

No universal purification in quantum mechanics

I. INTRODUCTION

The purification of noisy quantum systems, namely the reduction of quantum entropy in order to prepare a target pure system, is a fundamental primitive in quantum information science with both broad practical applications and deep theoretical significance. A standard and widely studied purification paradigm is to consume many copies of a noisy state or channel and concentrate their purity onto a single output copy [1–3]. On the practical side, such purification processes underlie a wide range of tasks, including quantum error correction and a number of variants of coherent quantum error mitigation [4–7], quantum cooling [8–10], and various resource distillation and concentration protocols [1, 11–13]. At the same time, purification also serves as a central conceptual framework in several foundational areas, including quantum Shannon theory [14], resource theories [15–17], and quantum thermodynamics [18]. Here, the number of input copies represents the cost of producing a pure output system, thus naturally characterizes the complexity of the task. A recurring theme in the literature is about analyzing the fundamental limitation of the number of input copies.

According to the relationship between the input and output of the purification process, purification tasks can be broadly divided into two categories. The first category consists of tasks whose target output pure state or unitary is fixed and independent of the input. Prominent examples include entanglement distillation [1] and magic state distillation [12], where the usual target outputs are respectively a Bell state and a single-qubit T state. Limitations for such tasks originate from restrictions on the allowed physical maps, commonly referred to as free operations [13, 17, 19–21], without which the task is trivialized by replacement maps that always output the desired target.

The second and more fundamental category consists of input-dependent purification tasks, i.e., universal purification tasks, where the target output pure state or unitary itself depends on the input state or channel. Key examples include quantum purity amplification [22, 23] and quantum princi-

pal component analysis [24]. Beyond these examples, input-dependent purification also fundamentally underlies notions of quantum learning and estimation. In particular, the preparation of a pure dilation state¹ of a mixed state ρ using many copies of it, namely a global pure state whose reduced density matrix equals ρ , has important applications in learning nonlinear quantum properties [25–28]. More generally, estimating physical properties of ρ from many copies can itself be viewed as an input-dependent purification task, where the target pure output encodes the correct classical prediction of the target property. Unlike the first category, the limitations of input-dependent purification tasks persist even without imposing any restrictions on the allowed operations. In other words, the obstruction is no longer simply a consequence of resource-theoretic constraints, and its physical and information-theoretic origin remains far from fully understood. Understanding the source of these limitations is therefore of fundamental importance, both for developing practical quantum purification protocols and for clarifying the ultimate limits of quantum learning.

In this work, we unveil a new physical principle: a general no-go principle on quantum purification, which rules out the possibility of purifying any finite copies of noisy states or channels into input-dependent pure states or unitaries. This conclusion is fully protocol-independent: it applies to any physically admissible strategy, including sequential, coherent, and adaptive protocols. With a short proof, we elucidate that this impossibility originates from two fundamental mathematical features of quantum mechanics—*linearity* and *positivity*—thus uncovering a new intrinsic limitation of quantum information processing, as demonstrated in Fig. 1. Beyond its theoretical significance, this limitation directly constrains the performance of practical quantum tasks. By relaxing the requirement of producing an exactly pure state to allowing nearly pure outputs, we derive a general sample complexity lower bound of approximate purification. The lower bound meets the performance of various tasks, including quantum purity amplification and quantum principal component analysis algorithms. Importantly, the performance of approximate purification implies a standard quantum limit for estimating arbitrary functions of unknown quantum systems, greatly enhancing the former result derived using quantum Fisher information and showing the deep connection between quantum purification and learning. Beyond these general conclusions, we show that these fundamental limitations become significantly more stringent when specific operational constraints are imposed in several important tasks. For the preparation of a pure dilation state, we prove an exponential sample complexity lower bound. For the purification of bosonic Gaussian




FIG. 1. The linearity and positivity of quantum mechanics fundamentally constrain the preparation of pure outputs that depend non-trivially on noisy inputs.

¹ In this work, we use “pure dilation state” to denote pure-state extension often referred to as “purification” in the literature, whereas “purification” represents a transformation process where the pure target state can be arbitrary.

states, we show the no-go result with passive Gaussian operations, even for approximate purification.

II. NO UNIVERSAL EXACT PURIFICATION

A universal quantum purification task involves finding a quantum process, represented by $\mathcal{M}$, that can map many copies of an arbitrary noisy input system into a pure state or unitary depending on the input. At the same time, the quantum process, $\mathcal{M}$, is universal and independent of the input system. The starting point of our framework can be summarized in the following observation, which very crisply showcases the fundamental obstructions to purification.

