


<https://doi.org/10.1038/s41534-026-01273-4>

Learning to erase quantum states: thermodynamic implications of quantum learning theory

Haimeng Zhao¹✉, Yuzhen Zhang²✉ & John Preskill^{1,3}✉

The energy cost of erasing quantum states depends on our knowledge of the states. We show that learning algorithms can acquire such knowledge to erase many copies of an unknown state at the optimal energy cost. This is proved by showing that learning can be made fully reversible and has no fundamental energy cost itself. With simple counting arguments, we relate the energy cost of erasing quantum states to their complexity, entanglement, and magic. We further show that the constructed erasure protocol is computationally efficient when learning is efficient. Conversely, under standard cryptographic assumptions, we prove that the optimal energy cost cannot be achieved efficiently in general. These results also enable efficient work extraction based on learning. Together, our results establish a concrete connection between quantum learning theory and thermodynamics, highlighting the physical significance of learning processes and enabling provably-efficient learning-based protocols for thermodynamic tasks.

Do abstract learning processes have tangible physical consequences? For example, does the (in)ability to learn impact the amount of physical resources required to perform certain tasks?

Learning is the process of acquiring information, which can be used to execute actions^{1–4}. Landauer’s principle^{5–10} illustrates how the energy cost of erasing a physical system depends on our knowledge of the system. In particular, given a system S in a state ρ with degenerate Hamiltonian $\mathcal{H} = 0$, the amount of work required to transform it to a reference pure state $|0\rangle$ in an environment at temperature T is $W \geq H(S)k_B T \ln 2$, where k_B is the Boltzmann constant and $H(S)$ is the entropy of the state. Entropy measures our ignorance of the state, and additional knowledge of the system reduces the optimal work cost to $H(S|M)k_B T \ln 2$, where $H(S|M)$ is the entropy of the system conditioned on the memory M that records the knowledge¹¹. However, this statement does not take into account the potential cost of acquiring the knowledge recorded in the memory.

To formalize this process of “learning to erase”, we consider the simple scenario in which a source repeatedly produces an unknown n -qubit state $|\psi_x\rangle$ from a class of m possible states $\mathcal{C} = \{|\psi_x\rangle\}_{x=1}^m$. At first, due to our ignorance of the state, we have to pay work to erase it. As we collect more copies of the state $|\psi_x\rangle$, we progressively learn enough information to identify the state. Once we have learned the state, we can erase additional copies without further work by reversing the state preparation unitary and uncomputing them, $U_x^\dagger |\psi_x\rangle = |0\rangle$ (By learning, we mean taking multiple copies of $|\psi_x\rangle$ as input and outputting a circuit description of U_x , the unitary

that prepares $|\psi_x\rangle$). The circuit implementing the state unpreparation unitary $U_x^\dagger$ can then be derived by reversing the order of the gates and taking their Hermitian conjugation. The requirement of outputting the circuit description of the target state is necessary for properly defining the computational complexity of the learning algorithms, as we detail in Supplementary Information B. This is very different from constructing $U_x^\dagger$ from quantum queries to U_x , which is generally hard¹². We note that all the learning algorithms considered in this work satisfy this definition).

In this paper, we rigorously study this process using tools recently developed in quantum learning theory^{13–21} and explore its thermodynamic implications. We explain a general method that lifts any learning algorithm to a reversible one, enabling it to acquire knowledge from multiple copies of the unknown state and erase many additional copies at the optimal energy cost, saturating Landauer’s limit (Fig. 1a). We further pinpoint the quantitative relation between the energy cost of erasing a class of quantum states and their complexity, as measured by circuit depth, entanglement, or magic (Fig. 1b).

After establishing the information-theoretical optimality of learning-to-erase protocols, we study their computational efficiency—whether they can be implemented in a number of elementary operations that scales polynomially with the system size n and number of copies N . Computational efficiency is crucial in that if a protocol requires exponential time to execute, it would take longer than the age of the universe to erase only a few hundreds of qubits. We show that whenever the learning algorithm and

¹Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA, USA. ²Department of Physics, University of California, Santa Barbara, CA, USA. ³AWS Center for Quantum Computing, Pasadena, CA, USA. ✉e-mail: haimengzhao@icloud.com; yuzhenzhang@ucsb.edu; preskill@caltech.edu




Fig. 1 | Overview of main results. a Reversible learning algorithm $\mathcal{L}$ can acquire knowledge of the unknown n -qubit state and erase N copies of it at the optimal work cost, saturating Landauer's limit. Here, S stores the copies used by the learning algorithm, R stores the rest, M is the memory of the learning algorithm, and M' is an auxiliary memory. **b** The work cost of erasing physically relevant classes of states

grows with the complexity of the states, as measured by circuit depth d , magic t , entanglement entropy $\mathcal{S}$, and degree k . When the complexity is bounded by a constant, Landauer's limit can be achieved efficiently by learning. In contrast, when the complexity grows poly-logarithmically, no polynomial-time quantum algorithm can erase the states without paying a nearly maximal amount of work.

state preparation are efficient, the constructed erasure protocol is also efficient. We illustrate this with a wide range of physically relevant classes of states, including shallow-circuit states, doped stabilizer states, matrix product states, and low-degree phase states. Conversely, under standard cryptographic assumptions, we prove that there exist classes of low-complexity states that are hard to learn and also hard to erase. For these states, Landauer's principle allows N copies to be erased by expending an amount of work independent of N , but for any efficient erasure protocol the cost scales nearly linearly with N . This large gap between the work cost of computationally-bounded agents versus what is information-theoretically possible is a genuine quantum many-body phenomenon and a computational no-go result complementary to the third law of thermodynamics^{22–25}. It also implies that the existing general purpose compress-to-erase methods for achieving Landauer's limit^{11,26–29} require resources superpolynomial in n (Note that the time efficiency claims in these previous works refer to a runtime of $\text{poly}(N, D)$, where $D = 2^n$ is the Hilbert space dimension of a single copy that grows exponentially with the system size).

Finally, we show that our learning-to-erase protocols can be used to extract work^{7,30,31} in Methods, where similar optimality and (in)efficiency results hold. In particular, the no-go result implies that in general one can extract only a small fraction of all the possible work in polynomial time. In contrast, given efficiently learnable states, we can efficiently extract the optimal amount of work by learning.

Results

General formulation

We first consider the general case where the class of possible states $\mathcal{C} = \{|\psi_x\rangle\}_{x=1}^m$ is arbitrary and known (along with circuit descriptions of their state preparation unitaries $\{U_x\}_{x=1}^m$). We assume that these states are distinct, in the sense that they are pairwise apart by a constant trace distance: $\forall x \neq x', d_{\text{tr}}(|\psi_x\rangle, |\psi_{x'}\rangle) = \frac{1}{2} \|\langle\psi_x| - \langle\psi_{x'}|\|_1 \geq \epsilon = \Theta(1)$. (In this work, we use the standard notations for asymptotics. For two positive functions $f(n)$ and $g(n)$, $f(n) = O(g(n))$ if $\exists n_0, C > 0$ such that $\forall n > n_0$, $f(n) \leq Cg(n)$. $f(n) = \Omega(g(n))$ if $g(n) = O(f(n))$. $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$). The state we aim to erase can be described by $\rho = \sum_{x=1}^m p_x (|\psi_x\rangle\langle\psi_x|)^{\otimes N}$, where $\{p_x\}$ are the prior probabilities. We consider the state-agnostic setting, where we do not have knowledge of $\{p_x\}$ and thus the state ρ that we want to erase, but we do have knowledge of the class $\mathcal{C}$ from which ρ is generated. For example, the distribution $\{p_x\}$ can be supported on any single state $|\psi_x\rangle$, or it can be uniform. Our goal is to design a protocol that can erase ρ with high success probability no matter what the distribution $\{p_x\}$ is. To do so, we describe a learning algorithm that identifies x , lift it to a reversible algorithm, and then use it to erase the state.

