

Article

Autoencoding-Assisted Quantum Cloning Machine

Qian Jun Beh ¹ , Moritz Straeter ¹, Zeen Sun ¹ , Leong Chuan Kwek ^{1,2,3,4,*}  and Yuancheng Zhan ^{1,2,*} 

¹ Centre for Quantum Technologies (CQT), National University of Singapore, Singapore 117543, Singapore; qjbeh@u.nus.edu (Q.J.B.); moritz.straeter@u.nus.edu (M.S.); zeen_sun@u.nus.edu (Z.S.)

² School of Electrical and Electronic Engineering (EEE), Nanyang Technological University, Singapore 639798, Singapore

³ National Institute of Education (NIE), Nanyang Technological University, Singapore 637616, Singapore

⁴ MajuLab, CNRS-UNS-NUS-NTU International Joint Research Unit, Singapore 117543, Singapore

* Correspondence: cqtklc@nus.edu.sg (L.C.K.); zhan0530@e.ntu.edu.sg (Y.Z.)

Abstract

Quantum cloning machines are essential in quantum information processing, finding applications in areas such as quantum communication and cryptographic protocols. However, the fidelity of universal quantum cloning machines diminishes as the dimension of the Hilbert space increases, resulting in significantly lower efficiency when cloning high-dimensional quantum states compared to qubits. In this study, we introduce a Hybrid Quantum Autocloning Machine (HQAM) that combines quantum autoencoding with universal quantum cloning. The core concept involves compressing a high-dimensional quantum state into a lower-dimensional effective subspace through a quantum autoencoder, conducting the cloning process within this reduced subspace, and then reconstructing the state in the original Hilbert space. Our results show that, for input states with a strong overlap with the effective qubit subspace, the HQAM achieves cloning fidelities exceeding the benchmark fidelity of direct qutrit universal cloning and approaching the optimal qubit cloning limit, while maintaining robustness under noise. These findings demonstrate that compression-assisted cloning provides a practical strategy for improving cloning performance in high-dimensional quantum systems and may enable more efficient quantum information processing protocols.

Keywords: quantum information; quantum cloning; quantum autoencoding

1. Introduction

Quantum information processing unlocks fundamentally new capabilities for computation [1,2], communication [3,4], and simulation [5–8]. However, the practical implementation of quantum technologies is hindered by limited quantum resources. In near-term platforms [9–12], there is a scarcity of quantum memory, coherence times are finite, and manipulating high-dimensional quantum states often requires significant experimental and computational resources. These challenges highlight the need for protocols that leverage the intrinsic structure of quantum states to reduce the effective dimensionality of the information processed. A notable example is quantum cloning; while the no-cloning theorem prohibits the perfect duplication of arbitrary unknown quantum states [13], approximate cloning is permissible and plays a crucial role in quantum information theory. Universal quantum cloning machines (UQCMs) can produce optimal approximate copies of any state [14], but the fidelity of cloning diminishes as the dimension of the Hilbert space increases. This means that cloning higher-dimensional states, like qutrits, is inherently less



Academic Editor: Jie-Qiao Liao

Received: 26 March 2026

Revised: 4 May 2026

Accepted: 15 May 2026

Published: 18 May 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

accurate than cloning qubits. This observation suggests that if a high-dimensional state primarily exists within a lower-dimensional effective subspace, cloning within that reduced subspace could offer a more resource-efficient approach.

Quantum autoencoders (QAEs) [15,16] offer a natural framework for such dimensionality reduction. As the quantum analogue of classical autoencoders, QAEs compress quantum states into lower-dimensional latent spaces while preserving the information required for faithful reconstruction [15]. This approach is particularly effective when the relevant states occupy only a restricted region of the Hilbert space. Previous work has demonstrated both theoretical and experimental implementations of quantum compression, including photonic schemes that map qutrit states into effective qubit subspaces [17]. From a structural viewpoint, the optimal compression of an ensemble $\{p_i, |\psi_i\rangle\}$ is the principal component projector of its average state $\bar{\rho}$ —that is, quantum principal component analysis (quantum PCA) on the ensemble. In a fully classical simulation, this projector can be computed in closed form by diagonalising the empirical $\bar{\rho}$. In a quantum setting, however, $\bar{\rho}$ is not available, and obtaining it would require full state tomography of the ensemble. Quantum autoencoders provide a sample-based, tomography-free realisation of the same projector by training a parametrised unitary on per-state measurements alone [15,16,18], and are therefore the natural method when the latent subspace is not known a priori. Meanwhile, approximate quantum cloning has been extensively studied in both theory and experiments. The Bužek–Hillery UQCM provides the optimal symmetric $1 \rightarrow 2$ cloning protocol for qubits [19]. Despite these advances, quantum autoencoding and quantum cloning have largely been studied independently, and the potential benefits of combining state compression with quantum cloning remain largely unexplored.

In this work, we investigate a compression-assisted cloning strategy for high-dimensional quantum states that possess significant overlap with a lower-dimensional subspace. Instead of directly applying a d -dimensional cloning protocol, the input state is first compressed into an effective qubit subspace, cloning is performed within this reduced space, and the resulting states are subsequently reconstructed in the original Hilbert space. To realise this idea, we introduce a *Hybrid Quantum Autocloning Machine* (HQAM) that integrates a quantum autoencoder with the Bužek–Hillery UQCM. In this protocol, the encoder compresses the input qutrit state into a qubit representation while suppressing the population in an auxiliary junk mode. Cloning is then performed within the encoded qubit subspace, after which the cloned states are decoded back into the qutrit space.

We analyse this protocol using both analytical modelling and numerical optimisation. Our results show that on Haar-scrambled qutrit ensembles—in which the latent two-dimensional subspace V is unknown to the encoder—the trained autoencoder converges to the analytical quantum-PCA optimum within statistical precision. The resulting cloning fidelity exceeds the universal qutrit cloning bound $F = 3/4$ when the overlap with V is sufficiently large ($\chi \geq 3/\sqrt{10}$) and approaches the qubit cloning limit $F = 5/6$ as $\chi \rightarrow 1$. By contrast, a fixed canonical projection—equivalent to assuming the latent subspace is known and aligned with the lab basis—cannot solve the problem when the subspace is unknown; its mean fidelity sits at the Haar-random scramble limit $5/9$, independent of χ . The protocol degrades gracefully under leakage noise that drives states away from the effective qubit subspace. These results demonstrate that compression-assisted cloning provides a subspace-aware strategy for improving cloning performance in high-dimensional quantum systems whose latent structure is not known a priori and offer a pathway towards practical implementations of hybrid quantum cloning protocols.

2. Hybrid Quantum Autocloning Machine

We focus here on the theoretical and numerical HQAM framework. An illustrative photonic realisation of the encoding process is in Appendix A.

In quantum cloning, the no-cloning theorem states that no physical operation can perfectly copy an arbitrary unknown quantum state [13]. Nevertheless, approximate cloning is possible, and optimal approximate cloning protocols have been extensively studied. The Bužek–Hillery universal quantum cloning machine generates optimal symmetric copies of arbitrary qubit states. For a general qubit input state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, the cloning transformation can be written as

$$|\psi\rangle|\mathcal{R}\rangle|\mathcal{M}\rangle \rightarrow \sqrt{\frac{2}{3}}|\psi\rangle|\psi\rangle|\psi^\perp\rangle - \sqrt{\frac{1}{6}}[|\psi\rangle|\psi^\perp\rangle + |\psi^\perp\rangle|\psi\rangle]|\psi\rangle. \quad (1)$$

The performance of a cloning protocol is characterised by the cloning fidelity F , which measures the similarity between the cloned state and the original input. For universal symmetric $N \rightarrow M$ cloning of a d -dimensional quantum system, the optimal fidelity is given by [14]

$$F_{N \rightarrow M}(d) = \frac{N}{M} + \frac{(M-N)(N+1)}{M(N+d)}. \quad (2)$$

Since the optimal fidelity decreases as the Hilbert-space dimension d increases, cloning higher-dimensional quantum states is intrinsically less accurate than cloning lower-dimensional states. This limitation motivates the search for strategies that can exploit the structure within the state space to improve cloning performance.

