Arbitrary state preparation in quantum harmonic oscillators using neural networks

Nicolas Parra-A[a], Vladimir Vargas-Calderón[b], Herbert Vinck-Posada[a]

[a]*Grupo de Superconductividad y Nanotecnología, Departamento de Física, Universidad Nacional de Colombia, Bogotá, Colombia*
[b]*D-Wave Systems, Burnaby, British Columbia, Canada*

Abstract

Preparing quantum states is a fundamental task in various quantum algorithms. In particular, state preparation in quantum harmonic oscillators (HOs) is crucial for the manipulation of qudits and the implementation of high-dimensional algorithms. In this work, we develop a general methodology for quantum state preparation in an HO coupled to an auxiliary qubit, guaranteeing that any target state is physically preparable. Both the qubit and the HO are driven by two lasers with time-dependent phase modulation. The modulation times and phase values are generated by a neural network whose input is the desired target state. In contrast to conventional quantum control approaches, this framework eliminates the need for per-instance optimization of the control protocol. Instead, the control parameters required to prepare an arbitrary quantum state of the HO are obtained directly from a single forward pass through the neural network. Specifically, we present results for preparing arbitrary qubit, qutrit, and qudit ($n = 4$) states in the HO, achieving average fidelities of 99.99%, 99.5%, and 98.9%, respectively, across random target states.

Keywords: Quantum state preparation, Harmonic oscillators, Neural networks, Qudits, Jaynes–Cummings model

1. Introduction

State preparation refers to the process of transitioning a system from an initial state to a target state through a series of interactions applied to the system. This process is a fundamental step in many quantum technologies and, in particular, in quantum algorithms [1]. Improperly preparing the initial states can significantly affect the algorithm's behavior [2]. For instance, in the simulation of physical systems, an initial state must be prepared before evolving it [3, 4]. In quantum machine learning algorithms, information needs to be encoded into an initial state, which is then used to make predictions [5, 6]. Other examples requiring state preparation include algorithms for solving equations [7, 8], quantum chemistry [9, 10], quantum metrology [11, 12], and other applications such as quantum communications and quantum memories [13, 14, 15].

The literature offers various approaches to determine the physical interactions needed to guide a system from an initial state to a target state. Some of these works employ Krotov's methods [16], quantum steering [17], quantum circuits [18, 19], reinforcement learning [20, 21] and other techniques [22, 23, 24, 25, 26]. However, these methods require the construction of an optimization problem for each target state–which makes them less suitable for applications requiring rapid or repeated state preparation across many different targets–or techniques that require considerable computational resources. Other, more analytical techniques aim to solve the dynamics required to reach a specific target state from an initial state [27, 28]. These methods often yield a separate family of equations for each target state, making them impractical for multiple applications. Even though approximate methods have been proposed as alternatives [29, 30, 31, 32], a truly practical solution demands a technique capable of rapidly receiving the target state on demand and predicting the interactions required in the physical system, ensuring efficient state preparation.

In this work, we propose using a neural network in which the input is the target state to be prepared, and the output consists of the interactions required to achieve that state. We focus on preparing states in n-dimensional systems, also called qudits. These systems have attracted growing interest in recent years due to their enhanced capacity for encoding information, robustness to noise, and applications in high-dimensional quantum computing [33, 34, 35, 36, 37]. The challenge of using these systems lies in the complexity of their manipulation, especially when it comes to state preparation [38]. In particular, we consider using harmonic oscillators (HOs) to realize control on a finite subset of an HO's levels.

In general, HOs cannot be fully controllable under a finite set of steps [39, 28]. To address this, an auxiliary system is employed. In the case of an HO, an auxiliary qubit is used to interact with the oscillator over a finite period. Under certain approximations, such as the rotating wave approximation [40], an interacting Jaynes-Cummings-type system is formed, which is known to be approximately controllable within a truncated n-level manifold [41].

Superconducting quantum electrodynamics circuits can realize such systems [42] with Rabi oscillation periods of tens of nanoseconds [43]. Both the HO and the qubit can be driven by microwave control lines where the phase can be modulated rapidly at a few nanoseconds or at the sub-nanosecond scale [44]. As a result, multiple phase switches can be executed within a single Rabi period. This substantially simplifies numerical simulations of the system dynamics: between successive switches, the control Hamiltonian reduces to that of a resonant drive with fixed amplitude and phase. The result-

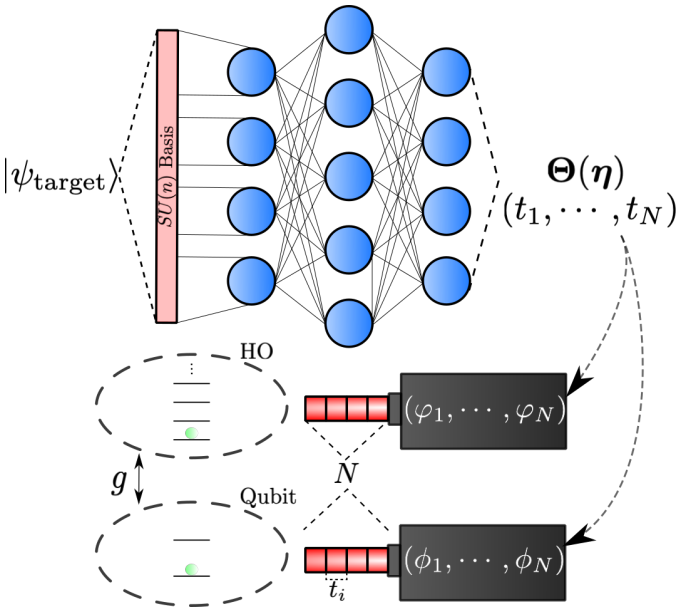

Figure 1: Scheme of the proposed methodology. The target state passes through a layer that computes the expectation values of the elements of the $SU(n)$ basis. These values are fed into a neural network with parameters $\boldsymbol{\eta}$. The network outputs the phases of a sequence of N pulses applied to the HO ($\boldsymbol{\varphi}$) and to the ancillary qubit ($\boldsymbol{\phi}$), as well as the switch times $\boldsymbol{t}$ at which phase changes occur. These pulses are applied to the system's initial state $|0,0\rangle$, leaving the harmonic oscillator prepared in the target state. More details are given in section 2.

ing piecewise time-independent form of the Jaynes-Cummings-type Hamiltonian can be straightforwardly integrated to obtain the system's dynamics.

Taking this into account, in our proposal, the neural network will receive the target state to be prepared in the HO as input, and the output will consist of the parameters of the drivings applied both to the HO and to the ancillary qubit to evolve the HO into the target state. In fig. 1, we show a schematic representation of the methodology.

The quality or precision of our model is evaluated using a state fidelity criterion, which is far superior to other neural-network-based proposals in the literature (cf. the noiseless case in [45]). Additionally, we investigate how the number of pulses in the sequence enhances the precision with which the state can be prepared. The term "arbitrary" refers to target states within the finite, truncated Hilbert space considered in the model. Its use is justified by the controllability result of Pinna and Panati [41], which shows that the Jaynes–Cummings-type dynamics is approximately controllable in such a truncated n-level manifold. Therefore, the neural network is trained to predict, without per-instance optimization, pulse sequences for representative cases ($n = 2, 3$, and 4) where this theoretical reachability applies.

The structure of this paper is as follows. In section 2, we describe the physical model, introduce the control parametrization, and detail the full training pipeline used to learn pulse sequences from target states. In section 3, we then present and analyze the numerical results for qubit, qutrit, and qudit ($n = 4$) state preparation, including a study of how the number of pulses affects the achievable fidelity. Finally, in section 4 we summarize the main findings, discuss experimental considerations and limitations, and outline possible directions for future work.

2. Method

In this section, we detail our proposal for using a neural network to predict the pulse parameters necessary to bring an HO coupled to a qubit to a separable state such that the HO state is any target state, when the HO basis is truncated to the first n levels.