To study fundamental principles, we make several technical simplifications. First, we focus on the energy cost incurred by the irreversibility in

logical information processing, and leave the cost from implementation details such as error correction and device maintenance to future work. Second, we allow the use of ancilla qubits as long as they are restored to the initial state, so that one cannot hide entropy or work cost in the ancilla qubits³¹. Third, we assume that learning algorithms output circuit descriptions of the target states so that their computational complexity can be properly defined. See Supplementary Information B for a formal description.

We adopt a learning algorithm¹³ based on hypothesis selection³² and Clifford classical shadows³³. Given s copies of the state $|\psi_x\rangle$, we apply a random n -qubit Clifford gate on each copy and measure in the computational basis. The measurement statistics allow us to classically calculate unbiased estimators for the expectation values of $\binom{m}{2} = O(m^2)$ Helstrom measurements³⁴ between all pairs of possible states in $\mathcal{C}$. We can then use quantum hypothesis selection to find out the candidate state in $\mathcal{C}$ that is closest to $|\psi_x\rangle$ in trace distance. When we have $s = O(\log(m)/\epsilon^2) = O(\log m)$ samples, we are guaranteed to find the correct x with high probability p_{succ} . A detailed proof can be found in ref. 13 Proposition 1.

The above learning algorithm involves randomness generation and irreversible measurements that may incur extra work cost. To avoid this, we describe a general method that lifts a learning algorithm to a reversible one. Formally speaking, a learning algorithm $\mathcal{L}_0$ takes the system $|\psi_x\rangle_S^{\otimes s}$ and a classical memory register $|0\rangle_M$ as input and writes the learning outcome into $|x\rangle_M$ with high probability. It is implemented using a set of quantum and classical gates, random number generators, and projective measurements. We first replace every quantum and classical gates by their reversible counterparts^{7,35}. Then we replace projective measurements by their coherent version: to measure a set of projectors $\{P_a\}$, we apply an unitary U on the system and an ancilla A that records the measurement outcome. U completes the isometry $\sum_a P_a \otimes |a\rangle\langle 0|$ such that $U|\psi\rangle_S|0\rangle_A = \sum_a P_a |\psi\rangle_S |a\rangle_A$ for any $|\psi\rangle_S$. When classical random numbers are needed, we coherently measure $|+\rangle$ states in the computational basis. All operations that depend on random numbers or measurement outcomes can be realized using reversible gates controlled by the ancilla. In this way, we have constructed a reversible learning algorithm $\mathcal{L}$ (a unitary) that satisfies

$$\mathcal{L}|\psi_x\rangle_S^{\otimes s}|0\rangle_M|0\rangle_A = \sum_{x'=1}^m c_{x'|x} |x'\rangle_M |\text{junk}_{x'}\rangle_{S,A}, \quad (1)$$

where the coefficients $c_{x'|x} \in \mathbb{C}$ are determined by the probability of predicting x' given the state $|\psi_x\rangle$. In particular, $|c_{x|x}|^2 \geq p_{\text{succ}}$ is close to one for any x .

Now we construct an erasure protocol $\mathcal{E}(\mathcal{L})$ using the reversible learning algorithm $\mathcal{L}$. The construction is summarized in Fig. 1a. Given

$N \geq s$ copies of $|\psi_x\rangle$, the protocol proceeds by first executing $\mathcal{L}$ on s copies. We call the s copies S and the rest R . Then we introduce an additional classical memory M' and use CNOT gates to copy the content of M into M' . Next, we apply the inverse $\mathcal{L}^{-1}$ to S, M, A to uncompute³⁶ and get rid of the junk. We erase each copy in S and R by applying the inverse state preparation unitary conditioned on the learning outcome recorded in M' and uncomputing the copy. Finally, we erase M' bit by bit using standard classical Landauer erasure⁷, costing energy $W = (\log_2 m)k_B T \ln 2$. We prove the correctness of this erasure protocol in Supplementary Information C and provide a proof sketch in Methods.

We note that the only step that costs energy in this erasure protocol is the step of erasing learning outcomes, which costs $W = (\log_2 m)k_B T \ln 2$ joules of work, independent of N . This confirms our intuition that once we have learned the state, we can erase many more copies without further work. This general method of constructing reversible learning algorithms illustrates that learning as a physical process does not itself have a fundamental energy cost^{37,38}. In principle, the energy cost only occurs when the learning agent needs to erase and recycle its memory. In practice, there may be an extra energy cost if the learning algorithm is not implemented in a fully coherent and reversible way.

The sample complexity s of the learning algorithm plays an interesting role in the erasure protocol. The above construction works as long as $N \geq s$ and we have sufficient quantum working memory (ancilla supply) so that we have enough information to learn and space to make everything reversible. Extra work may be needed when these assumptions are not met. This shows that the usual $N \rightarrow \infty$ limit taken in thermodynamics is not necessary and can be significantly relaxed to $N \geq s$. Assuming that the learning algorithm and the state preparation unitary have time complexity T_{learn} and T_{prep} , the total time complexity of $\mathcal{E}(\mathcal{L})$ is $O(T_{\text{learn}} + \log m + NT_{\text{prep}})$. Therefore, whenever the learning and state preparation steps are efficient, the erasure protocol is also efficient.

We show that the above work cost is optimal by calculating the lower bound given by Landauer's principle. Since we are only given one copy of ρ (which contains N copies of $|\psi_x\rangle$), the Landauer's limit is given by the one-shot max-entropy $W \geq H_{\text{max}}(\rho)k_B T \ln 2$, where $H_{\text{max}}(\rho) = \log_2(\text{rank}(\rho))$ is the max-entropy of ρ ^{9,11,31,39,40}. To calculate the rank of ρ , we construct the Gram matrix $G_{xx'} = \langle \psi_x | \psi_{x'} \rangle^N$. Assuming $N > 2 \log(m-1)/\log(1/(1-\epsilon^2)) = \Theta(\log m)$ (which is naturally satisfied when $N \geq s$ and we always assume this from now on), we have $\sum_{x' \neq x} |G_{xx'}| \leq \sum_{x' \neq x} (1 - \epsilon^2)^{N/2} < \sum_{x' \neq x} (m-1)^{-1} = 1 = G_{xx}$, $\forall x$. This means that G is diagonally dominant and thus non-singular, implying that $\{|\psi_x\rangle\}_{x=1}^m$ are linearly independent. Therefore, when we maximize over all possible $\{p_x\}$ in the state-agnostic setting, we have $\text{rank}(\rho) = m$ and Landauer's limit reads $(\log_2 m)k_B T \ln 2$. This coincides with the work cost of $\mathcal{E}(\mathcal{L})$, proving that the learning-to-erase protocol is information-theoretically optimal.