Observation 1 (No universal exact purification). *Let t be any finite integer, and let $\mathcal{M}$ be a nonzero positive map which maps positive semi-definite matrices to positive semi-definite matrices. Suppose that for any state ρ supported on a subspace of the space of all density matrices $\mathcal{D}_S \subseteq \mathcal{D}(\mathcal{H}_d)$ with $\dim(\mathcal{D}_S) \geq 2$, the output $\mathcal{M}(\rho^{\otimes t})$ is proportional to a pure state. Then, for any $\rho \in \mathcal{D}_S$, the output of $\mathcal{M}$ must always be proportional to a fixed pure state,*

$$\mathcal{M}(\rho^{\otimes t}) \propto \Psi. \quad (1)$$

Thus, any universal exact purification protocol must be trivial: whenever the output is required to be pure for all inputs, it cannot depend non-trivially on the input state. Although the assumption of producing an exactly pure output state is idealized and not directly achievable in realistic settings, this observation identifies the fundamental physical origin of the no-purification phenomenon. Specifically, the proof relies solely on the linearity and positivity of $\mathcal{M}$, two fundamental structural principles of quantum mechanics, which also underlie all later no-go results on approximate purification derived in this work. Moreover, since we do not require $\mathcal{M}$ to be trace-preserving, the result applies to all quantum operations, including post-selection and adaptive strategies.

Extending beyond quantum states, our no-purification result can be naturally generalized to quantum channels and higher-order maps [29]. In particular, there is no universal procedure that transforms t copies of an unknown input object into an output object with a pure Choi representation that depends non-trivially on the input. This excludes, for example, universal transformations from quantum channels to unitary channels, from quantum states to unitary channels, and from quantum channels to pure states. Moreover, the conclusion is independent of the specific protocol structure: it applies to arbitrary physically allowed strategies, including sequential, coherent, and adaptive protocols. The key reason is that any physically admissible supermap acting on channels can be represented, at the level of Choi operators, as a completely positive trace-preserving map acting on the corresponding Choi states [30–32]. Since the Choi state of a unitary channel is pure, any such hypothetical transformation would induce a quantum channel mapping input Choi states to input-dependent pure states, directly contradicting Observation 1.

III. APPROXIMATE PURIFICATION

In the above, we have addressed the fundamental restriction on perfect quantum purification, whose targets are pure states or unitaries. Nonetheless, approximate pure systems are often already useful in many tasks, thus the characterization of an approximate version of the principle is crucial for practical purposes. We analyzed the robustness of Observation 1, showing a clear trade-off relationship. Specifically, for any positive trace-preserving map $\mathcal{M}$ with a fixed t , if its output state $\mathcal{M}(\rho^{\otimes t})$ is close to a pure state, then the difference between its output states will be small. When the purification task demands the output state to vary significantly according to the input state, then one requires a large number of input states to obtain high purity.

Theorem 1 (Approximate purification). *Let $\mathcal{M}$ be a positive trace-preserving map. Suppose that there exists a parameter $0 \leq \epsilon < 1$ such that for all ρ supported on a subspace $\mathcal{D}_S \subseteq \mathcal{D}(\mathcal{H}_d)$ with $\dim(\mathcal{D}_S) \geq 2$*

$$\text{Tr}[\mathcal{M}(\rho^{\otimes t})^2] \geq 1 - \epsilon. \quad (2)$$

Then there exists a state σ_0 such that, for all $\rho \in \mathcal{D}_S$,

$$D(\sigma_0, \mathcal{M}(\rho^{\otimes t})) \leq (8\pi\sqrt{2t} + 12)\sqrt{\epsilon}. \quad (3)$$

Here $D(\cdot, \cdot)$ denotes the trace distance $D(\omega, \tau) = \frac{1}{2}\|\omega - \tau\|_1$.

By the same Choi-state reduction, the approximate purification bound in Theorem 1 also extends to all channel-involving settings and applies to arbitrary physically allowed strategies. Specifically, if one aims to find a universal quantum protocol that approximately transforms t copies of quantum channel into a unitary or pure quantum state, or transforms t copies of quantum states into a unitary, similar trade-off exists.

