7820

Table 8: Impact of increasing network width. Here h refers to the size of hidden units.

557 **B.3 Impact of Training Datasets**

558 In addition to refining the model architecture, the development of a specially designed dataset for
 559 pre-training SuperEncoder is essential. Currently, the dataset utilized is FractalDB [21], which is
 560 originally designed for computer vision tasks. However, given the wide range of applications of QSP,
 561 there is a need to accommodate diverse types of classical data from various domains. Therefore, how
 562 to create a comprehensive dataset that could fully unleash the potential of SuperEncoder remains an
 563 open question. While developing a pre-trained model that performs well in all kinds of applications
 564 may be challenging, we advocate for a strategy that combines pre-training with fine-tuning for the
 565 practical deployment of SuperEncoder, similar to the approach used with foundation models in
 566 classical machine learning. To substantiate this approach, we have compiled a separate dataset that
 567 encompasses a variety of statistical distributions not limited to those utilized for evaluation (but with
 568 different settings). As demonstrated in Table 9, after fine-tuning, the performance of SuperEncoder
 569 improves by approximately 0.03.

Dataset	Pre-training	Pre-training+Finetuning
Uniform	0.9731	0.9909
Normal	0.8201	0.8879
Log-normal	0.9421	0.9717
Exponential	0.9464	0.9729
Dirichlet	0.9737	0.9903
Avg	0.9310	0.9627

Table 9: Fidelity improvements after fine-tuning SuperEncoder using a dataset consisting of different distributions.

570 **C Runtime Estimation for Training on Real Devices**

571 Although we have theoretically analyzed the feasibility of training SuperEncoder using states from
 572 real devices (Section 3.2), its practical implementation poses significant challenges. Specifically,
 573 state-of-the-art quantum tomography techniques, such as classical shadow [55, 16], require numerous
 574 *snapshots*, each measuring a distinct observable.

575 To train SuperEncoder, each sample in the training dataset necessitates one classical shadow to obtain
 576 the prepared state. For instance, with the FractalDB-60 dataset, one training epoch requires 60,000
 577 classical shadows. Our experiments on the IBM cloud platform reveal an average runtime of 3.02

578 seconds per circuit job excluding queuing time. Suppose the number of snapshots is 1000, then the
579 total runtime to train SuperEncoder for 10 epochs is about 1,812,000,000 seconds³, roughly 57 years,
580 making the process prohibitively expensive and time-consuming.

581 However, quantum tomography is under active investigation, and we expect more efficient techniques
582 to emerge for acquiring noisy quantum states from real devices. Additionally, with the advancement
583 of quantum computing system, future systems may have tightly integrated quantum-classical hetero-
584 geneous architectures (shorter runtime per job) while being capable of executing numerous quantum
585 circuits in parallel (jobs within a classical shadow can execute in parallel). Hence, we anticipate the
586 training of SuperEncoder to be feasible in the future.

³ $10 \times 1000 \times 60000 \times 3.02$

587 **NeurIPS Paper Checklist**

588 **1. Claims**

589 Question: Do the main claims made in the abstract and introduction accurately reflect the
590 paper's contributions and scope?

591 Answer: [\[Yes\]](#)

592 Justification: This work aims at training-free approximate quantum state preparation. As
593 claimed in the abstract and introduction.

594 Guidelines:

- 595 • The answer NA means that the abstract and introduction do not include the claims
596 made in the paper.
- 597 • The abstract and/or introduction should clearly state the claims made, including the
598 contributions made in the paper and important assumptions and limitations. A No or
599 NA answer to this question will not be perceived well by the reviewers.
- 600 • The claims made should match theoretical and experimental results, and reflect how
601 much the results can be expected to generalize to other settings.
- 602 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
603 are not attained by the paper.