We remark that the above construction works for any valid learning algorithm, not just for the one described above. In fact, a simple compress-to-erase protocol commonly used in the literature^{9,11} can already achieve the optimal energy cost (see Supplementary Information B). However, this compress-to-erase protocol has exponential runtime, and in fact we will show that this computational inefficiency is inevitable. In contrast, we will see below that learning-to-erase protocols are provably-efficient and energy-optimal for physically relevant classes of states. Moreover, in regimes where our protocols are inefficient, our computational hardness results show that no efficient protocols can achieve the optimal energy cost. Therefore, learning-to-erase can be regarded as an operational way to instantiate compression-based erasure that achieves the same optimal energy cost but with provable computational efficiency guarantees.

Work and complexity

We now apply the general formulation to several physically relevant classes of states, and illustrate how the work cost of erasure grows with the complexity of the states in the target class. For these classes, we give time efficient erasure protocols built from efficient learning algorithms. The results are

summarized in Fig. 1b. Note that we are considering erasure under the assumption that the system Hamiltonian is trivial. If the Hamiltonian is nontrivial, then erasure might induce a change in internal energy that should be included in the thermodynamic cost.

Our first example is shallow-circuit states $\mathcal{C} = \{|\psi\rangle = U|0^n\rangle\}$, where U can be an arbitrary quantum circuit with depth d on a $O(1)$ -dimensional lattice built from a finite-size two-qubit gate set G ^{14,15,41}. This class naturally describes the output state of finite-time evolution of physical systems⁴² and noisy intermediate-scale quantum computation⁴³, offering quantum advantage over classical computation^{44–49}. The number of possible states in this class is $m = |G|^{\Theta(nd)}$. Thus the required work cost is

$$W = \Theta(nd)k_B T \ln 2, \quad (2)$$

growing linearly with circuit depth d , when N is sufficiently large. In Methods, we generalize this to the case where each gate is non-discrete and can be any two-qubit unitary. The same work cost scaling applies up to logarithmic factors. For constant d , this optimal work cost is achieved efficiently using $\mathcal{E}(\mathcal{L})$ with a polynomial-time learning algorithm $\mathcal{L}$ that learns the local inversions of the shallow circuit and sews them together (the discrete gate set version of ref. 15). In particular, we have $T_{\text{learn}} = \text{poly}(n) \cdot 2^{d^{O(1)}}$ and $T_{\text{prep}} = O(nd)$, and the total time complexity of erasure is $O(\text{poly}(n) \cdot 2^{d^{O(1)}} + ndN)$.

A second example is t -doped stabilizer states $\mathcal{C} = \{|\psi\rangle = U|0^n\rangle\}$, where U is composed of an arbitrary number of Clifford gates and at most t T -gates ($T = \text{diag}(1, e^{i\pi/4})$)^{16,17}. Here, t characterizes the amount of magic (or non-stabilizerness) in the state⁵⁰, which is closely related to universality^{51,52}, classical simulability^{53–55}, and error correction^{56,57} of quantum computation. The number of possible states in the class is $m = (2^{\Theta(n^2)})^{t+1} \times n^t = \exp(\Theta(n^2(t+1)))$ ⁵³ because there are $2^{\Theta(n^2)}$ possible Clifford unitaries in between each pair of T -gates. When N is large, the work cost for erasure is

$$W = \Theta(n^2(t+1))k_B T \ln 2, \quad (3)$$

growing linearly with the magic t . In particular, the work cost for erasing stabilizer states is $\Theta(n^2)k_B T \ln 2$. For $t = O(\log n)$, this is achieved efficiently by an efficient learning algorithm for t -doped stabilizer states, which identifies the state exactly by bell sampling⁵⁸ and estimating the stabilizer group (see ref. 16). We have $T_{\text{learn}} = \text{poly}(n, 2^t)$ and $T_{\text{prep}} = O(n^2 t)$; thus the total time complexity of erasure is $O(\text{poly}(n, 2^t) + n^2 t N)$.

Our third example is matrix product states (MPS) $\mathcal{C} = \{|\psi\rangle = \sum_{x \in \{0,1\}^n} A_1^{x_1} \cdots A_n^{x_n} |x\rangle\}$, where $A_i^x \in \mathbb{C}^{\chi_i \times \chi_{i+1}}$ with bond dimension $\chi_1 = \chi_{n+1} = 1$ and $\chi_i \leq 2^S$ ¹⁸. Here, S is the maximal entanglement entropy across every cut of the system. MPS representations with constant S exist for states with bounded entanglement, including ground states of one-dimensional gapped local Hamiltonians⁵⁹ and symmetry-protected topological phases⁶⁰, as well as important states like AKLT⁶¹, GHZ, W and cluster states that are useful for quantum metrology⁶² and universal quantum computation⁶³. In Methods, we show that the work cost for erasure

$$W = \exp(\tilde{O}(S))k_B T \ln 2 \quad (4)$$

grows exponentially with the entanglement entropy S . For constant S , this is achieved efficiently using an efficient learning algorithm for MPS that sequentially learns $n O(S)$ -qubit unitaries to disentangle the MPS (see ref. 18). We have $T_{\text{learn}} = \text{poly}(n, 2^S)$ and $T_{\text{prep}} = O(n4^S)$, and the total time complexity of erasure is $O(\text{poly}(n, 2^S) + n4^S N)$.

The last example is low-degree phase states $\mathcal{C} = \{|\psi\rangle = 2^{-n/2} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle\}$, where $f: \{0,1\}^n \rightarrow \{0,1\}$ is an arbitrary Boolean function with degree at most $k = O(1)$ ¹⁹. They correspond to hypergraph states that are generated by k -body multi-controlled Z gates on $|+\rangle^{\otimes n}$. The degree k is closely related to the level of Clifford hierarchy⁶⁴ and the amount of quantum contextuality⁶⁵. They are useful resources in measurement-based quantum computation^{63,66} and appear in quantum

advantage experiments^{67,68} and oracle query algorithms³⁶. By counting polynomial Boolean functions, we see that the number of such states is

$$m = \prod_{j=0}^k 2^{\binom{n}{j}} = \exp \Theta(n^k) \text{ and the required work cost is}$$

$$W = \Theta(n^k) k_B T \ln 2, \quad (5)$$

growing exponentially with degree k . For constant k , the erasure is efficient using an efficient learning algorithm that measures single-qubit X and Z observables and learns the gradient of f with respect to each argument (see ref. 19, Theorem 3). Here, $T_{\text{learn}} = O(n^{3k-2})$ and $T_{\text{prep}} = O(kn^k)$ ⁶⁹, and the total time complexity of erasure is $O(n^{3k-2} + kn^k N)$.