The system under consideration, i.e., an HO coupled to a qubit, can receive pulses on both the oscillator and the qubit. The basis of the quantum states of this system is $|\alpha\rangle \otimes |\beta\rangle$, where $|\alpha\rangle \in \{|0\rangle, \cdots, |n-1\rangle\}$ and $|\beta\rangle \in \{|0\rangle, |1\rangle\}$. Note that we truncate the HO levels up to n excitations to be able to represent states and operators in a computer as vectors and matrices. Sometimes we drop the tensor product notation to save space. Modeling the HO-qubit interaction as a Jaynes-Cummings-type system, the Hamiltonian describing this system is [46] ($\hbar = 1$):

$$\hat{H}(t) = g(\hat{a}\hat{\sigma}_+ + \hat{a}^\dagger\hat{\sigma}_-) + \zeta(t)(e^{i\phi(t)}\hat{\sigma}_- + e^{-i\phi(t)}\hat{\sigma}_+) + \xi(t)(e^{i\varphi(t)}\hat{a} + e^{-i\varphi(t)}\hat{a}^\dagger) + \Delta_c\hat{n} + \Delta_{\omega_{eg}}\frac{\hat{\sigma}_z}{2}, \quad (1)$$

where $\hat{a}$ is the HO annihilation operator ($\hat{a}^\dagger$ is the creation operator), $\hat{\sigma}_+$ and $\hat{\sigma}_-$ are the raising and lowering operators of the qubit, g is the oscillator-qubit interaction constant, ζ and ξ are the amplitudes of the driving pulses, and ϕ, φ are the phases of the pulses for the qubit and HO, respectively. Δ_c and $\Delta_{\omega_{eg}}$ are the detunings of the HO and the qubit with respect to the driving lasers, respectively.

In this work, we assume resonance of the qubit and the HO with the laser ($\Delta_{\omega_{eg}} = \Delta_c = 0$). Further, we assume constant and equal amplitudes for both the HO and the qubit, i.e., $\zeta(t) = \xi(t) = \frac{\Omega}{2}$, for some pumping constant Ω. We model the phases of the drivings with a composition of pulses [47]—a technique to increase the robustness of quantum control by relying on multiple constant pulses—using piecewise constant functions of the form $\phi(t) = \phi_i$ for $t_{i-1} < t < t_i$, defined by the phase values $\{\phi_i\}_{i=1}^N$ and by the switch times $\{t_i\}_{i=1}^N$, where $t_0 = 0$. We assume synchronization of both lasers, so that the switch times are shared for both lasers, but the phase values for the HO are independent parameters $\{\varphi_i\}_{i=1}^N$. The total evolution time to prepare a target state $|\psi_{\text{target}}\rangle$ in the HO is t_N. As the initial state, we always consider the state $|0,0\rangle$ because this is typically the state of the unexcited system at low temperatures.

With these considerations, the evolution to a target state is given by:

$$\begin{aligned}\rho &= \text{Tr}_q\left[|\psi(\boldsymbol{\Theta})\rangle\langle\psi(\boldsymbol{\Theta})|\right] \\ &\equiv \text{Tr}_q\left[\prod_{i=N}^{1} e^{i\hat{H}(\phi_i,\varphi_i)(t_{i-1}-t_i)}\,|0,0\rangle\langle 0,0|\prod_{i=1}^{N} e^{-i\hat{H}(\phi_i,\varphi_i)(t_{i-1}-t_i)}\right] \end{aligned} \quad (2)$$

where $\text{Tr}_q[\cdot]$ is the partial trace of the qubit degrees of freedom, $\boldsymbol{\Theta}$ is shorthand notation for the pulse parameters $\{\phi_i\}_{i=1}^N \cup$

$\{\varphi_i\}_{i=1}^N \cup \{t_i\}_{i=1}^N$, and $\hat{H}(\phi, \varphi)$ makes the parameterization of the Hamiltonian explicit. Here, $|\psi(\mathbf{\Theta})\rangle$ denotes the final pure state generated by the pulse sequence, $|\psi(\mathbf{\Theta})\rangle = \hat{U}(\mathbf{\Theta})\,|0,0\rangle$. Parameters satisfying $\rho = \rho_{\text{target}} := \left|\psi_{\text{target}}\right\rangle\!\left\langle\psi_{\text{target}}\right|$ do exist, as pointed out by Pinna and Panati [41], who show that this preparation can be achieved, although, in general, the parameters must be functions of time. Since we only consider the phase modulation of both drivings, optimal constant pulse parameters might not satisfy eq. (2).

To achieve high-fidelity preparation of target states in the HO, we maximize

$$F = \text{Tr}\left[\left|\psi_{\text{target}}\right\rangle\!\left\langle\psi_{\text{target}}\right| \text{Tr}_{\text{q}}\left(|\psi(\mathbf{\Theta})\rangle\langle\psi(\mathbf{\Theta})|\right)\right]. \tag{3}$$

From this point onward, we refer to the infidelity $(1 - F)$ for convenience in the subsequent analyses.

There are many state-of-the-art alternatives for determining the pulse sequence required to prepare a state [16, 17, 48, 19, 20, 21, 22, 23]. However, these methods rely on iterative techniques that optimize a cost function, whose dimension increases as larger states are targeted, for each target state. The fact that each target state requires a separate solving routine makes these algorithms less suitable for real-world implementations requiring rapid encoding and/or state preparation across many different targets, such as in quantum machine learning [5]. In contrast, the proposed neural-network approach shifts the computational burden to a one-time offline training phase, after which control parameters for any new target state are obtained in a single forward pass. This makes the per-state inference cost independent of the particular target state, which is advantageous in scenarios requiring repeated or rapid state preparation across many different targets. In this work, we propose using a neural network to predict the pulse parameters $\mathbf{\Theta}$ as a function of the input target state.

Our neural network is, thus, a parameterized function $f_{\boldsymbol{\eta}} : \mathbb{R}^d \to \mathbb{R}^{3N}$, where d real numbers characterize the information in $\left|\psi_{\text{target}}\right\rangle$ and $3N$ real numbers specify the Hamiltonian parameters. $\boldsymbol{\eta}$ are the neural network parameters (for more details, see appendix Appendix A). We can thus write the infidelity in eq. (3) explicitly as a function of these parameters as

$$1 - F(\psi_{\text{target}}, \boldsymbol{\eta}) = \left\langle\psi_{\text{target}}\middle|\rho_{\text{prepared}}(\boldsymbol{\eta})\middle|\psi_{\text{target}}\right\rangle, \tag{4}$$

where $\rho_{\text{prepared}}(\boldsymbol{\eta})$ is the state prepared in the HO, given by

$$\rho_{\text{prepared}}(\boldsymbol{\eta}) = \text{Tr}_{\text{q}}\left(\left|\psi(\mathbf{\Theta}(\boldsymbol{\eta}))\right\rangle\!\left\langle\psi(\mathbf{\Theta}(\boldsymbol{\eta}))\right|\right). \tag{5}$$

The target state is fed into the neural network by passing the expected values of the $d = n^2 - 1$ matrices that form a complete basis of $SU(n)$ [49] with respect to the target state $\left|\psi_{\text{target}}\right\rangle$ (for more details, see appendix Appendix B). A schematic of the information flow through the neural network can be seen in Figure 1. The neural network parameters are optimized using stochastic gradient descent [50]. In the machine learning jargon, carrying out this optimization is called "training" the neural network. Training is done by minimizing the average infidelity over random target states, i.e.,

$$C(\boldsymbol{\eta}) = 1 - \mathbb{E}\left[F(\psi_{\text{target}}, \boldsymbol{\eta})\right] \tag{6}$$

where the expected value in eq. (6) is taken over random states $\left|\psi_{\text{target}}\right\rangle$, sampled from $SU(n)$ according to the Haar measure [51]. Thus, we aim to minimize the average infidelity between a set of target states and the states prepared using the phases predicted by the neural network. As previously mentioned, this is done through stochastic gradient descent, which requires $\nabla_{\boldsymbol{\eta}} C(\boldsymbol{\eta})$ to be calculated. In modern software, this can be easily achieved through automatic differentiation [52].