The fundamental limitation of approximate universal purification has several important consequences. In particular, whenever the outputs corresponding to two different input states remain distinguishable, namely $D(\mathcal{M}(\rho_1^{\otimes t}), \mathcal{M}(\rho_2^{\otimes t})) = \Omega(1)$, Theorem 1 immediately implies the sample complexity lower bound

$$t = \Omega(\epsilon^{-1}). \quad (4)$$

Interestingly, the ϵ^{-1} scaling appears universally across a broad range of quantum information processing tasks. Examples include quantum purity amplification [22, 23], which aims to prepare the eigenstate of ρ corresponding to its largest eigenvalue, and quantum principal component analysis [24], which aims to transform many copies of ρ into the unitary evolution $e^{-i\rho}$. Notably, the same ϵ^{-1} behavior also appears in covariant codes [33, 34] and quantum channel purification [35], which aim to prepare the ideal unitary channel using several noisy ones. Together, these important examples suggest that approximate universal purification provides a common information-theoretic framework underlying several previously unrelated complexity phenomena in quantum information processing.

Another fundamental consequence of approximate purification arises in quantum learning and estimation. Indeed, estimating a property of an unknown quantum state ρ , represented by a function $f(\rho)$, from many copies of ρ can itself be viewed as a universal purification task.

Corollary 1 (Standard quantum limit for arbitrary function). *Given t copies of unknown quantum state ρ , which is supported on a subspace $\mathcal{D}_S \subseteq \mathcal{D}(\mathcal{H}_d)$ with $\dim(\mathcal{D}_S) \geq 2$, any target function $f(\rho)$, and an estimator $\hat{f}$ constructed by any measurement performed on t copies of ρ with*

$$\sup_{\rho} |f(\rho) - \mathbb{E}[\hat{f}]| \leq \frac{1}{10} \left[\sup_{\rho} f(\rho) - \inf_{\rho} f(\rho) \right] = \frac{1}{10} \Delta_f, \quad (5)$$

then the variance of $\hat{f}$ is lower bounded by

$$\sup_{\rho} \text{Var}(\hat{f}) = \Omega(\Delta_f^2 t^{-1}). \quad (6)$$

This result also directly generalizes to the estimation of channel properties and substantially generalizes previous derivations in quantum metrology based on the additivity of quantum Fisher information [36, 37], as it does not require the target function to satisfy any smoothness or differentiability assumptions and applies to global estimation. It shows that the standard quantum limit is an extremely fundamental principle for quantum learning, while also demonstrating the tightness of the lower bound established in Theorem 1.

IV. CASE STUDIES

A. Approximate pure dilation state preparation

We define $\Psi_{A,B}$ as the pure dilation state of a mixed state ρ so that $\text{Tr}_B(\Psi_{A,B}) = \rho$. Besides being a fundamental concept, the pure dilation state also carries broad practical significance in areas such as quantum algorithms, learning, and simulation. By definition, this task aims to find a quantum operation, $\mathcal{M}$, that acts on $\rho^{\otimes t}$ and the output state is some pure dilation state satisfying two requirements: (1) the output is a pure state; (2) the marginal of the output state is ρ . Theorem 1 shows that it is impossible to simultaneously achieve these two requirements, but in practice, it is also valuable to understand if we can do so up to some approximation.

Theorem 2. *Let $\mathcal{M} : \mathcal{D}(\mathcal{H}_A^{\otimes t}) \rightarrow \mathcal{D}(\mathcal{H}_{A,B})$ be a positive trace-preserving map with $d \geq C$ for a sufficiently large absolute constant C . Suppose there exist parameters $\epsilon, \eta \geq 0$ such that $\xi := 2(\sqrt{\epsilon} + \epsilon + \eta) < \frac{1}{3}$ and, for arbitrary ρ ,*

$$(1) \text{ (Near purity) } \text{Tr}[\mathcal{M}(\rho^{\otimes t})^2] \geq 1 - \epsilon,$$

$$(2) \text{ (Marginal accuracy) } D(\text{Tr}_B[\mathcal{M}(\rho^{\otimes t})], \rho) \leq \eta.$$

Then

$$t \geq \frac{ed}{128 \left[\xi + 13 \exp\left(-\frac{d}{32(\ln d)^2}\right) \right]^2}. \quad (7)$$

For qubit systems, the linear scaling with d in Eq. (7) implies that the sample complexity t scales exponentially with the system size n (where $d = 2^n$).

B. Gaussian purification

We investigate purification using passive Gaussian quantum channels [38–40]. Because these operations arise naturally in linear optical implementations [38], they provide a standard framework for studying bosonic quantum systems and quantum optical platforms. We will show that, when restricting to only these operations, even approximate universal purification encounters fundamental limitations.