In Fig. 1b, we summarize the above results and visualize how the optimal work cost of erasure grows with the four different complexity measures: circuit depth d , magic t , entanglement entropy S , and degree k . In particular, the work cost scales linearly with depth ($\sim nd$) and magic ($\sim nt$), where the growth with magic has a larger slope. Meanwhile, the work cost scales exponentially with entanglement entropy ($\sim e^S$) and degree ($\sim n^k$), where the growth with degree is faster. The computational complexity of erasure grows exponentially with all four measures of complexity. When the complexity is bounded by a constant, the optimal work cost can be achieved efficiently by learning.

Computational hardness

For the above special classes of states, we have constructed efficient erasure protocols with optimal work cost. However, we show that this is not possible in general. We consider a particular class of states, pseudorandom states $\{|\psi_x\rangle\}$ ^{70,71}, which cannot be efficiently distinguished from Haar random states with non-negligible ($1/\text{poly}(n)$) probability. Concretely, under the standard cryptographic conjecture that there exist one-way functions secure against any sub-exponential time quantum adversary (e.g., based on the learning with errors problems⁷²), pseudorandom states can be constructed with $d = \text{polylog}(n)$ circuit depth²⁰. Let $N = \text{poly}(n)$. Because the entropy of $\rho = \sum_x p_x (|\psi_x\rangle\langle\psi_x|)^{\otimes N}$ depends on m but not N , Landauer's principle asserts that the N copies of pseudorandom states can be erased with low work cost $\Theta(nd) k_B T \ln 2 = \Theta(n \text{polylog}(n)) k_B T \ln 2$, independent of N . On the other hand, for the Haar ensemble, the Landauer cost of erasure is

$$W_{\text{Haar}} = \log_2 \binom{N + 2^n - 1}{N} k_B T \ln 2, \quad (6)$$

because the density operator realized by the Haar ensemble has full rank over the totally symmetric subspace for N copies of a 2^n -dimensional Hilbert space, which has dimension $\binom{N + 2^n - 1}{N}$. In the following, we prove that any polynomial-time quantum algorithm that can erase a pseudorandom state $\rho = \sum_x p_x (|\psi_x\rangle\langle\psi_x|)^{\otimes N}$ with success probability close to one must require work cost W_{Haar} . The intuition is simple: if we can erase pseudorandom states with less work than W_{Haar} then we can distinguish them from Haar random states by measuring the work cost, violating the definition of pseudorandom states.

To prove the statement, we suppose for the sake of contradiction that there exists an efficient erasure protocol $\mathcal{E}$ that erases ρ with work cost less than W_{Haar} . We construct a distinguisher $\mathcal{D}$ that distinguishes pseudorandom states from Haar random states. The distinguisher proceeds by executing $\mathcal{E}$ on the given states and conducting two tests: (1) determine if the erasure is successful by measuring the projector $|0\rangle\langle 0|$; and (2) test if the work cost is less than W_{Haar} by measuring the energy change in the energy source (Here, we assume that the energy source is classical and can be measured deterministically. If the energy source is quantum, we can always measure the work cost to $1/\text{poly}(n)$ accuracy in polynomial time. Thus, the work cost lower bound in the no-go result will change to $W_{\text{Haar}} - \text{negl}(n)$ where $\text{negl}(n)$ is a function

that decays faster than any polynomial of n . In Supplementary Information D, we provide the proof details of the fully quantum case). $\mathcal{D}$ outputs 0 if both tests pass, and 1 otherwise.

When given pseudorandom states, we have $\Pr[\mathcal{D}((|\psi_x\rangle\langle\psi_x|)^{\otimes N}) = 0] = \Pr[\mathcal{D}(\rho) = 0]$ close to one, where the probability is over x and the randomness in $\mathcal{D}$; here we used linearity of quantum channels and the guarantee that $\mathcal{E}$ can erase pseudorandom states. When given Haar random states $\rho_{\text{Haar}} = \int d\psi (|\psi\rangle\langle\psi|)^{\otimes N}$, the probability of outputting 0 is equal to $\Pr[\mathcal{D}(|\psi\rangle\langle\psi|) = 0] = \Pr[\mathcal{D}(\rho_{\text{Haar}}) = 0]$, the probability of one-shot successful erasure of ρ_{Haar} with work cost less than W_{Haar} . Because the erasure protocol $\mathcal{E}$ is assumed to be efficient, this probability must also be close to one by the definition of pseudorandom states. But this contradicts Landauer's principle, which asserts that it is not possible to perform one-shot erasure with work cost less than $H_{\text{max}}(\rho_{\text{Haar}}) k_B T \ln 2 = \log_2(\text{rank}(\rho_{\text{Haar}})) k_B T \ln 2 = W_{\text{Haar}}$. This completes the proof of the no-go result. We provide more details of the proof in Supplementary Information D.

For a quantitative understanding, we note that $W_{\text{Haar}}/(k_B T \ln 2) \geq nN(1 - (\log_2 N)/n)$, which goes to the maximal value nN as n grows large with $N = \text{poly}(n)$. But Landauer's limit for this class of states is only $\Theta(nd) = n \text{polylog}(n)$, independent of N . This means that Landauer's limit cannot be achieved in polynomial time, and asymptotically maximal work must be paid.

We remark that this is a computational no-go result complementary to the third law of thermodynamics (We consider the unattainability formulation of the third law of thermodynamics, which applies not only to the cooling of thermal states but also the erasure/purification of arbitrary states. See e.g., ref. 23), which asserts that perfect erasure as a form of cooling requires diverging resources (time, size of the heat bath, etc.)^{22–25}. The third law does not rule out the possibility of achieving Landauer's limit in polynomial time by paying $1/\text{poly}(n)$ extra work and using an infinite size heat bath. In contrast, our result shows that even with an infinite size heat bath, any procedure that can erase a particular class of low-complexity states in polynomial time must pay a nearly maximal amount of work, under standard cryptographic assumptions. This is a many-body phenomenon arising from the complexity of large systems.

Discussion

We have established a concrete connection between quantum learning theory and thermodynamics. Our results illustrate that seemingly abstract learning processes have tangible physical consequences that determine the energy cost or gain in thermodynamic tasks. This allows us to study properties of learning itself from a thermodynamic perspective. In particular, we find that learning has no fundamental energy cost, if implemented in a fully coherent fashion.

This connection also enables provably-efficient and energy-optimal thermodynamic protocols based on learning, improving upon existing compression-based protocols that have exponential time complexity^{11,26–29}. This contributes to a growing literature on state-agnostic thermodynamic and compression protocols^{38,73–75}. In particular, we circumvent the potential energy cost of measurements and revive the use of learning algorithms for thermodynamic tasks³⁸. In a sense, learning is a form of compression — namely compressing quantum states into their classical descriptions — and efficient learning yields efficient compression⁷³. These efficient erasure protocols may find applications in the initialization of distributed quantum metrology systems and in recycling qubits for repeated use in quantum experiments. They also provide a viable way to experimentally demonstrate Landauer's limit^{76–78} and materialize Maxwell's demon as a learning agent^{79–82} in the quantum many-body regime. These efficient energy-optimal protocols may be a key component in building future large-scale, energy-efficient quantum computers⁸³.