3. Results

In this section, we present results on the performance of the neural network in minimizing the cost function defined in Eq. (6). From this point on, we fix the energy scale by choosing $g = 1$. The training datasets used in the results we present contain 4096 states, sampled using the Haar measure [51]. This number was chosen by characterizing how many states are required for the infidelity to converge, i.e., such that increasing the number of states in the training dataset does not yield any improvement. Additionally, in our numerical simulations we truncate the HO Hilbert space to a finite dimension $n_{\text{comp}} > n$ and embed the desired n-level target state in this larger computational basis. This accounts for truncation (edge) effects: during the pulse sequence, population can transiently leak into levels above n, even though we evaluate the final infidelity only within the target n-dimensional subspace. We set $n_{\text{comp}} = 6$, as previous convergence tests indicated that this cutoff is sufficient for state-preparation simulations with $n = 2$, $n = 3$, and $n = 4$.

3.1. Number of pulses

One of the most common aspects in the literature on pulse composition is that increasing the number of pulses in the sequence improves the quality of the preparation [22, 23]. To evaluate this effect, we trained a separate model for each number of pulses ($N = 2$–14) and qudit dimension ($n = 2$–4). For each setting, we sampled 1000 target states from the Haar measure and computed the average infidelity of the prepared states. The results are shown in fig. 2.

In general, we observe that increasing the number of pulses (or laser phase shifts) reduces the infidelity, in agreement with the literature. However, for each n there is an optimal point, since increasing the number of pulses beyond that does not further reduce the preparation infidelity. In the qubit case, the lowest infidelity occurs at $N = 9$, achieving an average preparation infidelity of 0.01%. This infidelity is highly competitive compared to different state-of-the-art techniques. For qutrits, the infidelity already starts to increase noticeably, reaching 0.5% with a sequence of $N = 11$, although it remains a low preparation infidelity. Finally, for qudits with $n = 4$, the infidelity increases slightly further, reaching 1.1% at $N = 13$. Although this result may not appear as strong, there are not many techniques that report performance for qudit ($n = 4$) preparation, so it provides an interesting starting baseline. On the other hand, the algorithm

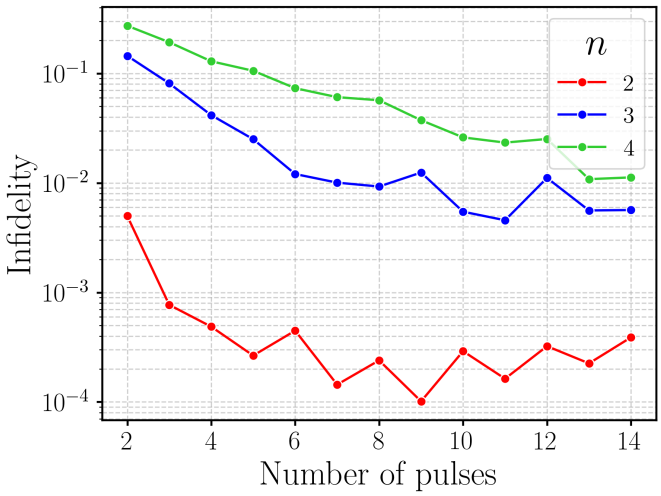

Figure 2: Preparation infidelity using the proposed neural network for 1000 different states as a function of the number of pulses. The red line corresponds to preparing qubit states in the HO, the blue line corresponds to preparing qutrit states in the HO, and the green line corresponds to preparing qudit states ($n = 4$).

appears to exhibit a strong dependence on the number of qudit levels, which is expected because the Hilbert space increases in size; however, given the relatively small increase from $n = 3$ to $n = 4$, the drop does not seem abrupt. This will require future scaling analyses.

3.2. Qubit performance

We now analyze the state-preparation performance in the qubit case ($n = 2$). As discussed in section 3.1, the best-performing model is obtained with a sequence of $N = 9$ pulses, reaching an average preparation infidelity of 0.01% over 1000 Haar-random target states.

To better understand how the model works to prepare a state, fig. 3 illustrates three representative target states—the Hadamard states $|+\rangle$ and $|-\rangle$, and the magic state $|T\rangle$—as they are prepared using the pulse sequence predicted by the neural network. The first set of panels shows the trajectory of the qubit prepared state on the Bloch sphere after each pulse. Notably, individual pulses do not necessarily move the state monotonically closer to the target; rather, the early pulses can drive the system through intermediate directions in state space, while the final pulses perform the fine corrections that align the trajectory with the desired target state.

The figure also reports the phase-switching pattern predicted for each pulse applied to the HO and to the ancillary qubit. No simple, easily recognizable pattern is expected in these phases: they emerge from the learned control strategy and are precisely the degrees of freedom that allow the protocol to compensate accumulated phase and population errors along the sequence.

Finally, the qubit purity is shown throughout the preparation. It typically decreases at intermediate steps, indicating transient entanglement between the qubit and the HO during the control process, and then returns close to unity at the end of the sequence, consistent with the intended final separable state. Achieving such infidelity preparation for states that are central to quantum information processing—particularly non-stabilizer resource states such as the magic state $|T\rangle$, which are often challenging for alternative preparation strategies—highlights the competitiveness of the proposed approach.

3.3. Qutrit performance

We now analyze the state-preparation performance in the qutrit case ($n = 3$). As discussed in section 3.1, the best-performing model is obtained with a sequence of $N = 11$ pulses, reaching an average preparation infidelity of 0.5% over 1000 Haar-random target states. In this setting, the dynamics cannot be visualized as directly as in the qubit case using a Bloch-sphere representation; instead, we characterize the preparation quality through the HO Wigner function,

$$W(x,p) = \frac{1}{\pi\hbar}\int_{-\infty}^{\infty} \langle x-y|\hat{\rho}|x+y\rangle\, e^{2ipy/\hbar}\, dy. \tag{7}$$

Figure 4 compares the HO Wigner function of the prepared state (first column) against that of the target state (second column) for two representative qutrit targets. Each row corresponds to a different target state, $\frac{1}{\sqrt{3}}(|0\rangle + |1\rangle + |2\rangle)$ and $\frac{1}{\sqrt{3}}(|0\rangle + |1\rangle - |2\rangle)$, respectively. We evaluated the absolute pointwise difference $|W_{\text{target}}(x,p) - W_{\text{prepared}}(x,p)|$, but the differences are practically negligible.

Overall, the prepared and target Wigner functions show excellent agreement: the model reproduces the dominant phase-space support as well as the finer nonclassical structure. The remaining discrepancies are small and typically localized, consistent with the slight increase in average infidelity compared to the qubit case. These results confirm that the proposed approach remains highly effective for qutrit state preparation.

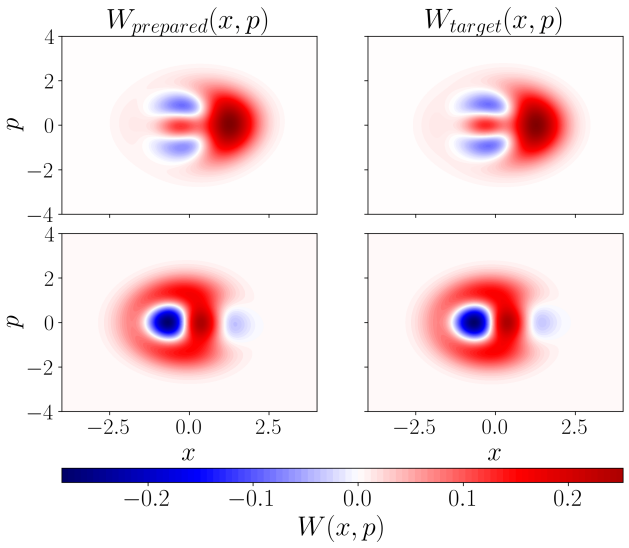

Figure 4: Representative qutrit-state preparation for the target states $\frac{1}{\sqrt{3}}(|0\rangle + |1\rangle - |2\rangle)$ and $\frac{1}{\sqrt{3}}(|0\rangle + |1\rangle + |2\rangle)$ with average infidelity 0.5% for a pulse sequence of $N = 11$. The first column shows the Wigner function of the prepared HO state and the second column shows the target-state Wigner functions.