Theorem 3 (No Gaussian universal approximate purification). *Let $\mathcal{M} : \mathcal{D}(\mathcal{H}^{\otimes t}) \rightarrow \mathcal{D}(\mathcal{H})$ be a passive Gaussian channel. Suppose that there exists a parameter $0 \leq \epsilon < 1$ such that for all Gaussian states ρ contained in a subspace $\mathcal{G}$ of Gaussian states with covariance matrices that are larger than $\mathbb{I}$ in matrix ordering, we have that*

$$\text{Tr}[\mathcal{M}(\rho^{\otimes t})^2] \geq 1 - \epsilon. \quad (8)$$

Then there exists a Gaussian state Ψ such that, for all Gaussian states $\rho \in \mathcal{G}$,

$$D(\mathcal{M}(\rho^{\otimes t}), \Psi) \leq 4(1 + \sqrt{3})\epsilon(1 + 4\epsilon). \quad (9)$$

Comparing this result with Theorem 1 or Theorem 2, we can notice the main difference: the r.h.s. of Eq. (9) is independent of t , indicating that coherent state manipulation does not bring an advantage. This is to an extent reminiscent of earlier findings in a different context that show that coherent state manipulation of Gaussian states finds obstructions [41–44]. This leads to an important consequence: passive Gaussian operation cannot perform an arbitrary input-dependent purification task, even approximately.

V. DISCUSSION AND OUTLOOK

Identifying the fundamental no-go principles of quantum information processing has been one of the central and defining goals of quantum information science. A common paradigm for establishing such results is to first prove an exact impossibility theorem in an idealized setting, then study the achievable performance and resource requirements in the approximate regime, and finally explore the implications in specific scenarios of practical relevance. In this work, we have completed this program for the fundamental problem of universal purification. Observation 1 establishes the no-go theorem in the exact setting, while Theorem 1 characterizes its asymptotic behavior in the approximate regime. Furthermore, Corollary 1 and Theorem 2 reveal deep connections between approximate universal purification and quantum learning. Finally, Observation ?? and Theorem 3 show that, under physically motivated restrictions on the allowed operations, namely passive Gaussian operations, universal purification can remain fundamentally impossible even in the approximate regime.