Our results also exemplify the drastic change of physical resources required by certain tasks when we are in the many-body regime. Ideas that

date back to the early days of information theory, like Landauer's principle, may require reexamination via the lens of complexity^{84,85}.

Our work opens up many interesting future directions. First, our setting can be extended to more realistic scenarios, where we want to erase many copies of mixed states $\rho = \sum_x p_x \rho_x^{\otimes N}$. A natural choice could be thermal states of unknown Hamiltonians that describe physical systems in thermal equilibrium. Efficient learning and erasing may be possible in this case^{86,87}, where we expect the optimal work cost to scale with $N: \sum_x p_x N H(\rho_x) + H(\{p_x\})$, because after we have identified ρ_x , we still need to pay work to erase each copy because it is mixed. Second, our results may be extended to continuous variable systems such as fermionic or bosonic systems and inspire efficient thermodynamic protocols there^{88,89}. Third, learning may yield efficient protocols for other thermodynamic tasks beyond erasure and work extraction.

Our computational hardness results may find cryptological applications related to energy storage, manipulation, and transfer. For example, one may envision an encrypted battery, where any computationally-bounded agent without the secret key cannot extract much work from the battery even when having access to the physical substrate of the battery; the full energy is only efficiently extractable when one has the secret key. Intriguingly, limitations on erasing information in or extracting work from black holes may apply if the evolution of these systems is modeled accurately by pseudorandom unitaries^{90–92}. Finally, it may be instructive to study other physical properties of learning processes beyond energy cost, such as their consequences for information dynamics⁹³, phase transitions⁹⁴, and noise robustness²¹.

During the revision of this manuscript, we became aware of related independent works by Watanabe and Takagi⁹⁵ and Lumbreras et al.⁹⁶, released subsequent to the initial posting of our preprint. Ref. 95 addresses the information-theoretic achievability of state-agnostic work extraction when infinitely many copies are given ($N \rightarrow \infty$). Ref. 96 applies a known regret bound from online learning of quantum states to online work extraction from qubits ($n = 1$). Both works start from the premise that standard tomography destroys many copies of the states and uses measurements that may incur additional energy cost. In contrast, our work shows that one can circumvent these costs by making tomography fully coherent and reversible, thereby achieving the optimal energy cost by learning.

Our work also studies the computational complexity aspect of erasure and work extraction, providing provably efficient protocols for structured states and proving no-go theorems for high-complexity states in the many-body regime ($n \gg 1$). This aspect is not considered in refs. 95,96. In particular, ref. 96 is restricted to single-qubit pure states ($n = 1$). Ref. 95 studies mixed states in general systems and has an excess energy cost per copy that scales as $\sim 2^{O(n)}/N$, which approaches zero only when the number of copies N is exponential in n . In fact, our no-go theorem proves that no general-purpose protocol can have runtime polynomial in n while dissipating energy sublinear in N , as doing so would break post-quantum cryptography. As a corollary, the polylog(N) dissipation scaling derived in ref. 96 for qubits ($n = 1$) cannot be generalized to computationally efficient protocols for many-body systems.

In addition, our work relates the energy cost of erasing quantum states to their structural properties, including circuit depth, entanglement, magic, and Boolean function degree. This relation between different resources has not been addressed before. We provide a detailed discussion of other related works in Supplementary Information A.

Methods

Proof of correctness

Here we provide the proof sketch for the correctness of the erasure protocol $\mathcal{E}(\mathcal{L})$. A more detailed proof can be found in Supplementary Information C.

The key is to calculate the overlap between $|x\rangle_{M'} |\psi_x\rangle_S^{\otimes N} |0\rangle_{M,A}$ and the intermediate state of $\mathcal{E}(\mathcal{L})$ after uncomputation (apart from R):

$$\begin{aligned} & (\langle x|_{M'} \langle \psi_x|_S^{\otimes N} \langle 0|_{M,A}) \left(\mathcal{L}^\dagger \sum_{x'} c_{x'|x} |x'\rangle_{M'} |x'\rangle_M |\text{junk}_{x'}\rangle_{S,A} \right) \\ &= \sum_{x'} c_{x'|x} \langle x|x'\rangle_{M'} \left(\langle \psi_x|_S^{\otimes N} \langle 0|_{M,A} \mathcal{L}^\dagger \right) \left(|x'\rangle_M |\text{junk}_{x'}\rangle_{S,A} \right) \\ &= c_{x|x} \left(\sum_{x'} c_{x''|x}^* \langle x''|_M \langle \text{junk}_{x''}|_{S,A} \right) |x\rangle_M |\text{junk}_x\rangle_{S,A} \\ &= |c_{x|x}|^2 \geq p_{\text{succ}}, \end{aligned} \quad (7)$$

which means that the trace distance between the two states is $\sqrt{1 - |c_{x|x}|^4} \leq \sqrt{1 - p_{\text{succ}}^2}$. Due to the data processing inequality satisfied by the trace distance⁹⁷, after executing the rest of the erasure protocol (a fixed quantum channel), we have $d_{\text{tr}}(\mathcal{E}(\mathcal{L})(|\psi_x\rangle_S^{\otimes N} |0\rangle), |0\rangle) \leq \sqrt{1 - p_{\text{succ}}^2}$. When the input is the mixture $\rho = \sum_x p_x (|\psi_x\rangle_S \langle \psi_x|)^{\otimes N}$, we invoke the triangle inequality and have trace distance $\frac{1}{2} \|\mathcal{E}(\mathcal{L})(\rho \otimes |0\rangle \langle 0|) - |0\rangle \langle 0|\|_1 \leq \sum_x p_x d_{\text{tr}}(\mathcal{E}(\mathcal{L})(|\psi_x\rangle_S^{\otimes N} |0\rangle), |0\rangle) \leq \sqrt{1 - p_{\text{succ}}^2}$, which is close to zero. This concludes the proof of the correctness.

Continuous classes of states

It is straightforward to generalize the results to erase a continuous class of states $\mathcal{C}$. We can use the erasure protocol for a discrete class in conjunction with a minimal ϵ -covering net of $\mathcal{C}$, similar to ref. 13, Theorem 1. The resulting erasure protocol will be able to approximately erase every copy to ϵ error in trace distance. The work cost will be $(\log m(\mathcal{C}, d_{\text{tr}}, \epsilon)) k_B T \ln 2$ where $m(\mathcal{C}, d_{\text{tr}}, \epsilon)$ is the size of the covering net. Meanwhile, this protocol can perfectly erase states drawn from a maximal $\Theta(\epsilon)$ -packing net of $\mathcal{C}$. By Landauer's principle, this requires work $(\log \tilde{m}(\mathcal{C}, d_{\text{tr}}, \Theta(\epsilon))) k_B T \ln 2$ when N is large, where $\tilde{m}(\mathcal{C}, d_{\text{tr}}, \Theta(\epsilon))$ is the size of the packing net. Since m and $\tilde{m}$ are equivalent up to a constant factor change in ϵ ⁹⁸, the erasure protocol is asymptotically optimal.