604 **2. Limitations**

605 Question: Does the paper discuss the limitations of the work performed by the authors?

606 Answer: [\[Yes\]](#)

607 Justification: SuperEncoder sacrifices fidelity, as discussed in Section 4.4.

608 Guidelines:

- 609 • The answer NA means that the paper has no limitation while the answer No means that
610 the paper has limitations, but those are not discussed in the paper.
- 611 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 612 • The paper should point out any strong assumptions and how robust the results are to
613 violations of these assumptions (e.g., independence assumptions, noiseless settings,
614 model well-specification, asymptotic approximations only holding locally). The authors
615 should reflect on how these assumptions might be violated in practice and what the
616 implications would be.
- 617 • The authors should reflect on the scope of the claims made, e.g., if the approach was
618 only tested on a few datasets or with a few runs. In general, empirical results often
619 depend on implicit assumptions, which should be articulated.
- 620 • The authors should reflect on the factors that influence the performance of the approach.
621 For example, a facial recognition algorithm may perform poorly when image resolution
622 is low or images are taken in low lighting. Or a speech-to-text system might not be
623 used reliably to provide closed captions for online lectures because it fails to handle
624 technical jargon.
- 625 • The authors should discuss the computational efficiency of the proposed algorithms
626 and how they scale with dataset size.
- 627 • If applicable, the authors should discuss possible limitations of their approach to
628 address problems of privacy and fairness.
- 629 • While the authors might fear that complete honesty about limitations might be used by
630 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
631 limitations that aren't acknowledged in the paper. The authors should use their best
632 judgment and recognize that individual actions in favor of transparency play an impor-
633 tant role in developing norms that preserve the integrity of the community. Reviewers
634 will be specifically instructed to not penalize honesty concerning limitations.

635 **3. Theory Assumptions and Proofs**

636 Question: For each theoretical result, does the paper provide the full set of assumptions and
637 a complete (and correct) proof?

638 Answer: [\[Yes\]](#)

639 Justification: All these necessary contents for theoretical results are included in Section 3.2.

640 Guidelines:

- 641 • The answer NA means that the paper does not include theoretical results.
- 642 • All the theorems, formulas, and proofs in the paper should be numbered and cross-
- 643 referenced.
- 644 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 645 • The proofs can either appear in the main paper or the supplemental material, but if
- 646 they appear in the supplemental material, the authors are encouraged to provide a short
- 647 proof sketch to provide intuition.
- 648 • Inversely, any informal proof provided in the core of the paper should be complemented
- 649 by formal proofs provided in appendix or supplemental material.
- 650 • Theorems and Lemmas that the proof relies upon should be properly referenced.

651 4. Experimental Result Reproducibility

652 Question: Does the paper fully disclose all the information needed to reproduce the main ex-

653 perimental results of the paper to the extent that it affects the main claims and/or conclusions

654 of the paper (regardless of whether the code and data are provided or not)?

655 Answer: [Yes]

656 Justification: Our code is open-source with instructions to reproduce our results, as described

657 in Section 4.1. We also describe the details of experiment settings in Appendix A.

658 Guidelines:

- 659 • The answer NA means that the paper does not include experiments.
- 660 • If the paper includes experiments, a No answer to this question will not be perceived
- 661 well by the reviewers: Making the paper reproducible is important, regardless of
- 662 whether the code and data are provided or not.
- 663 • If the contribution is a dataset and/or model, the authors should describe the steps taken
- 664 to make their results reproducible or verifiable.
- 665 • Depending on the contribution, reproducibility can be accomplished in various ways.
- 666 For example, if the contribution is a novel architecture, describing the architecture fully
- 667 might suffice, or if the contribution is a specific model and empirical evaluation, it may
- 668 be necessary to either make it possible for others to replicate the model with the same
- 669 dataset, or provide access to the model. In general, releasing code and data is often
- 670 one good way to accomplish this, but reproducibility can also be provided via detailed
- 671 instructions for how to replicate the results, access to a hosted model (e.g., in the case
- 672 of a large language model), releasing of a model checkpoint, or other means that are
- 673 appropriate to the research performed.
- 674 • While NeurIPS does not require releasing code, the conference does require all submis-
- 675 sions to provide some reasonable avenue for reproducibility, which may depend on the
- 676 nature of the contribution. For example
 - 677 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
 - 678 to reproduce that algorithm.
 - 679 (b) If the contribution is primarily a new model architecture, the paper should describe
 - 680 the architecture clearly and fully.
 - 681 (c) If the contribution is a new model (e.g., a large language model), then there should
 - 682 either be a way to access this model for reproducing the results or a way to reproduce
 - 683 the model (e.g., with an open-source dataset or instructions for how to construct
 - 684 the dataset).
 - 685 (d) We recognize that reproducibility may be tricky in some cases, in which case
 - 686 authors are welcome to describe the particular way they provide for reproducibility.
 - 687 In the case of closed-source models, it may be that access to the model is limited in
 - 688 some way (e.g., to registered users), but it should be possible for other researchers
 - 689 to have some path to reproducing or verifying the results.

