Toward the Goldilocks Blind Compression of Quantum States

Hyunho Cha[1] **Chae-Yeun Park**[2] **Jungwoo Lee**[1*]

[1]NextQuantum and Department of Electrical and Computer Engineering, Seoul National University
[2]School of Integrated Technology and Department of Quantum Information, Yonsei University

Abstract

Quantum autoencoders (QAEs) are learning architectures that compress quantum data into a low-dimensional latent state while preserving the information needed for reconstruction. We study blind single-copy compression of quantum states through a k-qubit bottleneck and investigate the minimal circuit width required to attain the information-theoretic optimum under average infidelity. Between the conventional architecture, which is narrow but nonuniversal, and fully general *completely positive and trace preserving* (CPTP) realizations, which are universal but overparameterized, we identify a balanced regime. We prove that for every distribution of pure n-qubit states, there exists a QAE with k encoder ancillas and n decoder ancillas that achieves the optimal fidelity over all CPTP encoder–decoder pairs. The encoder-side statement is sharp in that we construct source families for which every optimal scheme necessarily uses at least k encoder ancillas, thereby determining the universal encoder threshold exactly. On the decoder side, we show that isometric decoders are optimal for several analytically tractable source families, but we also exhibit an explicit counterexample demonstrating that decoder isometry is not universally sufficient. Nevertheless, numerical experiments indicate that the performance gap is practically negligible.

1 Introduction

The rapid development of controllable quantum platforms has brought into focus a central challenge in quantum physics, namely how researchers can manipulate and preserve information encoded in exponentially large Hilbert spaces with finite and noisy resources. Manipulation of quantum states serves as the structural foundation for diverse fields, ranging from quantum simulation of many-body dynamics and chemistry to quantum computation, communication, and sensing [Lloyd, 1996; Nielsen and Chuang, 2010; Bernien et al., 2017; Degen et al., 2017; Feynman, 2018; Preskill, 2018; Bharti et al., 2022]. Yet the same exponential growth of state space that fuels quantum advantage also makes quantum memory and coherent processing increasingly scarce commodities. For near-term devices in particular, the available number of high-fidelity qubits and coherent operations remains a dominant bottleneck [Preskill, 2018; Bharti et al., 2022]. This necessitates architectures that use the quantum hardware at an appropriate scale by being neither underpowered for the task nor overprovisioned in ways that amplify noise.

Compression is the canonical response to resource scarcity in classical information theory [Shannon, 1948], and quantum information theory has long established that quantum data can likewise be compressed. The quantum noiseless coding theorem shows that an i.i.d. source described by a density matrix ρ can be asymptotically and faithfully compressed to $S(\rho) := -\operatorname{Tr}(\rho \log \rho)$ qubits per signal [Jozsa and Schumacher, 1994; Schumacher, 1995; Dür et al., 2001; Koashi and Imoto, 2001a,b;

*Corresponding author: junglee@snu.ac.kr

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Winter, 2002; Khanian and Winter, 2022], with refinements that address converses and general fidelity limits [Barnum et al., 1996], universality with minimal prior knowledge [Jozsa et al., 1998; Hayashi and Matsumoto, 2002], and extensions to mixed-state sources and different notions of side information or visibility of the signal identity [Horodecki, 1998; Hayden et al., 2002]. Experimental demonstrations of this compression have also been realized [Mitsumori et al., 2003]. At the same time, many emerging use cases operate far from the asymptotic i.i.d. regime and often call for single-copy, distribution-dependent compression under experimentally meaningful distortion measures [Barnum, 2000; Datta et al., 2012, 2013; Wilde et al., 2013; Salek et al., 2018; Hadiashar and Nayak, 2020].

A parallel set of ideas has emerged from machine learning. Classical autoencoders learn low-dimensional latent representations by jointly training an encoder and decoder to minimize reconstruction error, enabling data compression, denoising, and feature extraction [Hinton and Salakhutdinov, 2006; Vincent et al., 2008; Bengio et al., 2013]. Variational autoencoders extend this paradigm into probabilistic latent-variable models that support sampling and regularization of representations [Kingma and Welling, 2013; Rezende et al., 2014; Zhao et al., 2017]. These concepts have begun to influence quantum information processing through variational quantum algorithms and quantum machine learning models that optimize physically motivated cost functions using parameterized quantum circuits [Schuld et al., 2015; Biamonte et al., 2017; McClean et al., 2018; Benedetti et al., 2019; Khoshaman et al., 2019; Sim et al., 2019; Cerezo et al., 2021; Holmes et al., 2022].

Within this landscape, quantum autoencoders (QAEs) were proposed as a route to compression of quantum data that may not admit efficient classical descriptions [Romero et al., 2017]. In a QAE, a trainable encoder maps an n-qubit input state into a k-qubit latent register ($k < n$) plus discarded degrees of freedom. A decoder then aims to reconstruct the original state from the latent register. QAEs have been studied theoretically and deployed in diverse physical contexts [Romero et al., 2017; Ding et al., 2019; Lamata et al., 2019; Pepper et al., 2019; Bondarenko and Feldmann, 2020; Huang et al., 2020; Cao and Wang, 2021; Ma et al., 2023]. These developments demonstrate that QAE-style compression is becoming an architectural primitive for quantum technologies.

A major open issue, however, concerns resource optimality. Most QAE implementations adopt a unitary–trace architecture with a circuit acting on the system and a small number of ancillas (or auxiliary quantum systems), which can severely restrict the set of realizable encoder and decoder channels. From the perspective of quantum channel theory, the most general encoder and decoder are *completely positive and trace preserving* (CPTP) maps [Stinespring, 1955; Kraus, 1971; Jamiołkowski, 1972; Jiang et al., 2013]. Indeed, recent fully quantum variational autoencoder frameworks have proposed universal constructions capable of representing arbitrary CPTP maps by introducing sufficiently large ancillas [Wang et al., 2025]. While attractive in principle, this "maximally expressive" design collides with noisy reality and high training costs [McClean et al., 2018; Preskill, 2018; Bharti et al., 2022]. This motivates a fundamental question: for a given objective, what is the minimal circuit width required to achieve the information-theoretic optimum?

In this work, we investigate this question for blind, single-copy quantum state compression under the *infidelity* loss [Mahler et al., 2013; Huang et al., 2020; Cerezo et al., 2021]. In essence, for the infidelity objective there is a minimally universal architecture wide enough to reach the information-theoretic optimum, yet narrow enough to avoid the redundant ancilla overhead.

On the encoder side, we determine the ancilla cost sharply. Specifically, we prove matching sufficient and necessary conditions on the encoder ancilla size. In this sense, the encoder ancilla requirement is characterized exactly. This singles out a natural minimal encoder architecture that is expressive enough to attain the optimum while avoiding unnecessary overhead.

We further provide both analytical and numerical evidence for a conceptually appealing practical principle. Although our counterexample demonstrates that, in blind single-copy compression with infidelity loss, an isometric decoder need not always attain the theoretical optimum, numerical results suggest that restricting the decoder to be isometric nevertheless yields empirically near-optimal performance. We prove optimality of isometric decoders for several sources with closed-form distributions, and we find consistent numerical support on datasets constructed from MNIST-encoded quantum states [LeCun et al., 1998] and additional physical benchmarks, where the observed performance gap is negligible. This makes isometric decoders particularly attractive in practice. They offer near-optimal performance while reducing model size and computational overhead.

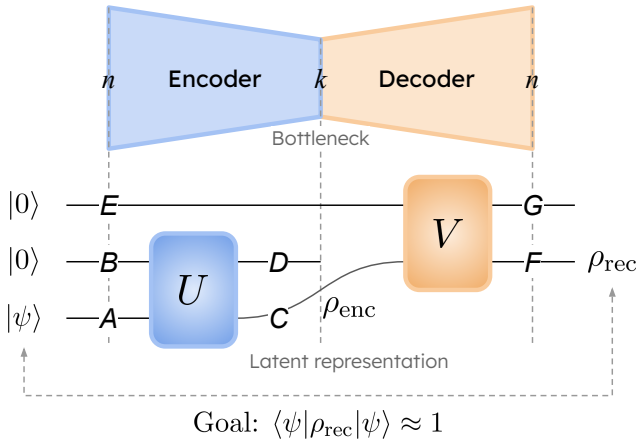

Figure 1: General structure of a QAE that compresses an n-qubit input state $|\psi\rangle$ into a k-qubit latent system and subsequently reconstructs the state using a decoder. Labels A–G identify the individual quantum registers.

2 Preliminaries and prior work

For a positive integer d, let $H_d \cong \mathbb{C}^d$ be a d-dimensional Hilbert space and let $L(H_d)$ denote the linear operators on H_d. A qubit is a two-dimensional quantum system, so an n-qubit register has Hilbert space $H_{2^n} \cong (\mathbb{C}^2)^{\otimes n}$ and hence dimension 2^n. For quantum registers $A, B, \ldots$, we write $H_A, H_B, \ldots$ for the associated Hilbert spaces and $H_{AB} := H_A \otimes H_B$ for the composite system. If the same physical space is viewed under two factorizations, we write for example $AB \equiv CD$. We use the computational basis $\{|x\rangle : x \in \{0,1\}^n\}$ of an n-qubit system (equivalently $\{|j\rangle : 0 \le j \le 2^n - 1\}$ after fixing the binary-to-integer identification). In Dirac notation, $|\psi\rangle$ is a *ket*, $\langle\psi| := |\psi\rangle^\dagger$ is the corresponding *bra*, and $\langle\phi|\psi\rangle$ is the inner product. State vectors are always taken to have unit norm.

A pure state is represented by a unit vector $|\psi\rangle$, or equivalently by the rank-one *density operator* $|\psi\rangle\langle\psi|$. More generally, a mixed state on a Hilbert space H is a density operator $\rho \in L(H)$, meaning $\rho \ge 0$ and $\mathrm{Tr}(\rho) = 1$. Tensor products are used for spaces, state vectors, and operators. For a bipartite state $\rho_{AB} \in L(H_{AB})$, the reduced state on A is $\rho_A := \mathrm{Tr}_B(\rho_{AB})$, where Tr_B is the partial trace. Operationally, this is the operation of discarding subsystem B. Likewise $\rho_B := \mathrm{Tr}_A(\rho_{AB})$.

A *superoperator* is a linear map between operator spaces. The physical maps considered here are quantum channels, namely CPTP maps $\Phi : L(H_A) \to L(H_B)$. Trace preservation means $\mathrm{Tr}[\Phi(X)] = \mathrm{Tr}[X]$, while complete positivity means that $\Phi \otimes \mathrm{id}_R$ sends positive operators to positive operators for every auxiliary register R. Closed-system evolution is unitary: if $U \in \mathrm{U}(d)$, the unitary group on H_d, then the induced channel is the conjugation map

$$\mathrm{Ad}_U(\rho) := U\rho U^\dagger.$$

More generally, every CPTP map admits a *Kraus representation*

$$\Phi(\rho) = \sum_{i=1}^{r} K_i \rho K_i^\dagger, \qquad \sum_{i=1}^{r} K_i^\dagger K_i = I,$$

and, equivalently, a Stinespring realization by adjoining an ancilla in a fixed state, applying a joint unitary, and partially tracing out an environment. The smallest possible number r of Kraus operators is the *Kraus rank* of Φ.

A more detailed mathematical formulation can be found in Appendix B.

Blind single-copy compression means that the encoder receives one copy of an unknown source state $|\psi\rangle$, drawn from a source distribution μ, and must compress it to a smaller k-qubit register without being given a classical label for $|\psi\rangle$. The decoder then attempts to reconstruct the original state from that single compressed quantum register. When a pure state is said to be Haar-random, it is drawn from the unitarily invariant probability measure on the unit sphere of the relevant Hilbert space (or subspace) [Haar, 1933]. A schematic depiction of the blind compression setting is provided in Figure 1.

Prior architectures. Among the many QAE constructions studied in the literature, two architectures are especially relevant for the resource question considered here. The *conventional QAE* introduced in [Romero et al., 2017; Huang et al., 2020; Bravo-Prieto, 2021; Cao and Wang, 2021; Ngairangbam et al., 2022; Srikumar et al., 2022; Wu et al., 2024] uses an n-qubit unitary on the input register, discards $n-k$ qubits to obtain a k-qubit latent register, and reconstructs by adjoining $n-k$ fresh ancillas followed by another n-qubit unitary. Anticipating the (n, k, n_B, n_E) notation introduced in Section 3, this is the $(n, k, 0, n-k)$ architecture. Its main virtue is its small circuit width, since both the encoder and decoder act on only n qubits. However, it constrains the encoder to be a unitary-on-system followed by a partial trace with no additional encoder ancilla, and it constrains the decoder to be an isometric channel. Accordingly, it is not universal for blind single-copy compression through a k-qubit bottleneck.

At the opposite end of the expressivity spectrum, [Wang et al., 2025] proposed the *ζ-QVAE*, a quantum variational autoencoder with mixed latent representations and an explicit latent-space regularization term. In that framework, both the encoder and decoder are modeled as general CPTP maps via *Stinespring dilations* [Stinespring, 1955] with sufficiently large ancillas, so the construction is universal. Specialized to the compression setting studied here, one convenient realization corresponds to $n_B = n + 2k$ encoder ancillas and $n_E = 2n + k$ decoder ancillas, so that both associated unitaries act on $2(n + k)$ qubits. Thus the conventional QAE and the ζ-QVAE represent the two extremes relevant to our analysis. The former is resource-efficient but nonuniversal, whereas the latter is universal but substantially wider than may be necessary if one is only interested in infidelity. This contrast motivates the main question of the present work, namely whether there exists a strictly smaller universal architecture for blind single-copy compression under the infidelity objective.

3 Main results

Overview. In this section, we establish four primary results. First, k encoder ancilla qubits and n decoder ancilla qubits always suffice to match the best possible CPTP encoder–decoder pair. Second, the encoder count k is unavoidable in the worst case. Third, the smallest physical decoder, an isometric decoder with $n-k$ ancilla qubits, is not always optimal. Fourth, despite this counterexample, isometric decoders are near-optimal when the source is concentrated near a 2^k-dimensional subspace.

For an n-qubit source distribution μ over pure states $|\psi\rangle \in H_{2^n}$ and integers $n > k \geq 1$, we evaluate an encoder $\mathcal{E} : L(H_{2^n}) \to L(H_{2^k})$ and a decoder $\mathcal{D} : L(H_{2^k}) \to L(H_{2^n})$ by the average fidelity

$$F_\mu(\mathcal{E}, \mathcal{D}) := \mathbb{E}_{|\psi\rangle\sim\mu}\big[\langle\psi|(\mathcal{D} \circ \mathcal{E})(|\psi\rangle\langle\psi|)|\psi\rangle\big]. \tag{1}$$

The corresponding average infidelity is $L_\mu(\mathcal{E}, \mathcal{D}) := 1 - F_\mu(\mathcal{E}, \mathcal{D})$.

Definition 3.1 ((n, k, n_B, n_E)-QAE)**.** *Fix $n > k \geq 1$. An (n, k, n_B, n_E)-QAE is specified by two unitaries (U, V) and acts as follows.*

(i) Prepare the input state $|\psi\rangle$ on an n-qubit register A and $|0\rangle^{\otimes n_B}$ on an n_B-qubit ancilla register B. Apply an $(n + n_B)$-qubit unitary U and reinterpret $AB \equiv CD$, where C has k qubits and D has $n + n_B - k$ qubits. The encoded state is

$$\rho_{\text{enc}} := \mathrm{Tr}_D\big(U(|\psi\rangle\langle\psi|_A \otimes |0^{n_B}\rangle\langle 0^{n_B}|_B)U^\dagger\big) \in L(H_C).$$

(ii) Prepare $|0\rangle^{\otimes n_E}$ on an n_E-qubit ancilla register E. Apply a $(k + n_E)$-qubit unitary V to CE and reinterpret $CE \equiv FG$, where F has n qubits and G has $k + n_E - n$ qubits. The reconstructed state is

$$\rho_{\text{rec}} := \mathrm{Tr}_G\big(V(\rho_{\text{enc}} \otimes |0^{n_E}\rangle\langle 0^{n_E}|_E)V^\dagger\big) \in L(H_F).$$

More generally, if the encoder ancilla B and decoder ancilla E are initialized to fixed density operators $\sigma_B \in L(H_B)$ and $\tau_E \in L(H_E)$, respectively, then the induced encoder and decoder channels are

$$\mathcal{E}_{U,\sigma_B}(\rho_A) := \mathrm{Tr}_D\big[U(\rho_A \otimes \sigma_B)U^\dagger\big], \qquad \mathcal{D}_{V,\tau_E}(\rho_C) := \mathrm{Tr}_G\big[V(\rho_C \otimes \tau_E)V^\dagger\big].$$

Lemma 3.2. *Fix $n > k \geq 1$, ancilla counts $n_B \geq 0$ and $n_E \geq n - k$, and a source distribution μ over pure n-qubit states. Then*

$$\max_{U,V,\sigma_B,\tau_E} F_\mu(\mathcal{E}_{U,\sigma_B}, \mathcal{D}_{V,\tau_E}) = \max_{U,V} F_\mu\left(\mathcal{E}_{U,|0^{n_B}\rangle\langle 0^{n_B}|}, \mathcal{D}_{V,|0^{n_E}\rangle\langle 0^{n_E}|}\right),$$

where on the left the maximum is over all unitaries U, V of the appropriate sizes and all density operators $\sigma_B \in L(H_B)$ and $\tau_E \in L(H_E)$, and on the right the maximum is over the ordinary (n, k, n_B, n_E)-QAEs of Definition 3.1.

Proof. See Appendix C.2. □

Lemma 3.2 implies that allowing arbitrary fixed mixed ancilla states does not change the optimal average fidelity. Thus, without loss of generality, we may fix the encoder and decoder ancilla states to all-zero states for the purpose of infidelity minimization. We maintain this convention throughout the remainder of this work.

Remark 3.3. *At first glance, [Ma et al., 2024] may appear to contradict Lemma 3.2, since they report an advantage from using mixed reference states. There is, however, no contradiction, because the optimization problem studied in [Ma et al., 2024] is different from the problem considered here. Lemma 3.2 applies to our separately linear objective $F_\mu(\mathcal{E}, \mathcal{D})$, whereas [Ma et al., 2024] optimizes a nonlinear cost function. Moreover, their protocol is not blind, because it uses input-dependent reference states, effectively employing state-dependent side information.*

Our first theorem identifies a universal resource-efficient architecture by showing that k ancillas on the encoder side and n ancillas on the decoder side always suffice.

Theorem 3.4 (Universal sufficiency of (k, n) ancillas)**.** *For every distribution μ of pure n-qubit states, there exists an (n, k, k, n)-QAE attaining $\max_{\mathcal{E},\mathcal{D}} F_\mu(\mathcal{E}, \mathcal{D})$, where the maximum ranges over all CPTP encoders $\mathcal{E} : L(H_{2^n}) \to L(H_{2^k})$ and decoders $\mathcal{D} : L(H_{2^k}) \to L(H_{2^n})$.*

Proof sketch. By a Choi-matrix reformulation, the fidelity functional is separately linear in $\mathcal{E}$ and $\mathcal{D}$. Hence an optimum can be chosen with both channels extreme. Extreme encoders $L(H_{2^n}) \to L(H_{2^k})$ have Choi rank at most 2^n, and extreme decoders $L(H_{2^k}) \to L(H_{2^n})$ have Choi rank at most 2^k. Padded Kraus lists of these lengths can be embedded into isometries

$$\widetilde{U} : H_{2^n} \to H_{2^k} \otimes H_{2^n}, \qquad \widetilde{V} : H_{2^k} \to H_{2^n} \otimes H_{2^k},$$

which extend to unitaries on $n + k$ qubits. The full proof is developed throughout Appendix C, concluding in Appendix C.4.

The universal sufficiency theorem is sharp on the encoder side. The source family below already forces k encoder ancillas.

Theorem 3.5 (Informal)**.** *For every choice of input size n and bottleneck size $k < n$, there is a source distribution of pure n-qubit states for which fewer than k encoder ancillas are insufficient to reach the best possible reconstruction fidelity. More precisely, for this source, no (n, k, n_B, n_E)-QAE with $n_B < k$ can attain the optimum value of $F_\mu(\mathcal{E}, \mathcal{D})$ over all CPTP encoders and decoders through the k-qubit bottleneck. Hence every globally optimal QAE for this source must satisfy $n_B \geq k$. In other words, the k encoder ancillas used in the universal construction are necessary in the worst case.*

Proof sketch. The proof uses a one-parameter source family $\{\mu_{1,\varepsilon}\}_{0<\varepsilon<1}$, whose construction is given in Appendix E. Intuitively, the states are all close to one fixed reference state, but their small differences point in many possible directions. A compression scheme that is optimal must therefore preserve the reference state perfectly while also handling these hidden directions as well as the k-qubit bottleneck allows. The optimal behavior has a simple operational form. The channel keeps the part of the input lying in a subspace of the same size as the bottleneck and sends the remaining part back to the reference state. In other words, the best channel for this source is a "reset" channel. The key point is that implementing this reset behavior through the bottleneck requires a sufficiently rich encoder. It can be shown that any architecture with $n_B < k$ cannot be optimal for this source. A detailed restatement and full proof are given in Appendix E.

Corollary 3.6 (Sharp universal encoder threshold)**.** *For blind single-copy compression under the infidelity objective, k encoder ancillas are universally sufficient and, in the worst case, necessary.*

For the source $\mu_{1,\varepsilon}$ with sufficiently small ε, as well as the sources discussed in Appendix D, the optimal fidelity is always achieved by an isometric decoder. Furthermore, as we will demonstrate in Section 4, isometric decoders also appear sufficient for the heuristic datasets constructed from real-world data (MNIST). In particular, there is no noticeable difference between the settings $(n_B, n_E) = (k, n-k)$ and $(n_B, n_E) = (k, n-k+1)$. Taken together, these observations might suggest that isometric decoders always suffice. A natural remaining question is whether the decoder ancilla count can *always* be reduced to $n-k$, equivalently, whether the decoder may always be restricted to an isometric channel. The answer is negative. In what follows, we establish this by constructing an explicit counterexample.

Proposition 3.7 (Decoder isometry is not sufficient)**.** *Let $n = 2$, $k = 1$,*

$$|\psi_\phi\rangle = \frac{|0\rangle + e^{i\phi}|1\rangle}{\sqrt{2}}, \qquad |\Psi_\phi\rangle = |\psi_\phi\rangle^{\otimes 2}, \qquad \phi \in [0, 2\pi), \tag{2}$$

and let μ_{ph} be the uniform distribution on ϕ. Then

$$\sup_{\mathcal{E},\, V^\dagger V = I_2} F_{\mu_{\mathrm{ph}}}(\mathcal{E}, \mathcal{D}_V) < \frac{3}{4} \leq \sup_{\mathcal{E},\mathcal{D}} F_{\mu_{\mathrm{ph}}}(\mathcal{E}, \mathcal{D}), \tag{3}$$

where $\mathcal{D}_V(\rho) = V\rho V^\dagger$.

Proof sketch. A concrete encoder–decoder pair with a decoder of Kraus rank 2 attains fidelity $3/4$. For an isometric decoder, averaging over ϕ produces a rank-2 projector optimization problem whose value is at most $3/4$. Equality would force a CPTP map sending the input family to a family of pure projected states. A Gram-matrix criterion shows that no such pure-state transformation exists. The full argument is given in Appendix F.

Because the gap is visually negligible, we employ two distinct methods for greater precision. The first uses a training set of 2000 states and a test set of 1000 states. The second is a more computationally efficient method that replaces the continuous integral over the phase distribution with an exact evaluation at a finite set of points, justified by the following observation.

Observation 3.8 (Informal)**.** *The average fidelity for the source μ_{ph} can be computed from only five phases.*

A detailed restatement and full proof are given in Appendix F.9.

In all cases, we adopt the exponential parameterization for the trainable unitaries, $U = \exp(iH)$, where H is a trainable Hermitian operator. The distinction between isometric and non-isometric decoders is mathematically clear but quite subtle, as illustrated in Figure 2.

Remark 3.9. *The negative result of Proposition 3.7 is genuinely source-dependent. Appendix D shows that for the Haar prior on pure states, and more generally for certain mixtures with perfectly compressible priors, a rank-one decoder is already optimal.*

The universal sufficiency theorem is sharp on the decoder side when $k = 1$. Specifically, we show that the phase-family counterexample of Proposition 3.7 extends from $(n, k) = (2, 1)$ to every pair $(n, 1)$ with $n \geq 2$.

For $n \geq 2$, let R denote an $(n-2)$-qubit register, let

$$|0_R\rangle := |0\rangle^{\otimes(n-2)}, \qquad |\Omega_\phi\rangle := |\Psi_\phi\rangle \otimes |0_R\rangle,$$

where $|\Psi_\phi\rangle$ is the two-qubit phase family from Eq. (2), and let $\mu_{\mathrm{ph}}^{(n)}$ be the distribution induced by the uniform choice of $\phi \in [0, 2\pi)$.

Theorem 3.10 (Sharp decoder threshold for $k = 1$)**.** *For every $n \geq 2$,*

$$\sup_{\mathcal{E},\, V^\dagger V = I_2} F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E}, \mathcal{D}_V) < \frac{3}{4} \leq \sup_{\mathcal{E},\mathcal{D}} F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E}, \mathcal{D}),$$

where $\mathcal{E} : L(H_{2^n}) \to L(H_2)$, $\mathcal{D} : L(H_2) \to L(H_{2^n})$, and $\mathcal{D}_V(\rho) = V\rho V^\dagger$. Consequently, every optimal $(n, 1, n_B, n_E)$-QAE for the source $\mu_{\mathrm{ph}}^{(n)}$ satisfies $n_E \geq n$.

Proof. See Appendix F.10. □

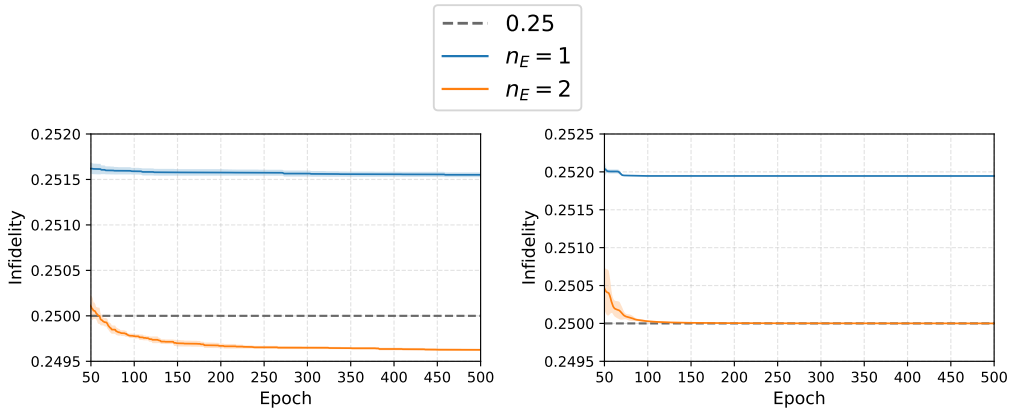

(a) (b)

Figure 2: Infidelity curves for an isometric decoder ($n_E = 1$) and a non-isometric decoder ($n_E = 2$) with $n_B = 1$ encoder ancilla qubit, under the source distribution in Eq. (2) (with 2-sigma confidence intervals). The dashed line represents the theoretical upper bound of $1/4$ for the minimum infidelity over all CPTP pairs $(\mathcal{E}, \mathcal{D})$, implied by Eq. (3). The displayed panels represent results obtained via (a) training with 2000 training states and 1000 test states, and (b) evaluation at discrete points as justified by Observation 3.8. While the converged losses are nearly indistinguishable on the original scale, the magnified view makes clear that the isometric decoder fails to reach the fidelity $3/4$ achieved by a non-isometric decoder.