Quantum autoencoders provide a natural framework for such dimensionality reduction. A quantum autoencoder compresses quantum states into lower-dimensional representations while preserving the information required for reconstruction. In a QAE, an encoder maps an input state in a high-dimensional Hilbert space into a lower-dimensional latent subspace, and a decoder subsequently reconstructs the original state from the compressed representation. Formally, the encoding process can be written as

$$|\psi\rangle_{\text{in}} \rightarrow U_e |\psi\rangle_{\text{in}} = |\psi\rangle_{\text{enc}},$$

where U_e is a unitary encoder, and the encoded state occupies an effective subspace of the original Hilbert space. The decoder reconstructs the input state through the inverse transformation $U_e^\dagger$. If the relevant quantum states predominantly occupy a lower-dimensional subspace, compressing the state prior to cloning may, therefore, provide a more resource-efficient strategy. This observation forms the basis of the compression-assisted cloning approach investigated in this work.

2.1. Analytical Model

Consider a general normalised qutrit state $|\psi\rangle_{\text{in}} = c_0|0\rangle + c_1|1\rangle + c_2|2\rangle$, with $|c_0|^2 + |c_1|^2 + |c_2|^2 = 1$ and $c_0, c_1, c_2 \in \mathbb{C}$. The quantum autoencoder first applies an encoding unitary, giving $U_e |\psi\rangle_{\text{in}}$, and then projects into a two-dimensional qutrit subspace, yielding the normalised logical qubit state

$$|\psi_{\text{enc}}\rangle = \frac{\hat{\Pi} U_e |\psi\rangle_{\text{in}}}{\sqrt{1-\varepsilon}}, \quad (3)$$

where $\hat{\Pi} = |1\rangle\langle 1| + |2\rangle\langle 2|$ and ε is the compression error, given by

$$\varepsilon = \text{tr}\left[(\mathbb{I} - \hat{\Pi}) U_e |\psi\rangle_{\text{in}} \langle\psi|_{\text{in}} U_e^\dagger\right]. \quad (4)$$

The projection is successful with probability $1 - \varepsilon$. Applying the Bužek–Hillery qubit quantum cloning machine to the encoded logical qubit state $|\psi_{\text{enc}}\rangle$ yields

$$\rho_{\text{enc}}^{(A)} = \rho_{\text{enc}}^{(B)} = \frac{5}{6} |\psi_{\text{enc}}\rangle \langle \psi_{\text{enc}}| + \frac{1}{6} |\psi_{\text{enc}}^{\perp}\rangle \langle \psi_{\text{enc}}^{\perp}|, \quad (5)$$

where $|\psi_{\text{enc}}^{\perp}\rangle$ is the state in the encoded two-dimensional subspace orthogonal to $|\psi_{\text{enc}}\rangle$. Decoding with $U_e^{\dagger}$ gives

$$\rho_{\text{out}}^{(A)} = \rho_{\text{out}}^{(B)} = U_e^{\dagger} \left(\frac{5}{6} |\psi_{\text{enc}}\rangle \langle \psi_{\text{enc}}| + \frac{1}{6} |\psi_{\text{enc}}^{\perp}\rangle \langle \psi_{\text{enc}}^{\perp}| \right) U_e. \quad (6)$$

The one-copy fidelity F with the input state is therefore

$$F = \langle \psi |_{\text{in}} U_e^{\dagger} \left(\frac{5}{6} |\psi_{\text{enc}}\rangle \langle \psi_{\text{enc}}| + \frac{1}{6} |\psi_{\text{enc}}^{\perp}\rangle \langle \psi_{\text{enc}}^{\perp}| \right) U_e | \psi \rangle_{\text{in}} \quad (7)$$

$$= \frac{5}{6} \left| \langle \psi |_{\text{in}} U_e^{\dagger} | \psi_{\text{enc}} \rangle \right|^2 + \frac{1}{6} \left| \langle \psi |_{\text{in}} U_e^{\dagger} | \psi_{\text{enc}}^{\perp} \rangle \right|^2. \quad (8)$$

Since $|\psi_{\text{enc}}\rangle$ is proportional to $\hat{\Gamma} U_e | \psi \rangle_{\text{in}}$, we have

$$\left| \langle \psi |_{\text{in}} U_e^{\dagger} | \psi_{\text{enc}} \rangle \right|^2 = 1 - \varepsilon. \quad (9)$$

Moreover, because $|\psi_{\text{enc}}^{\perp}\rangle$ lies in the range of $\hat{\Gamma}$ and is orthogonal to $|\psi_{\text{enc}}\rangle$,

$$\langle \psi |_{\text{in}} U_e^{\dagger} | \psi_{\text{enc}}^{\perp} \rangle = 0. \quad (10)$$

Hence

$$F = \frac{5}{6} (1 - \varepsilon). \quad (11)$$

The HQAM produces copies with higher fidelity than direct universal qutrit cloning if

$$\frac{5}{6} (1 - \varepsilon) > \frac{3}{4} \Rightarrow \varepsilon < \frac{1}{10}. \quad (12)$$

For an ensemble of qutrit quantum states $\mathcal{E} = \{p_i, |\psi_i\rangle\}$ with average state

$$\bar{\rho} = \sum_i p_i |\psi_i\rangle \langle \psi_i|, \quad (13)$$

the average compression error $\bar{\varepsilon}$ is given by

$$\bar{\varepsilon} = 1 - \text{tr} [\hat{\Gamma} U_e \bar{\rho} U_e^{\dagger}]. \quad (14)$$

This quantity is minimised when U_e diagonalises $\bar{\rho}$ and maps the two eigenvectors with the largest eigenvalues to $|1\rangle$ and $|2\rangle$. For an average state $\bar{\rho}$ with eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$, the minimum average compression error is

$$\bar{\varepsilon}_{\text{min}} = \lambda_3. \quad (15)$$

The corresponding maximum average cloning fidelity is

$$\bar{F}_{\text{max}} = \frac{5}{6} (\lambda_1 + \lambda_2) = \frac{5}{6} (1 - \lambda_3) = \frac{5}{6} (1 - \bar{\varepsilon}_{\text{min}}). \quad (16)$$

In the limit of perfect compressibility of the ensemble ($\bar{\epsilon} = 0$), the qubit cloning fidelity $F = 5/6$ is recovered, as expected.

Equation (16) characterises the optimal encoder for an ensemble $\{p_i, |\psi_i\rangle\}$: the unitary U_e diagonalises the average state $\bar{\rho}$ and maps its two leading eigenvectors onto the qubit subspace $\{|1\rangle, |2\rangle\}$, while the eigenvector with the smallest eigenvalue is mapped to the junk mode $|0\rangle$. This is the principal component projector of $\bar{\rho}$, i.e., quantum principal component analysis of the ensemble. The numerical implementation in Section 2.2 provides a sample-based realisation of this projector that does not require explicit access to $\bar{\rho}$.

2.2. Numerical Implementation

To complement the illustrative photonic discussion in Appendix A, we numerically evaluate the performance of the Hybrid Quantum Autocloning Machine. As illustrated in Figure 1, the protocol consists of three stages: the input qutrit state is first compressed into an effective qubit subspace by an optimised encoder, the encoded state is then cloned using the Bužek–Hillery universal quantum cloning machine, and the resulting two-clone state is finally decoded back into the qutrit space. The performance of the HQAM is quantified by the fidelity between the decoded output state and the original input state.




Figure 1. Illustration of the hybrid autocloning machine. The input is a qutrit state, the encoder outputs an encoded qubit state in the effective two-dimensional subspace, and $U^\dagger \otimes U^\dagger$ denotes the decoder. The labels “ 3×1 ” and “ 2×1 ” denote qutrit and qubit column vectors, respectively.