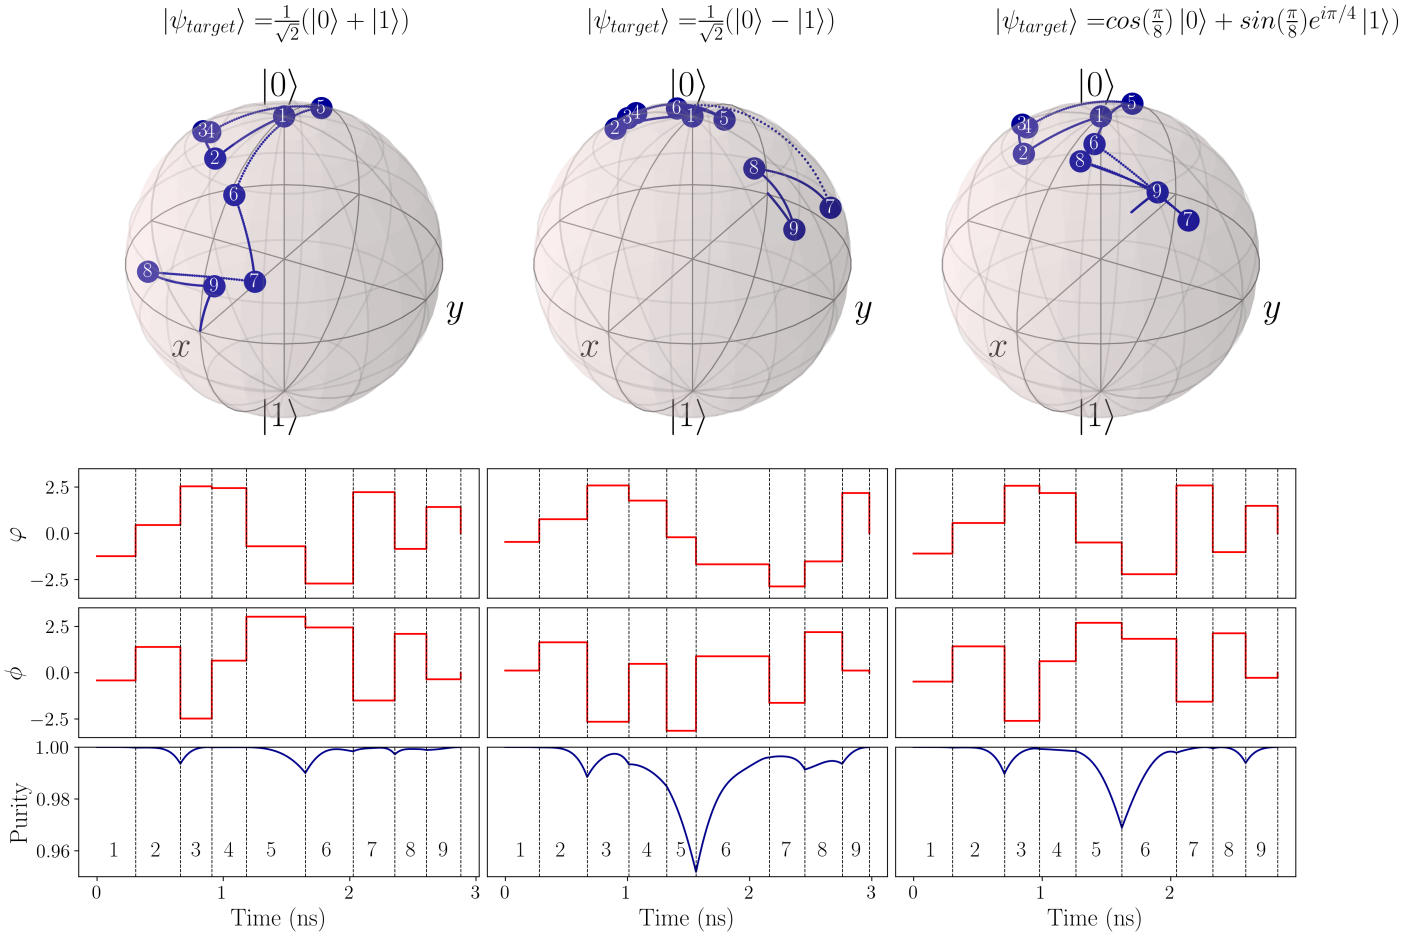

Figure 3: Three representative qubit-state preparation examples (average infidelity 0.01% for $N = 9$). The panels show the target state on the Bloch sphere, the neural-network-predicted phase-switching sequences, and the purity evolution during the preparation.

3.4. Qudit ($n = 4$) performance

We now turn to the qudit case with $n = 4$. As discussed in section 3.1, the best-performing model is obtained with $N = 13$ pulses, achieving an average preparation infidelity of 1.1% over 1000 Haar-random target states. Before discussing these examples, we clarify the sense in which two-qubit entangled states are encoded in the oscillator. The four-dimensional subspace spanned by the first oscillator levels is isomorphic to two logical qubits, and we use the identification $|0\rangle \equiv |00\rangle$, $|1\rangle \equiv |01\rangle$, $|2\rangle \equiv |10\rangle$, and $|3\rangle \equiv |11\rangle$. Under this tensor-product structure, entanglement is defined between the two logical qubits encoded in the oscillator subspace, as in standard bosonic-qubit encodings [42].

Figure 5 summarizes representative $n = 4$ preparation examples using the same Wigner-function diagnostics as in the qutrit case. With the encoding above, the oscillator states $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |3\rangle)$ and $|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|1\rangle - |2\rangle)$ correspond to the Bell states $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ and $\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, respectively. Each row corresponds to one of these target states. The first column shows the Wigner function of the prepared HO state, the second column shows the target-state Wigner function. Even at $n = 4$, the learned pulse sequences are able to capture the main phase-space structure of the target states, including interference patterns and negative regions that signal nonclassicality. The remaining infidelity is associated with small, spatially localized deviations rather than a global distortion of the Wigner distribution, suggesting that the protocol is already close to the target and primarily limited by fine feature resolution as the Hilbert-space dimension increases. Nevertheless, achieving this level of performance for $n = 4$ qudits is particularly relevant because it demonstrates the ability to prepare nontrivial two-qubit entangled states encoded within the oscillator subspace, a key resource for quantum-information processing. Bell states are especially useful benchmarks in this setting because they are simple enough to interpret under the two-logical-qubit mapping, while still requiring coherent superpositions between separated oscillator levels and well-defined relative phases. Their preparation therefore tests not only population transfer within the four-level manifold, but also the phase coherence needed to generate entanglement-like resources in a single bosonic mode.

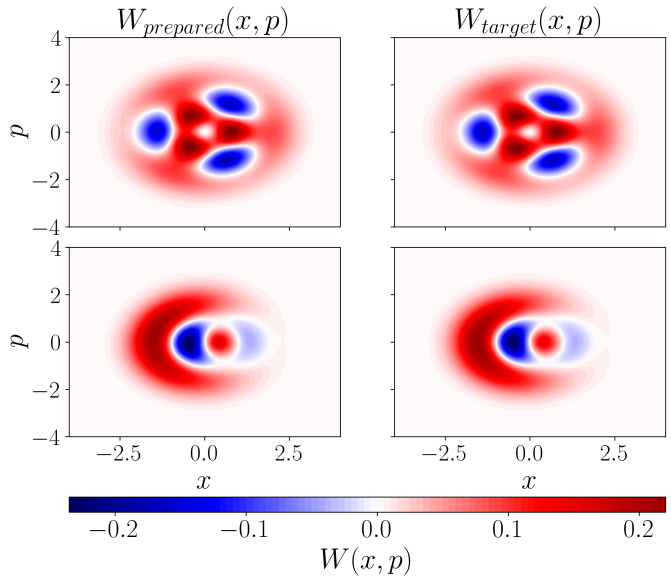

Figure 5: Representative qudit-state ($n = 4$) preparation for the target Bell states $|\Phi^+\rangle$ and $|\Psi^-\rangle$ with average infidelity 1.1% for a pulse sequence of $N = 13$. The first column shows the Wigner function of the prepared HO state and the second column shows the target-state Wigner function.