Table 1: Ancilla requirements and status of several architectures for blind single-copy compression under the infidelity objective. The row labeled "reset channel" refers to the explicit optimal construction for the family $\mu_{1,\varepsilon}$ from Theorem 3.5. Decoder ancillas $n-k$ correspond to an isometric decoder, because the decoder has just enough qubits to expand the k-qubit latent register to n output qubits and has no leftover environment to discard.

Scheme	Encoder ancillas	Decoder ancillas	Unitary widths	Status (infidelity)
Conventional QAE	0	$n-k$	n, n	not universal
ζ-QVAE [Wang et al., 2025]	$n+2k$	$2n+k$	$2(n+k), 2(n+k)$	universal
Theorem 3.4	k	n	$n+k, n+k$	universal
Reset channel for $\mu_{1,\varepsilon}$	k	$n-k$	$n+k, n$	optimal for small ε
Decoder isometry for μ_{ph}	–	$n-k$	–, n	suboptimal

Corollary 3.11 (Sharp universal decoder threshold for $k = 1$)**.** *For blind single-copy compression through a one-qubit bottleneck under the infidelity objective, n decoder ancillas are universally sufficient and, in the worst case, necessary.*

Table 1 summarizes the ancilla requirements and circuit dimensions of the compression architectures considered in this work.

The next theorem shows that, although isometric decoders need not be optimal in general, they remain near-optimal whenever the average source state is well concentrated on its top 2^k eigenspaces.

Theorem 3.12 (Informal)**.** *For blind single-copy compression, restricting the decoder to be an isometry is near-optimal whenever the source is close to an m-dimensional subspace, where $m = 2^k$ is the latent dimension. More precisely, if η_m denotes the weight of the average source state outside its top m eigenspaces, then the best isometric decoder achieves at least a fraction $1 - \left(1 - \frac{1}{m}\right)\eta_m$ of the unrestricted optimum average fidelity.*

A detailed restatement and full proof are given in Appendix H.

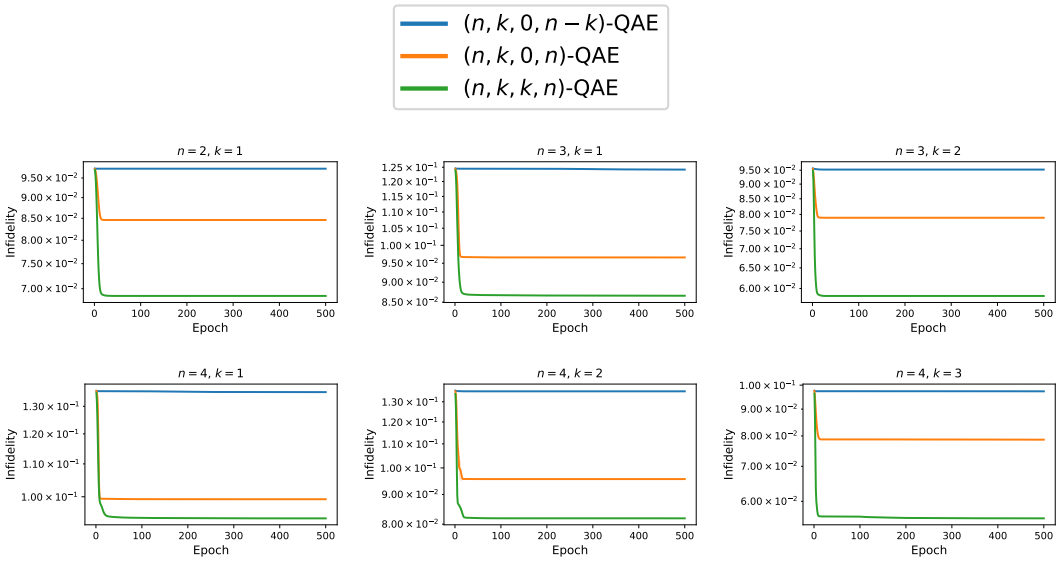

Figure 3: Test infidelity curves for source $\mu_{1,0.1}$, where $n \in \{2, 3, 4\}$ and $k \in \{1, \ldots, n-1\}$ (with 2-sigma confidence intervals). Across all settings, the performance differences among $(n, k, 0, n-k)$-QAE, $(n, k, 0, n)$-QAE, and (n, k, k, n)-QAE are clearly visible. The converged values for $(2, 1, 0, 1)$-QAE and $(2, 1, 1, 2)$-QAE are consistent with the first-order bounds derived in Appendix E.2.

4 Experiments

4.1 Setup

We perform experiments on $\mu_{1,0.1}$ to examine whether k encoder ancillas are necessary. The choice of $\mu_{1,0.1}$ is sufficient to illustrate Corollary 3.6. To this end, we compare (i) $(n, k, 0, n-k)$-QAE, (ii) $(n, k, 0, n)$-QAE, and (iii) (n, k, k, n)-QAE. If we observe a clear performance gain both from (i) to (ii) and from (ii) to (iii), this would suggest that when the number of encoder ancillas is smaller than k, even enlarging the decoder ancilla count to n (which is already sufficient for the decoder itself under the infidelity loss) remains insufficient to attain the optimum.

We also introduce μ_2, an empirical distribution induced by the MNIST dataset (version 3.0.1, CC BY-SA 3.0 license). More precisely, μ_2 is the distribution over quantum states obtained by encoding MNIST images into normalized state vectors in $\mathbb{C}^d$ via the data-preparation procedure described in Appendix I. This construction intentionally creates quantum states whose average source state has most of its mass in a low-dimensional subspace, matching the regime where Theorem 3.12 predicts that isometric decoders should be near-optimal.

For the experiments reported in Sections 4.2 and 4.3, we generated 2000 training samples and 1000 test samples. The encoder and decoder unitaries were trained for 500 epochs using the Adam optimizer with a learning rate of 10^{-3} and a batch size of 64. All computations were performed on a 32-core Intel Xeon Processor (Skylake, 2.5 GHz) and an NVIDIA H100 NVL GPU running CUDA 12.9, with 62 GB of memory. Training time information is provided in Appendix J. The experiments with uncertainty bands use ten independent training runs with random seeds $0, \ldots, 9$.

4.2 Numerical results for $\mu_{1,0.1}$

We compare three architectures using the $\mu_{1,0.1}$ dataset: (i) the $(n, k, 0, n-k)$-QAE (conventional), (ii) the $(n, k, 0, n)$-QAE, and (iii) the (n, k, k, n)-QAE. The results are shown in Figure 3. We observe clear performance gaps among (i), (ii), and (iii). This indicates that increasing the decoder ancilla count to n is insufficient to attain the optimum over all CPTP encoder–decoder pairs when the number of encoder ancillas is smaller than k.

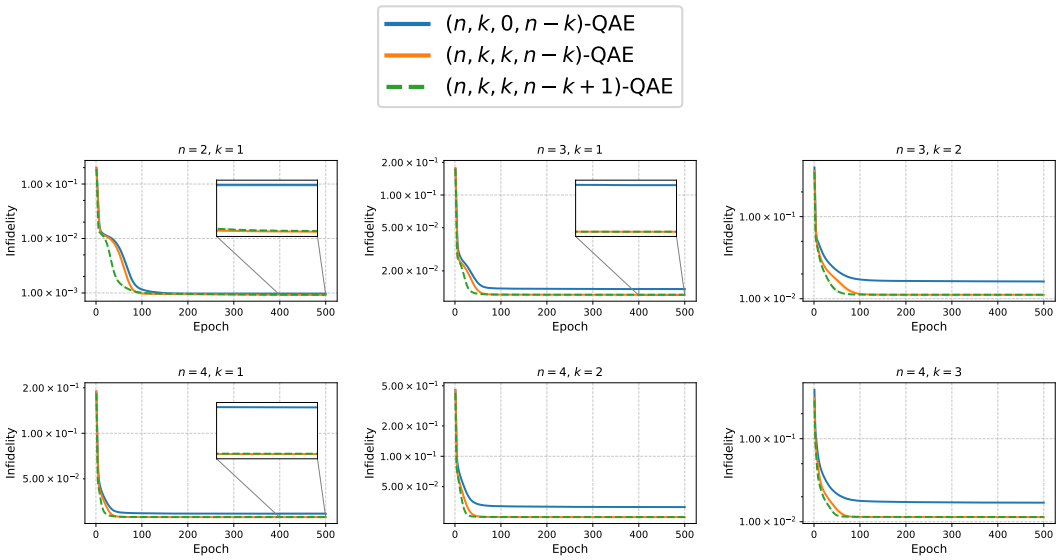

Figure 4: Test infidelity curves for MNIST source μ_2, where $n \in \{2, 3, 4\}$ and $k \in \{1, \ldots, n-1\}$ (with 2-sigma confidence intervals). Across all settings, the performance difference between $(n, k, k, n-k)$-QAE and $(n, k, k, n-k+1)$-QAE is negligible relative to the gap between $(n, k, k, n-k)$-QAE and $(n, k, 0, n-k)$-QAE.

4.3 Numerical results for μ_2

We compare three architectures using the μ_2 dataset: (i) the $(n, k, 0, n-k)$-QAE (conventional), (ii) the $(n, k, k, n-k)$-QAE, and (iii) the $(n, k, k, n-k+1)$-QAE. The results are shown in Figure 4. We observe a clear performance gap between (i) and (ii), whereas there is no noticeable difference between (ii) and (iii). This indicates that increasing the decoder ancilla from $n-k$ to $n-k+1$ provides no practical advantage, at least within the scope of our experiments. While this result may appear to contradict Proposition 3.7, we note again that the performance gap between $n_E = 1$ and $n_E = 2$ described in the counterexample is analytically valid but numerically negligible.

4.4 Additional distributions

Increased-tail MNIST. To test sensitivity to the concentration of the original MNIST source, we consider a larger-tail preparation setting. Its definition and numerical infidelities are given in Appendix K.1 and Table 3.

Coherently perturbed GHZ states. We consider pure states obtained by coherently perturbing the Greenberger–Horne–Zeilinger (GHZ) state [Greenberger et al., 1989]. The source definition and infidelities appear in Appendix K.2 and Table 4.

Output depolarization. To examine a departure from ideal reconstruction, we apply independent single-qubit depolarization to every reconstructed output qubit. Appendix K.3 specifies the channel and Table 5 reports infidelities.

Trace distance. To examine whether the empirical comparison depends on the distortion measure, we also report results based on trace distance in Appendix K.4, Table 6.

5 Discussion

Natural settings with $k \ll n$. A sufficient structural condition is that the source (almost) lies in a common subspace of dimension $r \ll 2^n$, which permits reliable compression to $\lceil \log_2 r \rceil$ qubits. Examples include:

- Identically prepared product states $|\phi\rangle^{\otimes n}$, which lie in an $(n+1)$-dimensional permutation-symmetric subspace [Plesch and Bužek, 2010; Rozema et al., 2014].
- A fixed-r_{ex} excitation sector of dimension $\binom{n}{r_{\mathrm{ex}}}$, requiring $k = \lceil \log_2 \binom{n}{r_{\mathrm{ex}}} \rceil = O(r_{\mathrm{ex}} \log n)$ for fixed excitation number.
- A code space for n physical qubits that carries q logical qubits, with $k = q \ll n$ when the code has large physical redundancy.
- A low-energy source supported on R relevant eigenstates, with $k = \lceil \log_2 R \rceil$.
- Nonorthogonal state families used in QAE-assisted group discrimination [Huang et al., 2020], when their joint span is small.

For a fixed source, k can be selected by a validation sweep with multiple initializations, choosing the smallest bottleneck meeting a target fidelity.

Limitations. This work restricts attention to blind single-copy compression of pure states, with performance measured solely by reconstruction infidelity. Accordingly, our results should not be interpreted as statements about other objectives, such as latent-space regularization, representation learning quality, robustness, or trainability. Such objectives may be valuable in practice, but they lie beyond the scope of the present theory and may modify the optimal resource counts identified here. The conclusions should also be distinguished from results in the literature that optimize different loss functions or assume different informational access. Furthermore, our empirical validation uses an artificially engineered quantum dataset, which is useful for probing the low-dimensional concentration regime but does not adequately represent the entanglement structures of physical quantum data, such as many-body ground states.

6 Conclusion

Our main results are resource-theoretic statements about blind single-copy compression under the infidelity objective. We identify an architecture that is wide enough to attain the optimum over all CPTP encoder–decoder pairs, yet narrower than fully general models. In particular, k encoder ancillas and n decoder ancillas are universally sufficient, and the encoder threshold is sharp. The decoder side remains less completely characterized. Specifically, the tight universal decoder ancilla lower bound as a function of n and k remains open. Because infidelity is the primary objective in many compression tasks, the ancilla counts identified here provide a natural baseline. A promising direction is to study how additional objectives, including explicit regularization and other desiderata for representation learning, alter the architecture or lead to genuinely different optimal tradeoffs.

Acknowledgments

This work is in part supported by the National Research Foundation of Korea (NRF, RS-2024-00451435 (20%), RS-2024-00413957 (20%)), Institute of Information & Communications Technology Planning & Evaluation (IITP, RS-2025-02305453 (15%), RS-2025-02273157 (15%), RS-2025-25442149 (15%), RS-2021-II211343 (15%)) grant funded by the Ministry of Science and ICT (MSIT), Institute of New Media and Communications (INMAC), and the BK21 FOUR program of the Education, Artificial Intelligence Graduate School Program (Seoul National University), and Research Program for Future ICT Pioneers, Seoul National University in 2026.

References

Howard Barnum. Quantum rate-distortion coding. *Physical Review A*, 62(4):042309, 2000.

Howard Barnum, Christopher A Fuchs, Richard Jozsa, and Benjamin Schumacher. General fidelity limit for quantum channels. *Physical Review A*, 54(6):4707, 1996.

Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.

Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S Zibrov, Manuel Endres, Markus Greiner, et al. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584, 2017.

Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermanni Heimonen, Jakob S Kottmann, Tim Menke, et al. Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1):015004, 2022.

Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.

Dmytro Bondarenko and Polina Feldmann. Quantum autoencoders to denoise quantum data. *Physical Review Letters*, 124(13):130502, 2020.

Carlos Bravo-Prieto. Quantum autoencoders with enhanced data encoding. *Machine Learning: Science and Technology*, 2(3):035028, 2021.

Chenfeng Cao and Xin Wang. Noise-assisted quantum autoencoder. *Physical Review Applied*, 15(5): 054012, 2021.

Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021.

Anthony Chefles, Richard Jozsa, and Andreas Winter. On the existence of physical transformations between sets of quantum states. *International Journal of Quantum Information*, 2(01):11–21, 2004.

Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Theoretical framework for quantum networks. *Physical Review A*, 80(2):022339, 2009.

Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975.

James W Cooley and John W Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19(90):297–301, 1965.

Nilanjana Datta, Min-Hsiu Hsieh, and Mark M Wilde. Quantum rate distortion, reverse Shannon theorems, and source-channel separation. *IEEE Transactions on Information Theory*, 59(1): 615–630, 2012.

Nilanjana Datta, Joseph M Renes, Renato Renner, and Mark M Wilde. One-shot lossy quantum data compression. *IEEE Transactions on Information Theory*, 59(12):8057–8076, 2013.

Christian L Degen, Friedemann Reinhard, and Paola Cappellaro. Quantum sensing. *Reviews of Modern Physics*, 89(3):035002, 2017.

Yongcheng Ding, Lucas Lamata, Mikel Sanz, Xi Chen, and Enrique Solano. Experimental implementation of a quantum autoencoder via quantum adders. *Advanced Quantum Technologies*, 2(7-8): 1800065, 2019.

W Dür, G Vidal, and JI Cirac. Visible compression of commuting mixed states. *Physical Review A*, 64(2):022308, 2001.

Richard P Feynman. Simulating physics with computers. *Feynman and Computation*, 22:133–153, 2018.

Shmuel Friedland and Raphael Loewy. On the extreme points of quantum channels. *Linear Algebra and its Applications*, 498:553–573, 2016.

Christopher A Fuchs and Jeroen Van De Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, 1999.

Daniel M Greenberger, Michael A Horne, and Anton Zeilinger. Going beyond Bell's theorem. In *Bell's theorem, quantum theory and conceptions of the universe*, pages 69–72. Springer, 1989.

Alfred Haar. Der Massbegriff in der Theorie der kontinuierlichen Gruppen. *Annals of Mathematics*, 34(1):147–169, 1933.

Shima Bab Hadiashar and Ashwin Nayak. On the Entanglement Cost of One-Shot Compression. *Quantum*, 4:286, 2020.

Masahito Hayashi and Keiji Matsumoto. Quantum universal variable-length source coding. *Physical Review A*, 66(2):022311, 2002.

Patrick Hayden, Richard Jozsa, and Andreas Winter. Trading quantum for classical resources in quantum data compression. *Journal of Mathematical Physics*, 43(9):4404–4444, 2002.

Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

Zoë Holmes, Kunal Sharma, Marco Cerezo, and Patrick J Coles. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum*, 3(1):010313, 2022.

Michał Horodecki. Limits for compression of quantum information carried by ensembles of mixed states. *Physical Review A*, 57(5):3364, 1998.

Chang-Jiang Huang, Hailan Ma, Qi Yin, Jun-Feng Tang, Daoyi Dong, Chunlin Chen, Guo-Yong Xiang, Chuan-Feng Li, and Guang-Can Guo. Realization of a quantum autoencoder for lossless compression of quantum data. *Physical Review A*, 102(3):032412, 2020.

Andrzej Jamiołkowski. Linear transformations which preserve trace and positive semidefiniteness of operators. *Reports on Mathematical Physics*, 3(4):275–278, 1972.

Min Jiang, Shunlong Luo, and Shuangshuang Fu. Channel-state duality. *Physical Review A*, 87(2): 022310, 2013.

Richard Jozsa and Benjamin Schumacher. A new proof of the quantum noiseless coding theorem. *Journal of Modern Optics*, 41(12):2343–2349, 1994.

Richard Jozsa, Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Universal quantum information compression. *Physical Review Letters*, 81(8):1714, 1998.

Zahra Baghali Khanian and Andreas Winter. General mixed-state quantum data compression with and without entanglement assistance. *IEEE Transactions on Information Theory*, 68(5):3130–3138, 2022.

Amir Khoshaman, Walter Vinci, Brandon Denis, Evgeny Andriyash, Hossein Sadeghi, and Mohammad H Amin. Quantum variational autoencoder. *Quantum Science and Technology*, 4(1):014001, 2019.

Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Masato Koashi and Nobuyuki Imoto. Compressibility of quantum mixed-state signals. *Physical Review Letters*, 87(1):017902, 2001a.

Masato Koashi and Nobuyuki Imoto. Teleportation cost and hybrid compression of quantum signals. *arXiv preprint quant-ph/0104001*, 2001b.

Karl Kraus. General state changes in quantum theory. *Annals of Physics*, 64(2):311–335, 1971.

Lucas Lamata, Unai Alvarez-Rodriguez, Jose David Martin-Guerrero, Mikel Sanz, and Enrique Solano. Quantum autoencoders via quantum adders with genetic algorithms. *Quantum Science and Technology*, 4(1):014007, 2019.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.

Hailan Ma, Chang-Jiang Huang, Chunlin Chen, Daoyi Dong, Yuanlong Wang, Re-Bing Wu, and Guo-Yong Xiang. On compression rate of quantum autoencoders: Control design, numerical and experimental realization. *Automatica*, 147:110659, 2023.

Hailan Ma, Gary J Mooney, Ian R Petersen, Lloyd CL Hollenberg, and Daoyi Dong. Quantum autoencoders using mixed reference states. *npj Quantum Information*, 10(1):86, 2024.

Dylan H Mahler, Lee A Rozema, Ardavan Darabi, Christopher Ferrie, Robin Blume-Kohout, and Aephraim M Steinberg. Adaptive quantum state tomography improves accuracy quadratically. *Physical Review Letters*, 111(18):183601, 2013.

Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.

Laleh Memarzadeh and Barry C Sanders. Group-covariant extreme and quasiextreme channels. *Physical Review Research*, 4(3):033206, 2022.

Yasuyoshi Mitsumori, John A Vaccaro, Stephen M Barnett, Erika Andersson, Atsushi Hasegawa, Masahiro Takeoka, and Masahide Sasaki. Experimental demonstration of quantum source coding. *Physical Review Letters*, 91(21):217902, 2003.

Ashwin Nayak and Pranab Sen. Invertible quantum operations and perfect encryption of quantum states. *Quantum Information & Computation*, 7(1):103–110, 2007. ISSN 1533-7146.

Vishal S Ngairangbam, Michael Spannowsky, and Michihisa Takeuchi. Anomaly detection in high-energy physics using a quantum autoencoder. *Physical Review D*, 105(9):095004, 2022.

Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2010.

Alex Pepper, Nora Tischler, and Geoff J Pryde. Experimental realization of a quantum autoencoder: The compression of qutrits via machine learning. *Physical Review Letters*, 122(6):060501, 2019.

Martin Plesch and Vladimír Bužek. Efficient compression of quantum information. *Physical Review A*, 81(3):032317, 2010.

John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.

Daniel Puzzuoli and John Watrous. Ancilla dimension in quantum channel discrimination. *Annales Henri Poincaré*, 18(4):1153–1184, 2017.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pages 1278–1286. PMLR, 2014.

Jonathan Romero, Jonathan P Olson, and Alan Aspuru-Guzik. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001, 2017.

Lee A Rozema, Dylan H Mahler, Alex Hayat, Peter S Turner, and Aephraim M Steinberg. Quantum data compression of a qubit ensemble. *Physical Review Letters*, 113(16):160504, 2014.

Mary Beth Ruskai. Some open problems in quantum information theory. *arXiv preprint arXiv:0708.1902*, 2007.

Sina Salek, Daniela Cadamuro, Philipp Kammerlander, and Karoline Wiesner. Quantum rate-distortion coding of relevant information. *IEEE Transactions on Information Theory*, 65(4): 2603–2613, 2018.

Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, 2015.

Benjamin Schumacher. Quantum coding. *Physical Review A*, 51(4):2738, 1995.

Claude Elwood Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948.

Sukin Sim, Peter D Johnson, and Alán Aspuru-Guzik. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12):1900070, 2019.

Maiyuren Srikumar, Charles D Hill, and Lloyd CL Hollenberg. Clustering and enhanced classification using a hybrid quantum autoencoder. *Quantum Science and Technology*, 7(1):015020, 2022.

W Forrest Stinespring. Positive functions on C*-algebras. *Proceedings of the American Mathematical Society*, 6(2):211–216, 1955.

Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine Learning*, pages 1096–1103, 2008.

Gaoyuan Wang, Jonathan Warrell, Prashant S Emani, and Mark Gerstein. Quantum variational autoencoder utilizing regularized mixed-state latent representations. *Physical Review A*, 111(4): 042416, 2025.

Mark M Wilde, Nilanjana Datta, Min-Hsiu Hsieh, and Andreas Winter. Quantum rate-distortion coding with auxiliary resources. *IEEE Transactions on Information Theory*, 59(10):6755–6773, 2013.

Andreas Winter. Compression of sources of probability distributions and density operators. *arXiv preprint quant-ph/0208131*, 2002.

Jun Wu, Hao Fu, Mingzheng Zhu, Haiyue Zhang, Wei Xie, and Xiang-Yang Li. Quantum circuit autoencoder. *Physical Review A*, 109(3):032623, 2024.

Shengjia Zhao, Jiaming Song, and Stefano Ermon. InfoVAE: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*, 2017.

A List of symbols and logical framework

Symbol	Meaning
Hilbert-space and channel notation	
$H_d \cong \mathbb{C}^d$	A d-dimensional Hilbert space.
$L(H)$	Linear operators on the Hilbert space H.
$H_A, H_B, \dots$	Hilbert spaces associated with registers $A, B, \dots$
$H_{AB} := H_A \otimes H_B$	Composite Hilbert space of registers A and B.
$AB \equiv CD$	The same physical system viewed under two different factorizations.
$\{\lvert x\rangle : x \in \{0,1\}^n\}$	Computational basis of an n-qubit register.
$\{\lvert j\rangle : 0 \le j \le 2^n - 1\}$	Computational basis of an n-qubit register.
$\lvert v\rangle^\perp$	Subspace orthogonal to the vector $\lvert v\rangle$.
Tr_B	Partial trace over subsystem B.
$\Phi : L(H_A) \to L(H_B)$	Generic quantum channel.
$\mathrm{Ad}_U(\rho) = U\rho U^\dagger$	Unitary channel induced by the unitary U.
$\{K_i\}$	Kraus operators of a channel.
$\mathrm{rank}_{\mathrm{Kraus}}(\Phi)$	Kraus rank (equivalently, Choi rank) of Φ.
$\Lambda(\Phi)$	Choi matrix of the linear map Φ.
$\Lambda_1 \star \Lambda_2$	Link product of two Choi matrices.
ρ^T, T_B	Transpose of ρ and partial transpose over subsystem B.
Compression setting and QAE architecture	
$d := 2^n$	Input/output Hilbert-space dimension.
$m := 2^k$	Latent-space dimension (bottleneck dimension).
$F_\mu(\mathcal{E}, \mathcal{D})$	Average reconstruction fidelity for source distribution μ.
$L_\mu(\mathcal{E}, \mathcal{D}) := 1 - F_\mu(\mathcal{E}, \mathcal{D})$	Average infidelity.
(n, k, n_B, n_E)-QAE	QAE architecture with n_B encoder ancillas and n_E decoder ancillas.
ρ_{enc}	Encoded state on the latent register C.
ρ_{rec}	Reconstructed state on the output register F.
Source families	
μ_{Haar}	Haar distribution on pure states in $\mathbb{C}^d$.
$H^\perp := \lvert 0\rangle^\perp \subset \mathbb{C}^d$	Orthogonal complement of the reference basis vector $\lvert 0\rangle$.
$\lvert\psi_\varepsilon(\varphi)\rangle$	Perturbed source state $\lvert\psi_\varepsilon(\varphi)\rangle = \sqrt{1-\varepsilon}\,\lvert 0\rangle + \sqrt{\varepsilon}\,\lvert\varphi\rangle$, $\lvert\varphi\rangle \in H^\perp$, $\lVert\varphi\rVert = 1$.
$\mu_{1,\varepsilon}$	Distribution induced by Haar-uniform $\lvert\varphi\rangle$ on the states $\lvert\psi_\varepsilon(\varphi)\rangle$.
$F_\varepsilon(E, D)$	Average fidelity for the source family $\mu_{1,\varepsilon}$.
$\Phi_P(\rho)$	Reset channel $\Phi_P(\rho) = P\rho P + \mathrm{Tr}\big[(I_d - P)\rho\big]\,\lvert 0\rangle\langle 0\rvert$.
Decoder isometry counterexample	
$\lvert\psi_\phi\rangle$	Single-qubit phase state $\lvert\psi_\phi\rangle = (\lvert 0\rangle + e^{i\phi}\lvert 1\rangle)/\sqrt{2}$.
$\lvert\Psi_\phi\rangle$	Two-qubit source state $\lvert\Psi_\phi\rangle = \lvert\psi_\phi\rangle^{\otimes 2}$.
μ_{ph}	Uniform distribution of ϕ on $[0, 2\pi)$.
$\mathcal{D}_V(\rho) = V\rho V^\dagger$	Isometric decoder generated by an isometry V.
$\lvert\Sigma\rangle$	Symmetric Bell state $\lvert\Sigma\rangle = (\lvert 01\rangle + \lvert 10\rangle)/\sqrt{2}$.
$P_{t,\gamma}$	Rank-2 projector $P_{t,\gamma} = \lvert\Sigma\rangle\langle\Sigma\rvert + \lvert\eta_{t,\gamma}\rangle\langle\eta_{t,\gamma}\rvert$.
$\lvert\eta_{t,\gamma}\rangle$	Vector $\lvert\eta_{t,\gamma}\rangle = \cos t\,\lvert 00\rangle + e^{i\gamma}\sin t\,\lvert 11\rangle$.
$\lvert\chi_\phi^{(t,\gamma)}\rangle$	Normalized projection of $\lvert\Psi_\phi\rangle$ onto $\mathrm{im}(P_{t,\gamma})$.
Source-dependent multiplicative guarantee	
$\lambda_i(X)$	Eigenvalues of a Hermitian operator X, listed in nonincreasing order.
$s_m := \sum_{i=1}^m \lambda_i(\bar\rho)$	Weight of the average source state on its top m eigenvalues.
$\eta_m := 1 - s_m$	Tail weight of the average source state outside its top m eigenvalues.
$p_P(\psi) := \langle\psi\rvert P\lvert\psi\rangle$	Probability mass of the source state $\lvert\psi\rangle$ inside $\mathrm{im}(P)$.
$\mathcal{E}_{P,\xi}$	Encoder in the isometric lower bound construction.
$\mathcal{D}_{V_P}$	Associated isometric decoder in the lower bound construction.