We model the realistic setting in which the encoder receives states drawn from a structured ensemble whose latent two-dimensional subspace is unknown a priori. Datasets are generated in two steps: (i) draw a state $|\varphi\rangle$ in canonical coordinates with the $|0\rangle$ component fixed in magnitude and a uniformly random phase and the $(|1\rangle, |2\rangle)$ sampled uniformly on the χ -radius sphere within $\text{span}\{|1\rangle, |2\rangle\}$,

$$|\varphi\rangle = c_0 |0\rangle + c_1 |1\rangle + c_2 |2\rangle, \quad |c_0|^2 = 1 - \chi^2, \quad |c_1|^2 + |c_2|^2 = \chi^2, \quad (17)$$

so that every state has overlap with $\text{span}\{|1\rangle, |2\rangle\}$ exactly equal to χ ; (ii) apply a fixed Haar-random unitary $U_{\text{scramble}} \in U(3)$ so that the actual input state is

$$|\psi_{\text{in}}\rangle = U_{\text{scramble}} |\varphi\rangle. \quad (18)$$

The resulting ensemble has fixed overlap χ with the rotated subspace $V = U_{\text{scramble}} \cdot \text{span}\{|1\rangle, |2\rangle\}$ (a fully two-dimensional plane in $\mathbb{C}^3$; the parameter χ controls how concentrated the data is within V , not the dimensionality of V itself). The encoder receives $|\psi_{\text{in}}\rangle$ without any access to U_{scramble} or to the empirical $\bar{\rho}$. Seven values $\chi \in \{0.85, 0.875, 0.90, 0.925, 0.95, 0.975, 0.99\}$ are considered, and ten independent draws of U_{scramble} are used to average over the choice of latent subspace. Each $(\chi, U_{\text{scramble}})$ pair contributes a training set of 2000 states and a held-out test set of 500 states. Details of the state generation procedure are provided in Appendix C. Because every state has overlap exactly χ with V and is uniform within V , the average state has eigenvalues $(\chi^2/2, \chi^2/2, 1 - \chi^2)$

with eigenvectors aligned with V and $V^\perp$. For $\chi > \sqrt{2/3}$, which covers the entire range considered here, the smallest eigenvalue is $\lambda_3 = 1 - \chi^2$, and Equation (16) reduces to

$$\bar{F}_{\max}(\chi) = \frac{5}{6} \chi^2. \quad (19)$$

This is the direct counterpart of Equation (16) for the structured ensembles benchmarked in Section 3.1; the empirical HQAM fidelities reported in Table 1 agree with Equation (19) to within sample noise ($\lesssim 5 \times 10^{-4}$). The encoder is parametrised as

$$U_e = e^{iH}, \quad H = \sum_{j=1}^8 w_j \sigma_j, \quad (20)$$

where σ_j are the Gell–Mann matrices and w_j are trainable parameters. The choice of a parametric encoder optimised from per-state measurements is what makes the protocol applicable in a quantum setting where the ensemble is provided as a stream of physical states rather than classical descriptions. As shown in Section 2.1, the optimal encoder for a given ensemble is the principal component projector of $\bar{\rho}$. In a fully classical simulation, this projector can be computed in closed form by diagonalising the empirical $\bar{\rho}$. In a quantum setting, however, $\bar{\rho}$ is not directly available; obtaining it would require full state tomography of the ensemble. The autoencoder side-steps this by training U_e on per-state measurements of the junk mode population alone, providing a sample-based, tomography-free realisation of quantum PCA on the input ensemble [15–17]. We verify in Section 3.1 that the trained encoder converges to the analytical optimum of Equation (16) within statistical precision. The encoded state is obtained as

$$|\psi_{\text{encoded}}\rangle = U_e |\psi_{\text{in}}\rangle, \quad \mathcal{L}_{\text{junk}} = |\langle 0 | \psi_{\text{encoded}} \rangle|^2. \quad (21)$$

Here $|\psi_{\text{encoded}}\rangle$ is still formally a qutrit-state vector in the three-dimensional Hilbert space. When the junk mode population is suppressed, the relevant information is concentrated in the effective two-dimensional subspace, and only this qubit component is passed to the cloning stage.

The encoded state is then processed by the Bužek–Hillery UQCM. In numerical form, the cloning stage is described by

$$|\psi_{\text{encoded}}\rangle_{2 \times 1} \xrightarrow{\text{cloning}} |\psi_{\text{cloned}}\rangle_{8 \times 1} \rightarrow \rho_{8 \times 8}^{\text{cloned}} \xrightarrow{\text{Tr}_M[\rho]} \rho_{4 \times 4}^{AB}, \quad (22)$$

where A and B denote the two clone subsystems, and the machine degree of freedom has been traced out. The subscript “ 2×1 ” here denotes the two-component column vector associated with the effective qubit subspace, not the full encoded qutrit representation. To decode the cloned state, we take the decoder unitary to be the adjoint of the encoder,

$$U_d = U_e^\dagger \otimes U_e^\dagger, \quad (23)$$

since the decoding acts on the tensor product space of the two clones. As the cloned state ρ^{AB} lies in a two-qubit space while the decoder acts in the two-qutrit space, we embed $\rho_{4 \times 4}^{AB}$ into the 9×9 qutrit tensor space using the isometry

$$\mathbb{I}_{3 \times 2} = \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbb{I}_{3 \times 2}^\dagger \mathbb{I}_{3 \times 2} = \mathbb{I}_{2 \times 2}. \quad (24)$$

The embedded cloned state is therefore

$$\rho_{9 \times 9}^{AB} = (\mathbb{I}_{3 \times 2} \otimes \mathbb{I}_{3 \times 2}) \rho_{4 \times 4}^{AB} (\mathbb{I}_{3 \times 2} \otimes \mathbb{I}_{3 \times 2})^\dagger, \quad (25)$$

which is then decoded as

$$\rho_{\text{decoded}} = U_d \rho_{9 \times 9}^{AB} U_d^\dagger. \quad (26)$$

The state ρ_{decoded} in Equation (26) is the full jointly decoded two-qutrit state. The single-clone fidelities reported in the numerical results are computed from its one-clone marginals,

$$\rho_{\text{decoded}}^A = \text{Tr}_B[\rho_{\text{decoded}}], \quad \rho_{\text{decoded}}^B = \text{Tr}_A[\rho_{\text{decoded}}], \quad (27)$$

with

$$F_A = \langle \psi_{\text{in}} | \rho_{\text{decoded}}^A | \psi_{\text{in}} \rangle, \quad F_B = \langle \psi_{\text{in}} | \rho_{\text{decoded}}^B | \psi_{\text{in}} \rangle. \quad (28)$$

Equivalently, because the decoder is local on the two clone systems, tracing out one subsystem after applying $U_e^\dagger \otimes U_e^\dagger$ gives the same one-clone marginal as decoding the corresponding reduced clone state separately. The separate reduced-clone implementation therefore evaluates the same F_A and F_B as the marginal evaluation of the full two-qutrit decoded state, although it does not store the full joint correlations.

The encoder parameters are optimised by minimising the total loss

$$L = \mathcal{L}_{\text{cloning}} + \beta \mathcal{L}_{\text{junk}}, \quad (29)$$

where $\mathcal{L}_{\text{cloning}} = 1 - F_A - F_B + (F_A - F_B)^2$ and $\mathcal{L}_{\text{junk}} = |\langle 0 | \psi_{\text{encoded}} \rangle|^2$, F_A and F_B denote the fidelities of the two decoded clones, and the penalty term $(F_A - F_B)^2$ enforces approximate symmetry between them. The parameter β controls the relative weight of the junk mode suppression term. For the symmetric Bužek–Hillery cloner of Section 2.1, the symmetry penalty $(F_A - F_B)^2$ vanishes identically, $F_A = F_B$, and so contributes nothing to the gradient; it is retained as a guard for non-ideal or asymmetric cloning channels.

We analyse the algebraic structure of the loss in the ideal pipeline. The cloning step in Section 2.1 is a Bužek–Hillery universal cloner; by universality, the per-state single-clone fidelities depend on the input only through its overlap $\chi(\psi)$ with the encoder’s qubit subspace, $F_A(\psi) = F_B(\psi) = \frac{5}{6} \chi^2(\psi)$. Substituting into the loss definitions gives $\mathcal{L}_{\text{cloning}}(\psi) = 1 - \frac{5}{3} \chi^2(\psi)$ and $\mathcal{L}_{\text{junk}}(\psi) = 1 - \chi^2(\psi)$, so the total loss reduces to

$$\mathcal{L}(\psi) = (1 + \beta) - \left(\frac{5}{3} + \beta\right) \chi^2(\psi), \quad (30)$$

an affine function of $\chi^2(\psi)$ alone. Averaging over the training ensemble, minimising $\langle \mathcal{L} \rangle$ is therefore equivalent to maximising $\langle \chi^2 \rangle$, which is exactly the criterion that defines the analytical optimum (16). The loss is structurally consistent with the quantum-PCA solution of Section 2.1: any choice of $\beta > 0$ drives the encoder to the same minimiser, and the role of β is only to rescale the gradient magnitude during training. The cloning loss term is retained because it makes the operational objective (cloning fidelity) explicit in the loss and because the reduction (30) relies on universality of the ideal cloner: as soon as the cloning channel deviates from the Bužek–Hillery form (for example, under cloning-stage noise), F_A ceases to be a function of χ^2 alone and the two loss terms become structurally distinct.