-
- [1] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Mixed-state entanglement and quantum error correction, *Phys. Rev. A* **54**, 3824 (1996).
 - [2] J. I. Cirac, A. K. Ekert, and C. Macchiavello, Optimal purification of single qubits, *Phys. Rev. Lett.* **82**, 4344 (1999).
 - [3] J. Fiurásek, Optimal probabilistic cloning and purification of quantum states, *Phys. Rev. A* **70**, 032308 (2004).
 - [4] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [5] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, *Rev. Mod. Phys.* **95**, 045005 (2023).
 - [6] W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, Virtual distillation for quantum error mitigation, *Phys. Rev. X* **11**, 041036 (2021).
 - [7] B. Koczor, Exponential error suppression for near-term quantum devices, *Phys. Rev. X* **11**, 031057 (2021).
 - [8] J. Cotler, S. Choi, A. Lukin, H. Gharibyan, T. Grover, M. E. Tai, M. Rispoli, R. Schittko, P. M. Preiss, A. M. Kaufman, M. Greiner, H. Pichler, and P. Hayden, Quantum virtual cooling, *Phys. Rev. X* **9**, 031013 (2019).
 - [9] P. Zeng, J. Sun, and X. Yuan, Universal quantum algorithmic cooling on a quantum computer, *arXiv preprint arXiv:2109.15304* (2021).
 - [10] D. An, J.-P. Liu, and L. Lin, Linear combination of hamiltonian simulation for nonunitary dynamics with optimal state preparation cost, *Phys. Rev. Lett.* **131**, 150603 (2023).
 - [11] J.-W. Pan, C. Simon, C. Brukner, and A. Zeilinger, Entanglement purification for quantum communication, *Nature* **410**, 1067 (2001).
 - [12] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, *Phys. Rev. A* **71**, 022316 (2005).
 - [13] K. Fang and Z.-W. Liu, No-go theorems for quantum resource purification, *Phys. Rev. Lett.* **125**, 060405 (2020).
 - [14] M. M. Wilde, From classical to quantum Shannon theory, *arXiv preprint arXiv:1106.1445* (2011).
 - [15] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Quantum entanglement, *Rev. Mod. Phys.* **81**, 865 (2009).
 - [16] E. Chitambar and G. Gour, Quantum resource theories, *Rev. Mod. Phys.* **91**, 025001 (2019).
 - [17] Z.-W. Liu, K. Bu, and R. Takagi, One-shot operational quantum resource theory, *Phys. Rev. Lett.* **123**, 020401 (2019).
 - [18] G. Gour, M. P. Müller, V. Narasimhachar, R. W. Spekkens, and N. Yunger Halpern, The resource theory of informational nonequilibrium in thermodynamics, *Physics Rep.* **583**, 1–58 (2015).
 - [19] K. Fang and Z.-W. Liu, No-go theorems for quantum resource purification: New approach and channel theory, *PRX Quantum* **3**, 010337 (2022).
 - [20] B. Regula and R. Takagi, Fundamental limitations on distillation of quantum channel resources, *Nature Comm.* **12** (2021).
 - [21] A. Zang, X. Chen, E. Chitambar, M. Suchara, and T. Zhong, No-go theorems for universal entanglement purification, *Phys. Rev. Lett.* **134**, 190803 (2025).
 - [22] Z. Li, H. Fu, T. Isogawa, and I. Chuang, Optimal quantum purity amplification, *arXiv preprint arXiv:2409.18167* (2024).
 - [23] A. M. Childs, H. Fu, D. Leung, Z. Li, M. Ozols, and V. Vyas, Streaming quantum state purification, *Quantum* **9**, 1603 (2025).
 - [24] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, *Nat. Phys.* **10**, 631 (2014).
 - [25] A. Gilyén and T. Li, Distributional property testing in a quantum world, *arXiv preprint arXiv:1902.00814* (2019).
 - [26] Q. Wang and Z. Zhang, Quantum lower bounds by sample-to-query lifting, *arXiv preprint arXiv:2308.01794* (2023).
 - [27] Z. Liu, W. Gong, Z. Du, and Z. Cai, Exponential separations between quantum learning with and without purification, *arXiv preprint arXiv:2410.17718* (2024).
 - [28] K. Chen, Q. Wang, and Z. Zhang, Local test for unitarily invariant properties of bipartite quantum states, *arXiv preprint arXiv:2404.04599* (2024).
 - [29] Y. Liu and X. Yuan, Operational resource theory of quantum channels, *Phys. Rev. Res.* **2**, 012035 (2020).
 - [30] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Transforming quantum operations: Quantum supermaps, *Europhys. Lett.* **83**, 30004 (2008).
 - [31] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Quantum circuit architecture, *Phys. Rev. Lett.* **101**, 060401 (2008).
 - [32] Z. Liu, Y. Xiao, and Z. Cai, Non-markovian noise suppression simplified through channel representation, *arXiv preprint arXiv:2412.11220* (2024).
 - [33] P. Faist, S. Nezami, V. V. Albert, G. Salton, F. Pastawski, P. Hayden, and J. Preskill, Continuous symmetries and approximate quantum error correction, *Phys. Rev. X* **10**, 041018 (2020).
 - [34] L. Kong and Z.-W. Liu, Near-optimal covariant quantum error-correcting codes from random unitaries with symmetries, *PRX Quantum* **3**, 020314 (2022).
 - [35] R. Niwa, S. Yoshida, K. Ono, T. Utsumi, Z. Li, Y. Yang, R. Takagi, and M. Mura, Scaling-optimal purification of noisy qubit unitary channels, *arXiv preprint arXiv:2606.12394* (2026).
 - [36] S. L. Braunstein and C. M. Caves, Statistical distance and the geometry of quantum states, *Phys. Rev. Lett.* **72**, 3439 (1994).
 - [37] M. Tsang, F. Albarelli, and A. Datta, Quantum semiparametric estimation, *Phys. Rev. X* **10**, 031023 (2020).
 - [38] N. J. Cerf, G. Leuchs, and E. S. Polzik, *Quantum information with continuous variables of atoms and light* (World Scientific, 2007).
 - [39] C. Weedbrook, S. Pirandola, R. Garcia-Patron, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, Gaussian quantum information, *Rev. Mod. Phys.* **84**, 621 (2012).
 - [40] F. Caruso, J. Eisert, V. Giovannetti, and A. S. Holevo, Optimal unitary dilation for bosonic gaussian channels, *Phys. Rev. A* **84**, 022306 (2011).
 - [41] J. Eisert, S. Scheel, and M. B. Plenio, Distilling Gaussian states with Gaussian operations is impossible, *Phys. Rev. Lett.* **89**, 137903 (2002).
 - [42] G. Giedke and J. Ignacio Cirac, Characterization of Gaussian operations and distillation of Gaussian states, *Phys. Rev. A* **66**, 032316 (2002).
 - [43] J. Fiurásek, Gaussian transformations and distillation of entangled Gaussian states, *Phys. Rev. Lett.* **89**, 137904 (2002).
 - [44] G. Giedke, J. Eisert, J. I. Cirac, and M. B. Plenio, Entanglement transformations of pure gaussian states, *Quant. Inf. Comp.* **3**, 211 (2003).