

Entanglement-assisted Hamiltonian dynamics learning

Ayaka Usui,^{1,*} Guillermo Abad-López,^{1,2,†} Hari Krishnan
SV,^{1,‡} Anna Sanpera,^{1,3,§} and Some Sankar Bhattacharya^{1,¶}

¹*Física Teòrica: Informació i Fenòmens Quàntics, Departament de Física,
Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain*

²*Qilimanjaro Quantum Tech, Carrer de Venèçuela 74, 08019, Barcelona, Spain*
³*ICREA, Pg. Lluís Companys 23, 08010 Barcelona, Spain*

Approximating the dynamics generated by complex many-body Hamiltonians with simpler, experimentally feasible models is a central challenge at the intersection of quantum simulation and quantum learning. Quantum generative adversarial networks (QGANs) have recently emerged as a powerful framework for learning unitary dynamics, often achieving significantly more compact representations than standard Trotter-based approaches. However, their practical performance is frequently limited by training instabilities, including optimization plateaus and convergence to sub-optimal local minima, especially as system size and circuit depth increase. In this work, we introduce an entanglement-assisted learning strategy that enhances the training of QGANs through the controlled use of an ancillary qubit. Our approach dynamically enlarges the generator’s Hilbert space by coupling a randomly initialized ancilla to the system mid-training. We show that the combination of entanglement, randomization, and adaptive circuit extension reshapes the optimization landscape, enabling the learning process to escape stagnation and achieve higher fidelities. We analyze these behaviors by computing the bipartite ancilla-system von Neumann entanglement entropy during training, and utilizing Common Random Numbers to demonstrate the amplification of gradient variance upon ancilla insertion. Through numerical simulations, we demonstrate that this method substantially improves the approximation of nontrivial target dynamics (including Heisenberg models) using hardware-efficient ansatzes with limited locality. We further analyze the roles of ancilla connectivity, timing of introduction, and parameter initialization, identifying the key ingredients responsible for the observed performance gains. Our results highlight a general, theoretically grounded, and potentially scalable strategy for mitigating optimization bottlenecks in variational quantum algorithms, with applications in quantum simulation, control, and machine learning [1].

As a proof of concept, we consider a simple 3-local target Hamiltonian, $H_T = \sigma_z^1 \sigma_z^2 \sigma_z^3$, with the corresponding operator, $U_T = e^{-i\sigma_z^1 \sigma_z^2 \sigma_z^3}$, acting on three qubits for a unit time. We remark that the specific choice of target does not impact the generality of the learning procedure presented here. As already pointed out, the generator consists of multiple layers, each with nearest-neighbor interactions and local rotations, while the discriminator consists of a generalized local measurement, parametrized by a linear combination of Pauli operators before a local computational basis measurement in each qubit. Both the discriminator and generator parameters are randomly initialized at the start of each training.

This generator architecture becomes universal in the limit of infinitely many layers. Consequently, the expressivity of our generator increases with the number of layers. Based on our numerical results, the shorter gate sequence that performed sufficiently well is obtained using three layers. We fix this configuration for the remainder of our study.

Each training sequence consists of up to 3000 iterations or until a target fidelity of 99% is reached at iteration

i^* , i.e., $F_{i^*}^{(j)} \geq 0.99$, where $F_i^{(j)}$ denotes the fidelity at iteration i for the j th training run. For each run, we take as a figure of merit the maximum fidelity attained during training, $F_{\max}^{(j)} = \max_i F_i^{(j)}$. When multiple independent runs are performed, we report the average value:

$$F_{\max}^{\text{avg}} = \langle F_{\max}^{(j)} \rangle_j = \langle \max_i F_i^{(j)} \rangle_j.$$

As we will see, the overarching role of the ancilla is to help the generator escape local minima and plateaus. We analyze its impact from three complementary perspectives: (i) its connectivity with the generator qubits, (ii) the timing of its introduction, and (iii) its initial parameters.

Although the interaction between the ancilla and the generator qubits is always implemented via a ZZ gate, the way the ancilla qubit connects to the generator qubits is crucial. To explore this aspect, we consider several connectivity configurations for each layer, as shown in Fig. 1(a): (a1) the ancilla interacts with a single qubit; (a2) the ancilla interacts with two neighbor qubits; (a3) the ancilla interacts with two distant qubits at the edges; and (a4) the ancilla interacts with all generator qubits.