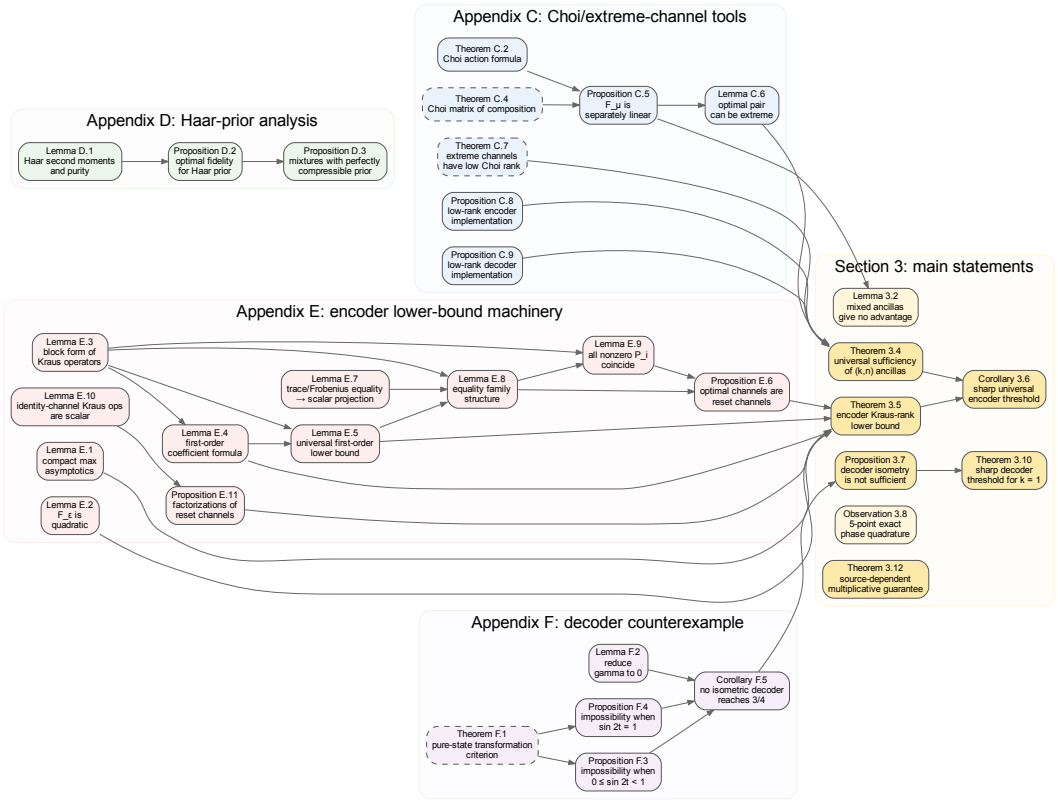

Figure 5: Logical dependency graph of the results. Dashed nodes are used as cited black boxes.

B Basic concepts

B.1 Hilbert spaces, qubits, registers, and Dirac notation

A finite-dimensional Hilbert space is a finite-dimensional complex vector space equipped with an inner product. Throughout, $H_d \cong \mathbb{C}^d$ denotes a d-dimensional Hilbert space and $L(H)$ denotes the linear operators on H. In Dirac notation, a vector is written as a ket $|\psi\rangle$, its conjugate transpose is the bra $\langle\psi| := |\psi\rangle^\dagger$, and the inner product of $|\phi\rangle$ and $|\psi\rangle$ is $\langle\phi|\psi\rangle$. A vector represents a physical pure state only when it has unit norm, $\langle\psi|\psi\rangle = 1$. For a vector $|v\rangle$, we write $|v\rangle^\perp$ for the subspace orthogonal to $|v\rangle$.

A qubit is a system with Hilbert space $H_2 \cong \mathbb{C}^2$ and computational basis $\{|0\rangle, |1\rangle\}$. An n-qubit register has Hilbert space

$$H_{2^n} \cong (\mathbb{C}^2)^{\otimes n},$$

so an n-qubit system has dimension 2^n. For quantum registers $A, B, \ldots$, we write $H_A, H_B, \ldots$ for the corresponding Hilbert spaces. Composite systems are modeled by tensor products, so $H_{AB} := H_A \otimes H_B$. The same physical space may be reinterpreted under a different factorization, written for example as $AB \equiv CD$.

Tensor products are used not only for Hilbert spaces but also for states and operators. Thus $|\psi\rangle_A \otimes |\phi\rangle_B$ is a joint pure state on AB, often abbreviated as $|\psi\rangle_A|\phi\rangle_B$, and $X_A \otimes Y_B$ is an operator on H_{AB}. The computational basis of an n-qubit system is $\{|x_1 \cdots x_n\rangle : x_i \in \{0,1\}\}$, equivalently $\{|j\rangle : 0 \le j \le 2^n - 1\}$ after identifying a binary string with its integer value.

B.2 Pure states, mixed states, density operators, and reduced states

A pure state may be described by a unit vector $|\psi\rangle$, but the object acted on by channels and partial traces is the corresponding density operator $|\psi\rangle\langle\psi| \in L(H)$. These are different mathematical

objects: $|\psi\rangle$ is a vector in H, whereas $|\psi\rangle\langle\psi|$ is an operator on H. In particular, $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ are different vectors but determine the same density operator.

A mixed state on H is a density operator $\rho \in L(H)$, meaning a positive semidefinite operator with unit trace, $\rho \geq 0$ and $\mathrm{Tr}(\rho) = 1$. Equivalently, a mixed state can be written as a convex combination of pure states,

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|, \qquad p_i \geq 0, \qquad \sum_i p_i = 1.$$

The same density operator may arise from many different ensembles $\{p_i, |\psi_i\rangle\}$.

For a bipartite state ρ_{AB}, the reduced state (or marginal) on subsystem A is defined by the partial trace over B:

$$\rho_A := \mathrm{Tr}_B(\rho_{AB}) = \sum_j (I_A \otimes \langle j|_B)\, \rho_{AB}\, (I_A \otimes |j\rangle_B),$$

where $\{|j\rangle_B\}$ is any orthonormal basis of H_B. The result is basis independent. Likewise $\rho_B := \mathrm{Tr}_A(\rho_{AB})$. On simple tensors one has

$$\mathrm{Tr}_B(X_A \otimes Y_B) = X_A \ \mathrm{Tr}(Y_B),$$

and linearity extends this to all operators. Operationally, taking the partial trace is the same as discarding the corresponding subsystem. Thus ρ_A is the state accessible to an observer who keeps only subsystem A.

B.3 Quantum channels, unitary evolution, CPTP maps, and channel representations

A superoperator is a linear map $\Phi : L(H_A) \to L(H_B)$. The physically allowed maps are quantum channels, namely CPTP maps. Complete positivity means that for every auxiliary Hilbert space H_R and every positive semidefinite operator $X \in L(H_A \otimes H_R)$,

$$(\Phi \otimes \mathrm{id}_R)(X) \geq 0.$$

This is the correct positivity condition because a system may be entangled with external reference systems. Trace preservation means

$$\mathrm{Tr}[\Phi(X)] = \mathrm{Tr}[X]$$

for every operator X.

Closed-system evolution is unitary. The unitary group on H_d is

$$\mathrm{U}(d) := \{U \in L(H_d) : U^\dagger U = UU^\dagger = I_d\}.$$

If $U \in \mathrm{U}(d)$, the corresponding unitary channel is conjugation by U,

$$\mathrm{Ad}_U(\rho) := U\rho U^\dagger.$$

General open-system evolution is modeled by CPTP maps rather than only by unitary channels.

Every CPTP map $\Phi : L(H_A) \to L(H_B)$ admits a Kraus representation

$$\Phi(\rho) = \sum_{i=1}^{r} K_i \rho K_i^\dagger, \qquad K_i : H_A \to H_B, \qquad \sum_{i=1}^{r} K_i^\dagger K_i = I_A.$$

The minimal number r of nonzero Kraus operators in such a representation is the Kraus rank of Φ.

Equivalently, Stinespring's theorem states that there exist an auxiliary Hilbert space H_E and an isometry $V : H_A \to H_B \otimes H_E$ such that

$$\Phi(\rho) = \mathrm{Tr}_E\big(V\rho V^\dagger\big).$$

After adjoining an ancilla in a fixed state and extending the isometry to a unitary on a larger space, this becomes the familiar unitary-plus-ancilla-plus-partial-trace realization

$$\Phi(\rho) = \mathrm{Tr}_E\big[U\big(\rho \otimes |0\rangle\langle 0|_E\big)U^\dagger\big].$$

Concretely, if we write Φ in Kraus form as

$$\Phi(\rho) = \sum_{i=0}^{r-1} K_i \rho K_i^\dagger,$$

choose an orthonormal family $\{|i\rangle_E\}_{i=0}^{r-1}$ in the auxiliary space H_E, and define

$$V = \sum_{i=0}^{r-1} K_i \otimes |i\rangle_E,$$

then $V : H_A \to H_B \otimes H_E$ is an isometry (because $\sum_i K_i^\dagger K_i = I_{H_A}$), and

$$\begin{aligned} V\rho V^\dagger &= \left(\sum_{i=0}^{r-1} K_i \otimes |i\rangle_E\right) \rho \left(\sum_{j=0}^{r-1} K_j^\dagger \otimes \langle j|_E\right) \\ &= \sum_{i,j=0}^{r-1} K_i \rho K_j^\dagger \otimes |i\rangle\langle j|_E. \end{aligned}$$

Taking the partial trace over H_E gives

$$\begin{aligned} \mathrm{Tr}_E\big(V\rho V^\dagger\big) &= \sum_{k=0}^{r-1} (I_{H_B} \otimes \langle k|_E)\Big(\sum_{i,j=0}^{r-1} K_i \rho K_j^\dagger \otimes |i\rangle\langle j|_E\Big)(I_{H_B} \otimes |k\rangle_E) \\ &= \sum_{i,j,k=0}^{r-1} K_i \rho K_j^\dagger \,\langle k|i\rangle\langle j|k\rangle \\ &= \sum_{i=0}^{r-1} K_i \rho K_i^\dagger \\ &= \Phi(\rho). \end{aligned}$$

B.4 Blind single-copy compression and Haar-random pure states

In blind compression, the encoder is not told which source state was supplied. It must apply the same encoding channel to every input state. In single-copy compression, only one copy of the source state is available. Thus blind single-copy compression from n to $k < n$ qubits asks for channels

$$\mathcal{E} : L(H_{2^n}) \to L(H_{2^k}), \qquad \mathcal{D} : L(H_{2^k}) \to L(H_{2^n}),$$

that reconstruct an unknown pure n-qubit state as well as possible after passing through a k-qubit bottleneck.

Haar measure on $\mathrm{U}(d)$ is the unique probability measure invariant under left and right multiplication by fixed unitaries. The induced measure on pure states is obtained by applying a Haar-random unitary to a fixed reference vector. For example, if $U \sim \mathrm{Haar}$ on $\mathrm{U}(d)$, then $U|0\rangle$ is Haar-random on the unit sphere of H_d. More generally, if $S \subseteq H$ is a subspace, a Haar-random pure state in S is one drawn from the unique unitarily invariant probability measure on the unit sphere of S.

C Choi calculus, extreme channels, and the proof of Theorem 3.4

C.1 Choi representation and link product

Definition C.1 (Choi representation)**.** *Let $\Phi : L(H_{d_1}) \to L(H_{d_2})$ be linear. Its Choi matrix is*

$$\Lambda(\Phi) := \sum_{i,j=0}^{d_1-1} |i\rangle\langle j| \otimes \Phi(|i\rangle\langle j|) \in L(H_{d_1} \otimes H_{d_2}). \tag{4}$$

The Choi rank (or Kraus rank) of a linear map Φ is the rank of its corresponding Choi matrix.

Theorem C.2. *Let $\Phi : L(H_A) \to L(H_B)$ be linear. Then for every $\rho \in L(H_A)$,*

$$\Phi(\rho) = \mathrm{Tr}_A\big(\Lambda(\Phi)(\rho^T \otimes I_B)\big), \tag{5}$$

where the transpose is taken in the computational basis.

Definition C.3 (Link product)**.** *For $\Lambda_1 \in L(H_B \otimes H_C)$ and $\Lambda_2 \in L(H_A \otimes H_B)$, define*

$$\Lambda_1 \star \Lambda_2 := \mathrm{Tr}_B\Big((I_A \otimes \Lambda_1)(\Lambda_2^{T_B} \otimes I_C)\Big), \tag{6}$$

where T_B denotes partial transpose over H_B.

Theorem C.4 ([Chiribella et al., 2009], Theorem 1)**.** *If $\Phi_1 : L(H_A) \to L(H_B)$ and $\Phi_2 : L(H_B) \to L(H_C)$ are linear maps, then*

$$\Lambda(\Phi_2 \circ \Phi_1) = \Lambda(\Phi_2) \star \Lambda(\Phi_1). \tag{7}$$

Proposition C.5. *For any source distribution μ, the functional $F_\mu(\mathcal{E}, \mathcal{D})$ from Eq.* (1) *extends to a separately linear continuous functional on the vector spaces of linear maps $\mathcal{E}$ and $\mathcal{D}$.*

Proof. For a pure state $|\psi\rangle$,

$$\begin{aligned}(\mathcal{D} \circ \mathcal{E})(|\psi\rangle\langle\psi|) &= \mathrm{Tr}_A\Big((\Lambda(\mathcal{D}) \star \Lambda(\mathcal{E}))(|\psi\rangle\langle\psi|^T \otimes I)\Big) \\ &= \mathrm{Tr}_A\Big(\mathrm{Tr}_C\big[(I_A \otimes \Lambda(\mathcal{D}))(\Lambda(\mathcal{E})^{T_C} \otimes I)\big](|\psi\rangle\langle\psi|^T \otimes I)\Big). \end{aligned} \tag{8}$$

The right-hand side is bilinear in $\Lambda(\mathcal{E})$ and $\Lambda(\mathcal{D})$, hence separately linear in $\mathcal{E}$ and $\mathcal{D}$. Averaging over μ preserves separate linearity and continuity. □

C.2 Proof of Lemma 3.2

We recall and prove Lemma 3.2.

Lemma 3.2. *Fix $n > k \geq 1$, ancilla counts $n_B \geq 0$ and $n_E \geq n - k$, and a source distribution μ over pure n-qubit states. Then*

$$\max_{U,V,\sigma_B,\tau_E} F_\mu(\mathcal{E}_{U,\sigma_B}, \mathcal{D}_{V,\tau_E}) = \max_{U,V} F_\mu\big(\mathcal{E}_{U,|0^{n_B}\rangle\langle 0^{n_B}|}, \mathcal{D}_{V,|0^{n_E}\rangle\langle 0^{n_E}|}\big)\,,$$

where on the left the maximum is over all unitaries U, V of the appropriate sizes and all density operators $\sigma_B \in L(H_B)$ and $\tau_E \in L(H_E)$, and on the right the maximum is over the ordinary (n, k, n_B, n_E)-QAEs of Definition 3.1.

Proof. Let

$$M_{\mathrm{mix}} := \max_{U,V,\sigma_B,\tau_E} F_\mu(\mathcal{E}_{U,\sigma_B}, \mathcal{D}_{V,\tau_E})\,, \qquad M_0 := \max_{U,V} F_\mu\big(\mathcal{E}_{U,|0^{n_B}\rangle\langle 0^{n_B}|}, \mathcal{D}_{V,|0^{n_E}\rangle\langle 0^{n_E}|}\big)\,.$$

Both maxima exist because the relevant unitary groups and state spaces are compact, and the objective is continuous in all parameters.

Since $|0^{n_B}\rangle\langle 0^{n_B}|$ and $|0^{n_E}\rangle\langle 0^{n_E}|$ are particular density operators, we immediately have

$$M_0 \leq M_{\mathrm{mix}}.$$

It remains to prove the reverse inequality.

Fix arbitrary unitaries U, V and arbitrary density operators σ_B, τ_E. Choose spectral decompositions

$$\sigma_B = \sum_r p_r |\beta_r\rangle\langle\beta_r|, \qquad \tau_E = \sum_s q_s |\eta_s\rangle\langle\eta_s|,$$

with $p_r, q_s \geq 0$ and $\sum_r p_r = \sum_s q_s = 1$. Define the corresponding pure-ancilla encoder and decoder channels by

$$\mathcal{E}_r := \mathcal{E}_{U,|\beta_r\rangle\langle\beta_r|}, \qquad \mathcal{D}_s := \mathcal{D}_{V,|\eta_s\rangle\langle\eta_s|}.$$

By linearity,

$$\mathcal{E}_{U,\sigma_B} = \sum_r p_r \mathcal{E}_r, \qquad \mathcal{D}_{V,\tau_E} = \sum_s q_s \mathcal{D}_s.$$

Applying separate linearity twice gives

$$F_\mu(\mathcal{E}_{U,\sigma_B}, \mathcal{D}_{V,\tau_E}) = \sum_{r,s} p_r q_s\, F_\mu(\mathcal{E}_r, \mathcal{D}_s).$$

Hence there exist indices r_*, s_* such that

$$F_\mu(\mathcal{E}_{U,\sigma_B}, \mathcal{D}_{V,\tau_E}) \le F_\mu(\mathcal{E}_{r_*}, \mathcal{D}_{s_*}).$$

Now choose unitaries W_B on H_B and W_E on H_E satisfying

$$W_B|0^{n_B}\rangle = |\beta_{r_*}\rangle, \qquad W_E|0^{n_E}\rangle = |\eta_{s_*}\rangle,$$

and define

$$\widetilde{U} := U(I_A \otimes W_B), \qquad \widetilde{V} := V(I_C \otimes W_E).$$

Then, for every input state ρ_A,

$$\begin{aligned}
\mathcal{E}_{\widetilde{U},|0^{n_B}\rangle\langle 0^{n_B}|}(\rho_A) &= \mathrm{Tr}_D\Big[\widetilde{U}(\rho_A \otimes |0^{n_B}\rangle\langle 0^{n_B}|)\widetilde{U}^\dagger\Big] \\
&= \mathrm{Tr}_D\Big[U(\rho_A \otimes W_B|0^{n_B}\rangle\langle 0^{n_B}|W_B^\dagger)U^\dagger\Big] \\
&= \mathrm{Tr}_D\big[U(\rho_A \otimes |\beta_{r_*}\rangle\langle\beta_{r_*}|)U^\dagger\big] \\
&= \mathcal{E}_{r_*}(\rho_A).
\end{aligned}$$

Likewise, for every input state ρ_C,

$$\mathcal{D}_{\widetilde{V},|0^{n_E}\rangle\langle 0^{n_E}|}(\rho_C) = \mathcal{D}_{s_*}(\rho_C).$$

Therefore

$$F_\mu(\mathcal{E}_{U,\sigma_B}, \mathcal{D}_{V,\tau_E}) \le F_\mu\Big(\mathcal{E}_{\widetilde{U},|0^{n_B}\rangle\langle 0^{n_B}|}, \mathcal{D}_{\widetilde{V},|0^{n_E}\rangle\langle 0^{n_E}|}\Big) \le M_0.$$

Since the mixed-ancilla implementation (U, V, σ_B, τ_E) was arbitrary, it follows that

$$M_{\mathrm{mix}} \le M_0.$$

Together with $M_0 \le M_{\mathrm{mix}}$, this proves $M_{\mathrm{mix}} = M_0$. □

C.3 Extreme points and low-rank implementations

Lemma C.6. *There exists an optimal pair $(\widehat{\mathcal{E}}, \widehat{\mathcal{D}})$ for Eq.* (1) *such that both channels are extreme points of their respective CPTP sets.*

Proof. The set of CPTP maps between fixed finite-dimensional spaces is convex and compact. By continuity, the maximum of F_μ is attained. Fix an optimal decoder $\mathcal{D}_0$. Because $\mathcal{E} \mapsto F_\mu(\mathcal{E}, \mathcal{D}_0)$ is continuous and linear by Proposition C.5, Bauer's maximum principle implies that it attains its maximum at an extreme encoder $\widehat{\mathcal{E}}$. Now fix this encoder. The same argument applied to $\mathcal{D} \mapsto F_\mu(\widehat{\mathcal{E}}, \mathcal{D})$ yields an extreme decoder $\widehat{\mathcal{D}}$, and $(\widehat{\mathcal{E}}, \widehat{\mathcal{D}})$ is still optimal. □

Theorem C.7 ([Ruskai, 2007], Theorem 1)**.** *Let $\Phi : L(H_{d_1}) \to L(H_{d_2})$ be an extreme CPTP channel. Then its Choi rank is at most d_1. More generally, the closure of the set of extreme CPTP channels equals the set of CPTP channels with Choi rank at most d_1 [Choi, 1975; Ruskai, 2007; Friedland and Loewy, 2016; Memarzadeh and Sanders, 2022].*

Proposition C.8 (Implementing a low-rank encoder)**.** *Let $\Phi : L(H_{2^n}) \to L(H_{2^k})$ be a CPTP channel of Kraus rank at most 2^n. Then there exists an $(n+k)$-qubit unitary U such that, under a relabeling $AB \equiv CD$ with $|A| = n$, $|B| = k$, $|C| = k$, and $|D| = n$,*

$$\Phi(\rho) = \mathrm{Tr}_D\Big(U(\rho_A \otimes |0^k\rangle\langle 0^k|_B)U^\dagger\Big). \tag{9}$$

Proof. Choose Kraus operators $K_0, \dots, K_{2^n-1} : H_{2^n} \to H_{2^k}$ for Φ, padding with zero operators if necessary. Define

$$\widetilde{U} := \sum_{i=0}^{2^n-1} K_i \otimes |i\rangle_D : H_{2^n} \to H_{2^k} \otimes H_{2^n}. \tag{10}$$

Then

$$\widetilde{U}^\dagger\widetilde{U} = \sum_{i=0}^{2^n-1} K_i^\dagger K_i = I_{2^n},$$

so $\widetilde{U}$ is an isometry. Identify the domain of $\widetilde{U}$ with the 2^n-dimensional subspace $H_A \otimes \mathrm{span}\{|0^k\rangle_B\}$ of $H_A \otimes H_B$ and extend $\widetilde{U}$ to a unitary U on $H_A \otimes H_B \cong H_C \otimes H_D$. Then

$$\mathrm{Tr}_D\big(\widetilde{U}\rho\widetilde{U}^\dagger\big) = \sum_{i=0}^{2^n-1} K_i \rho K_i^\dagger = \Phi(\rho),$$

which is Eq. (9). □

Proposition C.9 (Implementing a low-rank decoder)**.** *Let $\Phi : L(H_{2^k}) \to L(H_{2^n})$ be a CPTP channel of Kraus rank at most 2^k. Then there exists an $(n+k)$-qubit unitary V such that, under a relabeling $CE \equiv FG$ with $|C| = k$, $|E| = n$, $|F| = n$, and $|G| = k$,*

$$\Phi(\rho) = \mathrm{Tr}_G\Big(V(\rho_C \otimes |0^n\rangle\langle 0^n|_E)V^\dagger\Big). \tag{11}$$

Proof. Choose Kraus operators $L_0, \dots, L_{2^k-1} : H_{2^k} \to H_{2^n}$ for Φ, padding with zeros if necessary, and define

$$\widetilde{V} := \sum_{i=0}^{2^k-1} L_i \otimes |i\rangle_G : H_{2^k} \to H_{2^n} \otimes H_{2^k}. \tag{12}$$

Again $\widetilde{V}^\dagger\widetilde{V} = I_{2^k}$, so $\widetilde{V}$ is an isometry. Identify H_{2^k} with the subspace $H_C \otimes \mathrm{span}\{|0^n\rangle_E\}$ of $H_C \otimes H_E$ and extend $\widetilde{V}$ to a unitary V on $H_C \otimes H_E \cong H_F \otimes H_G$. Then

$$\mathrm{Tr}_G\big(\widetilde{V}\rho\widetilde{V}^\dagger\big) = \sum_{i=0}^{2^k-1} L_i \rho L_i^\dagger = \Phi(\rho),$$

which is Eq. (11). □

C.4 Proof of Theorem 3.4

We recall and prove Theorem 3.4.

Theorem 3.4 (Universal sufficiency of (k, n) ancillas)**.** *For every distribution μ of pure n-qubit states, there exists an (n, k, k, n)-QAE attaining $\max_{\mathcal{E},\mathcal{D}} F_\mu(\mathcal{E}, \mathcal{D})$, where the maximum ranges over all CPTP encoders $\mathcal{E} : L(H_{2^n}) \to L(H_{2^k})$ and decoders $\mathcal{D} : L(H_{2^k}) \to L(H_{2^n})$.*

Proof. Let $(\widehat{\mathcal{E}}, \widehat{\mathcal{D}})$ be an optimal pair as in Lemma C.6. By Theorem C.7,

$$\mathrm{rank}_{\mathrm{Kraus}}(\widehat{\mathcal{E}}) \le 2^n, \qquad \mathrm{rank}_{\mathrm{Kraus}}(\widehat{\mathcal{D}}) \le 2^k.$$

Therefore Propositions C.8 and C.9 realize $\widehat{\mathcal{E}}$ and $\widehat{\mathcal{D}}$ by the encoder and decoder parts of an (n, k, k, n)-QAE. The resulting QAE attains

$$F_\mu(\widehat{\mathcal{E}}, \widehat{\mathcal{D}}) = \max_{\mathcal{E},\mathcal{D}} F_\mu(\mathcal{E}, \mathcal{D}),$$

as claimed. □

D Haar analysis and mixtures with perfectly compressible priors

We write $d := 2^n$ and $m := 2^k$ for the source and bottleneck dimensions.