All simulations were implemented in Python 3.10 on an RTX 4090 GPU. The trainable parameters are the coefficients $\mathbf{w} = (w_1, \dots, w_8)$ appearing in the encoder Hamiltonian $H = \sum_{j=1}^8 w_j \sigma_j$. These parameters are trained once for an ensemble of qutrit states with

a prescribed overlap χ and then kept fixed. Accordingly, the HQAM is designed for structured families of input states with prior subspace information, rather than arbitrary unknown qutrit states.

For reproducibility and held-out evaluation, every $(\chi, U_{\text{scramble}})$ combination is checked in as separate pseudo-training and pseudo-test files. For the main setting $\beta = 8.0$, $\eta = 0.01$, training is repeated for five random weight initialisations (seeds 0–4) for each of ten independent draws of U_{scramble} (seeds 0–9), giving 50 trained encoders per χ . Each run uses 200 training epochs, all 2000 training states, and all 500 held-out test states. Aggregated statistics are reported in Table 1. The standard deviation captures both seed-to-seed variation in the trained weights and scramble-to-scramble variation in the latent subspace.

Table 1. Held-out fidelity statistics of the trained HQAM encoder at $\beta = 8.0$, $\eta = 0.01$, aggregated over 50 runs per χ (10 Haar-random scrambles $\times$ 5 weight initialisations). The fourth column gives the analytical prediction $(5/6)\chi^2$ from Equation (19); the gap column reports the residual $(5/6)\chi^2 - \bar{F}_{\text{HQAM}}$, consistent with finite-sample noise across all χ .

χ	$\bar{F}_{\text{HQAM}}$ (Mean)	Std	$(5/6)\chi^2$ (Equation (19))	Gap
0.850	0.60232	2×10^{-4}	0.60208	-2.4×10^{-4}
0.875	0.63792	9×10^{-5}	0.63802	1.0×10^{-4}
0.900	0.67504	1×10^{-4}	0.67500	-4×10^{-5}
0.925	0.71279	1×10^{-4}	0.71302	2.3×10^{-4}
0.950	0.75214	3×10^{-5}	0.75208	-6×10^{-5}
0.975	0.79214	2×10^{-5}	0.79219	5×10^{-5}
0.990	0.81677	1×10^{-5}	0.81675	-2×10^{-5}

The parameters were optimised using gradient descent to minimise the total loss function defined in Equation (29). At iteration k , the parameter vector is updated according to

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \nabla_{\mathbf{w}} \mathcal{L}(\mathbf{w}_k), \quad (31)$$

where η denotes the learning rate controlling the step size of the optimisation. All trained encoders reported below use $\eta = 0.01$ and $\beta = 8$, for which training converges to the analytical optimum within ~ 25 epochs; smaller learning rates (e.g., $\eta = 0.001$) reach the same optimum but require substantially more epochs.

3. Results

In this section, we present the performance of the hybrid autocloning machine with the corresponding code available in [20].

We first benchmark the cloning fidelity achieved by the HQAM against the theoretical limits of universal cloning on Haar-scrambled structured ensembles, then investigate the robustness of the protocol under leakage noise that drives the state away from the effective qubit subspace.

3.1. Cloning Fidelity Benchmark

We benchmark the cloning performance of the HQAM on the Haar-scrambled ensembles described in Section 2.2, in which the latent two-dimensional subspace $V = U_{\text{scramble}} \cdot \text{span}\{|1\rangle, |2\rangle\}$ is unknown to the encoder. The reference values are the universal symmetric $1 \rightarrow 2$ cloning fidelities of Equation (2): $F = 5/6$ for qubits and $F = 3/4$ for qutrits. The analytical maximum achievable on a structured ensemble with smallest eigenvalue λ_3 of $\bar{\rho}$ is $\bar{F}_{\text{max}} = (5/6)(1 - \lambda_3)$, realised by the quantum-PCA encoder of Section 2.1.

The benchmark comparison in this work is intentionally restricted to universal cloning benchmarks. In particular, the relevant direct qutrit benchmark is the universal qutrit cloning fidelity $F = 3/4$. For the structured ensembles considered here, the HQAM exceeds this direct universal qutrit cloning benchmark when the input states have sufficiently large overlap with the effective qubit subspace. This result should be interpreted as improved performance relative to universal qutrit cloning on structured input families, rather than as a claim of optimality among all state-dependent qutrit cloners.

Figure 2 reports the held-out cloning fidelity as a function of χ for two encoders: HQAM (trained on $|\psi_{\text{in}}\rangle$ via the procedure of Section 2.2) and a fixed canonical projection (the trivial isometry that maps $|\psi_{\text{in}}\rangle$ onto $\text{span}\{|1\rangle, |2\rangle\}$ of the lab basis without any learning). Three observations can be made:




Figure 2. Cloning fidelity of HQAM versus the fixed canonical projection on Haar-scrambled qutrit ensembles, as a function of the overlap χ . Each input state exactly overlaps χ with the latent subspace V . Each point aggregates 10 Haar-random draws of U_{scramble} ; HQAM additionally averages over 5 random weight initialisations per scramble. Shaded bands indicate $\pm\text{std}$ across scramble seeds; the HQAM band is smaller than the markers. The trained HQAM encoder (orange) saturates the analytical maximum $\bar{F}_{\text{max}}(\chi) = (5/6)\chi^2$ of Equation (19)—the quantum-PCA optimum on the empirical $\bar{\rho}$ —and exceeds the universal qutrit cloning bound $F = 3/4$ for $\chi > \chi_{\text{crit}} = 3/\sqrt{10} \approx 0.9487$, approaching the qubit cloning limit $F = 5/6$ as $\chi \rightarrow 1$. The fixed canonical projection (gray) lacks access to the unknown latent V ; its mean fidelity sits at the Haar-random scramble limit $5/9 \approx 0.556$, independent of χ .

(i) *HQAM saturates the analytical maximum.* For every χ , the trained encoder agrees with the analytical prediction $\bar{F}_{\text{max}}(\chi) = (5/6)\chi^2$ from Equation (19) to within $\sim 5 \times 10^{-4}$ (Table 1), and the seed-to-seed standard deviation is below 2×10^{-4} . We additionally verified numerically that the analytical quantum-PCA encoder—obtained by diagonalising the empirical $\bar{\rho}$ on the training set—agrees with the trained HQAM to within $\max |\bar{F}_{\text{HQAM}} - \bar{F}_{\text{PCA}}| < 2 \times 10^{-4}$ across all $(\chi, U_{\text{scramble}})$ combinations, consistent with the analytical prediction of Section 2.1: the autoencoder converges to the principal component projector of $\bar{\rho}$ from per-state measurements alone. The fidelities of the two clones coincide ($F_A = F_B$; see Figure A3), confirming the symmetry expected from the Bužek–Hillery protocol.

(ii) *HQAM exceeds the universal qutrit cloning bound for sufficient χ .* The trained encoder approaches the qubit cloning limit $5/6$ as $\chi \rightarrow 1$ ($\bar{F} = 0.817$ at $\chi = 0.99$). The crossover

with the universal qutrit bound $F = 3/4$ follows directly from Equation (19): setting $\frac{5}{6}\chi^2 = 3/4$ gives

$$\chi_{\text{crit}} = \sqrt{\frac{9}{10}} = \frac{3}{\sqrt{10}} \approx 0.9487. \quad (32)$$

For $\chi > \chi_{\text{crit}}$, the structured ensemble is sufficiently compressible that HQAM exceeds direct universal qutrit cloning; for $\chi < \chi_{\text{crit}}$, even the analytical optimum $\bar{F}_{\text{max}}$ remains below $3/4$, so subspace exploitation cannot beat the universal qutrit cloner. The empirical crossover between $\chi = 0.925$ ($\bar{F} = 0.713$) and $\chi = 0.95$ ($\bar{F} = 0.752$) in Figure 2 matches this bound.

(iii) *The fixed canonical projection fails on scrambled data.* Without access to U_{scramble} , projecting onto $\text{span}\{|1\rangle, |2\rangle\}$ of the lab basis achieves $\bar{F} \approx 0.556$ across all χ values, with a 10-scramble band whose width grows from ± 0.02 at $\chi = 0.85$ to ± 0.11 at $\chi = 0.99$. This is consistent with the analytically predicted Haar-random scramble limit $5/9 \approx 0.556$: in expectation over Haar-random U_{scramble} , the per-state cloning fidelity reduces to $(5/6)\langle\chi_{\text{lab}}^2\rangle = (5/6) \cdot (2/3) = 5/9$, independent of χ . The gap to HQAM grows from 0.05 at $\chi = 0.85$ to 0.26 at $\chi = 0.99$, confirming that learning is required when the latent subspace is unknown; a fixed canonical isometry suffices only when the input distribution is hand-aligned with the lab basis.