4. Discussion

4.1. Computational cost and practical advantages

A meaningful assessment of the practical value of the proposed method requires distinguishing three time scales that are often conflated, and which we summarize in Table 1. The first is the *offline training cost*: a one-time optimization of the neural-network parameters $\boldsymbol{\eta}$ for a fixed Hilbert-space dimension n, using the training set of 4096 Haar-random states described above. This cost is incurred once per dimension n and is fully amortized over all subsequent uses of the trained model. It should be noted that this training cost is incurred separately for each Hilbert-space dimension n, as the network architecture and training set are adapted to each case; this represents an overhead that is absent in per-instance optimal-control methods. However, once a model is trained for a given n, it serves as a universal controller for *all* target states in that dimension, so the training cost is amortized over the full set of subsequent inference calls, in contrast to GRAPE- or Krotov-type methods, where the dominant computational cost is incurred at inference time and repeated independently for every new target state. The practical advantage of the proposed approach therefore depends on the number of distinct target states required: for applications involving many different target states, the amortized cost per state of the neural-network approach is substantially lower. The second is the *online inference time*: once training is complete, the platform operates in an on-demand regime, where obtaining the pulse parameters $\boldsymbol{\Theta}$ for a new target state requires only a single forward pass through the network. In our implementation, this on-demand generation of controls takes approximately 100 ms on a standard CPU. This stands in sharp contrast with iterative optimal-control methods such as Krotov or GRAPE, where each new target state requires a separate optimization loop that typically demands 10^2–10^3 iterations, each involving a full simulation of the system dynamics, amounting to runtimes of the order of minutes to hours per target state depending on the system dimension and convergence criterion [16, 48]. The third is the *physical control time* t_N: the total duration of the pulse sequence applied to the hardware. In our dimensionless units ($g = 1$), a single Rabi period is $T_{\text{Rabi}} = 2\pi/g \approx 6.28$, and the full $N = 9$–13 pulse sequences span a total time t_N of a few Rabi periods. Mapping to circuit-QED hardware, where $g/2\pi \sim$ 10–100 MHz [43], this corresponds to physical control times of order 10–500 ns, which is comparable to pulse durations reported in conventional optimal-control implementations on similar platforms [16]. Thus, the primary advantage of the proposed method is not a reduction in physical pulse duration, which remains platform dependent and broadly comparable across approaches, but rather the ability to generate control parameters on demand for an arbitrary new target state in ∼100 ms. This on-demand capability makes the method particularly suitable for applications requiring rapid or repeated state preparation, such as quantum machine-learning workflows where the target state changes with each input datum [5].

4.2. Limitations and outlook

Although the proposed approach eliminates per-instance optimization, its scalability is constrained by two distinct factors that are worth distinguishing. The first is the capacity of the neural network itself: the observed increas in average infidelity as n increases from 2 to 4 reflects the growing complexity of the mapping from target state to pulse parameters, which requires a more expressive model to be approximated accurately. This is not a fundamental limitation of the framework. By the universal approximation theorem [53], a feedforward neural network with sufficient width or depth can approximate any continuous function on a compact domain to arbitrary precision. In practice, this means that increasing the network size and complexity—larger hidden layers, additional depth, or more expressive architectures—can in principle recover the fidelity lost as n grows. Indeed, the network architectures used here were already found to grow with n (hidden-layer widths of [209, 173, 222, 347] for $n = 2$, [197, 264, 304] for $n = 3$, and [321, 442, 425] for $n = 4$; see Appendix A), and systematic exploration of larger architectures at fixed n is a natural next step. The number of pulses N also plays a role analogous to network capacity: as shown in Fig. 2, increasing N consistently reduces the infidelity, and larger n requires more pulses to reach near-optimal performance ($N = 9, 11, 13$ for $n = 2, 3, 4$ respectively), suggesting that both network size and pulse number should be scaled jointly as n increases.

The second, and more fundamental, constraint is the cost of generating training data. Each training sample requires simulating the full quantum dynamics of the HO-qubit system in a Hilbert space of dimension $2n_{\text{comp}}$, where the computational cutoff $n_{\text{comp}} > n$ must be enlarged as n grows to avoid truncation artifacts. The cost of each simulation scales as $O(n_{\text{comp}}^3)$ per pulse segment due to matrix exponentiation, and the dataset size required for convergence is also expected to grow with

Method	Offline cost	Online cost per target	Physical control time	Typical fidelity
Optimal control (Krotov/GRAPE-type) [16, 48]	None, apart from method setup	Separate iterative optimization per target; 10^2–10^3 iterations, each requiring a full dynamics simulation; total runtime of order minutes to hours per target state	Platform-dependent; tens to hundreds of ns in circuit-QED implementations [16, 43]	Infidelities $\leq 10^{-4}$ in benchmark settings
Reinforcement-learning control [20, 21]	Training of a policy or value function; hours of computation	Policy evaluation after training; no full optimal-control loop; sub-second to second-scale inference	Platform-dependent	Task-dependent; high-fidelity benchmarks reported
This work	One-time neural-network training for each n using 4096 Haar-random states	Single forward pass; ∼100 ms per target state on a standard CPU	N = 9–13 phase segments; t_N spans a few Rabi periods, corresponding to ∼10–500 ns in circuit-QED with $g/2\pi \sim$ 10–100 MHz [43]	99.99% (n = 2), 99.5% (n = 3), 98.9% (n = 4)

Table 1: Comparison of the relevant time scales for the present approach and representative control strategies. The table separates the offline optimization or training cost, the online cost of producing controls for a new target state, and the physical duration of the applied pulse sequence. The physical control time is hardware dependent for all methods; for the present phase-switching protocol it is set by the number of pulse segments and the experimentally available Rabi scale.

n. For the cases studied here, $n_{\mathrm{comp}} = 6$ and 4096 training states were sufficient; however, both quantities will need to increase for larger n, making classical simulation the practical bottleneck rather than the neural network itself. Extending the method to the larger subspace dimensions required by bosonic-encoding applications—such as cat qubits, GKP codes, or binomial codes, which typically involve tens to hundreds of oscillator levels—will likely require strategies beyond direct simulation, such as tensor-network-based state evolution, hardware-in-the-loop training, or transfer learning across dimensions. These directions are left for future work.

A central experimental requirement of the proposed protocol is the ability to implement rapid and accurate phase updates of the driving fields. Our control Hamiltonian is modeled as a sequence of segments with constant amplitude and constant phase; thus, the switching time (and the bandwidth of the control chain) sets the shortest achievable segment duration and, consequently, the finest time resolution available to the controller.

This assumption is well aligned with circuit-QED implementations, where microwave sources combined with IQ mixers and modern arbitrary-waveform generators can realize phase switches on the nanosecond scale, enabling multiple phase switches within a single Rabi period. In this regime, the piecewise-constant model is realistic and the learned sequences can be deployed after standard calibration steps (e.g., correcting for IQ imbalance [54, 55], finite rise times [56], and static phase offsets [57]).

However, the technique becomes unsuitable in platforms where the required control timescales are substantially shorter than the achievable electronic phase-update time. Such is the case of excitonic dynamics in semiconductor quantum dots [58] or polaritonic systems driven by ultrafast lasers [59]. In such settings, a digitally switched, nanosecond-resolved phase protocol cannot resolve the dynamics, and the piecewise-constant approximation breaks down.

Beyond switching speed, additional non-idealities may reduce the achievable fidelity in an experiment: (i) decoherence of both the HO and the ancilla during the total preparation time T; (ii) amplitude and phase distortions due to finite bandwidth and filtering; (iii) imperfect measurement of Hamiltonian coupling strengths and detunings; and (iv) leakage outside the truncated HO subspace (especially as n increases) and residual contributions from terms neglected under the rotating-wave approximation. These effects suggest that the reported fidelities should be interpreted as an upper bound for a given hardware platform, and that incorporating realistic noise models and robust-training strategies will be important for future work.