Lemma D.1. *Let $|\psi\rangle$ be Haar-random on $\mathbb{C}^d$.*

(i) For all operators $X, Y \in L(\mathbb{C}^d)$,

$$\mathbb{E}_{\psi\sim\mathrm{Haar}}\big[\langle\psi|X|\psi\rangle\,\langle\psi|Y|\psi\rangle\big] = \frac{\mathrm{Tr}(XY) + \mathrm{Tr}(X)\,\mathrm{Tr}(Y)}{d(d+1)}. \tag{13}$$

(ii) If $|\psi\rangle$ is Haar-random on $H_A \otimes H_B$ and $\rho_A := \mathrm{Tr}_B(|\psi\rangle\langle\psi|)$, then

$$\mathbb{E}_{\psi\sim\mathrm{Haar}}\,\mathrm{Tr}(\rho_A^2) = \frac{\dim H_A + \dim H_B}{1 + \dim H_A \dim H_B}. \tag{14}$$

Proposition D.2 (Optimal fidelity for the Haar prior)**.** *For the Haar prior on pure states in $\mathbb{C}^d$, one has*

$$\sup_{\mathcal{E},\mathcal{D}} F_{\mu_{\mathrm{Haar}}}(\mathcal{E},\mathcal{D}) = \frac{d+m^2}{d(d+1)}. \tag{15}$$

An optimal encoder–decoder pair is obtained by identifying $\mathbb{C}^d \cong \mathbb{C}^m \otimes \mathbb{C}^{d/m}$ and setting

$$\mathcal{E}(\rho) = \mathrm{Tr}_B(\rho), \qquad \mathcal{D}(\sigma) = \sigma \otimes \frac{I_B}{d/m}. \tag{16}$$

By Haar symmetry one may replace $I_B/(d/m)$ in Eq. (16) *by any pure state on H_B. In particular, a rank-one decoder is optimal.*

Proof. Let $\Phi := \mathcal{D} \circ \mathcal{E}$. Choose Kraus operators $\{E_j\}_j$ for $\mathcal{E}$ and $\{D_i\}_i$ for $\mathcal{D}$, and write $K_{ij} := D_i E_j$. For a pure input $|\psi\rangle$,

$$\langle\psi|\Phi(|\psi\rangle\langle\psi|)|\psi\rangle = \sum_{i,j} |\langle\psi|K_{ij}|\psi\rangle|^2 .$$

Applying Eq. (13) with $X = K_{ij}$ and $Y = K_{ij}^\dagger$ gives

$$\begin{aligned}\mathbb{E}_{\psi\sim\mathrm{Haar}}\big[\langle\psi|\Phi(|\psi\rangle\langle\psi|)|\psi\rangle\big] &= \frac{1}{d(d+1)} \sum_{i,j} \Big(\mathrm{Tr}(K_{ij}K_{ij}^\dagger) + |\mathrm{Tr}\, K_{ij}|^2 \Big) \\ &= \frac{1}{d+1} + \frac{1}{d(d+1)} \sum_{i,j} |\mathrm{Tr}(D_i E_j)|^2 . \end{aligned} \tag{17}$$

Now $E_j D_i$ is an $m \times m$ matrix, so

$$|\mathrm{Tr}(D_i E_j)|^2 = |\mathrm{Tr}(E_j D_i)|^2 \le m\ \mathrm{Tr}(D_i^\dagger E_j^\dagger E_j D_i).$$

Summing over i, j yields

$$\sum_{i,j} |\mathrm{Tr}(D_i E_j)|^2 \le m \sum_i \mathrm{Tr}\left(D_i^\dagger \Big(\sum_j E_j^\dagger E_j \Big) D_i \right) = m \sum_i \mathrm{Tr}(D_i^\dagger D_i) = m^2.$$

Substituting into Eq. (17) gives the upper bound in Eq. (15).

For achievability, identify $\mathbb{C}^d \cong \mathbb{C}^m \otimes \mathbb{C}^{d/m} = H_A \otimes H_B$, take $\mathcal{E}(\rho) = \mathrm{Tr}_B(\rho)$, and $\mathcal{D}(\sigma) = \sigma \otimes I_B/(d/m)$. If $\rho_A := \mathrm{Tr}_B(|\psi\rangle\langle\psi|)$, then

$$\langle\psi|\mathcal{D}(\mathcal{E}(|\psi\rangle\langle\psi|))|\psi\rangle = \mathrm{Tr}\left(|\psi\rangle\langle\psi| \Big(\rho_A \otimes \frac{I_B}{d/m} \Big) \right) = \frac{1}{d/m}\, \mathrm{Tr}(\rho_A^2).$$

Averaging and using Eq. (14) with $\dim H_A = m$ and $\dim H_B = d/m$ gives

$$\mathbb{E}_{\psi\sim\mathrm{Haar}}\big[\langle\psi|\mathcal{D}(\mathcal{E}(|\psi\rangle\langle\psi|))|\psi\rangle\big] = \frac{1}{d/m} \cdot \frac{m + d/m}{d+1} = \frac{d+m^2}{d(d+1)}.$$

This proves Eq. (15). Finally, because the Haar measure is invariant under unitary rotations on H_B, replacing $I_B/(d/m)$ by any pure state on H_B leaves the average fidelity unchanged. □

Proposition D.3 (Mixtures with a perfectly compressible prior)**.** *Let $\mu' = \mu_{\mathrm{pc}}$ be a prior on pure states in $\mathbb{C}^d$ such that there exist a unitary $U \in \mathrm{U}(d)$, a fixed reference state $|\varphi_{\mathrm{ref}}\rangle$, and an m-dimensional subsystem C with*

$$U|\psi\rangle_A = |\psi'\rangle_C \otimes |\varphi_{\mathrm{ref}}\rangle_D \qquad \text{for all } |\psi\rangle \in \mathrm{supp}(\mu').$$

Let $\mu = c\,\mu_{\mathrm{Haar}} + (1-c)\mu'$ with $0 < c < 1$. Then the maps

$$\mathcal{E}(\rho) := \mathrm{Tr}_D(U\rho U^\dagger), \qquad \mathcal{D}(\sigma) := U^\dagger\big(\sigma \otimes |\varphi_{\mathrm{ref}}\rangle\langle\varphi_{\mathrm{ref}}|\big)U$$

simultaneously maximize the fidelity for the Haar part and reconstruct every state in $\mathrm{supp}(\mu')$ *perfectly. Consequently,*

$$F_\mu(\mathcal{E},\mathcal{D}) = 1 - c\,\frac{(d+m)(d-m)}{d(d+1)}. \tag{18}$$

Proof. The chosen pair reconstructs μ' perfectly by construction, so its contribution to the average fidelity is $1-c$. On the Haar component the same pair is of the form considered in Proposition D.2, hence contributes $c(d+m^2)/(d(d+1))$. Therefore

$$F_\mu(\mathcal{E},\mathcal{D}) = (1-c) + c\,\frac{d+m^2}{d(d+1)} = 1 - c\,\frac{d^2-m^2}{d(d+1)},$$

which is Eq. (18). □

E The source $\mu_{1,\varepsilon}$ and the encoder ancilla lower bound

Throughout this appendix we set

$$d := 2^n, \qquad m := 2^k, \qquad q := d-1, \qquad H_\perp := |0\rangle^\perp \subset \mathbb{C}^d, \qquad \dim H_\perp = q.$$

For $\varepsilon \in (0,1)$ and a unit vector $|\varphi\rangle \in H_\perp$, write

$$|\psi_\varepsilon(\varphi)\rangle := \sqrt{1-\varepsilon}|0\rangle + \sqrt{\varepsilon}|\varphi\rangle,$$

and define

$$F_\varepsilon(\mathcal{E},\mathcal{D}) := \mathbb{E}_\varphi\big[\langle\psi_\varepsilon(\varphi)|(\mathcal{D}\circ\mathcal{E})(|\psi_\varepsilon(\varphi)\rangle\langle\psi_\varepsilon(\varphi)|)|\psi_\varepsilon(\varphi)\rangle\big], \tag{19}$$

where $|\varphi\rangle$ is Haar-uniform on the unit sphere of $H_\perp$.

E.1 Two preliminary lemmas

Lemma E.1. *Let $K \neq \varnothing$ be compact, and let $A, B, C : K \to \mathbb{R}$ be bounded and continuous. Define*

$$f(\varepsilon) := \max_{x\in K}\big(A(x) + \varepsilon B(x) + \varepsilon^2 C(x)\big), \qquad \varepsilon \geq 0.$$

If

$$A_{\max} := \max_{x\in K} A(x), \qquad S := \{x \in K : A(x) = A_{\max}\},$$

then

$$\limsup_{\varepsilon\to 0^+} \frac{f(\varepsilon) - A_{\max}}{\varepsilon} \leq \sup_{x\in S} B(x). \tag{20}$$

Proof. For $\varepsilon > 0$, choose $x_\varepsilon \in K$ such that

$$f(\varepsilon) = A(x_\varepsilon) + \varepsilon B(x_\varepsilon) + \varepsilon^2 C(x_\varepsilon).$$

Since $A(x_\varepsilon) \leq A_{\max}$,

$$\frac{f(\varepsilon) - A_{\max}}{\varepsilon} \leq B(x_\varepsilon) + \varepsilon C(x_\varepsilon). \tag{21}$$

Choose a sequence $\varepsilon_r \to 0^+$ along which the left-hand side tends to the limsup. By compactness of K, some subsequence of x_{ε_r} converges to a point $x_* \in K$. Passing to that subsequence and using continuity of A, B, C gives

$$A(x_*) = A_{\max}, \qquad x_* \in S.$$

Taking the limit in Eq. (21) along the same subsequence yields

$$\limsup_{\varepsilon\to 0^+} \frac{f(\varepsilon) - A_{\max}}{\varepsilon} \leq B(x_*) \leq \sup_{x\in S} B(x),$$

which proves Eq. (20). □

Lemma E.2. *For every fixed CPTP map $\Phi : L(\mathbb{C}^d) \to L(\mathbb{C}^d)$, the function*

$$\varepsilon \longmapsto F_\varepsilon(\Phi) := \mathbb{E}_\varphi\big[\langle\psi_\varepsilon(\varphi)|\Phi(|\psi_\varepsilon(\varphi)\rangle\langle\psi_\varepsilon(\varphi)|)|\psi_\varepsilon(\varphi)\rangle\big]$$

is a polynomial of degree at most 2 in ε. More precisely, there exist real numbers $A(\Phi), B(\Phi), C(\Phi)$ such that

$$F_\varepsilon(\Phi) = A(\Phi) + \varepsilon B(\Phi) + \varepsilon^2 C(\Phi) \qquad \text{for all } \varepsilon \in [0,1],$$

and

$$A(\Phi) = \langle 0|\Phi(|0\rangle\langle 0|)|0\rangle.$$

Proof. Write

$$A_0 := |0\rangle\langle 0|, \qquad A_1 := |\varphi\rangle\langle\varphi|, \qquad C_+ := |0\rangle\langle\varphi|, \qquad C_- := |\varphi\rangle\langle 0|.$$

Then

$$|\psi_\varepsilon(\varphi)\rangle\langle\psi_\varepsilon(\varphi)| = (1-\varepsilon)A_0 + \varepsilon A_1 + \sqrt{\varepsilon(1-\varepsilon)}\,(C_+ + C_-).$$

Because $(X, Y) \mapsto \mathrm{Tr}(\Phi(X)Y)$ is bilinear, the quantity

$$\mathrm{Tr}\big(\Phi(|\psi_\varepsilon(\varphi)\rangle\langle\psi_\varepsilon(\varphi)|)|\psi_\varepsilon(\varphi)\rangle\langle\psi_\varepsilon(\varphi)|\big)$$

is a finite linear combination of monomials in $(1-\varepsilon)$, ε, and $\sqrt{\varepsilon(1-\varepsilon)}$. Any term containing unequal numbers of $|\varphi\rangle$ and $\langle\varphi|$ picks up a nontrivial phase under $|\varphi\rangle \mapsto e^{i\theta}|\varphi\rangle$ and hence has zero Haar average. The surviving terms have equal numbers of $|\varphi\rangle$ and $\langle\varphi|$, so they contribute only powers $\varepsilon^0, \varepsilon^1, \varepsilon^2$. At $\varepsilon = 0$ the state is $|0\rangle$, giving the formula for $A(\Phi)$. □

E.2 Asymptotic separation for the conventional architecture when $(n, k) = (2, 1)$

We first analyze the conventional architecture in the minimal nontrivial case. Here

$$|\psi\rangle = \sqrt{1-\varepsilon}|00\rangle + \sqrt{\varepsilon}|\varphi\rangle, \qquad |\varphi\rangle \in \mathrm{span}\{|01\rangle, |10\rangle, |11\rangle\}, \qquad \|\varphi\| = 1,$$

and $|\varphi\rangle$ is Haar-random in the three-dimensional orthogonal complement of $|00\rangle$.

Let $U, V \in \mathrm{U}(4)$ be the encoder and decoder unitaries of a conventional $(2, 1, 0, 1)$-QAE. Define the decoder isometry

$$W := V(I \otimes |0\rangle) : H_2 \to H_4, \qquad W^\dagger W = I_2,$$

and for $b \in \{0, 1\}$ set

$$\widetilde{K}_b := (I \otimes \langle b|)U : H_4 \to H_2, \qquad K_b := W\widetilde{K}_b : H_4 \to H_4.$$

Then

$$\Phi_{U,V}(\rho) = \sum_{b=0}^{1} K_b \rho K_b^\dagger. \tag{22}$$

Because U is unitary,

$$\widetilde{K}_b\widetilde{K}_{b'}^\dagger = (I \otimes \langle b|)UU^\dagger(I \otimes |b'\rangle) = \delta_{bb'} I_2,$$

so

$$K_b K_{b'}^\dagger = \delta_{bb'} P, \qquad P := WW^\dagger. \tag{23}$$

Thus P is the rank-2 projector onto $\mathrm{ran}(W)$. If we set

$$Q_b := K_b^\dagger K_b,$$

then each Q_b is a rank-2 projection, $Q_0 Q_1 = 0$, and trace preservation implies

$$Q_0 + Q_1 = I_4. \tag{24}$$

For a pure input $|\psi\rangle$,

$$\langle\psi|\Phi_{U,V}(|\psi\rangle\langle\psi|)|\psi\rangle = \sum_{b=0}^{1} |\langle\psi|K_b|\psi\rangle|^2. \tag{25}$$

For each fixed (U, V), define

$$f_{U,V}(\varepsilon) := \mathbb{E}_\varphi\big[\langle\psi|\Phi_{U,V}(|\psi\rangle\langle\psi|)|\psi\rangle\big].$$

The Kraus pair (K_0, K_1) may be mixed by an arbitrary unitary $S \in \mathrm{U}(2)$ without changing the channel, and under this mixing the vector

$$v(\psi) := (\langle\psi|K_0|\psi\rangle, \langle\psi|K_1|\psi\rangle) \in \mathbb{C}^2$$

transforms as $v(\psi) \mapsto Sv(\psi)$, leaving Eq. (25) unchanged. We may therefore choose the Kraus gauge so that

$$\langle 00|K_0|00\rangle = \sqrt{f_0} \geq 0, \qquad \langle 00|K_1|00\rangle = 0, \qquad f_0 := \sum_{b=0}^{1} |\langle 00|K_b|00\rangle|^2 \in [0, 1]. \tag{26}$$

If $f_0 = 0$, then the constant term of $f_{U,V}(\varepsilon)$ is 0, so the pair is clearly suboptimal. Henceforth assume $f_0 > 0$.

For $b \in \{0, 1\}$ define

$$f_b^{(1)}(\varphi) := \langle\psi|K_b|\psi\rangle.$$

Expanding,

$$f_b^{(1)}(\varphi) = (1-\varepsilon)\langle 00|K_b|00\rangle + \sqrt{\varepsilon(1-\varepsilon)}\big(\langle 00|K_b|\varphi\rangle + \langle\varphi|K_b|00\rangle\big) + \varepsilon\langle\varphi|K_b|\varphi\rangle.$$

Since $\sqrt{\varepsilon(1-\varepsilon)} = \sqrt{\varepsilon} + o(\varepsilon)$, we may write

$$f_b^{(1)}(\varphi) = f_b^{(2)} + \sqrt{\varepsilon}\, f_b^{(3)}(\varphi) + \varepsilon f_b^{(4)}(\varphi) + o(\varepsilon),$$

where

$$f_b^{(2)} := \langle 00|K_b|00\rangle, \qquad f_b^{(3)}(\varphi) := \langle 00|K_b|\varphi\rangle + \langle\varphi|K_b|00\rangle, \qquad f_b^{(4)}(\varphi) := \langle\varphi|K_b|\varphi\rangle - f_b^{(2)}.$$

Hence

$$\begin{aligned}\left|f_b^{(1)}(\varphi)\right|^2 = \left|f_b^{(2)}\right|^2 &+ 2\sqrt{\varepsilon}\,\Re\big((f_b^{(2)})^* f_b^{(3)}(\varphi)\big) \\ &+ \varepsilon\Big(\left|f_b^{(3)}(\varphi)\right|^2 + 2\Re\big((f_b^{(2)})^* f_b^{(4)}(\varphi)\big)\Big) + o(\varepsilon).\end{aligned}$$

Taking the Haar expectation removes the $\sqrt{\varepsilon}$ term because $\mathbb{E}[|\varphi\rangle] = 0$:

$$\mathbb{E}\left|f_b^{(1)}(\varphi)\right|^2 = \left|f_b^{(2)}\right|^2 + \varepsilon\left(\mathbb{E}\left|f_b^{(3)}(\varphi)\right|^2 + 2\Re\big((f_b^{(2)})^*\mathbb{E}[f_b^{(4)}(\varphi)]\big)\right) + o(\varepsilon).$$

Summing over b and using Eq. (26),

$$f_{U,V}(\varepsilon) = f_0 + \varepsilon\sum_{b=0}^{1}\left(\mathbb{E}\left|f_b^{(3)}(\varphi)\right|^2 + 2\Re\big((f_b^{(2)})^*\mathbb{E}[f_b^{(4)}(\varphi)]\big)\right) + o(\varepsilon). \tag{27}$$

Let

$$\Pi := I_4 - |00\rangle\langle 00|,$$

the projector onto $\mathrm{span}\{|01\rangle, |10\rangle, |11\rangle\}$. For Haar-random $|\varphi\rangle$ in this three-dimensional subspace,

$$\mathbb{E}[|\varphi\rangle\langle\varphi|] = \frac{\Pi}{3}, \qquad \mathbb{E}[\langle\varphi|A|\varphi\rangle] = \frac{\mathrm{Tr}(\Pi A)}{3}. \tag{28}$$

Moreover,

$$\mathbb{E}\big[\langle u\,|\,\varphi\rangle\langle v\,|\,\varphi\rangle\big] = 0 \qquad \text{for all } u, v,$$

by invariance under the global phase $|\varphi\rangle \mapsto e^{i\theta}|\varphi\rangle$. Therefore

$$\mathbb{E}\left|f_b^{(3)}(\varphi)\right|^2 = \frac{\left\|\Pi K_b^\dagger|00\rangle\right\|^2 + \|\Pi K_b|00\rangle\|^2}{3}. \tag{29}$$

Set

$$p_0 := \langle 00|P|00\rangle, \qquad q_b := \langle 00|Q_b|00\rangle, \qquad q_0 + q_1 = 1.$$

Then by Eqs. (23) and (24),

$$\left\|K_b^\dagger|00\rangle\right\|^2 = \langle 00|P|00\rangle = p_0, \qquad \|K_b|00\rangle\|^2 = \langle 00|Q_b|00\rangle = q_b.$$

Since $\|\Pi v\|^2 = \|v\|^2 - |\langle 00\,|\,v\rangle|^2$,

$$\left\|\Pi K_b^\dagger|00\rangle\right\|^2 = p_0 - \left|f_b^{(2)}\right|^2, \qquad \|\Pi K_b|00\rangle\|^2 = q_b - \left|f_b^{(2)}\right|^2.$$

Substituting into Eq. (29) and summing over b gives

$$\sum_{b=0}^{1}\mathbb{E}\left|f_b^{(3)}(\varphi)\right|^2 = \frac{2p_0 + 1 - 2f_0}{3}. \tag{30}$$

Next define

$$\tau_b := \mathrm{Tr}(\Pi K_b \Pi).$$

By Eq. (28),

$$\mathbb{E}[f_b^{(4)}(\varphi)] = \frac{\tau_b}{3} - f_b^{(2)}.$$

Hence

$$\begin{aligned}\sum_{b=0}^{1} 2\Re\big((f_b^{(2)})^* \mathbb{E}[f_b^{(4)}(\varphi)]\big) &= \frac{2}{3}\Re\left(\sum_{b=0}^{1}(f_b^{(2)})^*\tau_b\right) - 2f_0 \\ &= \frac{2}{3}\sqrt{f_0}\,\Re(\tau_0) - 2f_0, \end{aligned} \tag{31}$$

where the last equality uses the gauge choice in Eq. (26).

The set $\mathrm{U}(4) \times \mathrm{U}(4)$ is compact, and by Lemmas E.1 and E.2 the constant term of $f_{U,V}(\varepsilon)$ is maximized at pairs with $f_0 = 1$ (for instance $U = V = I_4$). Therefore

$$\limsup_{\varepsilon\to 0^+} \frac{\max_{U,V} f_{U,V}(\varepsilon) - 1}{\varepsilon} \le \sup_{(U,V):f_0=1} \left(\frac{2p_0 + 1 - 2f_0}{3} + \frac{2}{3}\sqrt{f_0}\,\Re(\tau_0) - 2f_0\right). \tag{32}$$

If $f_0 = 1$, then $K_0|00\rangle = |00\rangle$ and $K_1|00\rangle = 0$. Because K_0 has rank 2, there exists a unit vector $|u\rangle \in \mathrm{span}\{|01\rangle, |10\rangle, |11\rangle\}$ such that

$$\mathrm{ran}(K_0) = \mathrm{span}\{|00\rangle, |u\rangle\}.$$

Consequently $\Pi K_0 \Pi$ has rank at most 1. Also $K_0^\dagger K_0 = Q_0 \le I_4$, so $\|K_0\|_\infty \le 1$, and hence

$$\|\Pi K_0 \Pi\|_\infty \le 1, \qquad \|\Pi K_0 \Pi\|_{\mathrm{F}} \le 1.$$

Therefore

$$|\tau_0| = |\mathrm{Tr}(\Pi K_0 \Pi)| \le \sqrt{\mathrm{rank}(\Pi K_0 \Pi)}\ \|\Pi K_0 \Pi\|_{\mathrm{F}} \le 1. \tag{33}$$

Combining $p_0 \le 1$ with Eq. (33), the expression on the right-hand side of Eq. (32) is at most

$$\frac{2p_0 + 2\Re(\tau_0) - 7}{3} \le -1.$$

Thus

$$\liminf_{\varepsilon\to 0^+} \frac{1 - \max_{U,V} f_{U,V}(\varepsilon)}{\varepsilon} \ge 1. \tag{34}$$

The identity choice $U = V = I_4$ attains equality in Eq. (34).

For unrestricted CPTP encoder–decoder pairs, consider the projector

$$\Pi' := |00\rangle\langle 00| + |01\rangle\langle 01|$$

and the channel

$$\Phi(\rho) = \Pi'\rho\Pi' + \mathrm{Tr}\big((I - \Pi')\rho\big)|00\rangle\langle 00|.$$

It admits an encoder with Kraus operators

$$E_0 = |0\rangle\langle 00| + |1\rangle\langle 01|, \qquad E_1 = |0\rangle\langle 10|, \qquad E_2 = |0\rangle\langle 11|,$$

and a decoder with a single Kraus operator

$$D_0 = |00\rangle\langle 0| + |01\rangle\langle 1|.$$

Write

$$|\varphi\rangle = \varphi_1|01\rangle + \sqrt{1 - |\varphi_1|^2}\,|\chi\rangle, \qquad |\chi\rangle \in \mathrm{span}\{|10\rangle, |11\rangle\}, \qquad \|\chi\| = 1.$$

Then

$$\Pi'|\psi\rangle = \sqrt{1-\varepsilon}|00\rangle + \sqrt{\varepsilon}\,\varphi_1|01\rangle, \qquad (I - \Pi')|\psi\rangle = \sqrt{\varepsilon(1 - |\varphi_1|^2)}\,|\chi\rangle,$$

and therefore

$$\begin{aligned}\langle\psi|\Phi(|\psi\rangle\langle\psi|)|\psi\rangle &= \|\Pi'|\psi\rangle\|^4 + (1-\varepsilon)\,\|(I - \Pi')|\psi\rangle\|^2 \\ &= 1 - \varepsilon\big(1 - |\varphi_1|^2\big) + o(\varepsilon).\end{aligned}$$

Since $\mathbb{E}\,|\varphi_1|^2 = 1/3$,

$$\mathbb{E}\big[\langle\psi|\Phi(|\psi\rangle\langle\psi|)|\psi\rangle\big] = 1 - \frac{2}{3}\varepsilon + o(\varepsilon), \tag{35}$$

which is strictly better than the conventional-architecture coefficient in Eq. (34).

E.3 A universal first-order lower bound for all factorized channels

In the remainder of this appendix, let $\Phi = \mathcal{D} \circ \mathcal{E}$ be an arbitrary factorized channel with

$$\mathcal{E} : L(\mathbb{C}^d) \to L(\mathbb{C}^m), \qquad \mathcal{D} : L(\mathbb{C}^m) \to L(\mathbb{C}^d), \qquad \Phi(|0\rangle\langle 0|) = |0\rangle\langle 0|.$$

Choose Kraus operators $\{E_\ell\}_\ell$ for $\mathcal{E}$ and $\{D_j\}_j$ for $\mathcal{D}$, and relabel the Kraus family $\{D_j E_\ell\}_{j,\ell}$ of Φ as $\{A_i\}_{i=1}^r$. Each A_i has rank at most m.