3.2. Robustness Under Leakage Noise

The leakage noise analysis below uses the canonical-aligned ensemble (no Haar scrambling) in order to isolate the effect of leakage on the cloning fidelity from the subspace recovery question of Section 3.1. The noise strength ϵ in this section therefore plays the role of a channel parameter rather than the structural overlap threshold χ of Section 2.2. The overlap of the input qutrit with the effective qubit subspace, denoted by χ , can be interpreted as the action of a noise channel that drives the state away from the ideal two-dimensional encoding space. From the perspective of the encoded qubit, this behaviour resembles the effect of a depolarising process. For reference, the standard depolarising channel acting on a qubit density matrix ρ is given by

$$\rho \rightarrow (1 - \epsilon)\rho + \epsilon \frac{\mathbb{I}}{2}, \quad (33)$$

where ϵ denotes the noise probability. A general derivation of depolarising channels for d -dimensional systems and a special case $d = 2$ are provided in Appendices E.1 and E.2, respectively.

In the present work, the leakage from the qubit subspace is modelled as a channel that transfers population from the encoded qubit subspace into the fixed qutrit state $|0\rangle$. The corresponding mapping can be written as

$$\rho \rightarrow (1 - \epsilon)\rho + \epsilon |0\rangle\langle 0|, \quad (34)$$

where $\epsilon := 1 - \chi$ and the effective qubit subspace is spanned by $\{|1\rangle, |2\rangle\}$. A proof that this mapping constitutes a valid, completely positive, and trace-preserving (CPTP) operation is given in Appendix E.3.

We evaluate the robustness of the HQAM under this leakage model by training the protocol on noisy input states with $\epsilon \in [0.01, 0.20]$. The corresponding results are shown in Figure 3. As expected, the average cloning fidelity initially decreases as the noise level increases, reflecting the degradation of information encoded in the qubit subspace. However, a non-monotonic behaviour appears when ϵ approaches the upper end of this interval. This effect is analysed more systematically by extending the simulations to $\epsilon \in [0.18, 0.50]$, as shown in Figure A4.




Figure 3. Performance of the HQAM with the leakage channel, where $\epsilon \in [0.01, 0.20]$. **(a)** Average fidelity $\langle F \rangle$ versus ϵ . **(b)** Average loss versus ϵ . The values for both fidelity and loss are taken at the plateau during training for each ϵ .

The observed behaviour can be understood analytically by considering the linearity of the HQAM pipeline. The full protocol consists of a sequence of operations—encoding, cloning, embedding, partial tracing, and decoding—each of which is a valid CPTP map. Consequently, the overall HQAM transformation $\sigma(\cdot)$ is also linear. Applying this map to the noisy state in Equation (34) gives

$$\sigma(\rho) = \sigma((1 - \epsilon)\rho + \epsilon |0\rangle \langle 0|) = (1 - \epsilon)\sigma(\rho) + \epsilon\sigma(|0\rangle \langle 0|), \quad (35)$$

where linearity has been used. Since the input state $|\psi_{\text{in}}\rangle$ is pure, the fidelity between the input and output states can be written as

$$F(|\psi_{\text{in}}\rangle, \sigma(\rho)) = \langle \psi_{\text{in}} | \sigma(\rho) | \psi_{\text{in}} \rangle \quad (36)$$

$$= (1 - \epsilon)F_{\text{qubit}} + \epsilon F_{\text{qutrit}}, \quad (37)$$

where F_{qubit} denotes the fidelity associated with the encoded qubit subspace and F_{qutrit} corresponds to the contribution from the $|0\rangle$ component.

When the noise level is small ($\epsilon \leq 0.05$), the state remains largely confined to the qubit subspace and the overall fidelity is dominated by F_{qubit} , allowing the HQAM to outperform direct qutrit cloning. As the noise increases ($0.05 \leq \epsilon \leq 0.17$), the encoded qubit component becomes increasingly distorted, leading to a gradual decline in cloning fidelity. For sufficiently large noise ($\epsilon \geq 0.17$), the state becomes dominated by the $|0\rangle$ component, which the decoder reconstructs accurately. In this regime, the overall fidelity is therefore governed primarily by F_{qutrit} , explaining the observed increase in fidelity.

Figure 4 further illustrates the contributions from the qubit and qutrit components. As ϵ increases, the fidelity associated with the qubit subspace F_{qubit} deteriorates due to the increasing leakage from the encoded subspace. In contrast, the fidelity associated with the $|0\rangle$ component, F_{qutrit} , improves as the noise channel increasingly biases the state toward $|0\rangle$,

which the decoder reconstructs reliably. Importantly, the overall fidelity remains bounded by the theoretical limits imposed by universal quantum cloning machines. For qubits, the optimal fidelity is $F = 5/6$, while for qutrits, it is $F = 3/4$. The fact that the observed fidelities never exceed these limits confirms that the HQAM remains fully consistent with the no-cloning theorem.




Figure 4. Performance of the HQAM with the leakage channel, where $\epsilon \in [0.01, 1.00]$. The two horizontal dashed lines are the average fidelities achieved by UQCM for qubit and qutrit, at $\frac{5}{6}$ and $\frac{3}{4}$, respectively. (a) The overall fidelity of the HQAM. (b) The fidelity of the effective qubit subspace. (c) The fidelity of the qutrit subspace.

4. Discussions and Conclusions

In this work, we proposed a Hybrid Quantum Autocloning Machine that integrates quantum autoencoding with universal quantum cloning. By compressing qutrit states into an effective qubit subspace through a unitary encoder, U_e , the protocol enables cloning to be performed in a lower-dimensional space before reconstructing the state in the original Hilbert space. This compression-assisted strategy reduces the effective resource requirements while remaining fully consistent with the limits imposed by the no-cloning theorem. The encoder plays the role of a sample-based, tomography-free realisation of quantum principal component analysis on the input ensemble: it converges to the principal component projector of the empirical average state $\bar{\rho}$.

An illustrative photonic encoding model is provided in Appendix A. Numerical optimisation of the encoder confirms that the HQAM achieves high cloning fidelities when the input states have a strong overlap with the effective qubit subspace. In particular, for χ above the analytical crossover $\chi_{\text{crit}} \geq 3/\sqrt{10}$, the achieved fidelity exceeds the benchmark fidelity of direct qutrit UQCM cloning and approaches the optimal qubit cloning limit as $\chi \rightarrow 1$. On Haar-scrambled ensembles where the latent subspace is unknown to the encoder,

the trained HQAM saturates the analytical maximum $\bar{F}_{\max}(\chi) = (5/6)\chi^2$ predicted by the principal component projector, while a fixed canonical isometry that lacks knowledge of the latent subspace cannot solve the problem and sits at the Haar-random scramble limit $5/9$. We further analyse the robustness of the protocol under a leakage noise model, which shows that increasing noise degrades the fidelity associated with the encoded qubit component, while the contribution from the reconstructed $|0\rangle$ state becomes dominant at large ϵ . This behaviour explains the non-monotonic fidelity trends observed in the simulations and highlights the role of the encoded subspace in determining the overall cloning performance.

Overall, the HQAM demonstrates how compression techniques can be combined with approximate quantum cloning to improve cloning performance for structured high-dimensional quantum states whose latent subspace is not classically available. Future work could extend this framework to higher-dimensional systems, alternative cloning protocols such as phase-covariant cloning [21], and experimental implementations on photonic or near-term quantum computing platforms [6,22–27].

Author Contributions: L.C.K., Y.Z., and Q.J.B. jointly conceived the idea. M.S. and Z.S. assisted with the theory and simulations. L.C.K. supervised and coordinated all the work. Y.Z. and Q.J.B. wrote the manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: L.C.K., and Y.Z. acknowledge support from NRF, Singapore, Ministry of Education, Singapore.

Data Availability Statement: The data underlying the results presented in this paper are available from the corresponding authors upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Illustrative Photonic Encoding Model




Figure A1. Schematic of the compression-assisted encoding process. The photon state at each intermediate step is labelled. Diagram modified from Ref. [17].