A more realistic treatment that includes open-system dynamics and imperfect control—e.g., decoherence during the pulse sequence as well as errors and jitter in the target phases and switch times—will be the subject of future work.

5. Conclusions

In this work, we introduced a neural-network-assisted strategy for arbitrary state preparation in a quantum harmonic oscillator (HO) coupled to an ancillary qubit. The key idea is to train a single feed-forward model that receives a desired target state as input and outputs, in a single forward pass, the parameters of a piecewise-constant phase-modulated driving sequence acting on both the HO and the qubit. By building on the controllability of the effective Jaynes–Cummings-type dynamics within a truncated n-level manifold, the approach guarantees that the predicted control protocol corresponds to a physically preparable target within the adopted model and truncation.

Our numerical results show that the method achieves high-quality preparation across Haar-random targets without per-instance optimization. For qubit targets (n = 2), we obtained an average fidelity of 99.99% (infidelity 0.01%) using N = 9 pulses; for qutrit targets (n = 3), we reached 99.5% (infidelity of 0.5%) with N = 11 pulses; and for qudit targets with n = 4, we achieved 98.9% (infidelity 1.1%) with N = 13 pulses. These results, together with the n = 4 Bell-state examples, indicate that the method can serve as an on-demand primitive for single-qudit preparation and for generating logical entanglement resources within a bosonic mode. Future work will focus on incorporating open-system dynamics, hardware distortions, and robust-training strategies to assess performance under realistic experimental noise.

ACKNOWLEDGMENTS

N. P.-A. and V. V.-C. would like to acknowledge support from the project “Aprendizaje de Máquina para Sistemas Cuánticos" HERMES code 57792 and QUIPU code 201010040147, and H. V.-P. from the project “Ampliación del uso de la mecánica cuántica desde el punto de vista experimental y su relación con la teoría, generando desarrollos útiles para metrología y computación cuántica a nivel nacional" BPIN code 2022000100133. We sincerely thank D. Martínez-Tibaduiza and D. B. Anghelo-Rodriguez for their valuable discussions, comments, and suggestions on this work.

References

[1] M. Rosenkranz, E. Brunner, G. Marin-Sanchez, N. Fitzpatrick, S. Dilkes, Y. Tang, Y. Kikuchi, M. Benedetti, Quantum state preparation for multivariate functions (2024). arXiv:2405.21058.
URL https://arxiv.org/abs/2405.21058

[2] M. Mohseni, A. Scherer, K. G. Johnson, O. Wertheim, M. Otten, N. A. Aadit, K. M. Bresniker, K. Y. Camsari, B. Chapman, S. Chatterjee, G. A. Dagnew, A. Esposito, F. Fahim, M. Fiorentino, A. Khalid, X. Kong, B. Kulchytskyy, R. Li, P. A. Lott, I. L. Markov, R. F. McDermott, G. Pedretti, A. Gajjar, A. Silva, J. Sorebo, P. Spentzouris, Z. Steiner, B. Torosov, D. Venturelli, R. J. Visser, Z. Webb, X. Zhan, Y. Cohen, P. Ronagh, A. Ho, R. G. Beausoleil, J. M. Martinis, How to build a quantum supercomputer: Scaling challenges and opportunities (2024). arXiv:2411.10406.
URL https://arxiv.org/abs/2411.10406

[3] E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. A. Eriksson, K.-M. C. Fu, M. Greiner, K. R. Hazzard, R. G. Hulet, A. J. Kollár, B. L. Lev, M. D. Lukin, R. Ma, X. Mi, S. Misra, C. Monroe, K. Murch, Z. Nazario, K.-K. Ni, A. C. Potter, P. Roushan, M. Saffman, M. Schleier-Smith, I. Siddiqi, R. Simmonds, M. Singh, I. Spielman, K. Temme, D. S. Weiss, J. Vučković, V. Vuletić, J. Ye, M. Zwierlein, Quantum simulators: Architectures and opportunities, PRX Quantum 2 (2021) 017003. doi:10.1103/PRXQuantum.2.017003.
URL https://link.aps.org/doi/10.1103/PRXQuantum.2.017003

[4] C. W. Bauer, Z. Davoudi, A. B. Balantekin, T. Bhattacharya, M. Carena, W. A. de Jong, P. Draper, A. El-Khadra, N. Gemelke, M. Hanada, D. Kharzeev, H. Lamm, Y.-Y. Li, J. Liu, M. Lukin, Y. Meurice, C. Monroe, B. Nachman, G. Pagano, J. Preskill, E. Rinaldi, A. Roggero, D. I. Santiago, M. J. Savage, I. Siddiqi, G. Siopsis, D. Van Zanten, N. Wiebe, Y. Yamauchi, K. Yeter-Aydeniz, S. Zorzetti, Quantum simulation for high-energy physics, PRX Quantum 4 (2023) 027001. doi:10.1103/PRXQuantum.4.027001.
URL https://link.aps.org/doi/10.1103/PRXQuantum.4.027001

[5] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, P. J. Coles, Challenges and opportunities in quantum machine learning, Nature Computational Science 2 (9) (2022) 567–576. doi:10.1038/s43588-022-00311-3.
URL https://doi.org/10.1038/s43588-022-00311-3

[6] F. A. González, V. Vargas-Calderón, H. Vinck-Posada, Classification with quantum measurements, Journal of the Physical Society of Japan 90 (4) (2021) 044002. doi:10.7566/JPSJ.90.044002.
URL https://doi.org/10.7566/JPSJ.90.044002

[7] P. Brearley, S. Laizet, Quantum algorithm for solving the advection equation using hamiltonian simulation, Phys. Rev. A 110 (2024) 012430. doi:10.1103/PhysRevA.110.012430.
URL https://link.aps.org/doi/10.1103/PhysRevA.110.012430

[8] H. Krovi, Improved quantum algorithms for linear and nonlinear differential equations, Quantum 7 (2023) 913. doi:10.22331/q-2023-02-02-913.
URL https://doi.org/10.22331/q-2023-02-02-913

[9] S. Fomichev, K. Hejazi, M. S. Zini, M. Kiser, J. Fraxanet, P. A. M. Casares, A. Delgado, J. Huh, A.-C. Voigt, J. E. Mueller, J. M. Arrazola, Initial state preparation for quantum chemistry on quantum computers, PRX Quantum 5 (2024) 040339. doi:10.1103/PRXQuantum.5.040339.
URL https://link.aps.org/doi/10.1103/PRXQuantum.5.040339

[10] B. Bauer, S. Bravyi, M. Motta, G. K.-L. Chan, Quantum algorithms for quantum chemistry and quantum materials science, Chemical Reviews 120 (22) (2020) 12685–12717. doi:10.1021/acs.chemrev.9b00829.
URL https://doi.org/10.1021/acs.chemrev.9b00829

[11] C. Couteau, S. Barz, T. Durt, T. Gerrits, J. Huwer, R. Prevedel, J. Rarity, A. Shields, G. Weihs, Applications of single photons in quantum metrology, biology and the foundations of quantum physics, Nature Reviews Physics 5 (6) (2023) 354–363. doi:10.1038/s42254-023-00589-w.
URL https://doi.org/10.1038/s42254-023-00589-w

[12] Y. Li, Z. Ren, Quantum metrology with an n-qubit w superposition state under noninteracting and interacting operations, Phys. Rev. A 107 (2023) 012403. doi:10.1103/PhysRevA.107.012403.
URL https://link.aps.org/doi/10.1103/PhysRevA.107.012403

[13] S. Dooley, W. J. Munro, K. Nemoto, Quantum metrology including state preparation and readout times, in: Quantum Information and Measurement (QIM) 2017, Optica Publishing Group, 2017, p. QW3A.5. doi:10.1364/QIM.2017.QW3A.5.
URL https://opg.optica.org/abstract.cfm?URI=QIM-2017-QW3A.5

[14] J. S. Kollath-Bönig, L. Dellantonio, L. Giannelli, T. Schmit, G. Morigi, A. S. Sørensen, Fast storage of photons in cavity-assisted quantum memories, Physical Review Applied 22 (4) (Oct. 2024). doi:10.1103/physrevapplied.22.044038.
URL http://dx.doi.org/10.1103/PhysRevApplied.22.044038

[15] S. R. Hasan, M. Z. Chowdhury, M. Saiam, Y. M. Jang, Quantum communication systems: Vision, protocols, applications, and challenges, IEEE Access 11 (2023) 15855–15877. doi:10.1109/ACCESS.2023.3244395.