Lemma E.3 (Block form)**.** *For every i there exists $\alpha_i \in \mathbb{C}$ such that*

$$A_i|0\rangle = \alpha_i|0\rangle.$$

With respect to the decomposition $\mathbb{C}^d = \mathrm{span}\{|0\rangle\} \oplus H_\perp$, we may write

$$A_i = \begin{pmatrix} \alpha_i & v_i^\dagger \\ 0 & B_i \end{pmatrix}, \qquad v_i \in H_\perp, \qquad B_i : H_\perp \to H_\perp,$$

and the trace preserving condition becomes

$$\sum_{i=1}^r |\alpha_i|^2 = 1, \qquad \sum_{i=1}^r \alpha_i v_i = 0, \qquad \sum_{i=1}^r \left(v_i v_i^\dagger + B_i^\dagger B_i\right) = I_{H_\perp}. \tag{36}$$

In particular,

$$\sum_{i=1}^r \|v_i\|^2 + \sum_{i=1}^r \|B_i\|_{\mathrm{F}}^2 = d - 1. \tag{37}$$

Proof. Because

$$|0\rangle\langle 0| = \Phi(|0\rangle\langle 0|) = \sum_{i=1}^r A_i|0\rangle\langle 0|A_i^\dagger,$$

each positive semidefinite summand on the right has support contained in $\mathrm{span}\{|0\rangle\}$. Hence $A_i|0\rangle \in \mathrm{span}\{|0\rangle\}$, proving the first claim. Writing the block decomposition and comparing the matrix blocks in $\sum_i A_i^\dagger A_i = I_d$ gives Eq. (36). Taking traces yields Eq. (37). □

Lemma E.4. *Define*

$$S := \sum_{i=1}^r \|v_i\|^2, \qquad T := \sum_{i=1}^r \|B_i\|_{\mathrm{F}}^2.$$

Then

$$F_\varepsilon(\Phi) = 1 - 2\varepsilon + \frac{\varepsilon}{d-1}\left(S + 2\Re \sum_{i=1}^r \alpha_i^* \,\mathrm{Tr}(B_i)\right) + o(\varepsilon). \tag{38}$$

Equivalently,

$$c(\Phi) := \lim_{\varepsilon \to 0^+} \frac{1 - F_\varepsilon(\Phi)}{\varepsilon} = 2 - \frac{1}{d-1}\left(S + 2\Re \sum_{i=1}^r \alpha_i^* \,\mathrm{Tr}(B_i)\right). \tag{39}$$

Proof. Fix a unit vector $|\varphi\rangle \in H_\perp$ and set

$$a := \sqrt{1-\varepsilon}, \qquad b := \sqrt{\varepsilon}, \qquad x_i(\varphi) := \langle v_i \,|\, \varphi\rangle, \qquad y_i(\varphi) := \langle \varphi | B_i | \varphi \rangle.$$

Then

$$\langle \psi_\varepsilon(\varphi) | A_i | \psi_\varepsilon(\varphi)\rangle = a^2 \alpha_i + ab\, x_i(\varphi) + b^2 y_i(\varphi).$$

Since

$$\langle \psi | \Phi(|\psi\rangle\langle\psi|) | \psi \rangle = \sum_{i=1}^r |\langle \psi | A_i | \psi \rangle|^2,$$

we obtain

$$F_\varepsilon(\Phi) = \sum_{i=1}^{r} \mathbb{E}_\varphi \Big[\big| a^2\alpha_i + ab\, x_i(\varphi) + b^2 y_i(\varphi) \big|^2 \Big]$$
$$= \sum_{i=1}^{r} \Big(a^4 \,|\alpha_i|^2 + a^2b^2\, \mathbb{E}\,|x_i(\varphi)|^2 + 2a^2b^2 \Re\big(\alpha_i^* \mathbb{E}[y_i(\varphi)]\big) \Big) + o(\varepsilon).$$

The terms of order a^3b and ab^3 vanish after Haar averaging: $\mathbb{E}[x_i(\varphi)] = 0$ by symmetry, and $\mathbb{E}[x_i(\varphi)^* y_i(\varphi)] = 0$ because $x_i(\varphi)^* y_i(\varphi)$ acquires a nontrivial phase under $|\varphi\rangle \mapsto e^{i\theta}|\varphi\rangle$. By the first moments of Haar measure on the unit sphere of $H_\perp$,

$$\mathbb{E}[x_i(\varphi)] = 0, \qquad \mathbb{E}\,|x_i(\varphi)|^2 = \frac{\|v_i\|^2}{d-1}, \qquad \mathbb{E}[y_i(\varphi)] = \frac{\mathrm{Tr}(B_i)}{d-1}.$$

Using $a^4 = (1-\varepsilon)^2 = 1 - 2\varepsilon + O(\varepsilon^2)$, $a^2b^2 = \varepsilon(1-\varepsilon) = \varepsilon + O(\varepsilon^2)$, and $\sum_i |\alpha_i|^2 = 1$, we obtain Eq. (38). Equation (39) is immediate. □

Lemma E.5 (Universal first-order lower bound)**.** *For every factorized channel $\Phi = \mathcal{D} \circ \mathcal{E}$ as above,*

$$c(\Phi) \geq \frac{d-m}{d-1}. \tag{40}$$

Proof. Let

$$I_+ := \{i : \alpha_i \neq 0\}.$$

If $i \in I_+$, then $|0\rangle \in \mathrm{im}(A_i)$ and $\mathrm{rank}(A_i) \leq m$, so $\mathrm{rank}(B_i) \leq m-1$. Hence

$$|\mathrm{Tr}(B_i)|^2 \leq \mathrm{rank}(B_i)\ \|B_i\|_{\mathrm{F}}^2 \leq (m-1)\,\|B_i\|_{\mathrm{F}}^2 \qquad (i \in I_+).$$

Therefore,

$$\left| \sum_{i=1}^{r} \alpha_i^* \,\mathrm{Tr}(B_i) \right| \leq \sum_{i \in I_+} |\alpha_i|\ |\mathrm{Tr}(B_i)|$$
$$\leq \left(\sum_{i \in I_+} |\alpha_i|^2 \right)^{1/2} \left(\sum_{i \in I_+} |\mathrm{Tr}(B_i)|^2 \right)^{1/2}$$
$$\leq \sqrt{(m-1) \sum_{i \in I_+} \|B_i\|_{\mathrm{F}}^2} \leq \sqrt{(m-1)T}.$$

Using $S + T = d - 1$ from Eq. (37),

$$S + 2\Re \sum_{i=1}^{r} \alpha_i^* \,\mathrm{Tr}(B_i) \leq d - 1 - T + 2\sqrt{(m-1)T}.$$

The right-hand side is maximized at $T = m-1$, where it equals $d+m-2$. Substituting into Eq. (39) gives

$$c(\Phi) \geq 2 - \frac{d+m-2}{d-1} = \frac{d-m}{d-1},$$

which proves Eq. (40). □

Define the optimal first-order coefficient

$$c_* := \frac{d-m}{d-1}. \tag{41}$$

The reset construction below attains this bound.

E.4 Classification of first-order-optimal channels

Proposition E.6. *Let $\Phi = \mathcal{D} \circ \mathcal{E}$ factor through $\mathbb{C}^m$ and satisfy $\Phi(|0\rangle\langle 0|) = |0\rangle\langle 0|$. If $c(\Phi) = c_*$, then*

$$\Phi(\rho) = \Phi_P(\rho) := P\rho P + \mathrm{Tr}\big((I_d - P)\rho\big)|0\rangle\langle 0| \tag{42}$$

for some rank-m orthogonal projection P with $|0\rangle \in \mathrm{im}(P)$.

Lemma E.7. *Let $B \neq 0$ be a linear operator on a finite-dimensional Hilbert space. If*

$$|\mathrm{Tr}(B)|^2 = \mathrm{rank}(B)\ \|B\|_{\mathrm{F}}^2\,,$$

then there exist a scalar $\beta \neq 0$ and an orthogonal projection Q of rank $\mathrm{rank}(B)$ such that

$$B = \beta Q.$$

Proof. Let $r := \mathrm{rank}(B)$, and let Q be the orthogonal projection onto $(\ker B)^\perp$. Then

$$\mathrm{rank}(Q) = \dim\big((\ker B)^\perp\big) = \mathrm{rank}(B) = r.$$

Moreover, $B = BQ$. Now equip the operator space with the Frobenius inner product

$$\langle X, Y\rangle_{\mathrm{F}} := \mathrm{Tr}(X^\dagger Y).$$

Since $Q = Q^\dagger$, and since the trace is cyclic in finite dimensions, we get

$$\mathrm{Tr}(B) = \mathrm{Tr}(BQ) = \mathrm{Tr}(QB) = \langle Q, B\rangle_{\mathrm{F}}.$$

Hence, by the Cauchy–Schwarz inequality,

$$|\mathrm{Tr}(B)|^2 = |\langle Q, B\rangle_{\mathrm{F}}|^2 \le \|Q\|_{\mathrm{F}}^2\ \|B\|_{\mathrm{F}}^2\,.$$

Because Q is an orthogonal projection,

$$\|Q\|_{\mathrm{F}}^2 = \mathrm{Tr}(Q^\dagger Q) = \mathrm{Tr}(Q^2) = \mathrm{Tr}(Q) = \mathrm{rank}(Q) = r.$$

Therefore

$$|\mathrm{Tr}(B)|^2 \le r\ \|B\|_{\mathrm{F}}^2\,.$$

By hypothesis, equality holds. So equality holds in the Cauchy–Schwarz inequality, and therefore B is a scalar multiple of Q:

$$B = \beta Q$$

for some $\beta \in \mathbb{C}$. Since $B \neq 0$, we must have $\beta \neq 0$. □

Lemma E.8 (The equality family)**.** *Assume $c(\Phi) = c_*$. After multiplying each Kraus operator A_i by a phase of modulus 1, we may assume:*

(i) each α_i is real and nonnegative.

(ii) if $\alpha_i = 0$, then $B_i = 0$.

(iii) if $\alpha_i > 0$, then there exists a rank-$(m-1)$ orthogonal projection P_i on $H_\perp$ such that

$$B_i = \alpha_i P_i.$$

Moreover,

$$\sum_i \alpha_i^2 = 1, \qquad \sum_i \alpha_i v_i = 0, \qquad \sum_i \big(v_i v_i^\dagger + \alpha_i^2 P_i\big) = I_{H_\perp}. \tag{43}$$

Proof. Because $c(\Phi) = c_*$, every inequality in the proof of Lemma E.5 is saturated. In particular,

$$d - 1 - T + 2\sqrt{(m-1)T}$$

must attain its maximum at the actual value of T, so $T = m - 1$.

Equality in

$$\sum_{i\in I_+} |\mathrm{Tr}(B_i)|^2 \le (m-1)\sum_{i\in I_+} \|B_i\|_{\mathrm{F}}^2 \le (m-1)T$$

forces $B_i = 0$ whenever $\alpha_i = 0$, and for every $i \in I_+$,

$$|\mathrm{Tr}(B_i)|^2 = (m-1)\,\|B_i\|_{\mathrm{F}}^2\,, \qquad \mathrm{rank}(B_i) = m-1.$$

By Lemma E.7, for each such i there exist $\beta_i \neq 0$ and a rank-$(m-1)$ projection P_i such that $B_i = \beta_i P_i$.

Equality in the Cauchy–Schwarz step

$$\sum_i |\alpha_i|\;|\mathrm{Tr}(B_i)| \le \left(\sum_i |\alpha_i|^2\right)^{1/2} \left(\sum_i |\mathrm{Tr}(B_i)|^2\right)^{1/2}$$

shows that there exists $c \ge 0$ with $|\mathrm{Tr}(B_i)| = c\,|\alpha_i|$ for every i. Using $\sum_i |\alpha_i|^2 = 1$ and $\sum_i |\mathrm{Tr}(B_i)|^2 = (m-1)^2$, we find $c = m-1$. Since $|\mathrm{Tr}(B_i)| = (m-1)\,|\beta_i|$, this implies $|\beta_i| = |\alpha_i|$ for every $i \in I_+$.

Finally, equality throughout

$$\Re \sum_i \alpha_i^* \,\mathrm{Tr}(B_i) \le \left|\sum_i \alpha_i^* \,\mathrm{Tr}(B_i)\right| \le \sum_i |\alpha_i|\;|\mathrm{Tr}(B_i)|$$

implies that $\sum_i \alpha_i^* \,\mathrm{Tr}(B_i)$ is real and nonnegative, and that all complex numbers $\alpha_i^* \,\mathrm{Tr}(B_i)$ have a common phase. Multiplying each A_i by an individual unit-modulus scalar does not change the channel, so we may choose those phases to make every α_i real and nonnegative. Then $\mathrm{Tr}(B_i)$ has the same phase, hence is also real and nonnegative, which forces $\beta_i = \alpha_i$. This proves the lemma. □

Lemma E.9. *Under the hypotheses of Lemma E.8, all nonzero projections P_i are equal.*

Proof. Set $P_i := 0$ when $\alpha_i = 0$, and define

$$U_i := \mathrm{im}(P_i) \subset H_\perp, \qquad N := \sum_i \alpha_i^2 P_i, \qquad W := \mathrm{im}(N) = \mathrm{span}\Big(\bigcup_i U_i\Big).$$

If the nonzero projections are not all equal, then at least two distinct $(m-1)$-dimensional subspaces occur among the U_i, and therefore

$$\dim W \ge m. \tag{44}$$

We will show that this contradicts the bottleneck dimension.

Step 1: the image of a factorized channel has dimension at most m^2. Since $\Phi = \mathcal{D} \circ \mathcal{E}$, the image of Φ as a linear map on the operator space is contained in the linear subspace $\mathcal{D}(L(\mathbb{C}^m))$. Hence

$$\dim \mathrm{im}(\Phi) \le \dim L(\mathbb{C}^m) = m^2. \tag{45}$$

Step 2: a block formula for Φ. Write an arbitrary $\rho \in L(\mathbb{C}^d)$ in block form as

$$\rho = \begin{pmatrix} a & x^\dagger \\ y & Z \end{pmatrix}, \qquad a \in \mathbb{C}, \quad x, y \in H_\perp, \quad Z \in L(H_\perp).$$

Using the block decomposition from Lemma E.3 together with $B_i = \alpha_i P_i$, a direct multiplication yields

$$A_i \rho A_i^\dagger = \begin{pmatrix} \alpha_i^2 a + \alpha_i v_i^\dagger y + \alpha_i x^\dagger v_i + v_i^\dagger Z v_i & \alpha_i^2 x^\dagger P_i + \alpha_i v_i^\dagger Z P_i \\ \alpha_i^2 P_i y + \alpha_i P_i Z v_i & \alpha_i^2 P_i Z P_i \end{pmatrix}.$$

Summing over i, using $\sum_i \alpha_i v_i = 0$ and

$$\sum_i v_i v_i^\dagger = I_{H_\perp} - N,$$

we obtain

$$\Phi(\rho) = \begin{pmatrix} a + \mathrm{Tr}\big((I_{H_\perp} - N)Z\big) & x^\dagger N + \sum_i \alpha_i v_i^\dagger Z P_i \\ N y + \sum_i \alpha_i P_i Z v_i & T(Z) \end{pmatrix}, \tag{46}$$

where

$$T(Z) := \sum_i \alpha_i^2 P_i Z P_i. \tag{47}$$

Step 3: a lower bound on $\dim \mathrm{im}(\Phi)$**.** From Eq. (46) we introduce the four linear subspaces

$$V_{00} := \mathrm{span}\{|0\rangle\langle 0|\},$$
$$V_{01} := \left\{ \begin{pmatrix} 0 & x^\dagger N \\ 0 & 0 \end{pmatrix} : x \in H_\perp \right\} = \{|0\rangle\langle w| : w \in W\},$$
$$V_{10} := \left\{ \begin{pmatrix} 0 & 0 \\ Ny & 0 \end{pmatrix} : y \in H_\perp \right\} = \{|w\rangle\langle 0| : w \in W\},$$
$$V_{11} := \left\{ \begin{pmatrix} 0 & 0 \\ 0 & T(Z) \end{pmatrix} : Z \in L(H_\perp) \right\}.$$

These subspaces live in disjoint matrix blocks, so their sum is direct. In particular,

$$\dim \mathrm{im}(\Phi) \geq 1 + \dim V_{01} + \dim V_{10} + \dim V_{11} \tag{48}$$

once we know that all four subspaces are contained in $\mathrm{im}(\Phi)$.

The inclusions $V_{00}, V_{01}, V_{10} \subseteq \mathrm{im}(\Phi)$ are immediate from Eq. (46). The inclusion $V_{11} \subseteq \mathrm{im}(\Phi)$ is slightly less immediate. Fix $Z \in L(H_\perp)$ and set

$$\rho_Z := \begin{pmatrix} 0 & 0 \\ 0 & Z \end{pmatrix}.$$

Then

$$\Phi(\rho_Z) = \begin{pmatrix} \mathrm{Tr}\big((I_{H_\perp} - N)Z\big) & \sum_i \alpha_i v_i^\dagger Z P_i \\ \sum_i \alpha_i P_i Z v_i & T(Z) \end{pmatrix}.$$

Define

$$c_Z := \sum_i \alpha_i P_i Z v_i.$$

For each i, $P_i Z v_i \in \mathrm{im}(P_i) \subseteq W$, hence $c_Z \in W$. Since N is positive semidefinite and $\mathrm{im}(N) = W$, the restriction $N|_W : W \to W$ is invertible. Therefore there exists $y_Z \in W$ with $N y_Z = c_Z$, and so

$$\begin{pmatrix} 0 & 0 \\ c_Z & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ N y_Z & 0 \end{pmatrix} \in V_{10}.$$

Next define

$$r_Z^\dagger := \sum_i \alpha_i v_i^\dagger Z P_i.$$

If $u \in W^\perp$, then $P_i u = 0$ for every i, hence $r_Z^\dagger u = 0$. Equivalently, there exists $u_Z \in W$ such that

$$r_Z^\dagger = u_Z^\dagger \Pi_W,$$

where Π_W is the orthogonal projection onto W. Again using invertibility of $N|_W$, choose $x_Z \in W$ with $N x_Z = u_Z$. Since $N = N\Pi_W = \Pi_W N$ and N is self-adjoint,

$$x_Z^\dagger N = u_Z^\dagger \Pi_W = r_Z^\dagger.$$

Therefore

$$\begin{pmatrix} 0 & r_Z^\dagger \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & x_Z^\dagger N \\ 0 & 0 \end{pmatrix} \in V_{01}.$$

Also,

$$\begin{pmatrix} \mathrm{Tr}((I_{H_\perp} - N)Z) & 0 \\ 0 & 0 \end{pmatrix} \in V_{00}.$$

Subtracting these three elements from $\Phi(\rho_Z) \in \mathrm{im}(\Phi)$ shows that

$$\begin{pmatrix} 0 & 0 \\ 0 & T(Z) \end{pmatrix} \in \mathrm{im}(\Phi),$$

and therefore $V_{11} \subseteq \mathrm{im}(\Phi)$.

Now $\dim V_{01} = \dim V_{10} = \dim W$. Choose an index i_0 with $\alpha_{i_0} > 0$. For every $X \in L(U_{i_0})$,

$$\langle X, T(X)\rangle_{\mathrm{HS}} = \sum_i \alpha_i^2 \left\| P_i X P_i \right\|_{\mathrm{F}}^2 \geq \alpha_{i_0}^2 \left\| P_{i_0} X P_{i_0} \right\|_{\mathrm{F}}^2 = \alpha_{i_0}^2 \left\| X \right\|_{\mathrm{F}}^2.$$

Hence T is injective on $L(U_{i_0})$, so

$$\dim V_{11} = \operatorname{rank}(T) \geq \dim L(U_{i_0}) = (m-1)^2.$$

Substituting into Eq. (48),

$$\dim \operatorname{im}(\Phi) \geq 1 + 2\dim W + (m-1)^2.$$

If the P_i are not all equal, then by Eq. (44)

$$\dim \operatorname{im}(\Phi) \geq 1 + 2m + (m-1)^2 = m^2 + 2,$$

contradicting Eq. (45). Therefore all nonzero P_i coincide. □

Proof of Proposition E.6. By Lemmas E.8 and E.9, there exists a single rank-$(m-1)$ orthogonal projection $P_\perp$ on $H_\perp$ such that

$$B_i = \alpha_i P_\perp \qquad \text{for all } i.$$

Set

$$P := |0\rangle\langle 0| + P_\perp.$$

Then P is a rank-m orthogonal projection with $|0\rangle \in \operatorname{im}(P)$.

Choose a unitary matrix $U = (u_{ji})$ on the Kraus index space such that the new coefficients

$$\alpha'_j := \sum_i u_{ji}\alpha_i$$

satisfy $(\alpha'_1, \alpha'_2, \ldots, \alpha'_r) = (1, 0, \ldots, 0)$. Define the mixed Kraus family

$$A'_j := \sum_i u_{ji} A_i.$$

The channel is unchanged. Because $B_i = \alpha_i P_\perp$, the lower-right blocks transform as

$$B'_j = \sum_i u_{ji} B_i = \left(\sum_i u_{ji}\alpha_i\right) P_\perp = \alpha'_j P_\perp.$$

Hence

$$B'_1 = P_\perp, \qquad B'_j = 0 \quad (j \geq 2).$$

Also,

$$v'_1 = \sum_i u_{1i} v_i = \sum_i \alpha_i v_i = 0$$

by Eq. (43). Therefore

$$A'_1 = P, \qquad A'_j = |0\rangle\langle v'_j| \quad (j \geq 2).$$

Using again Eq. (43),

$$\sum_{j\geq 2} v'_j {v'_j}^\dagger = I_{H_\perp} - P_\perp = I_d - P$$

on $H_\perp$. Thus

$$\begin{aligned}\Phi(\rho) &= P\rho P + \sum_{j\geq 2} |0\rangle\langle v'_j|\,\rho\,|v'_j\rangle\langle 0| \\ &= P\rho P + \operatorname{Tr}\big((I_d - P)\rho\big)|0\rangle\langle 0|,\end{aligned}$$

which is Eq. (42). □

E.5 Factorizations of reset channels

Lemma E.10 (Kraus operators of the identity channel)**.** *Let C be finite dimensional, and let $\{K_t\}_t$ be Kraus operators of the identity channel on $L(C)$:*

$$\sum_t K_t X K_t^\dagger = X \qquad \text{for all } X \in L(C).$$

Then every K_t is a scalar multiple of the identity on C.

Proof. Fix a unit vector $|\psi\rangle \in C$ and apply the channel identity to $X = |\psi\rangle\langle\psi|$:

$$|\psi\rangle\langle\psi| = \sum_t K_t|\psi\rangle\langle\psi|K_t^\dagger.$$

The right-hand side is a sum of positive semidefinite rank-one operators, while the left-hand side has one-dimensional support $\mathrm{span}\{|\psi\rangle\}$. Therefore every vector $K_t|\psi\rangle$ belongs to $\mathrm{span}\{|\psi\rangle\}$. As this holds for every unit vector, each K_t is a scalar multiple of the identity. □

Proposition E.11 (Factorizations of a reset channel)**.** *Let P be a rank-m orthogonal projection on $\mathbb{C}^d$ with $|0\rangle \in \mathrm{im}(P)$, and define*

$$\Phi_P(\rho) := P\rho P + \mathrm{Tr}\big((I_d - P)\rho\big)|0\rangle\langle 0|.$$

Assume

$$\Phi_P = \mathcal{D} \circ \mathcal{E}, \qquad \mathcal{E} : L(\mathbb{C}^d) \to L(\mathbb{C}^m), \qquad \mathcal{D} : L(\mathbb{C}^m) \to L(\mathbb{C}^d)$$

are CPTP. Then:

(i) there exists a unitary $U : \mathrm{im}(P) \to \mathbb{C}^m$ such that

$$\mathcal{E}(X) = UXU^\dagger \qquad \text{for all } X \in L(\mathrm{im}(P)).$$

(ii) the decoder is the inverse isometry channel,

$$\mathcal{D}(Y) = U^\dagger Y U \qquad \text{for all } Y \in L(\mathbb{C}^m).$$

(iii) the encoder has Kraus rank

$$\mathrm{rank}_{\mathrm{Kraus}}(\mathcal{E}) = d - m + 1.$$

Proof. Let $C := \mathrm{im}(P) \subset \mathbb{C}^d$. For every $X \in L(C)$ we have $PXP = X$ and $(I_d - P)X = 0$, so $\Phi_P(X) = X$. Thus $\mathcal{D} \circ \mathcal{E}$ acts as the identity channel on the full matrix algebra $L(C)$.

Choose Kraus operators $\{F_a\}_a$ for $\mathcal{E}$ and $\{G_\mu\}_\mu$ for $\mathcal{D}$. Then $\{G_\mu F_a P\}_{\mu,a}$ is a Kraus family for the identity channel on $L(C)$. By Lemma E.10, for every pair (μ, a) there is a scalar $c_{\mu a}$ such that

$$G_\mu F_a P = c_{\mu a} P. \tag{49}$$

Using $\sum_\mu G_\mu^\dagger G_\mu = I_m$, we obtain

$$PF_a^\dagger F_b P = \sum_\mu P F_a^\dagger G_\mu^\dagger G_\mu F_b P = \left(\sum_\mu c_{\mu a}^* c_{\mu b}\right) P. \tag{50}$$

So there is a positive semidefinite matrix $\Lambda = (\lambda_{ab})$ such that $PF_a^\dagger F_b P = \lambda_{ab} P$. Diagonalize Λ by a unitary change of Kraus basis. After this change we may assume

$$P\widetilde{F}_j^\dagger \widetilde{F}_\ell P = \lambda_j \delta_{j\ell} P, \qquad \lambda_j \geq 0, \qquad \sum_j \lambda_j = 1.$$

If $\lambda_j > 0$, define

$$U_j := \lambda_j^{-1/2} \widetilde{F}_j P : C \to \mathbb{C}^m.$$

Then $U_j^\dagger U_j = P$, so U_j is an isometry from the m-dimensional space C into the m-dimensional space $\mathbb{C}^m$, hence a unitary. If two coefficients λ_j and λ_ℓ were positive with $j \neq \ell$, then

$$U_j^\dagger U_\ell = (\lambda_j \lambda_\ell)^{-1/2} P \widetilde{F}_j^\dagger \widetilde{F}_\ell P = 0,$$

which is impossible because the product of two unitaries cannot be the zero operator. Therefore exactly one λ_j is nonzero. Since their sum is 1, it must equal 1. Renaming that Kraus operator as F_1, we conclude

$$F_1 P = U, \qquad F_j P = 0 \quad (j \geq 2),$$

for some unitary $U : C \to \mathbb{C}^m$. This proves part (i).

Now let $Y \in L(\mathbb{C}^m)$ be arbitrary. Because U is unitary, there exists $X \in L(C)$ with $Y = UXU^\dagger = \mathcal{E}(X)$. Hence

$$\mathcal{D}(Y) = \mathcal{D}(\mathcal{E}(X)) = \Phi_P(X) = X = U^\dagger Y U,$$

which proves part (ii).

Finally, part (ii) implies that for every $\rho \in L(\mathbb{C}^d)$,

$$\Phi_P(\rho) = \mathcal{D}(\mathcal{E}(\rho)) = U^\dagger \mathcal{E}(\rho) U, \qquad \text{so} \qquad \mathcal{E}(\rho) = U \Phi_P(\rho) U^\dagger.$$

Thus $\mathcal{E}$ and Φ_P have the same Kraus rank. To compute the latter, choose an orthonormal basis $\{e_1, \ldots, e_{d-m}\}$ of $(\operatorname{im} P)^\perp$. Then

$$\Phi_P(\rho) = P\rho P + \sum_{j=1}^{d-m} |0\rangle\langle e_j| \, \rho \, |e_j\rangle\langle 0|.$$

So Φ_P has the Kraus family

$$P, \qquad |0\rangle\langle e_1|, \ldots, |0\rangle\langle e_{d-m}|,$$

which is Hilbert–Schmidt orthogonal and hence linearly independent. Therefore $\operatorname{rank}_{\mathrm{Kraus}}(\Phi_P) = d - m + 1$, and the same is true for $\mathcal{E}$. □

E.6 Alternative proof that the decoder in Proposition E.11 is isometric

We provide an alternative proof of the isometric-decoder part of Proposition E.11. This proof is not needed elsewhere, but it is conceptually useful because it isolates the decoder from the Kraus-rank computation.