The main text focuses on the abstract HQAM model and its numerical implementation, while the photonic realisation is presented only as background material. Building on the information previously, we now analyse the compression-assisted cloning protocol in detail. The overall process of the analytical model for photonic implementation of the quantum autoencoder is illustrated in Figure A1. For convenience, we express the photon state in the fast ($|f\rangle$) and slow ($|s\rangle$) polarisation bases. The relation to the horizontal ($|H\rangle$) and vertical ($|V\rangle$) bases is given in Appendix B. The initial state produced by the single-photon source can therefore be written as

$$|i\rangle = \cos \theta |f\rangle + \sin \theta |s\rangle, \quad (\text{A1})$$

where θ is the polarisation parameter that determines the superposition between the fast and slow axes, and $|f\rangle$ and $|s\rangle$ denote the fast and slow polarisation axes, respectively.

The encoding procedure consists of a sequence of optical transformations implemented by waveplates and beam displacers, whose detailed operations are summarised in Appendix B.

After passing through a half-wave plate (HWP) and a quarter-wave plate (QWP), the photon state becomes

$$|1\rangle = \cos \theta |f\rangle_2 - i \sin \theta |s\rangle_2,$$

where subscripts 1 and 2 denote the upper and lower paths. The beam displacer then separates the polarisation components into two spatial paths, resulting in the state,

$$|2\rangle = \cos \theta |f\rangle_1 - i \sin \theta |s\rangle_2,$$

A polarisation scrambler is subsequently applied to the lower path, producing

$$|3\rangle = \cos \theta |f\rangle_1 - i \sin \theta (\cos \phi |f\rangle_2 + \sin \phi |s\rangle_2),$$

where ϕ denotes the scrambler rotation angle. The encoding stage then recombines and rotates the polarisation components using additional waveplates. After the action of an HWP at 45° , the photon state becomes

$$|4\rangle = \cos \theta |s\rangle_1 - i \sin \theta (\cos \phi |f\rangle_2 - i \sin \phi |s\rangle_2).$$

Combining the polarisation components in the upper path yields

$$|5\rangle = (-i \sin \theta \cos \phi) |f\rangle_1 + \cos \theta |s\rangle_1 - \sin \theta \sin \phi |s\rangle_2.$$

New spatial paths, a and b , are introduced, where path a corresponds to the recombined upper spatial mode and path b corresponds to the lower spatial mode. These paths represent physically distinct output ports after the polarisation-to-path encoding process. The photon state can be written as

$$|6\rangle = (-i \sin \theta \cos \phi) |f\rangle_a - i \cos \theta |s\rangle_a - \sin \theta \sin \phi |f\rangle_b.$$

Finally, the mode corresponding to the discarded dimension is routed into the junk mode. The remaining encoded state is therefore

$$|7\rangle = -i \sin \theta \cos \phi |f\rangle_c - \sin \theta \sin \phi |f\rangle_a - i \cos \theta |s\rangle_a.$$

Here, c denotes the discarded spatial mode corresponding to the auxiliary dimension, which is routed into the junk mode.

The sequence of optical operations described above can be represented by a set of unitary transformations $U_1, U_2, \dots, U_7$ acting on the photon state. The states $|i\rangle$ to $|4\rangle$ describe a single-photon polarisation state together with path labels, while the four-component vector below is only a bookkeeping representation in the combined path–polarisation basis, rather than a generic four-level input state. Writing the input state in vector form,

$$|\psi_{\text{input}}\rangle = \begin{bmatrix} 0 \\ 0 \\ \cos \theta \\ -i \sin \theta \end{bmatrix}, \quad (\text{A2})$$

Here, $|\psi_{\text{input}}\rangle$ represents an effective state vector description of the photonic state in a chosen computational basis. Each optical element corresponds to a matrix operation on

this state vector. The overall encoder transformation is therefore obtained by multiplying the individual unitaries,

$$U_e = U_7 U_6 U_5 U_4 = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -i \\ -i & 0 & 0 & 0 \end{bmatrix}. \quad (\text{A3})$$

To minimise the population of the junk mode (corresponding to path c), one may set $\theta = 0$, which suppresses the amplitude in that mode. Under this condition, the encoded state reduces to the qubit state

$$|\psi_{\text{out}}\rangle = -i|s\rangle. \quad (\text{A4})$$

Applying the Bužek–Hillery universal quantum cloning machine (Equation (1)) to this encoded state yields

$$|\psi_{\text{out}}\rangle |R\rangle |M\rangle \rightarrow -i\sqrt{\frac{2}{3}}|s\rangle |s\rangle |f\rangle + i\sqrt{\frac{1}{6}}(|s\rangle |f\rangle + |f\rangle |s\rangle)|s\rangle. \quad (\text{A5})$$

Expressing the resulting state in vector form gives

$$|\psi_{\text{cloned}}\rangle = i\sqrt{\frac{1}{6}} \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 & -2 & 0 \end{bmatrix}^\top. \quad (\text{A6})$$

The density matrix of the output state is $\rho_{\text{out}} = |\psi_{\text{cloned}}\rangle \langle \psi_{\text{cloned}}|$.

Tracing out the machine subsystem yields the reduced density matrices of the two clones

$$\rho_{\text{out}}^{(1)} = \rho_{\text{out}}^{(2)} = \begin{bmatrix} \frac{1}{6} & 0 \\ 0 & \frac{5}{6} \end{bmatrix}. \quad (\text{A7})$$

The cloning fidelity is therefore $F = \frac{5}{6}$, which corresponds to the optimal fidelity for symmetric qubit cloning, as predicted by Equation (2).

Appendix B. Transformation of Basis

The standard polarisation basis is $\{|H\rangle, |V\rangle\}$, while the alternative basis is defined by the fast and slow axes, denoted by $\{|f\rangle, |s\rangle\}$. Expressing the photon state in the principal basis simplifies the description of optical elements such as waveplates.

The transformation from the $\{|H\rangle, |V\rangle\}$ basis to the $\{|f\rangle, |s\rangle\}$ basis corresponds to a rotation (θ is the basis rotation angle)

$$\begin{bmatrix} \hat{f} \\ \hat{s} \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \hat{H} \\ \hat{V} \end{bmatrix}. \quad (\text{A8})$$

The inverse transformation is therefore

$$\begin{bmatrix} \hat{H} \\ \hat{V} \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \hat{f} \\ \hat{s} \end{bmatrix}. \quad (\text{A9})$$

Figure A2 illustrates this basis transformation and the corresponding photon state $|\psi\rangle$ in the rotated coordinate system. In our model, the fast-axis component can be represented as

$$|f\rangle = \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix}, \quad (\text{A10})$$

while the slow-axis component corresponds to the orthogonal state.




Figure A2. Visualisation of the basis transformation from the standard polarisation basis $\{|H\rangle, |V\rangle\}$ to the principal basis $\{|f\rangle, |s\rangle\}$. The solid black horizontal and vertical axes denote the original $\{|H\rangle, |V\rangle\}$ basis directions. The dashed black diagonal axes denote the rotated principal directions associated with the slow and fast axes, $\{|s\rangle, |f\rangle\}$. The angle θ indicates the rotation relating the two bases. The red vector represents the photon state $|\psi\rangle$, shown in the same plane so that its components can be resolved in the principal basis.

When the principal axes of a waveplate align with the chosen $\{|f\rangle, |s\rangle\}$ basis, the corresponding transformations take particularly simple matrix forms. A half-wave plate (HWP) introduces a relative phase shift of π between the slow and fast components and is represented by

$$\text{HWP}_{\text{principal}} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (\text{A11})$$

Similarly, a quarter-wave plate (QWP) introduces a phase shift of $\pi/2$ on the slow component, corresponding to

$$\text{QWP}_{\text{principal}} = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}. \quad (\text{A12})$$

A half-wave plate oriented at 45° exchanges the fast and slow components and can therefore be written as

$$\text{HWP}_{45} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \quad (\text{A13})$$

Finally, the scrambler used in the encoding circuit mixes the polarisation components and can be modelled by

$$\text{Scrambler} = \begin{bmatrix} 1 & 0 \\ \cos \phi & \sin \phi \end{bmatrix}, \quad (\text{A14})$$

where ϕ denotes the rotation angle of the scrambler.

Appendix C. State Generation Procedure




Figure A3. Plots for fidelities of both copies, F_A and F_B , for different χ at $\beta = 8.0$ and $\eta = 0.001$. (a) $\chi = 0.99$. (b) $\chi = 0.95$. (c) $\chi = 0.90$. (d) $\chi = 0.85$. The F_A and F_B are overlap, which means they are identical, aligning with the theory.