As an example, we bound the size of the ϵ -covering and packing nets of matrix product states. Every MPS with entanglement entropy S can be generated by a sequential quantum circuit with n gates each acting on $O(S)$ qubits⁹⁹. Utilizing existing covering number bounds on circuits with bounded gate complexity and fixed circuit structure¹³, Theorem 8, we have $\log m(\mathcal{C}, d_{\text{tr}}, \epsilon) \leq O(n 4^{O(S)} \log(n 4^{O(S)} / \epsilon)) = \exp \tilde{O}(S)$, where $\tilde{O}$ omits $\log n$, $\log \log(1/\epsilon)$ factors. To lower bound the packing number $\tilde{m}$, we note that an MPS with entanglement entropy S can represent any $\Theta(S)$ -qubit state tensored with zeroes on the remaining qubits. The set of general $\Theta(S)$ -qubit states have log-packing number $\exp \tilde{\Omega}(S)$. Thus we have $\log \tilde{m}(\mathcal{C}, d_{\text{tr}}, \epsilon) \geq \exp \tilde{\Omega}(S)$. This means that the work cost of erasure is $W = \exp(\Theta(S)) k_B T \ln 2$. Similarly, one can bound the covering and packing numbers of depth- d shallow circuit states with arbitrary two-qubit gates as $\exp \tilde{\Theta}(nd)$ ¹³, which implies the work cost of erasure $W = \Theta(nd) k_B T \ln 2$ when each gate is not assumed to be discrete.

The classical case

It helps to compare with the classical version of the problem. Classically, the state is a mixture of $m \leq 2^n$ possible computational basis states copied N times. The entropy is $\log_2 m \leq n$, which does not increase with N , even when the bit-strings are generated by pseudorandom functions. We can always erase N copies with work cost at most $n k_B T \ln 2$. Therefore, it is a genuine quantum feature to have a work cost W_{Haar} that scales with N . This originates from the fact that there are many more states when we allow superposition. The optimal classical erasure protocol is straightforward: we condition on the first copy of the bit-string, apply transversal CNOT gates to erase the other copies, and finally erase the first copy. From a learning perspective, this protocol learns the bitstring by looking at the first copy, and the knowledge of this one copy is sufficient for erasing the rest. In contrast, in the fully quantum case, many copies are needed to learn a state.

Work extraction

The thermodynamic implications of quantum learning theory extend beyond erasure. In the task of work extraction, it is information-theoretically possible to extract $W = (nN - H(\rho))k_B T \ln 2$ work from the state $\rho^{7,10,30,31}$. When learning is efficient, our learning-to-erase protocols give efficient work extraction protocols with the maximal work yield: we simply erase the state first, and then extract work from each refreshed qubit. In contrast, there exist ensembles such that no polynomial-time quantum algorithm can extract more work than $nNk_B T \ln 2 - W_{\text{Haar}} \leq N \log_2(N)k_B T \ln 2$. Otherwise, one could construct an efficient erasure protocol with less than W_{Haar} work cost for pseudorandom states by first extracting work and then erasing each qubit (which costs $nNk_B T \ln 2$ work and $O(nN) = \text{poly}(n)$ time). Therefore, for general states, one can extract in polynomial time only a fraction $\log_2(N)/n = O(\log(n)/n)$ of the available work, which becomes negligibly small for n large.

Data availability

No data are generated or analyzed in this theoretical work.

Received: 9 January 2026; Accepted: 7 May 2026;

Published online: 27 May 2026

References

- Jaynes, E. T. Information theory and statistical mechanics. *Phys. Rev.* **106**, 620–630 (1957).
- Jaynes, E. T. Information theory and statistical mechanics. II. *Phys. Rev.* **108**, 171–190 (1957).
- Jaynes, E. T. *Probability Theory: The Logic of Science*. (Cambridge University Press, Cambridge, 2003).
- Landauer, R. Information is physical. *Phys. Today* **44**, 23–29 (1991).
- Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* **5**, 183–191 (1961).
- Bennett, C. H. Logical reversibility of computation. *IBM J. Res. Dev.* **17**, 525–532 (1973).
- Bennett, C. H. The thermodynamics of computation—a review. *Int. J. Theor. Phys.* **21**, 905–940 (1982).
- Reeb, D. & Wolf, M. M. An improved Landauer principle with finite-size corrections. *N. J. Phys.* **16**, 103011 (2014).
- Dahlsten, O. C. O., Renner, R., Rieper, E. & Vedral, V. Inadequacy of von Neumann entropy for characterizing extractable work. *N. J. Phys.* **13**, 053015 (2011).
- Feynman, R. P. *Feynman Lectures on Computation: Anniversary Edition* 2nd edn (CRC Press, 2023).
- del Rio, L., Aberg, J., Renner, R., Dahlsten, O. & Vedral, V. The thermodynamic meaning of negative entropy. *Nature* **474**, 61–63 (2011).
- Yoshida, S., Soeda, A. & Murao, M. Reversing unknown qubit-unitary operation, deterministically and exactly. *Phys. Rev. Lett.* **131**, 120602 (2023).
- Zhao, H. et al. Learning quantum states and unitaries of bounded gate complexity. *PRX Quantum* **5**, 40306 (2024).
- Huang, H.-Y. et al. Learning shallow quantum circuits. In *Proc. 56th Annual ACM Symposium on Theory of Computing*, 1343–1351 <http://arxiv.org/abs/2401.10095> (2024).
- Landau, Z. & Liu, Y. Learning quantum states prepared by shallow circuits in polynomial time. In *Proc. 57th Annual ACM Symposium on Theory of Computing*, 1828–1838 <http://arxiv.org/abs/2410.23618> (2025).
- Leone, L., Oliviero, S. F. E. & Hama, A. Learning t-doped stabilizer states. *Quantum* **8**, 1361 (2024).
- Grewal, S., Iyer, V., Kretschmer, W. & Liang, D. Efficient learning of quantum states prepared with few non-clifford gates. *Quantum*, **9**, 1907 (2025).
- Cramer, M. et al. Efficient quantum state tomography. *Nat. Commun.* **1**, 149 (2010).
- Arunachalam, S., Bravyi, S., Dutt, A. & Yoder, T. J. Optimal algorithms for learning quantum phase states. In *18th Conference on the Theory of Quantum Computation, Communication and Cryptography* <http://arxiv.org/abs/2208.07851> (2023).
- Schuster, T., Haferkamp, J. & Huang, H.-Y. Random unitaries in extremely low depth. *Science* **389**, 92–96 (2025).
- Anshu, A. & Arunachalam, S. A survey on the complexity of learning quantum states. *Nat. Rev. Phys.* **6**, 59–69 (2024).
- Wilming, H. & Gallego, R. Third Law of Thermodynamics as a Single Inequality. *Phys. Rev. X* **7**, 041033 (2017).
- Masanes, L. & Oppenheim, J. A general derivation and quantification of the third law of thermodynamics. *Nat. Commun.* **8**, 14538 (2017).
- Scharlau, J. & Mueller, M. P. Quantum Horn’s lemma, finite heat baths, and the third law of thermodynamics. *Quantum* **2**, 54 (2018).
- Taranto, P. et al. Landauer versus Nernst: what is the true cost of cooling a quantum system? *PRX Quantum* **4**, 010332 (2023).
- Plesch, M. & Bužek, V. Efficient compression of quantum information. *Phys. Rev. A* **81**, 032317 (2010).
- Yang, Y., Chiribella, G. & Ebler, D. Efficient quantum compression for ensembles of identically prepared mixed states. *Phys. Rev. Lett.* **116**, 080501 (2016).
- Hayashi, M. *Quantum Information Theory: Mathematical Foundation*. Graduate Texts in Physics. <http://link.springer.com/10.1007/978-3-662-49725-8> (Springer, 2017).
- Bennett, C. H., Harrow, A. W. & Lloyd, S. Universal quantum data compression via nondestructive tomography. *Phys. Rev. A-At. Mol. Opt. Phys.* **73**, 032336 (2006).
- Alicki, R., Horodecki, M., Horodecki, P. & Horodecki, R. Thermodynamics of quantum informational systems - Hamiltonian description. *Open Systems & Information Dynamics*, **11**, 205–217 (2004).
- Faist, P., Dupuis, F., Oppenheim, J. & Renner, R. The minimal work cost of information processing. *Nat. Commun.* **6**, 7669 (2015).
- Bădescu, C. & O’Donnell, R. Improved quantum data analysis. *TheoretCS* **3**, 10924 (2024).
- Huang, H.-Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.* **16**, 1050–1057 (2020).
- Helstrom, C. W. Quantum detection and estimation theory. *J. Stat. Phys.* **1**, 231–252 (1969).
- Fredkin, E. & Toffoli, T. Conservative logic. *Int. J. Theor. Phys.* **21**, 219–253 (1982).
- Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information*. <https://books.google.com/books?hl=zh-CN&lr=&id=s4DEy7o-a0C&oi=fnd&pg=PR17&dq=Nielsen+chuang&ots=NJ5FdmvxWo&sig=mALXBRaW2lvAUhu70yB4hzH7ePQ> (Cambridge university press, 2010).
- Šafránek, D., Rosa, D. & Binder, F. C. Work extraction from unknown quantum sources. *Phys. Rev. Lett.* **130**, 210401 (2023).
- Xuereb, J., Junior, A., Clivaz, F., Bakhshinezhad, P. & Huber, M. What resources do agents need to acquire knowledge in quantum thermodynamics. Preprint at <https://doi.org/10.48550/arXiv.2410.18167> (2024).
- Renner, R. & Wolf, S. Smooth Renyi entropy and applications. In *International Symposium on Information Theory, 2004. ISIT 2004. Proceedings.*, 233– <https://ieeexplore.ieee.org/document/1365269> (2004).
- König, R., Renner, R. & Schaffner, C. The operational meaning of min- and max-entropy. *IEEE Trans. Inf. Theory* **55**, 4337–4347 (2009).
- Kim, H.-S., Kim, I. H. & Ranard, D. Learning state preparation circuits for quantum phases of matter <http://arxiv.org/abs/2410.23544> (2024).
- Poulin, D., Qarry, A., Somma, R. & Verstraete, F. Quantum simulation of time-dependent hamiltonians and the convenient illusion of hilbert space. *Phys. Rev. Lett.* **106**, 170501 (2011).

43. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
44. Terhal, B. M. & DiVincenzo, D. P. Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin Games. *Quantum Information and Computation*, **4**, 134–145 (2004).
45. Gao, X., Wang, S.-T. & Duan, L.-M. Quantum supremacy for simulating a translation-invariant ising spin model. *Phys. Rev. Lett.* **118**, 040502 (2017).
46. Giovannetti, V., Lloyd, S. & Maccone, L. Architectures for a quantum random access memory. *Phys. Rev. A* **78**, 52310 (2008).
47. Haferkamp, J. et al. Closing gaps of a quantum advantage with short-time hamiltonian dynamics. *Phys. Rev. Lett.* **125**, 250501 (2020).
48. Hangleiter, D. & Eisert, J. Computational advantage of quantum random sampling. *Rev. Mod. Phys.* **95**, 035001 (2023).
49. Bravyi, S., Gosset, D. & König, R. Quantum advantage with shallow circuits. *Science* **362**, 308–311 (2018).
50. Liu, Z.-W. & Winter, A. Many-Body Quantum Magic. *PRX Quantum* **3**, 020333 (2022).
51. Shi, Y. Both Toffoli and Controlled-NOT need little help to do universal quantum computation. *Quantum Information and Computation*, **3**(1), 84–92 <http://arxiv.org/abs/quant-ph/0205115> (2003).
52. Bravyi, S. & Kitaev, A. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Phys. Rev. A* **71**, 022316 (2005).
53. Aaronson, S. & Gottesman, D. Improved simulation of stabilizer circuits. *Phys. Rev. A* **70**, 52328 <http://arxiv.org/abs/quant-ph/0406196> (2008).
54. Bravyi, S. & Gosset, D. Improved classical simulation of quantum circuits dominated by clifford gates. *Phys. Rev. Lett.* **116**, 250501 (2016).
55. Bravyi, S. et al. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum* **3**, 181 (2019).
56. Knill, E. Fault-Tolerant Postselected Quantum Computation: Schemes. Preprint at <http://arxiv.org/abs/quant-ph/0402171> (2004).
57. Knill, E., Laflamme, R. & Viola, L. Theory of quantum error correction for general noise. *Phys. Rev. Lett.* **84**, 2525–2528 (2000).
58. Montanaro, A. Learning stabilizer states by Bell sampling. *Proceedings of the Royal Society A*, **463**, 2088 <http://arxiv.org/abs/1707.04012> (2017).
59. Hastings, M. B. An area law for one dimensional quantum systems. *J. Stat. Mech.* **2007**, P08024 <http://arxiv.org/abs/0705.2024> (2007).
60. Zeng, B., Chen, X., Zhou, D.-L. & Wen, X.-G. *Quantum Information Meets Quantum Matter: From Quantum Entanglement to Topological Phases of Many-Body Systems*. Quantum Science and Technology. <http://link.springer.com/10.1007/978-1-4939-9084-9> (Springer, 2019).
61. Affleck, I., Kennedy, T., Lieb, E. H. & Tasaki, H. Rigorous results on valence-bond ground states in antiferromagnets. *Phys. Rev. Lett.* **59**, 799–802 (1987).
62. Degen, C. L., Reinhard, F. & Cappellaro, P. Quantum sensing. *Rev. Mod. Phys.* **89**, 035002 (2017).
63. Jozsa, R. An introduction to measurement based quantum computation. In *Quantum Information Processing*, 137–158 <http://arxiv.org/abs/quant-ph/0508124> (2006).
64. Hu, J., Liang, Q. & Calderbank, R. Climbing the diagonal clifford hierarchy. *Academia Quantum*, **2**, 4 <https://doi.org/10.48550/arXiv.2110.11923> (2025).
65. Raussendorf, R. Contextuality in measurement-based quantum computation. *Phys. Rev. A At. Mol. Opt. Phys.* **88**, 022322 (2013).
66. Raussendorf, R. & Briegel, H. J. A one-way quantum computer. *Phys. Rev. Lett.* **86**, 5188–5191 (2001).
67. Bremner, M. J., Jozsa, R. & Shepherd, D. J. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **467**, 459–472 (2010).
68. Bremner, M. J., Montanaro, A. & Shepherd, D. J. Achieving quantum supremacy with sparse and noisy commuting quantum computations. *Quantum* **1**, 8 (2017).
69. Gidney, C. Constructing Large Controlled Nots <https://algassert.com/circuits/2015/06/05/Constructing-Large-Controlled-Nots.html> (2015).
70. Ji, Z., Liu, Y.-K. & Song, F. Pseudorandom Quantum States. In Shacham, H. & Boldyreva, A. (eds.) *Advances in Cryptology – CRYPTO 2018*, 126–152 (Springer International Publishing, Cham, 2018).
71. Ma, F. & Huang, H.-Y. How to construct random unitaries. In *Proc. 57th Annual ACM Symposium on Theory of Computing*, 806–809 <http://arxiv.org/abs/2410.10116> (2025).
72. Regev, O. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* **56**, 34:1–34:40 (2009).
73. Yang, Y. Compression of quantum shallow-circuit states. *Phys. Rev. Lett.* **134**, 010603 (2025).
74. Watanabe, K. & Takagi, R. Black box work extraction and composite hypothesis testing. *Phys. Rev. Lett.* **133**, 250401 (2024).
75. Chakraborty, S., Das, S., Ghorui, A., Hazra, S. & Singh, U. Sample complexity of black box work extraction. *Quantum Science and Technology*, **10**, 045070 (2025).
76. Bérut, A. et al. Experimental verification of Landauer's principle linking information and thermodynamics. *Nature* **483**, 187–189 (2012).
77. Gaudenzi, R., Burzurí, E., Maegawa, S., van der Zant, H. S. J. & Luis, F. Quantum Landauer erasure with a molecular nanomagnet. *Nat. Phys.* **14**, 565–568 (2018).
78. Scandi, M. et al. Minimally dissipative information erasure in a quantum dot via thermodynamic length. *Phys. Rev. Lett.* **129**, 270601 (2022).
79. Koski, J. V., Maisi, V. F., Pekola, J. P. & Averin, D. V. Experimental realization of a Szilard engine with a single electron. *Proc. Natl. Acad. Sci. USA* **111**, 13786–13789 (2014).
80. Vidrighin, M. D. et al. Photonic Maxwell's Demon. *Phys. Rev. Lett.* **116**, 050401 (2016).
81. Camati, P. A. et al. Experimental rectification of entropy production by Maxwell's demon in a quantum system. *Phys. Rev. Lett.* **117**, 240502 (2016).
82. Cottet, N. et al. Observing a quantum Maxwell demon at work. *Proc. Natl. Acad. Sci. USA* **114**, 7561–7564 (2017).
83. Auffèves, A. Quantum Technologies Need a Quantum Energy Initiative. *PRX Quantum* **3**, 020101 (2022).
84. Munson, A. et al. Complexity-constrained quantum thermodynamics. *PRX Quantum*, **6**(1), 010346 <http://arxiv.org/abs/2403.04828> (2024).
85. Leone, L., Rizzo, J., Eisert, J. & Jerbi, S. Entanglement theory with limited computational resources. *Nature Physics*, 1–8 <http://arxiv.org/abs/2502.12284> (2025).
86. Bakshi, A., Liu, A., Moitra, A. & Tang, E. Learning quantum hamiltonians at any temperature in polynomial time. In *Proc. of the 56th Annual ACM Symposium on Theory of Computing*, 1470–1477 (2024).
87. Chen, C.-F., Anshu, A. & Nguyen, Q. T. Learning quantum gibbs states locally and efficiently. Preprint at <https://doi.org/10.48550/arXiv.2504.02706> (2025).
88. Mele, A. A. & Herasymenko, Y. Efficient learning of quantum states prepared with few fermionic non-gaussian gates. *PRX Quantum* **6**, 010319 (2025).
89. Mele, F. A. et al. Learning quantum states of continuous variable systems. *Nature Physics*, 1–7 <http://arxiv.org/abs/2405.01431> (2024).
90. Hayden, P. & Preskill, J. Black holes as mirrors: Quantum information in random subsystems. *J. High. Energy Phys.* **2007**, 120–120 (2007).
91. Bouland, A., Fefferman, B. & Vazirani, U. Computational Pseudorandomness, the Wormhole Growth Paradox, and Constraints on the AdS/CFT Duality. In Vidick, T. (ed.) *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, vol. 151 of