Let $|0_m\rangle$ denote the first standard basis vector of $\mathbb{C}^m$. Because $|0\rangle \in \operatorname{im}(P)$ and $\operatorname{rank}(P) = m$, choose an orthonormal basis $|u_1\rangle, \ldots, |u_m\rangle$ of $\operatorname{im}(P)$ with $|u_1\rangle = |0\rangle$, and define an isometry $V : \mathbb{C}^m \to \mathbb{C}^d$ by

$$V|j-1\rangle = |u_j\rangle \qquad (1 \leq j \leq m).$$

Then $VV^\dagger = P$, $V^\dagger V = I_m$, and $V|0_m\rangle = |0\rangle$. Define the embedding channel

$$\mathcal{V}(X) := VXV^\dagger, \qquad X \in L(\mathbb{C}^m),$$

and the channel $R : L(\mathbb{C}^d) \to L(\mathbb{C}^m)$ by

$$R(Y) := V^\dagger Y V + \operatorname{Tr}\big((I_d - P)Y\big)|0_m\rangle\langle 0_m|.$$

Choose an orthonormal basis $|f_1\rangle, \ldots, |f_{d-m}\rangle$ of $\operatorname{im}(I_d - P)$. Then R has Kraus operators

$$K_0 := V^\dagger, \qquad K_j := |0_m\rangle\langle f_j| \quad (1 \leq j \leq d-m),$$

so R is CPTP. Moreover,

$$(R \circ \mathcal{V})(X) = X \qquad \text{for all } X \in L(\mathbb{C}^m).$$

Now set

$$\widetilde{\mathcal{E}} := \mathcal{E} \circ \mathcal{V} : L(\mathbb{C}^m) \to L(\mathbb{C}^m), \qquad \widetilde{\mathcal{D}} := R \circ \mathcal{D} : L(\mathbb{C}^m) \to L(\mathbb{C}^m).$$

Both maps are CPTP, and

$$\widetilde{\mathcal{D}} \circ \widetilde{\mathcal{E}} = R \circ \mathcal{D} \circ \mathcal{E} \circ \mathcal{V} = R \circ \Phi_P \circ \mathcal{V} = R \circ \mathcal{V} = \mathrm{id}_{L(\mathbb{C}^m)}.$$

Thus $\widetilde{\mathcal{E}}$ is a CPTP channel with a CPTP left inverse. By the standard reversibility characterization of left-invertible channels ([Puzzuoli and Watrous, 2017, Theorem 18 and Remark 19], [Nayak and Sen, 2007, Theorem 2.1], and [Nielsen and Chuang, 2010, Theorem 10.1]), such a channel must be of the form

$$\widetilde{\mathcal{E}}(X) = WXW^\dagger$$

for some unitary $W \in \mathrm{U}(m)$. On the other hand,

$$\mathcal{D} \circ \widetilde{\mathcal{E}} = \mathcal{D} \circ \mathcal{E} \circ \mathcal{V} = \Phi_P \circ \mathcal{V} = \mathcal{V}.$$

Therefore, for every $X \in L(\mathbb{C}^m)$,

$$\mathcal{D}(X) = \mathcal{D}\big(\widetilde{\mathcal{E}}(W^\dagger XW)\big) = \mathcal{V}(W^\dagger XW) = VW^\dagger XWV^\dagger.$$

Setting $A := VW^\dagger$ gives an isometry $A^\dagger A = I_m$ with $AA^\dagger = P$ and

$$\mathcal{D}(X) = AXA^\dagger.$$

This is the isometric form asserted in Proposition E.11.

E.7 From asymptotics to optimality for small ε

We now give a detailed restatement and proof of Theorem 3.5.

Theorem E.12 (Encoder Kraus-rank lower bound)**.** *Set $d := 2^n$, $m := 2^k$, and $H_\perp := |0\rangle^\perp \subset \mathbb{C}^d$. For $\varepsilon \in (0,1)$ define*

$$|\psi_\varepsilon(\varphi)\rangle := \sqrt{1-\varepsilon}|0\rangle + \sqrt{\varepsilon}|\varphi\rangle, \qquad \|\varphi\| = 1, \quad |\varphi\rangle \in H_\perp,$$

and let $\mu_{1,\varepsilon}$ be the distribution induced by Haar-uniform $|\varphi\rangle$ on the unit sphere of $H_\perp$. Then for every $R < d - m + 1$ there exists $\varepsilon_0(R) > 0$ such that, for every $0 < \varepsilon < \varepsilon_0(R)$, no optimal pair $(\mathcal{E}, \mathcal{D})$ maximizing $F_{\mu_{1,\varepsilon}}(\mathcal{E}, \mathcal{D})$ can have encoder Kraus rank at most R.

In particular, for all sufficiently small $\varepsilon > 0$, every optimal (n, k, n_B, n_E)-QAE satisfies

$$n_B \geq k.$$

Proof. Fix $R < d - m + 1$, and let K_R be the set of all pairs $(\mathcal{E}, \mathcal{D})$ such that

- $\mathcal{E} : L(\mathbb{C}^d) \to L(\mathbb{C}^m)$ is CPTP and $\mathrm{rank}_{\mathrm{Kraus}}(\mathcal{E}) \leq R$.
- $\mathcal{D} : L(\mathbb{C}^m) \to L(\mathbb{C}^d)$ is CPTP.

This set is compact: finite-dimensional CPTP sets are compact, and the rank condition is closed because it is the condition that the Choi matrix have rank at most R.

For $(\mathcal{E}, \mathcal{D}) \in K_R$, write $\Phi = \mathcal{D} \circ \mathcal{E}$ and view $F_\varepsilon(\mathcal{E}, \mathcal{D}) := F_\varepsilon(\Phi)$. By Lemma E.2, there are continuous coefficient functions A_R, B_R, C_R on K_R such that

$$F_\varepsilon(\mathcal{E}, \mathcal{D}) = A_R(\mathcal{E}, \mathcal{D}) + \varepsilon B_R(\mathcal{E}, \mathcal{D}) + \varepsilon^2 C_R(\mathcal{E}, \mathcal{D}).$$

Moreover,

$$A_R(\mathcal{E}, \mathcal{D}) = \langle 0|\Phi(|0\rangle\langle 0|)|0\rangle \leq 1.$$

If $K_R = \varnothing$, there is nothing to prove. Otherwise, pairing any admissible encoder with the replacement decoder $\mathcal{D}_{\mathrm{rep}}(Y) := \mathrm{Tr}(Y)|0\rangle\langle 0|$ gives a pair in K_R with $A_R = 1$. Hence

$$A_{\max} := \max_{(\mathcal{E},\mathcal{D}) \in K_R} A_R(\mathcal{E}, \mathcal{D}) = 1,$$

and the maximizing set

$$S_R := \{(\mathcal{E}, \mathcal{D}) \in K_R : A_R(\mathcal{E}, \mathcal{D}) = 1\}$$

is nonempty and compact.

Now fix $(\mathcal{E}, \mathcal{D}) \in S_R$. Then $\Phi(|0\rangle\langle 0|) = |0\rangle\langle 0|$, so Lemma E.5 applies and gives

$$-B_R(\mathcal{E}, \mathcal{D}) = c(\Phi) \geq c_*.$$

Suppose, for contradiction, that

$$\sup_{(\mathcal{E},\mathcal{D})\in S_R} B_R(\mathcal{E},\mathcal{D}) = -c_*.$$

Because S_R is compact and B_R is continuous, some pair $(\mathcal{E}_*,\mathcal{D}_*) \in S_R$ attains this supremum. Then $c(\mathcal{D}_*\circ\mathcal{E}_*) = c_*$. By Proposition E.6, the overall channel must be a reset channel, and then Proposition E.11 implies $\mathrm{rank}_{\mathrm{Kraus}}(\mathcal{E}_*) = d-m+1$, contradicting the assumption $\mathrm{rank}_{\mathrm{Kraus}}(\mathcal{E}_*) \le R < d-m+1$.

Therefore there exists $\delta_R > 0$ such that

$$\sup_{(\mathcal{E},\mathcal{D})\in S_R} B_R(\mathcal{E},\mathcal{D}) \le -c_* - \delta_R. \tag{51}$$

Applying Lemma E.1 to the compact set K_R and the functions A_R, B_R, C_R gives

$$\limsup_{\varepsilon\to 0^+} \frac{\max_{(\mathcal{E},\mathcal{D})\in K_R} F_\varepsilon(\mathcal{E},\mathcal{D}) - 1}{\varepsilon} \le -c_* - \delta_R.$$

Consequently, there exists $\varepsilon_1 > 0$ such that for all $0 < \varepsilon < \varepsilon_1$,

$$\max_{(\mathcal{E},\mathcal{D})\in K_R} F_\varepsilon(\mathcal{E},\mathcal{D}) \le 1 - \left(c_* + \frac{\delta_R}{2}\right)\varepsilon. \tag{52}$$

On the other hand, fix any rank-m projection P with $|0\rangle \in \mathrm{im}(P)$, choose an isometry $V : \mathbb{C}^m \to \mathbb{C}^d$ with $VV^\dagger = P$ and $V|0_m\rangle = |0\rangle$, and define

$$\mathcal{E}_P(\rho) := V^\dagger P\rho P V + \mathrm{Tr}\big((I_d - P)\rho\big)|0_m\rangle\langle 0_m|, \qquad \mathcal{D}_P(Y) := VYV^\dagger.$$

Then $\mathcal{D}_P\circ\mathcal{E}_P = \Phi_P$, and one checks directly from Eqs. (38) and (41) that $c(\Phi_P) = c_*$. Hence

$$F_\varepsilon(\Phi_P) = 1 - c_*\varepsilon + O(\varepsilon^2),$$

so there exists $\varepsilon_2 > 0$ such that for all $0 < \varepsilon < \varepsilon_2$,

$$F_\varepsilon(\Phi_P) \ge 1 - \left(c_* + \frac{\delta_R}{4}\right)\varepsilon. \tag{53}$$

Comparing Eqs. (52) and (53), we conclude that for all

$$0 < \varepsilon < \varepsilon_0(R) := \min\{\varepsilon_1, \varepsilon_2\},$$

no pair in K_R can be globally optimal. This proves the Kraus-rank statement.

Finally, if an (n, k, n_B, n_E)-QAE has encoder ancilla count n_B, then its encoder is implemented by a unitary on $n + n_B$ qubits followed by tracing out $n + n_B - k$ qubits, so its encoder Kraus rank is at most 2^{n+n_B-k}. If $n_B < k$, then

$$2^{n+n_B-k} \le 2^{n-1} < 2^n - 2^k + 1 = d - m + 1,$$

so such a QAE cannot be optimal for $\mu_{1,\varepsilon}$ once ε is sufficiently small. Therefore every optimal (n, k, n_B, n_E)-QAE must satisfy $n_B \ge k$. □

F A counterexample to sufficiency of isometric decoders

We recall and prove Proposition 3.7.

Proposition 3.7 (Decoder isometry is not sufficient)**.** *Let $n = 2$, $k = 1$,*

$$|\psi_\phi\rangle = \frac{|0\rangle + e^{i\phi}|1\rangle}{\sqrt{2}}, \qquad |\Psi_\phi\rangle = |\psi_\phi\rangle^{\otimes 2}, \qquad \phi \in [0, 2\pi), \tag{2}$$

and let μ_{ph} be the uniform distribution on ϕ. Then

$$\sup_{\mathcal{E},\, V^\dagger V = I_2} F_{\mu_{\mathrm{ph}}}(\mathcal{E},\mathcal{D}_V) < \frac{3}{4} \le \sup_{\mathcal{E},\mathcal{D}} F_{\mu_{\mathrm{ph}}}(\mathcal{E},\mathcal{D}), \tag{3}$$

where $\mathcal{D}_V(\rho) = V\rho V^\dagger$.

We prove Proposition 3.7. The source ensemble is the one-parameter family $\{|\Psi_\phi\rangle\}_\phi$ with $\phi \in [0, 2\pi)$ uniformly distributed. We consider blind single-copy compression from 2 qubits to 1 qubit and back, with arbitrary CPTP encoder

$$\mathcal{E} : L((\mathbb{C}^2)^{\otimes 2}) \to L(\mathbb{C}^2)$$

and arbitrary CPTP decoder

$$\mathcal{D} : L(\mathbb{C}^2) \to L((\mathbb{C}^2)^{\otimes 2}).$$

The performance functional is

$$F(\mathcal{E}, \mathcal{D}) := \frac{1}{2\pi} \int_0^{2\pi} \langle \Psi_\phi | \mathcal{D}(\mathcal{E}(|\Psi_\phi\rangle\langle\Psi_\phi|))|\Psi_\phi\rangle \, d\phi. \tag{54}$$

F.1 A Kraus-rank-2 decoder attaining $3/4$

Define the symmetric Bell state

$$|\Sigma\rangle := \frac{|01\rangle + |10\rangle}{\sqrt{2}}.$$

Then

$$|\Psi_\phi\rangle = \frac{1}{2}|00\rangle + \frac{e^{i\phi}}{\sqrt{2}}|\Sigma\rangle + \frac{e^{2i\phi}}{2}|11\rangle. \tag{55}$$

Take the encoder to be the partial trace over the second qubit,

$$\mathcal{E}(\rho) = \mathrm{Tr}_2(\rho),$$

so that $\mathcal{E}(|\Psi_\phi\rangle\langle\Psi_\phi|) = |\psi_\phi\rangle\langle\psi_\phi|$. Define a decoder with two Kraus operators $A_0, A_1 : \mathbb{C}^2 \to (\mathbb{C}^2)^{\otimes 2}$,

$$A_0 = \frac{1}{\sqrt{3}}|00\rangle\langle 0| + \sqrt{\frac{2}{3}}|\Sigma\rangle\langle 1|, \qquad A_1 = \sqrt{\frac{2}{3}}|\Sigma\rangle\langle 0| + \frac{1}{\sqrt{3}}|11\rangle\langle 1|. \tag{56}$$

Since $A_0^\dagger A_0 + A_1^\dagger A_1 = I_2$, this is a CPTP decoder.

For every ϕ,

$$A_0|\psi_\phi\rangle = \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{3}}|00\rangle + e^{i\phi}\sqrt{\frac{2}{3}}|\Sigma\rangle\right), \qquad A_1|\psi_\phi\rangle = \frac{1}{\sqrt{2}}\left(\sqrt{\frac{2}{3}}|\Sigma\rangle + e^{i\phi}\frac{1}{\sqrt{3}}|11\rangle\right).$$

Using Eq. (55), a direct calculation gives

$$|\langle\Psi_\phi \,|\, A_0\psi_\phi\rangle|^2 = \frac{3}{8}, \qquad |\langle\Psi_\phi \,|\, A_1\psi_\phi\rangle|^2 = \frac{3}{8}.$$

Therefore

$$\langle \Psi_\phi | \mathcal{D}(\mathcal{E}(|\Psi_\phi\rangle\langle\Psi_\phi|))|\Psi_\phi\rangle = \frac{3}{4} \qquad \text{for every } \phi. \tag{57}$$

In particular,

$$\sup_{\mathcal{E},\mathcal{D}} F(\mathcal{E}, \mathcal{D}) \geq \frac{3}{4}. \tag{58}$$

F.2 A universal upper bound for isometric decoders

Now restrict the decoder to Kraus rank 1, so

$$\mathcal{D}_V(\rho) = V\rho V^\dagger, \qquad V^\dagger V = I_2,$$

for some isometry $V : \mathbb{C}^2 \to (\mathbb{C}^2)^{\otimes 2}$. Let

$$P := VV^\dagger,$$

so P is a rank-2 orthogonal projector onto the image of V. Fix an arbitrary encoder $\mathcal{E}$ and set

$$\rho_\phi := V\mathcal{E}(|\Psi_\phi\rangle\langle\Psi_\phi|)V^\dagger.$$

Because $\mathcal{E}(|\Psi_\phi\rangle\langle\Psi_\phi|)$ is a density matrix on a two-dimensional space,

$$0 \leq \mathcal{E}(|\Psi_\phi\rangle\langle\Psi_\phi|) \leq I_2.$$

Multiplying by V and $V^\dagger$ gives

$$0 \le \rho_\phi \le P. \tag{59}$$

Hence

$$\langle \Psi_\phi | \rho_\phi | \Psi_\phi \rangle \le \langle \Psi_\phi | P | \Psi_\phi \rangle.$$

Averaging over ϕ yields

$$F(\mathcal{E}, \mathcal{D}_V) \le \mathrm{Tr}(P\rho), \qquad \rho := \frac{1}{2\pi} \int_0^{2\pi} |\Psi_\phi\rangle\langle\Psi_\phi| \, d\phi. \tag{60}$$

From Eq. (55), averaging kills the oscillatory cross-terms and gives

$$\rho = \frac{1}{4}|00\rangle\langle 00| + \frac{1}{2}|\Sigma\rangle\langle\Sigma| + \frac{1}{4}|11\rangle\langle 11|. \tag{61}$$

Thus ρ has eigenvalues $1/2, 1/4, 1/4, 0$, with eigenspaces spanned by $|\Sigma\rangle$, by $\mathrm{span}\{|00\rangle, |11\rangle\}$, and by the antisymmetric Bell state

$$|A\rangle := \frac{|01\rangle - |10\rangle}{\sqrt{2}}.$$

Therefore the maximal value of $\mathrm{Tr}(P\rho)$ over rank-2 projectors P is the sum of the two largest eigenvalues:

$$\max_{\mathrm{rank}(P)=2} \mathrm{Tr}(P\rho) = \frac{1}{2} + \frac{1}{4} = \frac{3}{4}. \tag{62}$$

So every isometric decoder satisfies

$$F(\mathcal{E}, \mathcal{D}_V) \le \frac{3}{4}. \tag{63}$$

Moreover, every maximizing projector has the form

$$P_{t,\gamma} = |\Sigma\rangle\langle\Sigma| + |\eta_{t,\gamma}\rangle\langle\eta_{t,\gamma}|, \qquad |\eta_{t,\gamma}\rangle := \cos t|00\rangle + e^{i\gamma}\sin t|11\rangle, \quad 0 \le t \le \frac{\pi}{2}, \quad \gamma \in [0, 2\pi). \tag{64}$$

F.3 Implications of equality in the upper bound

Assume, for contradiction, that some encoder $\mathcal{E}$ and isometry V satisfy

$$F(\mathcal{E}, \mathcal{D}_V) = \frac{3}{4}.$$

Then equality must hold in both Eqs. (60) and (62). Consequently:

(a) the projector $P = VV^\dagger$ must be one of the maximizers in Eq. (64), say $P = P_{t,\gamma}$.

(b) for every ϕ, the pointwise inequality

$$\langle \Psi_\phi | \rho_\phi | \Psi_\phi \rangle \le \langle \Psi_\phi | P_{t,\gamma} | \Psi_\phi \rangle$$

must actually be an equality.

The second claim follows because the integrand is continuous and nonnegative. If its integral is 0, then it vanishes identically.

Let

$$s_\phi := \langle \Psi_\phi | P_{t,\gamma} | \Psi_\phi \rangle = \| P_{t,\gamma} | \Psi_\phi \rangle \|^2.$$

Since every $P_{t,\gamma}$ contains $|\Sigma\rangle$ and Eq. (55) has a nonzero $|\Sigma\rangle$ component, $s_\phi > 0$ for all ϕ. Define the normalized projections

$$|\chi_\phi^{(t,\gamma)}\rangle := \frac{P_{t,\gamma}|\Psi_\phi\rangle}{\sqrt{s_\phi}}. \tag{65}$$

Because ρ_ϕ is supported on $\mathrm{im}(P_{t,\gamma})$,

$$\langle \Psi_\phi | \rho_\phi | \Psi_\phi \rangle = \langle \Psi_\phi | P_{t,\gamma} \rho_\phi P_{t,\gamma} | \Psi_\phi \rangle = s_\phi \, \langle \chi_\phi^{(t,\gamma)} | \rho_\phi | \chi_\phi^{(t,\gamma)} \rangle.$$

Using Eq. (59), the restriction of ρ_ϕ to $\mathrm{im}(P_{t,\gamma})$ is bounded above by the identity on that two-dimensional subspace. Hence

$$\langle \chi_\phi^{(t,\gamma)} | \rho_\phi | \chi_\phi^{(t,\gamma)} \rangle \le 1.$$

Equality of the fidelities therefore forces

$$\rho_\phi = |\chi_\phi^{(t,\gamma)}\rangle\langle\chi_\phi^{(t,\gamma)}| \qquad \text{for every } \phi.$$

Thus equality in Eq. (63) would imply the existence of a CPTP map

$$\Phi := \mathcal{D}_V \circ \mathcal{E}$$

that sends each input pure state $|\Psi_\phi\rangle\langle\Psi_\phi|$ to the pure state $|\chi_\phi^{(t,\gamma)}\rangle\langle\chi_\phi^{(t,\gamma)}|$. We next show that no such channel exists.

F.4 A pure-state transformation criterion

Theorem F.1 ([Chefles et al., 2004], Theorem 2)**.** *Let $\{|\alpha_j\rangle\}_{j=1}^m \subset H_A$ and $\{|\beta_j\rangle\}_{j=1}^m \subset H_B$ be finite families of unit vectors. There exists a CPTP map $\Phi : L(H_A) \to L(H_B)$ satisfying*

$$\Phi(|\alpha_j\rangle\langle\alpha_j|) = |\beta_j\rangle\langle\beta_j| \qquad (j = 1, \dots, m)$$

if and only if there exists a positive semidefinite matrix $M \in M_m(\mathbb{C})$ with diagonal entries $M_{jj} = 1$ such that

$$G_\alpha = M \circ G_\beta, \tag{66}$$

where G_α and G_β are the Gram matrices

$$(G_\alpha)_{ij} = \langle \alpha_i \,|\, \alpha_j \rangle, \qquad (G_\beta)_{ij} = \langle \beta_i \,|\, \beta_j \rangle,$$

and $\circ$ denotes Hadamard (entrywise) product.

F.5 Reduction of γ to 0

Lemma F.2. *Let*

$$U_\beta := |0\rangle\langle 0| + e^{i\beta}|1\rangle\langle 1|, \qquad W_\beta := U_\beta^{\otimes 2}.$$

Then:

(i) $W_\beta |\Psi_\phi\rangle = |\Psi_{\phi+\beta}\rangle$ *for every* ϕ, β.

(ii) $W_{\gamma/2} P_{t,0} W_{\gamma/2}^\dagger = P_{t,\gamma}$ *for every* t, γ.

(iii) therefore

$$W_{\gamma/2} |\chi_\phi^{(t,0)}\rangle = |\chi_{\phi+\gamma/2}^{(t,\gamma)}\rangle \qquad \text{for every } \phi. \tag{67}$$

Consequently, there exists a channel sending every $|\Psi_\phi\rangle\langle\Psi_\phi|$ to $|\chi_\phi^{(t,\gamma)}\rangle\langle\chi_\phi^{(t,\gamma)}|$ if and only if there exists a channel sending every $|\Psi_\phi\rangle\langle\Psi_\phi|$ to $|\chi_\phi^{(t,0)}\rangle\langle\chi_\phi^{(t,0)}|$.

Proof. Part (i) is immediate from

$$U_\beta |\psi_\phi\rangle = \frac{|0\rangle + e^{i(\phi+\beta)}|1\rangle}{\sqrt{2}} = |\psi_{\phi+\beta}\rangle.$$

For part (ii), recall that

$$P_{t,0} = |\Sigma\rangle\langle\Sigma| + |\eta_{t,0}\rangle\langle\eta_{t,0}|, \qquad |\eta_{t,0}\rangle = \cos t|00\rangle + \sin t|11\rangle.$$

Now

$$W_{\gamma/2}|\Sigma\rangle = e^{i\gamma/2}|\Sigma\rangle, \qquad W_{\gamma/2}|\eta_{t,0}\rangle = \cos t|00\rangle + e^{i\gamma}\sin t|11\rangle = |\eta_{t,\gamma}\rangle,$$

which proves part (ii). For part (iii), apply part (ii) and then part (i):

$$W_{\gamma/2} P_{t,0} |\Psi_\phi\rangle = P_{t,\gamma} W_{\gamma/2} |\Psi_\phi\rangle = P_{t,\gamma} |\Psi_{\phi+\gamma/2}\rangle.$$

Since $W_{\gamma/2}$ is unitary, the norms on both sides agree, proving Eq. (67).

If a channel Φ_γ sends $|\Psi_\phi\rangle\langle\Psi_\phi|$ to $|\chi_\phi^{(t,\gamma)}\rangle\langle\chi_\phi^{(t,\gamma)}|$ for all ϕ, then

$$\Phi_0 := \mathrm{Ad}_{W_{\gamma/2}^\dagger} \circ \Phi_\gamma \circ \mathrm{Ad}_{W_{\gamma/2}}$$

sends $|\Psi_\phi\rangle\langle\Psi_\phi|$ to $|\chi_\phi^{(t,0)}\rangle\langle\chi_\phi^{(t,0)}|$ by Eq. (67). The converse is obtained by reversing the conjugation. □

By Lemma F.2, it is enough to treat the case $\gamma = 0$. Write

$$P_t := P_{t,0}, \qquad |\eta_t\rangle := |\eta_{t,0}\rangle = \cos t|00\rangle + \sin t|11\rangle.$$

F.6 The case $0 \le \sin 2t < 1$

Set

$$s := \sin 2t, \qquad c := \cos 2t.$$

For the projector $P_t = |\Sigma\rangle\langle\Sigma| + |\eta_t\rangle\langle\eta_t|$, define the normalized projected outputs

$$|\chi_\phi^{(t)}\rangle := \frac{P_t|\Psi_\phi\rangle}{\|P_t|\Psi_\phi\rangle\|}.$$

Using Eq. (55),

$$P_t|\Psi_\phi\rangle = \frac{\cos t + e^{2i\phi}\sin t}{2}|\eta_t\rangle + \frac{e^{i\phi}}{\sqrt{2}}|\Sigma\rangle,$$

so

$$\|P_t|\Psi_\phi\rangle\|^2 = \frac{3 + s\cos(2\phi)}{4}. \tag{68}$$

In particular,

$$|\chi_0^{(t)}\rangle = \frac{(\cos t + \sin t)|\eta_t\rangle + \sqrt{2}|\Sigma\rangle}{\sqrt{3+s}},$$

$$|\chi_{\pi/2}^{(t)}\rangle = \frac{(\cos t - \sin t)|\eta_t\rangle + i\sqrt{2}|\Sigma\rangle}{\sqrt{3-s}},$$

$$|\chi_\pi^{(t)}\rangle = \frac{(\cos t + \sin t)|\eta_t\rangle - \sqrt{2}|\Sigma\rangle}{\sqrt{3+s}},$$

$$|\chi_{3\pi/2}^{(t)}\rangle = \frac{(\cos t - \sin t)|\eta_t\rangle - i\sqrt{2}|\Sigma\rangle}{\sqrt{3-s}}.$$

From these formulas one computes

$$\langle\chi_0^{(t)}\,|\,\chi_{\pi/2}^{(t)}\rangle = \frac{c+2i}{\sqrt{9-s^2}}, \qquad \langle\chi_0^{(t)}\,|\,\chi_\pi^{(t)}\rangle = \frac{s-1}{s+3}, \qquad \langle\chi_0^{(t)}\,|\,\chi_{3\pi/2}^{(t)}\rangle = \frac{c-2i}{\sqrt{9-s^2}}. \tag{69}$$

Now choose the four phases

$$\phi_1 = 0, \qquad \phi_2 = \frac{\pi}{2}, \qquad \phi_3 = \pi, \qquad \phi_4 = \frac{3\pi}{2}.$$

For the inputs,

$$\langle\Psi_0\,|\,\Psi_{\pi/2}\rangle = \frac{i}{2}, \qquad \langle\Psi_0\,|\,\Psi_\pi\rangle = 0, \qquad \langle\Psi_0\,|\,\Psi_{3\pi/2}\rangle = -\frac{i}{2}, \tag{70}$$

and cyclic symmetry determines the rest of the 4×4 Gram matrix. If a channel mapped the four input states $|\Psi_{\phi_j}\rangle\langle\Psi_{\phi_j}|$ to the four pure output states $|\chi_{\phi_j}^{(t)}\rangle\langle\chi_{\phi_j}^{(t)}|$, then Theorem F.1 would force a positive semidefinite matrix M_t with diagonal entries 1 such that

$$G_\Psi = M_t \circ G_{\chi^{(t)}}.$$

For these four phases all off-diagonal entries are uniquely determined, and one obtains

$$M_t = \begin{pmatrix} 1 & h & 0 & \overline{h} \\ \overline{h} & 1 & h & 0 \\ 0 & \overline{h} & 1 & h \\ h & 0 & \overline{h} & 1 \end{pmatrix}, \qquad h = \frac{\frac{i}{2}}{\frac{c+2i}{\sqrt{9-s^2}}} = \frac{\sqrt{9-s^2}(2+ic)}{2(5-s^2)}. \tag{71}$$

Now apply M_t to the vector

$$v := (1, -1, 1, -1)^\top.$$

A direct multiplication gives

$$M_t v = (1 - h - \overline{h})v = \left(1 - \frac{2\sqrt{9-s^2}}{5-s^2}\right) v.$$

For $0 \le s < 1$ we have

$$2\sqrt{9-s^2} \ge 2\sqrt{8} > 5 \ge 5 - s^2,$$

so

$$1 - \frac{2\sqrt{9-s^2}}{5-s^2} < 0.$$

Hence M_t has a negative eigenvalue and cannot be positive semidefinite, contradicting Theorem F.1. We have proved the following.