In Fig. 1(b), we depict the average fidelity F_i versus the iteration for a typical run where the training stagnates. After a few iterations, the fidelity reaches a steady value around $F \simeq 0.5$. Increasing the number of iterations does not improve the maximum fidelity, signaling a learning plateau. This is overcome by adding an ancilla in the

* ayaka.usui@uab.cat

† guillermo.abad@uab.cat

‡ harikrishnan.sasidharan@uab.cat

§ anna.sanpera@uab.cat

¶ somesankar.bhattacharya@uab.cat




FIG. 1. Ancilla-assisted learning when the time evolution operator, $e^{-i\sigma_z^1\sigma_z^2\sigma_z^3}$, is generated with a hardware-efficient ansatz. (a) Ancilla configurations for each layer. (b) Fidelity F_i versus training iteration. In a typical run, a learning plateau is signaled by a fidelity saturating at $F_i \ll 1$. The plateau is overcome after adding an ancilla at iteration $i = 3000$ (vertical dashed line) and letting the learning proceed. (c) The average maximum fidelity, F_{max}^{avg} , is plotted against the ancilla configurations. The dotted horizontal line indicates the reference fidelity achieved in the absence of an ancilla. (c1) Ancilla is added in the middle of training. The green dots (brown squares) denote the fidelity when the parameters of the ancilla are initiated randomly (set to zero). (c2) Ancilla is added from the beginning of training with random initial parameters.

middle of training with random initial parameters and configuration (a3).

Crucially, this escape from stagnation is physically driven by transient bipartite entanglement. By evaluating the von Neumann entropy of the ancilla's reduced density matrix, $S(\rho_a) = -\text{Tr}(\rho_a \log_2 \rho_a)$ with $\rho_a = \text{Tr}_{A,B}[\Psi_{Ga}\rangle\langle\Psi_{Ga}]$, we observe that the mid-training insertion injects maximal entanglement ($S(\rho_a) \approx 1$), which is subsequently consumed as the optimizer locates the global minimum. The asymptotic decay $S(\rho_a) \rightarrow 0$ acts as the exact algebraic signature of the generator block-diagonalizing to isolate the target unitary. Furthermore, to quantitatively prove the reshaping of the cost-function landscape, we utilize Common Random Numbers to evaluate the gradient vectors. By computing the Median Absolute Deviation (MAD), we demonstrate that this dynamical Hilbert space expansion effectively raises the gradient variance floor ($\text{MAD}(\nabla\mathcal{L}_{Ga}) \gg \text{MAD}(\nabla\mathcal{L}_G)$), mitigating the barren plateau.

The timing at which the ancilla is introduced also strongly influences the learning process. We examine two scenarios: in the first, the ancilla is present from the start of training ($i = 0$); in the second, it is added midway through the training ($i = 3000$).

Entanglement-assisted learning works by opening new directions in the optimization landscape, allowing the system to escape stagnation. For this mechanism to work, the ancilla must be introduced during the middle of training when stagnation occurs. Random initial parameters are crucial to effectively increase the phase space minimization landscape. Randomization produces a “kick” to the system, which makes it jump out of the learning plateau. Conversely, when the ancilla parameters are initially set to zero (no rotation), only marginal improvement is observed, indicating that a nontrivial initialization—or “kick”—is essential. Although randomization alone is not enough, when some parameters of the generator are randomized without expanding the phase space with an ancillary qubit, no improvement is achieved, as analyzed in our extended study [1].

Our numerical results confirm such a description. When the ancilla is added at the start of the training ($i = 0$), the improvement in fidelity is practically irrelevant ($\leq 3\%$), irrespective of the connectivity of the ancillary qubit (a1-a4), as shown in Fig. 1(c2).

On the other hand, clear improvement is observed in Fig. 1(c1), when the ancilla is added in the middle of the training process ($i = 3000$), with randomized parameters. In this case, the average maximum fidelity increases by over 11% of its original value if the ancilla is connected to all the qubits (a4), although fidelity also increases substantially ($> 8.5\%$) if connected to two, not previously connected, qubits (a3). The reason for this improvement is formalized through a dimensional expressivity analysis, $\mathcal{E}_n := \text{rank}(J_n)$, where J_n is the Jacobian matrix of the circuit configuration (an) with respect to its parameters. As rigorously demonstrated in our ex-

tended study [1] via *Proposition 1* and numerical SVD thresholding, the expressivity follows the strict hierarchy $\mathcal{E}_4 > \mathcal{E}_3 > \mathcal{E}_2 > \mathcal{E}_1$. This highlights a key topological consequence: while (a2) creates redundant pathways between adjacent qubits, configuration (a3) bridges distant qubits, shrinking the graph diameter and increasing expressivity. Importantly, we have verified that this entanglement-assisted recovery and its variance ampli-

fication perfectly generalize to complex, non-commuting target dynamics, including Heisenberg Hamiltonian models [1].

REFERENCES

- [1] G. A.-L. et al., [Entanglement-assisted hamiltonian dynamics learning](#) (2026), [arXiv:2602.15931 \[quant-ph\]](#).