In the numerical simulations, four ensembles of 2000 normalised qutrit states were generated at random. The parameter χ denotes the overlap between the qutrit state and the effective qubit subspace used by the encoder.

Four datasets were generated corresponding to different overlap thresholds, namely

$$\chi \geq 0.99, \quad \chi \geq 0.95, \quad \chi \geq 0.90, \quad \chi \geq 0.85.$$

Random qutrit states satisfying a given overlap constraint were generated by sampling three complex coefficients

$$c_k = x_k + iy_k, \quad x_k, y_k \sim N(0, 1), \quad k = 0, 1, 2, \quad (\text{A15})$$

followed by normalisation

$$|\psi_{\text{in}}\rangle = \frac{1}{\sqrt{|c_0|^2 + |c_1|^2 + |c_2|^2}} \begin{bmatrix} c_0 \\ c_1 \\ c_2 \end{bmatrix}. \quad (\text{A16})$$

For simulations requiring a fixed overlap value $\chi = c$, a modified procedure was used. First, the amplitude of the non-qubit component was chosen as

$$A = \sqrt{1 - \chi^2}. \quad (\text{A17})$$

A random phase $\phi \sim U(0, 2\pi)$ was sampled, and the first coefficient was set to

$$c_0 = Ae^{i\phi}. \quad (\text{A18})$$

Next, two complex numbers

$$z_k = x_k + iy_k, \quad x_k, y_k \sim N(0, 1), \quad k = 1, 2 \quad (\text{A19})$$

were generated and normalised to obtain

$$\tilde{z}_k = \frac{z_k}{\sqrt{|z_1|^2 + |z_2|^2}}. \quad (\text{A20})$$

Finally, the coefficients in the qubit subspace were scaled according to the prescribed overlap,

$$c_1 = \chi \tilde{z}_1, \quad c_2 = \chi \tilde{z}_2, \quad (\text{A21})$$

which yields the state

$$|\psi_{\text{in}}\rangle = \begin{bmatrix} c_0 \\ c_1 \\ c_2 \end{bmatrix}. \quad (\text{A22})$$

By construction, the state is normalised and satisfies

$$|c_0|^2 + |c_1|^2 + |c_2|^2 = 1. \quad (\text{A23})$$

Appendix D. Optimised Encoder Unitaries

In this appendix, we report the encoder unitaries obtained from the optimisation procedure. The encoder U_e was trained under four different overlap conditions between the qutrit Hilbert space and the effective qubit subspace. In all trials, the hyperparameters were fixed to $\beta = 8.0$ and $\eta = 0.01$. The optimisation was performed separately for datasets with overlap thresholds $\chi \geq 0.99$, $\chi \geq 0.95$, $\chi \geq 0.90$, and $\chi \geq 0.85$.

For the case $\chi \geq 0.99$, the optimised encoder unitary is

$$U_e^{(0.99)} = \begin{bmatrix} -0.961721 - 0.2739745i & 0.004417 + 0.001720511i & -0.001957 + 0.001891414i \\ -0.004914 - 0.001463557i & -0.92631 - 0.2890803i & 0.238638 - 0.03752783i \\ -0.001791 - 0.0006179848i & 0.192974 - 0.1453321i & 0.43669 - 0.8665668i \end{bmatrix}.$$

For the case $\chi \geq 0.95$, the optimised encoder unitary is

$$U_e^{(0.95)} = \begin{bmatrix} -0.309344 + 0.9509026i & -0.005232 + 0.00350046i & 0.007071 - 0.0007754599i \\ -0.002113 + 1.688474 \times 10^{-4}i & -0.156878 - 0.5922081i & -0.024709 - 0.7899771i \\ 0.001987 - 0.009043409i & -0.665418 + 0.4264479i & 0.577064 - 0.2056022i \end{bmatrix}.$$

For the case $\chi \geq 0.90$, the optimised encoder unitary is

$$U_e^{(0.90)} = \begin{bmatrix} -0.9954852 - 0.092547i & 0.01183747 - 0.001949i & 0.009575159 + 0.014441i \\ -0.01932928 - 0.005484i & -0.4346251 - 0.234607i & -0.4781819 - 0.725947i \\ -0.005968226 + 0.002199i & -0.0007486880 + 0.869434i & 0.2344573 - 0.434825i \end{bmatrix}.$$

For the case $\chi \geq 0.85$, the optimised encoder unitary is

$$U_e^{(0.85)} = \begin{bmatrix} -0.9802706 + 0.1965275i & -0.002942577 + 0.02091174i & 0.0004458398 - 0.0004020995i \\ 0.01452004 - 0.008657617i & 0.1423112 + 0.7797590i & -0.3457613 - 0.5018831i \\ 0.003714416 + 0.01211348i & -0.6093131 - 0.003839584i & -0.7250697 + 0.3206806i \end{bmatrix}.$$

These matrices represent the encoder transformations learned by the HQAM for different input state distributions characterised by the overlap parameter χ .

Appendix E. Depolarising and Leakage Channels




Figure A4. Performance of the HQAM under the leakage channel for $\epsilon \in [0.01, 0.50]$ in a representative fixed-seed run. (a) Total loss as a function of ϵ , showing the overall optimisation objective used in training. (b) Cloning loss as a function of ϵ , isolating the cloning-related contribution to the objective. (c) Average fidelity $\langle F \rangle$ as a function of ϵ , showing the non-monotonic dependence of cloning performance on the leakage strength over this interval.

In this appendix, we briefly review the depolarising channel for a general d -dimensional quantum system and show that the leakage model used in the main text corresponds to a valid, completely positive, and trace-preserving (CPTP) quantum operation.

Appendix E.1. Depolarising Channel in d Dimensions

For a d -dimensional quantum system with density matrix ρ , the depolarising channel is defined as

$$\mathcal{E}(\rho) = (1 - p)\rho + p \frac{\mathbb{I}}{d}, \quad (\text{A24})$$

where $\mathbb{I}$ denotes the identity operator and $\mathbb{I}/d$ is the maximally mixed state.

This channel can be expressed using the operator sum representation

$$\mathcal{E}(\rho) = \sum_k K_k \rho K_k^\dagger, \quad (\text{A25})$$

where the Kraus operators satisfy the completeness relation

$$\sum_k K_k^\dagger K_k = \mathbb{I}. \quad (\text{A26})$$

A convenient choice of Kraus operators is

$$K_0 = \sqrt{1-p} \mathbb{I}, \quad (\text{A27})$$

and

$$K_i = \sqrt{\frac{p}{d^2-1}} E_i, \quad i = 1, \dots, d^2-1, \quad (\text{A28})$$

where $\{E_i\}$ denotes an orthonormal basis of traceless operators satisfying

$$\text{Tr}(E_i^\dagger E_j) = \delta_{ij}. \quad (\text{A29})$$

With this choice, one obtains

$$\mathcal{E}(\rho) = (1-p)\rho + \frac{p}{d^2-1} \sum_{i=1}^{d^2-1} E_i \rho E_i^\dagger. \quad (\text{A30})$$

Using Schur's lemma, we have

$$\sum_{i=1}^{d^2-1} E_i \rho E_i^\dagger = \mathbb{I} - \frac{1}{d} \rho. \quad (\text{A31})$$

Then we obtain

$$\mathcal{E}(\rho) = (1-p)\rho + \frac{p}{d^2-1} \left(\mathbb{I} - \frac{1}{d} \rho \right) = \left(1-p - \frac{p}{d(d^2-1)} \right) \rho + \frac{p}{d^2-1} \mathbb{I}. \quad (\text{A32})$$

Since this expression does not recover the standard depolarising channel, we instead adopt the Kraus operators

$$K_0 = \sqrt{1-p-\frac{p}{d^2}} \mathbb{I}, \quad K_i = \sqrt{\frac{p}{d}} E_i, \quad i = 1, \dots, d^2-1, \quad (\text{A33})$$

Appendix E.2. Special Case: Qubit Depolarising Channel

For a qubit system ($d = 2$), the channel reduces to

$$\mathcal{E}(\rho) = (1-p)\rho + p \frac{\mathbb{I}}{2}, \quad (\text{A34})$$

which is the form used in Equation (33) of the main text.