Proposition F.3. *If $0 \le \sin 2t < 1$, then no CPTP map sends every $|\Psi_\phi\rangle\langle\Psi_\phi|$ to $|\chi_\phi^{(t,0)}\rangle\langle\chi_\phi^{(t,0)}|$.*

F.7 The case $\sin 2t = 1$

Now suppose $\sin 2t = 1$, i.e. $t = \pi/4$. Then

$$|\eta_{\pi/4}\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}.$$

A short calculation shows that

$$P_{\pi/4}|\Psi_\phi\rangle = \frac{e^{i\phi}}{\sqrt{2}}\big(\cos\phi\,|\eta_{\pi/4}\rangle + |\Sigma\rangle\big),$$

so the normalized projected outputs are

$$|\chi_\phi^{(\pi/4)}\rangle = e^{i\phi}\,\frac{\cos\phi\,|\eta_{\pi/4}\rangle + |\Sigma\rangle}{\sqrt{1+\cos^2\phi}}. \tag{72}$$

Choose the three phases

$$\phi_1 = 0, \qquad \phi_2 = \frac{\pi}{6}, \qquad \phi_3 = \frac{5\pi}{6}.$$

If a channel mapped the corresponding input states to the three pure outputs from Eq. (72), then Theorem F.1 would again force a positive semidefinite matrix M with unit diagonal satisfying

$$G_\Psi = M \circ G_{\chi^{(\pi/4)}}.$$

A direct computation of the forced entrywise quotients gives

$$M = \begin{pmatrix} 1 & \sqrt{14}/4 & \sqrt{14}/4 \\ \sqrt{14}/4 & 1 & 7/4 \\ \sqrt{14}/4 & 7/4 & 1 \end{pmatrix}. \tag{73}$$

Applying M to $(0, 1, -1)^\top$ gives

$$M(0, 1, -1)^\top = -\frac{3}{4}(0, 1, -1)^\top,$$

so M has a negative eigenvalue and is not positive semidefinite. Therefore:

Proposition F.4. *If $\sin 2t = 1$, then no CPTP map sends every $|\Psi_\phi\rangle\langle\Psi_\phi|$ to $|\chi_\phi^{(t,0)}\rangle\langle\chi_\phi^{(t,0)}|$.*

F.8 Conclusion of the counterexample

Combining Lemma F.2 and Propositions F.3 and F.4 gives the following impossibility statement.

Corollary F.5. *No encoder $\mathcal{E}$ and isometry V can satisfy*

$$F(\mathcal{E}, \mathcal{D}_V) = \frac{3}{4}.$$

Equivalently, the upper bound in Eq. (63) *is never attained.*

The set of encoders is compact in finite dimensions (for example via Choi matrices), and the set of isometries $V : \mathbb{C}^2 \to (\mathbb{C}^2)^{\otimes 2}$ is compact as well. Since $(\mathcal{E}, V) \mapsto F(\mathcal{E}, \mathcal{D}_V)$ is continuous, the maximum over all encoder-isometric-decoder pairs is attained. By Corollary F.5, that maximum is strictly smaller than $3/4$. Together with Eq. (58), this proves Proposition 3.7.

F.9 Detailed restatement and proof of Observation 3.8

Let $\Phi_{U,V} := \mathcal{D}_V \circ \mathcal{E}_U$,

$$f_{U,V}(\phi) := \langle \Psi_\phi | \Phi_{U,V}(|\Psi_\phi\rangle\langle\Psi_\phi|) | \Psi_\phi \rangle, \qquad \tilde{F}_{\mu_{\mathrm{ph}}}(U,V) := \frac{1}{2\pi} \int_0^{2\pi} f_{U,V}(\phi)\, d\phi. \tag{74}$$

Observation F.6. *Set $\phi_j := 2\pi j/5$ for $j = 0, 1, 2, 3, 4$. Then*

$$\tilde{F}_{\mu_{\mathrm{ph}}}(U,V) = \frac{1}{5} \sum_{j=0}^{4} f_{U,V}(\phi_j). \tag{75}$$

In particular, the phase integral in Eq. (74) *can be replaced by five evaluations of the fidelity integrand.*

Proof. Write

$$|\Sigma\rangle := \frac{|01\rangle + |10\rangle}{\sqrt{2}}, \qquad |v_0\rangle := \frac{1}{2}|00\rangle, \qquad |v_1\rangle := \frac{1}{\sqrt{2}}|\Sigma\rangle, \qquad |v_2\rangle := \frac{1}{2}|11\rangle. \tag{76}$$

From Eq. (2) and the definition of $|\Sigma\rangle$,

$$|\Psi_\phi\rangle = \frac{1}{2}|00\rangle + \frac{e^{i\phi}}{\sqrt{2}}|\Sigma\rangle + \frac{e^{2i\phi}}{2}|11\rangle = \sum_{r=0}^{2} e^{ir\phi}|v_r\rangle. \tag{77}$$

Therefore

$$|\Psi_\phi\rangle\langle\Psi_\phi| = \sum_{a,b=0}^{2} e^{i(a-b)\phi} |v_a\rangle\langle v_b|. \tag{78}$$

By linearity of $\Phi_{U,V}$,

$$\Phi_{U,V}(|\Psi_\phi\rangle\langle\Psi_\phi|) = \sum_{a,b=0}^{2} e^{i(a-b)\phi}\, \Phi_{U,V}(|v_a\rangle\langle v_b|)\,. \tag{79}$$

Also,

$$\langle \Psi_\phi| = \sum_{c=0}^{2} e^{-ic\phi} \langle v_c|, \qquad |\Psi_\phi\rangle = \sum_{d=0}^{2} e^{id\phi}|v_d\rangle. \tag{80}$$

Substituting Eqs. (79) and (80) into the definition of $f_{U,V}(\phi)$ gives

$$f_{U,V}(\phi) = \sum_{a,b,c,d=0}^{2} e^{i(a-b-c+d)\phi}\, \langle v_c | \Phi_{U,V}(|v_a\rangle\langle v_b|)\, |v_d\rangle. \tag{81}$$

Now $a, b, c, d \in \{0, 1, 2\}$, so the integer

$$m := a - b - c + d$$

always lies in $\{-4,-3,\dots,4\}$. Hence $f_{U,V}$ is a trigonometric polynomial of degree at most 4:

$$f_{U,V}(\phi)=\sum_{m=-4}^{4} c_m(U,V)e^{im\phi}$$

for suitable coefficients $c_m(U,V)$.

Averaging over $[0,2\pi)$ keeps only the zero Fourier mode. Thus

$$\tilde{F}_{\mu_{\rm ph}}(U,V)=c_0(U,V)=\sum_{a,b,c,d=0}^{2}\delta_{a-b-c+d,\,0}\,\langle v_c|\Phi_{U,V}(|v_a\rangle\langle v_b|)\,|v_d\rangle. \tag{82}$$

Now let $\omega:=e^{2\pi i/5}$ and $\phi_j:=2\pi j/5$. The root-of-unity identity gives

$$\frac{1}{5}\sum_{j=0}^{4}\omega^{jm}=\begin{cases}1, & m\equiv 0 \pmod 5,\\ 0, & m\not\equiv 0 \pmod 5.\end{cases}$$

Because $m\in\{-4,-3,\dots,4\}$, the congruence $m\equiv 0 \pmod 5$ is equivalent here to $m=0$. Therefore

$$\begin{aligned}\frac{1}{5}\sum_{j=0}^{4} f_{U,V}(\phi_j)&=\sum_{a,b,c,d=0}^{2}\left(\frac{1}{5}\sum_{j=0}^{4}e^{2\pi ij(a-b-c+d)/5}\right)\langle v_c|\Phi_{U,V}(|v_a\rangle\langle v_b|)\,|v_d\rangle\\ &=\sum_{a,b,c,d=0}^{2}\delta_{a-b-c+d,\,0}\,\langle v_c|\Phi_{U,V}(|v_a\rangle\langle v_b|)\,|v_d\rangle\\ &=\tilde{F}_{\mu_{\rm ph}}(U,V).\end{aligned} \tag{83}$$

This proves Eq. (75). □

Remark F.7. *The proof uses only the phase expansion in Eq.* (77) *and the linearity of* $\Phi_{U,V}$*. Hence the same argument works for any fixed linear map* $\Phi: L\big((\mathbb{C}^2)^{\otimes 2}\big)\to L\big((\mathbb{C}^2)^{\otimes 2}\big)$.

F.10 Proof of Theorem 3.10

We recall and prove Theorem 3.10.

Theorem 3.10 (Sharp decoder threshold for $k=1$)**.** *For every* $n\geq 2$*,*

$$\sup_{\mathcal{E},\,V^\dagger V=I_2} F_{\mu_{\rm ph}^{(n)}}(\mathcal{E},\mathcal{D}_V)<\frac{3}{4}\leq \sup_{\mathcal{E},\mathcal{D}} F_{\mu_{\rm ph}^{(n)}}(\mathcal{E},\mathcal{D}),$$

where $\mathcal{E}: L(H_{2^n})\to L(H_2)$*,* $\mathcal{D}: L(H_2)\to L(H_{2^n})$*, and* $\mathcal{D}_V(\rho)=V\rho V^\dagger$*. Consequently, every optimal* $(n,1,n_B,n_E)$*-QAE for the source* $\mu_{\rm ph}^{(n)}$ *satisfies* $n_E\geq n$*.*

Proof. We first prove the lower bound on the unrestricted optimum. Let $\mathcal{E}_{\rm ph}: L(H_4)\to L(H_2)$ and $\mathcal{D}_{\rm ph}: L(H_2)\to L(H_4)$ be the encoder and decoder constructed in Appendix F.1 for the source $\mu_{\rm ph}$, so that

$$F_{\mu_{\rm ph}}(\mathcal{E}_{\rm ph},\mathcal{D}_{\rm ph})=\frac{3}{4}.$$

Define

$$\mathcal{E}^{(n)}:=\mathcal{E}_{\rm ph}\circ{\rm Tr}_R,\qquad \mathcal{D}^{(n)}(\rho):=\mathcal{D}_{\rm ph}(\rho)\otimes|0_R\rangle\langle 0_R|.$$

Then $\mathcal{E}^{(n)}: L(H_{2^n})\to L(H_2)$ and $\mathcal{D}^{(n)}: L(H_2)\to L(H_{2^n})$ are CPTP, and for every ϕ we have

$$(\mathcal{D}^{(n)}\circ\mathcal{E}^{(n)})(|\Omega_\phi\rangle\langle\Omega_\phi|)=(\mathcal{D}_{\rm ph}\circ\mathcal{E}_{\rm ph})(|\Psi_\phi\rangle\langle\Psi_\phi|)\otimes|0_R\rangle\langle 0_R|.$$

Therefore

$$F_{\mu_{\rm ph}^{(n)}}(\mathcal{E}^{(n)},\mathcal{D}^{(n)})=F_{\mu_{\rm ph}}(\mathcal{E}_{\rm ph},\mathcal{D}_{\rm ph})=\frac{3}{4},$$

which proves

$$\sup_{\mathcal{E},\mathcal{D}} F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E},\mathcal{D}) \geq \frac{3}{4}.$$

We next bound the performance of isometric decoders. Fix an encoder $\mathcal{E} : L(H_{2^n}) \to L(H_2)$ and an isometry $V : H_2 \to H_{2^n}$, and let $P := VV^\dagger$, so P is a rank-two projector on H_{2^n}. For each ϕ, set

$$\rho_\phi := V\,\mathcal{E}(|\Omega_\phi\rangle\langle\Omega_\phi|)V^\dagger.$$

Since $\mathcal{E}(|\Omega_\phi\rangle\langle\Omega_\phi|)$ is a qubit density operator, it is bounded above by I_2, and hence

$$0 \leq \rho_\phi \leq VI_2V^\dagger = P.$$

Thus

$$\langle\Omega_\phi|\rho_\phi|\Omega_\phi\rangle \leq \langle\Omega_\phi|P|\Omega_\phi\rangle.$$

Averaging over ϕ gives

$$F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E},\mathcal{D}_V) \leq \mathrm{Tr}(P\,\overline{\rho}_{\mathrm{ph}}^{(n)}), \qquad \overline{\rho}_{\mathrm{ph}}^{(n)} := \mathbb{E}_\phi[|\Omega_\phi\rangle\langle\Omega_\phi|].$$

By the calculation already carried out in Eq. (61),

$$\overline{\rho}_{\mathrm{ph}} := \mathbb{E}_\phi[|\Psi_\phi\rangle\langle\Psi_\phi|] = \frac{1}{4}|00\rangle\langle 00| + \frac{1}{2}|\Sigma\rangle\langle\Sigma| + \frac{1}{4}|11\rangle\langle 11|,$$

so

$$\overline{\rho}_{\mathrm{ph}}^{(n)} = \overline{\rho}_{\mathrm{ph}} \otimes |0_R\rangle\langle 0_R|.$$

Hence the nonzero eigenvalues of $\overline{\rho}_{\mathrm{ph}}^{(n)}$ are $1/2$, $1/4$, and $1/4$. Since P has rank two,

$$\mathrm{Tr}(P\,\overline{\rho}_{\mathrm{ph}}^{(n)}) \leq \frac{1}{2} + \frac{1}{4} = \frac{3}{4}.$$

Therefore

$$\sup_{\mathcal{E},\,V^\dagger V = I_2} F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E},\mathcal{D}_V) \leq \frac{3}{4}.$$

To rule out equality, assume that some encoder $\mathcal{E}$ and isometry V satisfy

$$F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E},\mathcal{D}_V) = \frac{3}{4}.$$

Then the preceding inequalities are tight, so in particular

$$\mathrm{Tr}(P\,\overline{\rho}_{\mathrm{ph}}^{(n)}) = \frac{3}{4}.$$

Because the only nonzero eigenvalues of $\overline{\rho}_{\mathrm{ph}}^{(n)}$ are $1/2$, $1/4$, and $1/4$, equality implies that $\mathrm{im}(P)$ is spanned by an eigenvector of eigenvalue $1/2$ together with a unit vector from the $1/4$-eigenspace. In particular,

$$\mathrm{im}(P) \subseteq \mathrm{supp}(\overline{\rho}_{\mathrm{ph}}^{(n)}) \subseteq (\mathbb{C}^2)^{\otimes 2} \otimes \mathrm{span}\{|0_R\rangle\}.$$

Hence there exists an isometry $W : H_2 \to (\mathbb{C}^2)^{\otimes 2}$ such that

$$V = W \otimes |0_R\rangle.$$

Now define a two-qubit encoder $\mathcal{E}^\sharp : L(H_4) \to L(H_2)$ by

$$\mathcal{E}^\sharp(\rho) := \mathcal{E}(\rho \otimes |0_R\rangle\langle 0_R|).$$

Then for every ϕ,

$$(\mathcal{D}_V \circ \mathcal{E})(|\Omega_\phi\rangle\langle\Omega_\phi|) = (\mathcal{D}_W \circ \mathcal{E}^\sharp)(|\Psi_\phi\rangle\langle\Psi_\phi|) \otimes |0_R\rangle\langle 0_R|,$$

and therefore

$$F_{\mu_{\mathrm{ph}}}(\mathcal{E}^\sharp,\mathcal{D}_W) = F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E},\mathcal{D}_V) = \frac{3}{4}.$$

This contradicts Corollary F.5. Hence no encoder and isometry can attain fidelity $3/4$ for $\mu_{\mathrm{ph}}^{(n)}$.

The sets of encoders $\mathcal{E} : L(H_{2^n}) \to L(H_2)$ and isometries $V : H_2 \to H_{2^n}$ are compact in finite dimensions, and the map $(\mathcal{E}, V) \mapsto F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E}, \mathcal{D}_V)$ is continuous. Therefore the maximum over encoder-isometric-decoder pairs is attained. Since the value $3/4$ is impossible, that maximum is strictly smaller than $3/4$:

$$\sup_{\mathcal{E},\, V^\dagger V = I_2} F_{\mu_{\mathrm{ph}}^{(n)}}(\mathcal{E}, \mathcal{D}_V) < \frac{3}{4}.$$

Finally, consider an optimal $(n, 1, n_B, n_E)$-QAE for the source $\mu_{\mathrm{ph}}^{(n)}$. By Definition 3.1 we must have $n_E \geq n - 1$. If $n_E < n$, then necessarily $n_E = n - 1$. But then the discarded decoder register G has $1 + n_E - n = 0$ qubits, so the decoder channel is isometric. This is impossible for an optimal QAE, because the globally achievable fidelity is at least $3/4$ while all isometric decoders achieve fidelity strictly below $3/4$. Therefore every optimal $(n, 1, n_B, n_E)$-QAE for $\mu_{\mathrm{ph}}^{(n)}$ must satisfy $n_E \geq n$. □

G Non-isometric decoders can outperform isometric ones for a fixed suboptimal encoder

The counterexample in Appendix F is about *globally optimal* encoder–decoder pairs. A much simpler phenomenon already appears if the encoder is held fixed but suboptimal.

Let $n = 2$ and $k = 1$, and consider the source distribution consisting of

$$|\psi_0\rangle = |00\rangle, \qquad |\psi_1\rangle = |11\rangle,$$

sampled with probabilities $1 - \varepsilon$ and ε, where $0 < \varepsilon < 1/2$. With an optimal encoder one can achieve perfect reconstruction using, for example, the encoder whose Kraus operators are

$$K_0 = |0\rangle\langle 00| + |1\rangle\langle 11|, \qquad K_1 = |0\rangle\langle 01| + |1\rangle\langle 10|,$$

and the isometric decoder

$$\mathcal{D}(\rho) = W\rho W^\dagger, \qquad W = |00\rangle\langle 0| + |11\rangle\langle 1|.$$

Now instead fix the suboptimal encoder that discards the input and always outputs the maximally mixed qubit state $I_2/2$. If the decoder is allowed to be non-isometric, it can ignore its input and deterministically output $|00\rangle$, giving expected fidelity

$$(1 - \varepsilon) \cdot 1 + \varepsilon \cdot 0 = 1 - \varepsilon > \frac{1}{2}.$$

Suppose instead that the decoder is restricted to be an isometry W. Then the reconstructed output state is always

$$W \frac{I_2}{2} W^\dagger = \frac{1}{2} P, \qquad P := WW^\dagger,$$

where P is a rank-2 projector. For any pure state $|\psi\rangle$,

$$\langle\psi| \frac{1}{2} P |\psi\rangle \leq \frac{1}{2}.$$

Hence the expected fidelity is at most $1/2$. So for this fixed suboptimal encoder, non-isometric decoders strictly outperform all isometric decoders. This is weaker than Proposition 3.7, but it pinpoints the failure of a tempting proof strategy already noted in spirit by [Barnum et al., 1996].

H A source-dependent multiplicative guarantee for isometric decoders

Isometric decoders are not universally sufficient for optimality, but one can still prove a source-dependent approximation guarantee in the same blind single-copy setting. Let $d := 2^n$ and $m := 2^k$. For a source distribution μ over pure states $|\psi\rangle \in \mathbb{C}^d$, define

$$F^\star(\mu) := \sup_{\mathcal{E}, \mathcal{D}} F_\mu(\mathcal{E}, \mathcal{D}), \qquad F_{\mathrm{iso}}^\star(\mu) := \sup_{\mathcal{E},\, V^\dagger V = I_m} F_\mu(\mathcal{E}, \mathcal{D}_V).$$

Also let

$$\bar{\rho} := \mathbb{E}_{|\psi\rangle\sim\mu}[|\psi\rangle\langle\psi|], \qquad s_m := \sum_{i=1}^{m} \lambda_i(\bar{\rho}), \qquad \eta_m := 1 - s_m,$$

where all eigenvalues are listed in nonincreasing order and padded with zeros when necessary. For each rank-m orthogonal projector P on $\mathbb{C}^d$, define

$$p_P(\psi) := \langle\psi|P|\psi\rangle, \qquad Y_P := \mathbb{E}_{|\psi\rangle\sim\mu}\Big[(1 - p_P(\psi))P|\psi\rangle\langle\psi|P\Big].$$

Definition H.1. *For finite-dimensional Hilbert spaces A and B, we say that a state $\rho \in L(A \otimes B)$ is separable if it can be written as*

$$\rho = \sum_{x=1}^{N} p_x \, \sigma_x \otimes \tau_x, \qquad p_x \geq 0, \qquad \sum_{x=1}^{N} p_x = 1,$$

where each σ_x is a density operator on A and each τ_x is a density operator on B.

Finally, define

$$B_\mu := \sup_{\substack{P=P^\dagger=P^2 \\ \operatorname{rank}(P)=m}} \big(\mathbb{E}_{|\psi\rangle\sim\mu}\big[p_P(\psi)^2\big] + \lambda_{\max}(Y_P)\big)$$

and

$$U_\mu := \sup\left\{ \sum_{i=1}^{m} \lambda_i(\Omega) : \begin{array}{l} \Omega \in L(\mathbb{C}^d \otimes H_E),\ \Omega \geq 0,\ \operatorname{Tr}(\Omega) = 1, \\ \operatorname{Tr}_E(\Omega) = \bar{\rho},\ \Omega \text{ separable}, \\ H_E \text{ finite dimensional} \end{array} \right\}.$$

We now give a detailed restatement and proof of Theorem 3.12.

Theorem H.2. *With the notation above,*

$$\frac{F^\star_{\text{iso}}(\mu)}{F^\star(\mu)} \geq \frac{B_\mu}{U_\mu}.$$

Moreover,

$$U_\mu \leq s_m, \qquad B_\mu \geq \left(1 - \frac{1}{m}\right) s_m^2 + \frac{s_m}{m},$$

and therefore

$$\frac{F^\star_{\text{iso}}(\mu)}{F^\star(\mu)} \geq \frac{B_\mu}{s_m} \geq 1 - \left(1 - \frac{1}{m}\right)\eta_m.$$

Proof. We divide the proof into four steps.

Step 1: the unrestricted optimum is at most U_μ. Fix any encoder $\mathcal{E} : L(\mathbb{C}^d) \to L(\mathbb{C}^m)$ and decoder $\mathcal{D} : L(\mathbb{C}^m) \to L(\mathbb{C}^d)$. Choose a Stinespring isometry $W : \mathbb{C}^m \to \mathbb{C}^d \otimes H_E$ such that

$$\mathcal{D}(\sigma) = \operatorname{Tr}_E\big(W\sigma W^\dagger\big),$$

where H_E is finite dimensional, and set

$$P := WW^\dagger.$$

Then P is a rank-m orthogonal projector on $\mathbb{C}^d \otimes H_E$. For a pure input state $|\psi\rangle\langle\psi|$ and a latent state $\sigma \in L(\mathbb{C}^m)$,

$$\langle\psi|\mathcal{D}(\sigma)|\psi\rangle = \operatorname{Tr}\big[(|\psi\rangle\langle\psi| \otimes I_E)W\sigma W^\dagger\big].$$

As σ ranges over all density operators on $\mathbb{C}^m$, the state $W\sigma W^\dagger$ ranges over all density operators supported on $\operatorname{im}(P)$. Hence

$$\sup_{\sigma\geq 0,\, \operatorname{Tr}\sigma=1} \langle\psi|\mathcal{D}(\sigma)|\psi\rangle = \lambda_{\max}\big(P(|\psi\rangle\langle\psi| \otimes I_E)P\big).$$

Now define

$$R_\psi := (\langle\psi| \otimes I_E)P.$$

Then

$$P(|\psi\rangle\langle\psi| \otimes I_E)P = R_\psi^\dagger R_\psi, \qquad (\langle\psi| \otimes I_E)P(|\psi\rangle \otimes I_E) = R_\psi R_\psi^\dagger.$$

Therefore these two operators have the same nonzero eigenvalues, and so

$$\lambda_{\max}\big(P(|\psi\rangle\langle\psi| \otimes I_E)P\big) = \lambda_{\max}\big((\langle\psi| \otimes I_E)P(|\psi\rangle \otimes I_E)\big).$$

Since $(\langle\psi| \otimes I_E)P(|\psi\rangle \otimes I_E)$ is positive semidefinite on H_E,

$$\lambda_{\max}\big((\langle\psi| \otimes I_E)P(|\psi\rangle \otimes I_E)\big) = \sup_{\tau \geq 0,\, \mathrm{Tr}\,\tau = 1} \mathrm{Tr}\big[P(|\psi\rangle\langle\psi| \otimes \tau)\big].$$

Consequently,

$$\langle\psi|(\mathcal{D} \circ \mathcal{E})(|\psi\rangle\langle\psi|)|\psi\rangle \leq \sup_{\tau \geq 0,\, \mathrm{Tr}\,\tau = 1} \mathrm{Tr}\big[P(|\psi\rangle\langle\psi| \otimes \tau)\big].$$

Averaging over μ and taking the supremum over all measurable state-valued maps $|\psi\rangle \mapsto \tau_\psi$ gives

$$F_\mu(\mathcal{E}, \mathcal{D}) \leq \sup_{|\psi\rangle \mapsto \tau_\psi} \mathrm{Tr}(P\Omega), \qquad \Omega := \mathbb{E}_{|\psi\rangle \sim \mu}[|\psi\rangle\langle\psi| \otimes \tau_\psi].$$

Every such Ω is separable and satisfies $\mathrm{Tr}_E(\Omega) = \bar{\rho}$. Since P has rank m,

$$\mathrm{Tr}(P\Omega) \leq \sum_{i=1}^{m} \lambda_i(\Omega) \leq U_\mu.$$

Thus $F_\mu(\mathcal{E}, \mathcal{D}) \leq U_\mu$ for every CPTP pair $(\mathcal{E}, \mathcal{D})$, and therefore

$$F^\star(\mu) \leq U_\mu.$$

Step 2: the auxiliary upper bound U_μ is at most s_m. Fix a finite-dimensional auxiliary space H_E and a separable state $\Omega \in L(\mathbb{C}^d \otimes H_E)$ with $\mathrm{Tr}_E(\Omega) = \bar{\rho}$. Because the ambient spaces are finite dimensional, Ω admits a finite pure-product decomposition,

$$\Omega = \sum_{x=1}^{N} p_x\, |a_x\rangle\langle a_x| \otimes |b_x\rangle\langle b_x|, \qquad p_x \geq 0, \quad \sum_{x=1}^{N} p_x = 1.$$

Set

$$u_x := \sqrt{p_x}|a_x\rangle \in \mathbb{C}^d, \qquad w_x := u_x \otimes |b_x\rangle \in \mathbb{C}^d \otimes H_E.$$

Let A be the $d \times N$ matrix whose x-th column is u_x, and let B be the $(d \dim H_E) \times N$ matrix whose x-th column is w_x. Then

$$\bar{\rho} = AA^\dagger, \qquad \Omega = BB^\dagger.$$

Define the Gram matrices

$$H := A^\dagger A, \qquad G := B^\dagger B, \qquad K_{xy} := \langle b_x \,|\, b_y\rangle.$$