- Leibniz International Proceedings in Informatics (LIPIcs)*, 63:1–63:2 (Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2020). <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2020.63>.
92. Akers, C. et al. Holographic pseudoentanglement and the complexity of the AdS/CFT dictionary. Preprint at <http://arxiv.org/abs/2411.04978> (2024).
 93. Swingle, B., Bentsen, G., Schleier-Smith, M. & Hayden, P. Measuring the scrambling of quantum information. *Phys. Rev. A* **94**, 040302 (2016).
 94. Zdeborová, L. & Krzakala, F. Statistical physics of inference: Thresholds and algorithms. *Adv. Phys.* **65**, 453–552 (2016).
 95. Watanabe, K. & Takagi, R. Universal work extraction in quantum thermodynamics. *Nature Communications*, **17**(1), 1857 <https://doi.org/10.48550/arXiv.2504.12373> (2026).
 96. Lumbreras, J., Huang, R. C., Hu, Y., Gu, M. & Tomamichel, M. Quantum state-agnostic work extraction (almost) without dissipation. Preprint at <https://doi.org/10.48550/arXiv.2505.09456> (2025).
 97. Wilde, M. M. *Quantum Information Theory*. <https://www.cambridge.org/core/books/quantum-information-theory/9DC2CA59F45636D4F0F30D971B677623>. Cambridge University Press, 2013).
 98. Vershynin, R. High-Dimensional Probability: An Introduction with Applications in Data Science. in *Cambridge Series in Statistical and Probabilistic Mathematics*. <https://www.cambridge.org/core/books/highdimensional-probability/797C466DA29743D2C8213493BD2D2102>. (Cambridge University Press, 2018).
 99. Schön, C., Solano, E., Verstraete, F., Cirac, J. I. & Wolf, M. M. Sequential generation of entangled multiqubit states. *Phys. Rev. Lett.* **95**, 110503 (2005).

Acknowledgements

We thank Soonwon Choi, Bartek Czech, Jens Eisert, Philippe Faist, Yingfei Gu, Hsin-Yuan Huang, Jarrod R. McClean, Jinzhao Wang, and Yuxiang Yang for helpful discussions. We gratefully acknowledge support from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator, and the National Science Foundation (PHY-2317110). The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

Author contributions

H.Z. and Y.Z. conceived this project. H.Z., Y.Z., and J.P. contributed to the development of this project. H.Z. wrote the manuscript with inputs from all authors.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41534-026-01273-4>.

Correspondence and requests for materials should be addressed to Haimeng Zhao, Yuzhen Zhang or John Preskill.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026