690 5. Open access to data and code

691 Question: Does the paper provide open access to the data and code, with sufficient instruc-

692 tions to faithfully reproduce the main experimental results, as described in supplemental

693 material?

694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745

Answer: [Yes]

Justification: See Section 4.1.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We illustrate the experimental settings in Section 4.1, and provides additional details in Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Throughout our experiments, we set the random seed to be fixed for all libraries we used.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- 746
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- 747
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- 748
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- 749
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.
- 750
- 751
- 752
- 753
- 754
- 755

756 8. Experiments Compute Resources

757 Question: For each experiment, does the paper provide sufficient information on the com-
758 puter resources (type of compute workers, memory, time of execution) needed to reproduce
759 the experiments?

760 Answer: [Yes]

761 Justification: We describe the computer resources used in this paper in Section 4.1.

762 Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

771 9. Code Of Ethics

772 Question: Does the research conducted in the paper conform, in every respect, with the
773 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

774 Answer: [Yes]

775 Justification: We have read the code of ethics and followed its requirements.

776 Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

782 10. Broader Impacts

783 Question: Does the paper discuss both potential positive societal impacts and negative
784 societal impacts of the work performed?

785 Answer: [NA]

786 Justification: This work has no societal impact.

787 Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- 795
- 796
- 797
- 798
- 799
- 800
- 801
- 802
- 803
- 804
- 805
- 806
- 807
- 808
- 809
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
 - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
 - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

810 11. Safeguards

811 Question: Does the paper describe safeguards that have been put in place for responsible
812 release of data or models that have a high risk for misuse (e.g., pretrained language models,
813 image generators, or scraped datasets)?

814 Answer: [NA]

815 Justification: This paper poses no such risks as our released model and datasets are only
816 able to be used for quantum state preparation.

817 Guidelines:

- 818
- 819
- 820
- 821
- 822
- 823
- 824
- 825
- 826
- 827
- The answer NA means that the paper poses no such risks.
 - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
 - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
 - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

828 12. Licenses for existing assets

829 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
830 the paper, properly credited and are the license and terms of use explicitly mentioned and
831 properly respected?

832 Answer: [Yes]

833 Justification: We use an open-source dataset FractalDB, we cite the original paper and
834 indicates the version we use in Section 4.1.

835 Guidelines:

- 836
- 837
- 838
- 839
- 840
- 841
- 842
- 843
- 844
- 845
- 846
- The answer NA means that the paper does not use existing assets.
 - The authors should cite the original paper that produced the code package or dataset.
 - The authors should state which version of the asset is used and, if possible, include a URL.
 - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
 - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
 - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- 847
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- 848
- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.
- 849
- 850

851 **13. New Assets**

852 Question: Are new assets introduced in the paper well documented and is the documentation
853 provided alongside the assets?

854 Answer: [Yes]

855 Justification: We submit our assets in zip file and also put them on the anonymous github
856 repository, we have included a README file with detailed descriptions.

857 Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

866 **14. Crowdsourcing and Research with Human Subjects**

867 Question: For crowdsourcing experiments and research with human subjects, does the paper
868 include the full text of instructions given to participants and screenshots, if applicable, as
869 well as details about compensation (if any)?

870 Answer: [NA]

871 Justification: This paper does not involve crowdsourcing nor research with human subjects.

872 Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

881 **15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human
882 Subjects**

883 Question: Does the paper describe potential risks incurred by study participants, whether
884 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
885 approvals (or an equivalent approval/review based on the requirements of your country or
886 institution) were obtained?

887 Answer: [NA]

888 Justification: This paper does not involve crowdsourcing nor research with human subjects.

889 Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
 - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
 - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
 - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.
- 890
- 891
- 892
- 893
- 894
- 895
- 896
- 897
- 898
- 899