The nonzero eigenvalues of $AA^\dagger$ and $A^\dagger A$ coincide, and likewise for $BB^\dagger$ and $B^\dagger B$. Hence the nonzero eigenvalues of $\bar{\rho}$ are the eigenvalues of H, and the nonzero eigenvalues of Ω are the eigenvalues of G. Moreover,

$$G_{xy} = \langle w_x \,|\, w_y\rangle = \langle u_x \,|\, u_y\rangle\langle b_x \,|\, b_y\rangle = H_{xy}K_{xy},$$

so

$$G = H \circ K = K \circ H,$$

where $\circ$ denotes the entrywise product. Because K is a Gram matrix of unit vectors, it is positive semidefinite and has diagonal entries equal to 1. Thus the Schur multiplier

$$\Phi_K(X) := K \circ X$$

is completely positive by the Schur product theorem, and it is both unital and trace preserving because $K_{xx} = 1$ for all x. Its Hilbert–Schmidt adjoint is

$$\Phi_K^\dagger(X) = \overline{K} \circ X,$$

and $\overline{K}$ is again a correlation matrix, so $\Phi_K^\dagger$ is also positive, unital, and trace preserving. Now let X be any Hermitian matrix. By the Ky Fan variational principle,

$$\sum_{i=1}^{m} \lambda_i(X) = \max_{\substack{0 \le Q \le I \\ \mathrm{Tr}(Q)=m}} \mathrm{Tr}(QX).$$

Applying this to $G = \Phi_K(H)$ yields

$$\begin{aligned}\sum_{i=1}^{m} \lambda_i(G) &= \max_{\substack{0 \le Q \le I \\ \mathrm{Tr}(Q)=m}} \mathrm{Tr}\big(Q\,\Phi_K(H)\big) \\ &= \max_{\substack{0 \le Q \le I \\ \mathrm{Tr}(Q)=m}} \mathrm{Tr}\big(\Phi_K^\dagger(Q)\,H\big) \\ &\le \max_{\substack{0 \le Y \le I \\ \mathrm{Tr}(Y)=m}} \mathrm{Tr}(YH) = \sum_{i=1}^{m} \lambda_i(H).\end{aligned}$$

Because the eigenvalues of H are the nonzero eigenvalues of $\bar{\rho}$, the right-hand side equals s_m. Therefore

$$\sum_{i=1}^{m} \lambda_i(\Omega) = \sum_{i=1}^{m} \lambda_i(G) \le s_m.$$

As this holds for every feasible Ω, we conclude that

$$U_\mu \le s_m.$$

Step 3: the isometric optimum is at least B_μ. Fix a rank-m orthogonal projector P on $\mathbb{C}^d$ and a unit vector $|\xi\rangle \in \mathrm{im}(P)$. Choose an isometry $V_P : \mathbb{C}^m \to \mathbb{C}^d$ such that

$$V_P V_P^\dagger = P, \qquad V_P|0\rangle = |\xi\rangle.$$

Define an encoder and decoder by

$$\mathcal{E}_{P,\xi}(\rho) := V_P^\dagger P \rho P V_P + \mathrm{Tr}\big((I-P)\rho\big)|0\rangle\langle 0|, \qquad \mathcal{D}_{V_P}(\sigma) := V_P \sigma V_P^\dagger.$$

The map $E_{P,\xi}$ is CPTP because it is the sum of two completely positive maps and

$$\mathrm{Tr}\big(\mathcal{E}_{P,\xi}(\rho)\big) = \mathrm{Tr}(P\rho P) + \mathrm{Tr}\big((I-P)\rho\big) = \mathrm{Tr}(\rho).$$

For a pure input state $|\psi\rangle\langle\psi|$, write $p := p_P(\psi) = \langle\psi|P|\psi\rangle$. Using $V_P V_P^\dagger = P$ and $V_P|0\rangle\langle 0|V_P^\dagger = |\xi\rangle\langle\xi|$, we obtain

$$(\mathcal{D}_{V_P} \circ \mathcal{E}_{P,\xi})(|\psi\rangle\langle\psi|) = P|\psi\rangle\langle\psi|P + (1-p)|\xi\rangle\langle\xi|.$$

Therefore

$$\langle\psi|(\mathcal{D}_{V_P} \circ \mathcal{E}_{P,\xi})(|\psi\rangle\langle\psi|)|\psi\rangle = p^2 + (1-p)\,|\langle\xi\,|\,\psi\rangle|^2\,.$$

Averaging over μ gives

$$F_\mu(\mathcal{E}_{P,\xi}, \mathcal{D}_{V_P}) = \mathbb{E}_{|\psi\rangle\sim\mu}[p_P(\psi)^2] + \langle\xi|Y_P|\xi\rangle.$$

Since Y_P is supported on $\mathrm{im}(P)$,

$$\sup_{\substack{|\xi\rangle \in \mathrm{im}(P) \\ \|\xi\|=1}} F_\mu(\mathcal{E}_{P,\xi}, \mathcal{D}_{V_P}) = \mathbb{E}_{|\psi\rangle\sim\mu}[p_P(\psi)^2] + \lambda_{\max}(Y_P).$$

Taking the supremum over all rank-m projectors P proves that

$$F_{\mathrm{iso}}^\star(\mu) \ge B_\mu.$$

Step 4: a lower bound on B_μ. Let $P_\star$ be the orthogonal projector onto the span of the eigenvectors of $\bar{\rho}$ corresponding to its m largest eigenvalues. Then

$$\mathbb{E}_{|\psi\rangle\sim\mu}[p_{P_\star}(\psi)] = \mathrm{Tr}(P_\star\bar{\rho}) = s_m.$$

Because $Y_{P_\star} \geq 0$ and $Y_{P_\star}$ is supported on the m-dimensional space $\mathrm{im}(P_\star)$,

$$\lambda_{\max}(Y_{P_\star}) \geq \frac{1}{m}\,\mathrm{Tr}(Y_{P_\star}) = \frac{1}{m}\mathbb{E}_{|\psi\rangle\sim\mu}\big[p_{P_\star}(\psi) - p_{P_\star}(\psi)^2\big].$$

Hence

$$\begin{aligned} B_\mu &\geq \mathbb{E}_{|\psi\rangle\sim\mu}\big[p_{P_\star}(\psi)^2\big] + \lambda_{\max}(Y_{P_\star}) \\ &\geq \left(1 - \frac{1}{m}\right)\mathbb{E}_{|\psi\rangle\sim\mu}\big[p_{P_\star}(\psi)^2\big] + \frac{1}{m}\mathbb{E}_{|\psi\rangle\sim\mu}[p_{P_\star}(\psi)]. \end{aligned} \tag{84}$$

By Jensen's inequality,

$$\mathbb{E}_{|\psi\rangle\sim\mu}\big[p_{P_\star}(\psi)^2\big] \geq \big(\mathbb{E}_{|\psi\rangle\sim\mu}[p_{P_\star}(\psi)]\big)^2 = s_m^2.$$

Substituting this and $\mathbb{E}[p_{P_\star}(\psi)] = s_m$ into Eq. (84) yields

$$B_\mu \geq \left(1 - \frac{1}{m}\right)s_m^2 + \frac{s_m}{m}.$$

Combining Steps 1–4, we obtain

$$\frac{F^\star_{\mathrm{iso}}(\mu)}{F^\star(\mu)} \geq \frac{B_\mu}{U_\mu} \geq \frac{B_\mu}{s_m} \geq \frac{\left(1-\frac{1}{m}\right)s_m^2 + \frac{s_m}{m}}{s_m} = \left(1 - \frac{1}{m}\right)s_m + \frac{1}{m} = 1 - \left(1 - \frac{1}{m}\right)\eta_m.$$

□

Remark H.3. *The converse $F^\star(\mu) \leq s_m$ is optimal as a function of $\bar{\rho}$ alone. Indeed, if*

$$\bar{\rho} = \sum_{i=1}^{d} \lambda_i(\bar{\rho})|i\rangle\langle i|$$

in an orthonormal eigenbasis and μ is the orthogonal eigenensemble that draws $|i\rangle$ with probability $\lambda_i(\bar{\rho})$, then a blind encoder can transmit the states $|1\rangle, \ldots, |m\rangle$ perfectly through the m-dimensional bottleneck and map every $|i\rangle$ with $i > m$ to a fixed latent basis state. With the obvious decoder, the resulting average fidelity is $\sum_{i=1}^{m}\lambda_i(\bar{\rho}) = s_m$. Thus any further improvement of the bound

$$\frac{F^\star_{\mathrm{iso}}(\mu)}{F^\star(\mu)} \geq 1 - \left(1 - \frac{1}{m}\right)\eta_m$$

must use information beyond the spectrum of $\bar{\rho}$ or must strengthen the lower bound on $F^\star_{\mathrm{iso}}(\mu)$.

I MNIST data preparation

Objects and dimensions. Each (rescaled) MNIST image is represented as a real matrix

$$x \in [0,1]^{28\times 28}.$$

Fix integers $n > k \geq 1$, and let $d = 2^n < 784$ and $m = 2^k$. We map each image to a pure n-qubit state represented by a complex vector

$$\psi \in \mathbb{C}^d, \qquad \|\psi\|_2 = 1.$$

Our goal is to concentrate most of the probability mass $\sum_{j=1}^{d}|\psi_j|^2$ into the first m entries. This data preparation scheme is designed to facilitate near-perfect reconstruction via a QAE with a bottleneck dimension of 2^k. By concentrating the signal power into the low-frequency subspace, we ensure that the reconstruction infidelity is small. This approach enhances the sensitivity of our benchmarks, as small variations in model performance become more discernible when the total error magnitude is low.

Unitary-scaled 2D FFT feature map. We compute a 2D discrete Fourier transform of the image with scaling [Cooley and Tukey, 1965]:

$$F := \frac{1}{\sqrt{784}}\texttt{FFT2}(x) \ \in \mathbb{C}^{28\times 28}.$$

We then apply an FFT shift

$$\tilde{F} := \texttt{fftshift}(F),$$

which permutes frequency bins so that the DC (zero-frequency) component is moved to the center of the 28×28 grid.

Low-frequency-first ordering. Let the center index be $(c_y, c_x) = (14, 14)$. For each frequency bin $(y, x) \in \{0, \ldots, 27\}^2$, define the squared radius from the center

$$d(y, x)^2 := (y - c_y)^2 + (x - c_x)^2.$$

We define an ordering (permutation) π of the 784 grid points by sorting bins in increasing order of $d(y, x)^2$ (with stable tie-breaking). Flattening $\tilde{F}$ into a length-784 vector and reordering by π yields a sequence

$$(a_1, a_2, \ldots, a_{784}) \in \mathbb{C}^{784},$$

where a_1 is the DC component, and increasing indices correspond to increasing spatial frequency. This low-frequency-first rule reflects the empirical fact that low spatial frequencies often capture most of the reconstruction-relevant content of images. We retain only the first d coefficients $(a_1, \ldots, a_d)$. This discards higher-frequency coefficients (lossy truncation).

Head/tail placement and energy allocation. Split the retained coefficients into a "head" and "tail":

$$u := (a_1, \ldots, a_m) \in \mathbb{C}^m, \qquad v := (a_{m+1}, \ldots, a_d) \in \mathbb{C}^{d-m}.$$

We build an unnormalized state $\widehat{\psi} \in \mathbb{C}^d$ by placing the head coefficients in the first m amplitudes and the tail coefficients next (optionally rescaled):

$$\widehat{\psi}_j = \begin{cases} u_j, & 1 \le j \le m, \\ \alpha\, v_{j-m}, & m+1 \le j \le d. \end{cases}$$

Given a target head energy fraction $e \in (0, 1)$, we choose $\alpha \ge 0$ so that, after normalization, the fraction of probability mass in the head equals e. Specifically, we choose $e = 0.9$ and set

$$\alpha = \sqrt{\frac{1-e}{e}}\sqrt{\frac{\|u\|^2}{\|v\|^2}}.$$

Normalization and output state. Finally we normalize to obtain a valid quantum state

$$\psi := \frac{\widehat{\psi}}{\|\widehat{\psi}\|_2}.$$

In non-degenerate cases, the construction ensures $\sum_{j=1}^{m} |\psi_j|^2 = e$, so a projective measurement onto the head subspace (spanned by the first 2^k computational basis vectors) succeeds with probability e. Moreover, because the head entries store the lowest-frequency Fourier coefficients, truncating to the first 2^k amplitudes corresponds to keeping a low-frequency approximation of the image, which is a simple proxy for preserving reconstruction-relevant information.

J Training time

Table 2: Training time (in seconds) across different configurations. For $\mu_{1,0.1}$, cfg1–cfg3 denote $(n, k, 0, n-k)$, $(n, k, 0, n)$, and (n, k, k, n)-QAEs, respectively. For μ_2, they denote $(n, k, 0, n-k)$, $(n, k, k, n-k)$, and $(n, k, k, n-k+1)$-QAEs.

Dataset	n	k	cfg1	cfg2	cfg3
$\mu_{1,0.1}$	2	1	129	102	115
	3	1	108	121	126
	3	2	105	127	145
	4	1	121	134	146
	4	2	121	152	182
	4	3	121	198	276
μ_2	2	1	154	119	122
	3	1	121	126	128
	3	2	120	134	136
	4	1	126	137	148
	4	2	126	156	166
	4	3	126	194	204

K Additional numerical experiments

We report the additional experiments in Section 4.4. Each table compares $(n, k, 0, n-k)$-, $(n, k, k, n-k)$-, and $(n, k, k, n-k+1)$-QAEs for $n \in \{2, 3, 4, 5\}$ and $1 \leq k < n$. The middle architecture has an isometric decoder, and the final architecture permits a decoder of Kraus rank at most two.

For an evaluation ensemble $\{|\psi_j\rangle\}_{j=1}^{N_{\mathrm{ev}}}$ and a reconstructed state ρ_j^{rec}, define the empirical average infidelity by

$$\widehat{L}_{\mathrm{inf}} := \frac{1}{N_{\mathrm{ev}}} \sum_{j=1}^{N_{\mathrm{ev}}} \left(1 - \langle\psi_j|\rho_j^{\mathrm{rec}}|\psi_j\rangle\right). \tag{85}$$

The first three tables report this metric. The fourth reports average trace distance.

K.1 Increased-tail MNIST

The preparation in Appendix I allocates a prescribed mass to a fixed coordinate subspace. Write $d = 2^n$, $m = 2^k$, and $\tau := 1 - e$. The original preparation uses $\tau = 0.1$. Increasing this allocation to $\tau = 0.2$ corresponds to $e = 0.8$. Table 3 tests sensitivity to the larger-tail preparation.

Table 3: Final average infidelity for the increased-tail MNIST setting.

		Ancilla counts (n_B, n_E)		
n	k	$(0, n-k)$	$(k, n-k)$	$(k, n-k+1)$
2	1	0.002000	0.001903	0.001904
3	1	0.025293	0.022532	0.022548
3	2	0.022232	0.016264	0.016277
4	1	0.057153	0.054477	0.054491
4	2	0.053883	0.044734	0.044748
4	3	0.022449	0.015791	0.015802
5	1	0.090231	0.087808	0.087839
5	2	0.098216	0.084887	0.084977
5	3	0.077549	0.060488	0.060556
5	4	0.027539	0.021319	0.021364

K.2 Coherently perturbed GHZ states

Let

$$|\mathrm{GHZ}_n\rangle := \frac{|0\rangle^{\otimes n} + |1\rangle^{\otimes n}}{\sqrt{2}}, \qquad |G_{n,\xi}\rangle := U_{n,\xi}|\mathrm{GHZ}_n\rangle, \qquad U_{n,\xi} \sim \nu_{\mathrm{pert},n}, \tag{86}$$

where $\nu_{\mathrm{pert},n}$ denotes the distribution of coherent perturbation unitaries used to prepare the samples. The resulting source distribution is the pushforward of $\nu_{\mathrm{pert},n}$ under $U \mapsto U|\mathrm{GHZ}_n\rangle$. Specifically, we set

$$U_{n,\xi} = \bigotimes_{i=1}^{n} U_i, \quad U_i = \exp\left[-\frac{i\theta_i}{2}(r_{i,x}X + r_{i,y}Y + r_{i,z}Z)\right],$$
$$\mathbf{r}_i = \frac{\mathbf{g}_i}{\|\mathbf{g}_i\|_2}, \quad \mathbf{g}_i \sim \mathcal{N}(\mathbf{0}, I_3), \quad \theta_i \sim \mathcal{N}(0, 0.3^2).$$

Table 4 provides the results.

Table 4: Final average infidelity for the coherently perturbed GHZ source of Eq. (86).

		Ancilla counts (n_B, n_E)		
n	k	$(0, n-k)$	$(k, n-k)$	$(k, n-k+1)$
2	1	0.005028	0.003407	0.003337
3	1	0.007157	0.004795	0.004803
3	2	0.005545	0.003195	0.003214
4	1	0.007575	0.006745	0.006759
4	2	0.007603	0.005081	0.005084
4	3	0.002508	0.001718	0.001728
5	1	0.008179	0.008176	0.008196
5	2	0.007452	0.006587	0.006605
5	3	0.004200	0.003324	0.003374
5	4	0.000046	0.000036	0.000037

K.3 Output depolarization

Let $\mu_{\mathrm{dep}} = \mu_2$ denote the MNIST source used in this experiment. For $\rho_\psi^{\mathrm{rec}} := (\mathcal{D} \circ \mathcal{E})(|\psi\rangle\langle\psi|)$, output noise acts as

$$\rho_{\psi,p}^{\mathrm{out}} := \mathcal{N}_p^{\otimes n}(\rho_\psi^{\mathrm{rec}}), \qquad p = 0.01, \tag{87}$$

where

$$\mathcal{N}_p(X) = (1-p)X + p\,\mathrm{Tr}(X)\frac{I_2}{2} \tag{88}$$

for a single-qubit operator X. The evaluation loss is

$$L_{\mathrm{dep}}(\mathcal{E}, \mathcal{D}) := 1 - \mathbb{E}_{\psi \sim \mu_{\mathrm{dep}}} \langle\psi|\rho_{\psi,p}^{\mathrm{out}}|\psi\rangle. \tag{89}$$

The target remains the original pure source state. Table 5 measures reconstruction after this output channel.

K.4 Trace distance

Let $\mu_{\mathrm{tr}} = \mu_2$ denote the MNIST source used in this experiment. Define

$$L_{\mathrm{tr},\mu_{\mathrm{tr}}}(\mathcal{E}, \mathcal{D}) := \mathbb{E}_{\psi \sim \mu_{\mathrm{tr}}} \frac{1}{2} \left\| (\mathcal{D} \circ \mathcal{E})(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi| \right\|_1, \tag{90}$$

where $\|X\|_1 := \mathrm{Tr}\sqrt{X^\dagger X}$. The empirical evaluation is the average of the individual trace distances:

$$\widehat{L}_{\mathrm{tr}} := \frac{1}{N_{\mathrm{ev}}} \sum_{j=1}^{N_{\mathrm{ev}}} \frac{1}{2} \|\rho_j^{\mathrm{rec}} - |\psi_j\rangle\langle\psi_j|\|_1. \tag{91}$$

Table 6 examines the architecture comparison under this alternative metric. The relationship to infidelity is explained in Appendix L.

Table 5: Final average infidelity after independent single-qubit output depolarization.

		Ancilla counts (n_B, n_E)		
n	k	$(0, n-k)$	$(k, n-k)$	$(k, n-k+1)$
2	1	0.012446	0.012387	0.012388
3	1	0.029223	0.027760	0.027769
3	2	0.035319	0.030458	0.030459
4	1	0.049466	0.047943	0.047964
4	2	0.054067	0.048149	0.048167
4	3	0.041732	0.036295	0.036307
5	1	0.070251	0.068832	0.068845
5	2	0.080682	0.072489	0.072538
5	3	0.073797	0.062757	0.062785
5	4	0.050734	0.046040	0.046065

Table 6: Final average trace distance using the normalization $\frac{1}{2}\|\rho - \sigma\|_1$.

		Ancilla counts (n_B, n_E)		
n	k	$(0, n-k)$	$(k, n-k)$	$(k, n-k+1)$
2	1	0.024574	0.024565	0.024562
3	1	0.100797	0.100062	0.100103
3	2	0.100240	0.097757	0.097708
4	1	0.159591	0.158924	0.158906
4	2	0.150254	0.147490	0.147421
4	3	0.097383	0.094570	0.094630
5	1	0.206730	0.206119	0.206158
5	2	0.207809	0.204181	0.204241
5	3	0.176198	0.170996	0.171051
5	4	0.102333	0.099928	0.100112

L Quantitative trace-distance transfer

L.1 Pointwise and average trace-distance bounds

Set

$$\rho_\psi := (\mathcal{D} \circ \mathcal{E})(|\psi\rangle\langle\psi|), \qquad \ell_{\inf}(\psi) := 1 - \langle\psi|\rho_\psi|\psi\rangle, \qquad T(\rho, \sigma) := \frac{1}{2}\|\rho - \sigma\|_1.$$

The upper bound below is the Fuchs–van de Graaf bound [Fuchs and Van De Graaf, 1999], and the lower bound follows directly from the variational characterization of the trace distance [Nielsen and Chuang, 2010]:

$$\ell_{\inf}(\psi) \le T(\rho_\psi, |\psi\rangle\langle\psi|) \le \sqrt{\ell_{\inf}(\psi)}. \tag{92}$$

Let $L_{\inf,\mu} := L_\mu$ and $L_{\mathrm{tr},\mu} := \mathbb{E}_{\psi\sim\mu} T(\rho_\psi, |\psi\rangle\langle\psi|)$. Averaging Eq. (92) and applying Jensen's inequality yields

$$L_{\inf,\mu}(\mathcal{E}, \mathcal{D}) \le L_{\mathrm{tr},\mu}(\mathcal{E}, \mathcal{D}) \le \sqrt{L_{\inf,\mu}(\mathcal{E}, \mathcal{D})}. \tag{93}$$

Equation (93) also holds for empirical averages evaluated on the same samples and the same reconstructed states.

L.2 Transfer to the width-constrained architecture

Let $\mathfrak{C}$ be the set of all CPTP encoder–decoder pairs through the bottleneck, and let $\mathfrak{G} \subseteq \mathfrak{C}$ be the pairs realizable by (n, k, k, n)-QAEs. For $a \in \{\inf, \mathrm{tr}\}$ and a class $\mathfrak{A}$, write

$$L^\star_{a,\mathfrak{A}} := \min_{(\mathcal{E},\mathcal{D})\in\mathfrak{A}} L_{a,\mu}(\mathcal{E}, \mathcal{D}).$$

Theorem 3.4 supplies a pair $(\mathcal{E}_F, \mathcal{D}_F) \in \mathfrak{G}$ satisfying $L_{\mathrm{inf},\mu}(\mathcal{E}_F, \mathcal{D}_F) = L^\star_{\mathrm{inf},\mathfrak{C}}$. Consequently,

$$\begin{aligned} L^\star_{\mathrm{tr},\mathfrak{C}} &\le L^\star_{\mathrm{tr},\mathfrak{G}} \le L_{\mathrm{tr},\mu}(\mathcal{E}_F, \mathcal{D}_F) \\ &\le \sqrt{L^\star_{\mathrm{inf},\mathfrak{C}}} \le \sqrt{L^\star_{\mathrm{tr},\mathfrak{C}}}. \end{aligned} \tag{94}$$

Thus the infidelity-optimal architecture provides a quantitative trace-distance guarantee, but the inequalities do not establish $L^\star_{\mathrm{tr},\mathfrak{G}} = L^\star_{\mathrm{tr},\mathfrak{C}}$. An approximately infidelity-optimal trained pair satisfying

$$L_{\mathrm{inf},\mu}(\mathcal{E}, \mathcal{D}) \le L^\star_{\mathrm{inf},\mathfrak{C}} + \delta$$

also obeys

$$L_{\mathrm{tr},\mu}(\mathcal{E}, \mathcal{D}) \le \sqrt{L^\star_{\mathrm{inf},\mathfrak{C}} + \delta} \le \sqrt{L^\star_{\mathrm{tr},\mathfrak{C}} + \delta}.$$

L.3 A source-dependent bound with an isometric decoder

Let $\mathfrak{G}_{\mathrm{iso}}$ denote the $(n, k, k, n-k)$-QAE class. The construction used in Theorem H.2 belongs to this class and achieves at least

$$b_m = \left(1 - \frac{1}{m}\right)(1 - \eta_m)^2 + \frac{1 - \eta_m}{m}, \qquad m = 2^k.$$

Applying Eq. (93) to that construction gives

$$L^\star_{\mathrm{tr},\mathfrak{G}_{\mathrm{iso}}} \le \sqrt{1 - b_m} = \sqrt{\left(2 - \frac{1}{m}\right)\eta_m - \left(1 - \frac{1}{m}\right)\eta_m^2}. \tag{95}$$

NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: This work studies blind single-copy compression of pure n-qubit states under average infidelity, proves universal sufficiency of the (n, k, k, n)-QAE architecture, establishes a sharp encoder-side ancilla threshold, and presents an explicit counterexample showing that isometric decoders are not universally sufficient. The claimed scope is also described appropriately, because the paper explicitly limits the theory to the infidelity objective and notes that a general sharp decoder threshold for arbitrary k remains open.

   Guidelines:

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

2. **Limitations**

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

   Answer: [Yes]

   Justification: The paper explicitly discusses its limitations in the Discussion section by stating that the analysis is restricted to blind single-copy compression of pure n-qubit states under the average infidelity objective, and that it does not address additional objectives such as latent-space regularization.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

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

   Answer: [Yes]

   Justification: The Main results section specifies the source model, bottleneck size, ancilla counts, and channel classes. We provide short proof sketches in the main text together with full proofs in the appendices.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The paper specifies the experimental setups that support its main claims, including which QAE architectures are compared, the source distributions used, the training/test sample sizes, the number of epochs, etc. The paper also gives the MNIST state-preparation pipeline in the appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

(c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The code for model training and evaluation is included in the supplemental material (qae_submit.zip).

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.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://neurips.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, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The Experiments section states the concrete experimental details needed to understand the results, including the train/test split, number of epochs, optimizer, learning rate, batch size, and hardware.

   Guidelines:

   - The answer [N/A] 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: [Yes]

Justification: The main experimental figures report uncertainty bands explicitly as 2-sigma confidence intervals.

Guidelines:

- The answer [N/A] 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).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- 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.
- 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).
- 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.

8. **Experiments compute resources**

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

   Answer: [Yes]

   Justification: The Setup section reports the hardware used for the experiments, and the Training time section in the appendix reports training times across the reported dataset/configuration settings. This provides the type of compute workers, memory, and execution-time information.

   Guidelines:

   - The answer [N/A] 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).

9. **Code of ethics**

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

   Answer: [Yes]

   Justification: The paper reports standard theoretical and numerical research on quantum-state compression, preserves submission anonymity, and does not involve human subjects, crowdsourcing, or high-risk model/data release.

   Guidelines:

   - The answer [N/A] 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).

10. **Broader impacts**

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

    Answer: [N/A]

    Justification: There is no societal impact of the work performed.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] 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.
    - 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).

11. **Safeguards**

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

    Answer: [N/A]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer [N/A] 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.

12. **Licenses for existing assets**

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

Answer: [Yes]

Justification: We use the MNIST handwritten digit dataset and credit its creators in the paper. We also state that MNIST (version 3.0.1) is available under the CC BY-SA 3.0 license and use the standard dataset for benchmarking.

Guidelines:

- The answer [N/A] 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.
- 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.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

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

Answer: [N/A]

Justification: The paper does not release new assets.

Guidelines:

- The answer [N/A] 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.

14. **Crowdsourcing and research with human subjects**

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

Answer: [N/A]

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

Guidelines:

- The answer [N/A] 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.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

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

Answer: [N/A]

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

Guidelines:

- The answer [N/A] 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.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.