[16] K. Rojan, D. M. Reich, I. Dotsenko, J.-M. Raimond, C. P. Koch, G. Morigi, Arbitrary-quantum-state preparation of a harmonic oscillator via optimal control, Phys. Rev. A 90 (2014) 023824. doi:10.1103/PhysRevA.90.023824.
URL https://link.aps.org/doi/10.1103/PhysRevA.90.023824

[17] D. Volya, P. Mishra, State preparation on quantum computers via quantum steering, IEEE Transactions on Quantum Engineering 5 (2024) 1–14. doi:10.1109/TQE.2024.3358193.

[18] W. Cai, J. Han, L. Hu, Y. Ma, X. Mu, W. Wang, Y. Xu, Z. Hua, H. Wang, Y. P. Song, J.-N. Zhang, C.-L. Zou, L. Sun, High-efficiency arbitrary quantum operation on a high-dimensional quantum system, Phys. Rev. Lett. 127 (2021) 090504. doi:10.1103/PhysRevLett.127.090504.
URL https://link.aps.org/doi/10.1103/PhysRevLett.127.090504

[19] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, M. Sarovar, From pulses to circuits and back again: A quantum optimal control perspective on variational quantum algorithms, PRX Quantum 2 (2021) 010101. doi:10.1103/PRXQuantum.2.010101.
URL https://link.aps.org/doi/10.1103/PRXQuantum.2.010101

[20] R. Porotti, A. Essig, B. Huard, F. Marquardt, Deep Reinforcement Learning for Quantum State Preparation with Weak Nonlinear Measurements, Quantum 6 (2022) 747. doi:10.22331/q-2022-06-28-747.
URL https://doi.org/10.22331/q-2022-06-28-747

[21] I. Khalid, C. A. Weidner, E. A. Jonckheere, S. G. Schirmer, F. C. Langbein, Sample-efficient model-based reinforcement learning for quantum control, Phys. Rev. Res. 5 (2023) 043002. doi:10.1103/PhysRevResearch.5.043002.
URL https://link.aps.org/doi/10.1103/PhysRevResearch.5.043002

[22] B. T. Torosov, N. V. Vitanov, Composite pulses with errant phases, Phys. Rev. A 100 (2019) 023410. doi:10.1103/PhysRevA.100.023410.
URL https://link.aps.org/doi/10.1103/PhysRevA.100.023410

[23] Z.-C. Shi, J.-H. Wang, C. Zhang, J. Song, Y. Xia, Universal composite pulses for robust quantum state engineering in four-level systems, Phys. Rev. A 109 (2024) 022441. doi:10.1103/PhysRevA.109.022441.
URL https://link.aps.org/doi/10.1103/PhysRevA.109.022441

[24] J. Zylberman, F. Debbasch, Efficient quantum state preparation with walsh series, Phys. Rev. A 109 (2024) 042401. doi:10.1103/PhysRevA.109.042401.
URL https://link.aps.org/doi/10.1103/PhysRevA.109.042401

[25] J. Iaconis, S. Johri, E. Y. Zhu, Quantum state preparation of normal distributions using matrix product states, npj Quantum Information 10 (1) (2024) 15. doi:10.1038/s41534-024-00805-0.
URL https://doi.org/10.1038/s41534-024-00805-0

[26] M. Zhou, F. A. Cárdenas-López, S. Dominique, X. Chen, Optimal control for open quantum system in circuit quantum electrodynamics (2024). arXiv:2412.20149.
URL https://arxiv.org/abs/2412.20149

[27] C. K. Law, J. H. Eberly, Arbitrary control of a quantum electromagnetic field, Phys. Rev. Lett. 76 (1996) 1055–1058. doi:10.1103/PhysRevLett.76.1055.
URL https://link.aps.org/doi/10.1103/PhysRevLett.76.1055

[28] K. Vogel, V. M. Akulin, W. P. Schleich, Quantum state engineering of the radiation field, Phys. Rev. Lett. 71 (1993) 1816–1819. doi:10.1103/PhysRevLett.71.1816.
URL https://link.aps.org/doi/10.1103/PhysRevLett.71.1816

[29] Y. R. Sanders, G. H. Low, A. Scherer, D. W. Berry, Black-box quantum state preparation without arithmetic, Phys. Rev. Lett. 122 (2019) 020502. doi:10.1103/PhysRevLett.122.020502.
URL https://link.aps.org/doi/10.1103/PhysRevLett.122.020502

[30] J. Bausch, Fast Black-Box Quantum State Preparation, Quantum 6 (2022) 773. doi:10.22331/q-2022-08-04-773.
URL https://doi.org/10.22331/q-2022-08-04-773

[31] T. M. L. de Veras, L. D. da Silva, A. J. da Silva, Double sparse quantum state preparation, Quantum Information Processing 21 (6) (2022) 204. doi:10.1007/s11128-022-03549-y.
URL https://doi.org/10.1007/s11128-022-03549-y

[32] D. Ramacciotti, A. I. Lefterovici, A. F. Rotundo, Simple quantum algorithm to efficiently prepare sparse states, Phys. Rev. A 110 (2024) 032609. doi:10.1103/PhysRevA.110.032609.
URL https://link.aps.org/doi/10.1103/PhysRevA.110.032609

[33] M. Luo, X. Wang, Universal quantum computation with qudits, Science China: Physics, Mechanics and Astronomy 57 (9) (2014) 1712–1717. doi:10.1007/s11433-014-5551-9.

[34] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, T. Monz, A universal qudit quantum processor with trapped ions, Nature Physics 18 (9) (2022) 1053–1057. arXiv:2109.06903, doi:10.1038/s41567-022-01658-0.

[35] W. Chen, J. Gan, J.-N. Zhang, D. Matuskevich, K. Kim, Quantum computation and simulation with vibrational modes of trapped ions, Chinese Physics B 30 (6) (2021) 060311. doi:10.1088/1674-1056/ac01e3.
URL https://dx.doi.org/10.1088/1674-1056/ac01e3

[36] V. Vargas-Calderón, N. Parra-A., H. Vinck-Posada, F. A. González, Many-qudit representation for the travelling salesman problem optimisation, Journal of the Physical Society of Japan 90 (11) (2021) 114002. arXiv:https://doi.org/10.7566/JPSJ.90.114002, doi:10.7566/JPSJ.90.114002.
URL https://doi.org/10.7566/JPSJ.90.114002

[37] K. Goswami, P. Schmelcher, R. Mukherjee, Qudit-based scalable quantum algorithm for solving the integer programming problem (2025). arXiv:2508.13906.
URL https://arxiv.org/abs/2508.13906

[38] D. Cozzolino, B. Da Lio, D. Bacco, L. K. Oxenløwe, High-dimensional quantum communication: Benefits, progress, and future challenges, Advanced Quantum Technologies 2 (12) (2019) 1900038. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/qute.201900038, doi:https://doi.org/10.1002/qute.201900038.
URL https://onlinelibrary.wiley.com/doi/abs/10.1002/qute.201900038

[39] A. M. Bloch, R. W. Brockett, C. Rangan, Finite controllability of infinite-dimensional quantum systems, IEEE Transactions on Automatic Control 55 (8) (2010) 1797–1805. doi:10.1109/TAC.2010.2044273.

[40] B. W. Shore, Two-state coherent excitation, Cambridge University Press, 2011, pp. 97–136.

[41] L. Pinna, G. Panati, Approximate controllability of the jaynes-cummings dynamics, Journal of Mathematical Physics 59 (7) (2018).

[42] J. J. García Ripoll, Quantum Information and Quantum Optics with Superconducting Circuits, Cambridge University Press, 2022. doi:10.1017/9781316779460.