A common Kraus representation employs the Pauli operators $\{\sigma_x, \sigma_y, \sigma_z\}$,

$$K_0 = \sqrt{1-p} \mathbb{I}, \quad K_i = \sqrt{\frac{p}{3}} \sigma_i, \quad i = x, y, z. \quad (\text{A35})$$

The completeness condition is verified as

$$\sum_i K_i^\dagger K_i = \mathbb{I}. \quad (\text{A36})$$

Appendix E.3. Validity of the Leakage Channel

The leakage channel introduced in Equation (34) of the main text is




Figure A5. Diagnostic held-out fidelity summary for the structured pseudo-test ensembles. The figure compares the learned HQAM repeated-seed evaluation with internal fixed-subspace and clone-free reconstruction diagnostics, together with the universal qutrit and qubit cloning reference fidelities. This diagnostic figure is provided to contextualise the structured-prior setting and should not be interpreted as evidence of superiority over all fixed-subspace or state-dependent qutrit cloning strategies.

$$\mathcal{E}(\rho) = (1 - \epsilon)\rho + \epsilon |0\rangle\langle 0|. \quad (\text{A37})$$

This channel describes population leakage from the encoded qubit subspace into a fixed qutrit state $|0\rangle$. To show that this mapping is a valid quantum operation, we construct the Kraus operators

$$R_i = |0\rangle\langle i|, \quad i = 0, \dots, d-1. \quad (\text{A38})$$

For any density matrix ρ ,

$$\sum_{i=1}^d R_i \rho R_i^\dagger = |0\rangle\langle 0| \text{Tr}(\rho). \quad (\text{A39})$$

Since $\text{Tr}(\rho) = 1$, this map replaces the input state by $|0\rangle\langle 0|$. The Kraus operators satisfy

$$\sum_i R_i^\dagger R_i = \sum_i |i\rangle\langle i| = \mathbb{I}. \quad (\text{A40})$$

Therefore, the leakage channel can be written as the convex combination

$$\mathcal{E}(\rho) = (1 - \epsilon)\rho + \epsilon |0\rangle\langle 0|, \quad (\text{A41})$$

which is a mixture of the identity channel and a replacement channel. Since both maps are CPTP, their convex combination is also CPTP. Hence, the leakage model used in the main text represents a valid quantum operation.

References

1. Bennett, C.H.; DiVincenzo, D.P. Quantum information and computation. *Nature* **2000**, *404*, 247–255. [[CrossRef](#)]
2. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information*; Cambridge University Press: Cambridge, UK, 2010.
3. Gisin, N.; Thew, R. Quantum communication. *Nat. Photonics* **2007**, *1*, 165–171. [[CrossRef](#)]
4. Couteau, C.; Barz, S.; Durt, T.; Gerrits, T.; Huwer, J.; Prevedel, R.; Rarity, J.; Shields, A.; Weihs, G. Applications of single photons to quantum communication and computing. *Nat. Rev. Phys.* **2023**, *5*, 326–338. [[CrossRef](#)]

5. Burger, A.; Kwek, L.C.; Poletti, D. Digital quantum simulation of the spin-Boson model under Markovian open-system dynamics. *Entropy* **2022**, *24*, 1766. [\[CrossRef\]](#)
6. Zhan, Y.; Zhang, H.; Erbanni, R.; Burger, A.; Wan, L.; Jiang, X.; Chae, S.; Liu, A.Q.; Poletti, D.; Kwek, L.C. Loop Quantum Photonic Chip for Coherent Multi-Time-Step Evolution. *Laser Photonics Rev.* **2026**, e02699. [\[CrossRef\]](#)
7. Gross, C.; Bloch, I. Quantum simulations with ultracold atoms in optical lattices. *Science* **2017**, *357*, 995–1001. [\[CrossRef\]](#)
8. Daley, A.J.; Bloch, I.; Kokail, C.; Flannigan, S.; Pearson, N.; Troyer, M.; Zoller, P. Practical quantum advantage in quantum simulation. *Nature* **2022**, *607*, 667–676. [\[CrossRef\]](#) [\[PubMed\]](#)
9. Lvovsky, A.I.; Sanders, B.C.; Tittel, W. Optical quantum memory. *Nat. Photonics* **2009**, *3*, 706–714. [\[CrossRef\]](#)
10. Julsgaard, B.; Sherson, J.; Cirac, J.I.; Fiurášek, J.; Polzik, E.S. Experimental demonstration of quantum memory for light. *Nature* **2004**, *432*, 482–486. [\[CrossRef\]](#)
11. Lau, J.W.Z.; Lim, K.H.; Shrotriya, H.; Kwek, L.C. NISQ computing: Where are we and where do we go? *AAPPS Bull.* **2022**, *32*, 27. [\[CrossRef\]](#)
12. Bharti, K.; Cervera-Lierta, A.; Kyaw, T.H.; Haug, T.; Alperin-Lea, S.; Anand, A.; Degroote, M.; Heimonen, H.; Kottmann, J.S.; Menke, T.; et al. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.* **2022**, *94*, 015004. [\[CrossRef\]](#)
13. Wootters, W.K.; Zurek, W.H. A single quantum cannot be cloned. *Nature* **1982**, *299*, 802–803. [\[CrossRef\]](#)
14. Scarani, V.; Iblisdir, S.; Gisin, N.; Acín, A. Quantum cloning. *Rev. Mod. Phys.* **2005**, *77*, 1225–1256. [\[CrossRef\]](#)
15. Romero, J.; Olson, J.P.; Aspuru-Guzik, A. Quantum autoencoders for efficient compression of quantum data. *Quantum Sci. Technol.* **2017**, *2*, 045001. [\[CrossRef\]](#)
16. Bravo-Prieto, C. Quantum autoencoders with enhanced data encoding. *Mach. Learn. Sci. Technol.* **2021**, *2*, 035028. [\[CrossRef\]](#)
17. Pepper, A.; Tischler, N.; Pryde, G.J. Experimental Realization of a Quantum Autoencoder: The Compression of Qutrits via Machine Learning. *Phys. Rev. Lett.* **2019**, *122*, 060501. [\[CrossRef\]](#)
18. Lloyd, S.; Mohseni, M.; Rebentrost, P. Quantum principal component analysis. *Nat. Phys.* **2014**, *10*, 631–633. [\[CrossRef\]](#)
19. Bužek, V.; Hillery, M. Quantum copying: Beyond the no-cloning theorem. *Phys. Rev. A* **1996**, *54*, 1844. [\[CrossRef\]](#) [\[PubMed\]](#)
20. Jun, B.Q. Quantum Autocloning Machine. 2026. Available online: <https://github.com/01QJ10/ACM> (accessed on 23 March 2026).
21. D’Ariano, G.M.; Macchiavello, C. Optimal phase-covariant cloning for qubits and qutrits. *Phys. Rev. A* **2003**, *67*, 042306. [\[CrossRef\]](#)
22. Hoch, F.; Rodari, G.; Caruccio, E.; Polacchi, B.; Carvacho, G.; Giordani, T.; Doosti, M.; Nicolau, S.; Pentangelo, C.; Piacentini, S.; et al. Variational quantum cloning machine on an integrated photonic interferometer. *Opt. Quantum* **2025**, *3*, 351–359. [\[CrossRef\]](#)
23. Zhan, Y.; Zhang, H.; Burger, A.; Cai, H.; Polett, D.; Kwek, L.; Liu, A. Highly Efficient Photonic Multiplexing Chip Using Low-Loss Silicon Nitride Recirculating Loops. In *Conference on Lasers and Electro-Optics (CLEO)*; IEEE: Charlotte, NC, USA, 2024; pp. 1–2.
24. Wang, H.; Ralph, T.C.; Renema, J.J.; Lu, C.Y.; Pan, J.W. Scalable photonic quantum technologies. *Nat. Mater.* **2025**, *24*, 1883–1897. [\[CrossRef\]](#) [\[PubMed\]](#)
25. Zhang, R.; Lou, Y.; Liu, X.; Wang, J.; Liu, S.; Jing, J. Experimental demonstration of deterministic all-optical quantum cloning with phase-conjugate inputs. *Photonics Res.* **2025**, *14*, 286–296. [\[CrossRef\]](#)
26. Pelofske, E.; Bärtschi, A.; Eidenbenz, S.; Garcia, B.; Kiefer, B. Probing quantum telecloning on superconducting quantum processors. *IEEE Trans. Quantum Eng.* **2024**, *5*, 1–19. [\[CrossRef\]](#)
27. Zhan, Y.; Zhang, H.; Wan, L.; Karim, M.; Cai, H.; Kwek, L.; Liu, A. A hybrid quantum computer for quantum finance computation of Monte-Carlo distribution. In *Proceedings of the Novel Optical Materials and Applications*; Optica Publishing Group: Maastricht, The Netherlands, 2022; p. JTh4A–4.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.