[43] A. Blais, A. L. Grimsmo, S. M. Girvin, A. Wallraff, Circuit quantum electrodynamics, Rev. Mod. Phys. 93 (2021) 025005. doi:10.1103/RevModPhys.93.025005.
URL https://link.aps.org/doi/10.1103/RevModPhys.93.025005

[44] W. D. Kalfus, D. F. Lee, G. J. Ribeill, S. D. Fallek, A. Wagner, B. Donovan, D. Ristè, T. A. Ohki, High-fidelity control of superconducting qubits using direct microwave synthesis in higher nyquist zones, IEEE Transactions on Quantum Engineering 1 (2020) 1–12. doi:10.1109/TQE.2020.3042895.

[45] Z.-W. Wang, H.-Y. Ma, Y.-A. Yan, L.-A. Wu, Z.-M. Wang (2025). arXiv:2508.20377, [link].
URL https://arxiv.org/abs/2508.20377

[46] I. A. Bocanegra-Garay, L. Hernández-Sánchez, I. Ramos-Prieto, F. Soto-Eguibar, H. M. Moya-Cessa, Invariant approach to the driven jaynes-cummings model, SciPost Physics 16 (1) (2024) 007. doi:10.21468/SciPostPhys.16.1.007.

[47] M. H. Levitt, Composite pulses, Progress in Nuclear Magnetic Resonance Spectroscopy 18 (2) (1986) 61–122. doi:https://doi.org/10.1016/0079-6565(86)80005-X.
URL https://www.sciencedirect.com/science/article/pii/007965658680005X

[48] M. Plesch, i. c. v. Brukner, Quantum-state preparation with universal gate decompositions, Phys. Rev. A 83 (2011) 032302. doi:10.1103/PhysRevA.83.032302.
URL https://link.aps.org/doi/10.1103/PhysRevA.83.032302

[49] W. Pfeifer, The Lie Algebras su(N), Birkhäuser Basel, Basel, 2003. doi:10.1007/978-3-0348-8097-8.
URL http://link.springer.com/10.1007/978-3-0348-8097-8

[50] M. V. Narkhede, P. P. Bartakke, M. S. Sutaone, A review on weight initialization strategies for neural networks, Artificial Intelligence Review 55 (1) (2022) 291–322. doi:10.1007/s10462-021-10033-z.
URL https://doi.org/10.1007/s10462-021-10033-z

[51] F. Mezzadri, How to generate random matrices from the classical compact groups (2007). arXiv:math-ph/0609050.
URL https://arxiv.org/abs/math-ph/0609050

[52] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, JAX: composable transformations of Python+NumPy programs (2018).
URL http://github.com/jax-ml/jax

[53] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, Neural Networks 2 (5) (1989) 359–366. doi:10.1016/0893-6080(89)90020-8.
URL https://doi.org/10.1016/0893-6080(89)90020-8

[54] S. W. Jolin, R. Borgani, M. O. Tholén, D. Forchheimer, D. B. Haviland, Calibration of mixer amplitude and phase imbalance in superconducting circuits, Review of Scientific Instruments 91 (12) (2020) 124707. arXiv:https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/5.0025836/14886296/124707_1_online.pdf, doi:10.1063/5.0025836.
URL https://doi.org/10.1063/5.0025836

[55] N. Wu, J. Lin, C. Xie, Z. Guo, W. Huang, L. Zhang, Y. Zhou, X. Sun, J. Zhang, W. Guo, et al., In situ mixer calibration for superconducting quantum circuits, Applied Physics Letters 125 (20) (2024).

[56] S. Oh, Errors due to finite rise and fall times of pulses in superconducting charge qubits, Phys. Rev. B 65 (2002) 144526. doi:10.1103/PhysRevB.65.144526.
URL https://link.aps.org/doi/10.1103/PhysRevB.65.144526

[57] M. Gely, J. Litarowicz, A. Leu, D. Lucas, In situ characterization of qubit-drive phase distortions, Phys. Rev. Appl. 22 (2024) 024001. doi:10.1103/PhysRevApplied.22.024001.
URL https://link.aps.org/doi/10.1103/PhysRevApplied.22.024001

[58] Z. Zhang, J. Sung, D. T. Toolan, S. Han, R. Pandya, M. P. Weir, J. Xiao, S. Dowland, M. Liu, A. J. Ryan, et al., Ultrafast exciton transport at early times in quantum dot solids, Nature materials 21 (5) (2022) 533–539.

[59] H. Jia, J. Cao, F. Chen, F. Peng, Y. Li, Y. Xu, L. Chen, Z. Ye, X. Zhao, S. Zhang, J. Jing, H. Xu, Z. Chen, T. Byrnes, H. Li, A. Kavokin, J. Wu, Femtosecond coherence dynamics of exciton–polaritons, National Science Review 13 (1) (2025) nwaf493. arXiv:https://academic.oup.com/nsr/article-pdf/13/1/nwaf493/65388832/nwaf493.pdf, doi:10.1093/nsr/nwaf493.
URL https://doi.org/10.1093/nsr/nwaf493

Appendix A. Neural network

In this work, we consider a feed-forward neural network, which results from the consecutive application of neural network layers. A neural network layer is a parameterized function $l : \mathbb{R}^{\text{in}} \to \mathbb{R}^{\text{out}}$ defined as

$$l(\boldsymbol{x}) = a(W\boldsymbol{x} + \mathbf{b}), \tag{A.1}$$

where a is an element-wise function, called the activation function, $W \in \mathbb{R}^{\text{out}} \times \mathbb{R}^{\text{in}}$ is called the weight matrix, and $\boldsymbol{b} \in \mathbb{R}^{\text{out}}$ is the bias vector. Both the weight matrix and the bias vector contain the parameters of the layer. Layers can be composed to form a feed-forward neural network, which is the type considered in this work. Such a neural network is defined as

$$f_{\boldsymbol{\eta}}(\boldsymbol{x}) = l_L \circ l_{L-1} \circ \cdots \circ l_1(\boldsymbol{x}) \tag{A.2}$$

for a neural network of L layers. We have made it explicit that $\boldsymbol{\eta}$ are the neural network's parameters, which are all the weight matrices and bias vectors from the neural network's layers. The last layer's output size is $3N$, to accommodate the $3N$ pulse parameters $\boldsymbol{\Theta} = \{\phi_i\}_{i=1}^N \cup \{\varphi_i\}_{i=1}^N \cup \{t_i\}_{i=1}^N$. The first layer's input is d, where d refers to the components that specify the input target state. The neural-network architecture depends on the target dimension n and was found by optimizing the network structure with Optuna. For $n = 2$ we used hidden-layer widths $[209, 173, 222, 347]$; for $n = 3$ we used $[197, 264, 304]$; and for $n = 4$ we used $[321, 442, 425]$.

Appendix B. Basis of SU(n)

The SU(n) basis that we consider is the union of the following $n \times n$ families of matrices [49]:

$$\begin{pmatrix} 0 & 1 & \cdots & 0 & 0 \\ 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 1 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \cdots \begin{pmatrix} 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix} \tag{B.1}$$

$$\begin{pmatrix} 0 & -i & \cdots & 0 & 0 \\ i & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & -i & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ i & 0 & 0 & \cdots & 0 \\ 0 & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \cdots \begin{pmatrix} 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & -i \\ 0 & 0 & \cdots & i & 0 \end{pmatrix} \tag{B.2}$$

$$\begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & -1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix}, \sqrt{\frac{1}{3}} \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & -2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix}, \cdots, \sqrt{\frac{2}{n(n-1)}} \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -n+1 \end{pmatrix} \tag{B.3}$$

This basis corresponds to $n^2 - 1$ matrices. Using this basis, for a target state $|\psi_{\text{target}}\rangle$, the inputs for the neural network are the $n^2 - 1$ expectation values $\langle\psi_{\text{target}}|G_k|\psi_{\text{target}}\rangle$, where $\{G_k\}_{k=1}^{n^2-1}$ denotes the SU(n) basis listed above.