

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Robust quantum gates using smooth pulses and physics-informed neural networks

Utkan Güngördü^{1,*} and J. P. Kestner¹

¹*Department of Physics, University of Maryland Baltimore County, Baltimore, MD 21250, USA*

The presence of decoherence in quantum computers necessitates the suppression of noise. Dynamically corrected gates via specially designed control pulses offer a path forward, but hardware-specific experimental constraints can cause complications. Here, we present a widely applicable method for obtaining smooth pulses which is not based on a sampling approach and does not need any assumptions with regards to the underlying statistics of the experimental noise. We demonstrate the capability of our approach by finding smooth shapes which suppress the effects of noise within the logical subspace as well as leakage out of that subspace.

I. INTRODUCTION

There has been significant progress in the past two decades towards the realization of a physical quantum computer. The greatest obstacle presently hampering the current efforts is the presence of unwanted couplings between the qubit and its hosting environment. This is particularly true for solid-state qubits such as Josephson-junction based qubits or semiconductor spin qubits, where decoherence is caused by charge noise, Overhauser effect, stray or drifting magnetic fields, quasiparticle poisoning, etc. An established way of suppressing the effects of such non-Markovian noise since the early days of NMR is to apply carefully designed control fields, which is also effective against calibration errors. For a robust quantum control protocol to be practical, however, it is important to take the limitations of the control hardware into account. Experimentally realistic control fields must have bounded bandwidth (and hence be smooth pulses) and amplitude.

For an ideal qubit with no noise, smooth pulses can be obtained analytically¹ or numerically². In many realizations, however, noise is detrimental, and smooth pulses that are robust against errors within a qubit's logical subspace have been proposed to address quasistatic noise in two-level systems^{3–5}, achieved by designing control protocols with reduced sensitivity to noise in a perturbative manner. Extending these (semi-)analytic approaches beyond two-level systems remains a challenge. This is an important problem because although two-level systems are useful for basic demonstrations, a useful quantum computer necessitates multiple qubits. Furthermore, systems such as the transmon⁶ and resonator-coupled spin qubits⁷ in which a pair of low lying energy states are used to encode a qubit suffer from leakage to excited states and also warrant a treatment that falls outside the scope of these robust smooth pulses.

Other existing approaches numerically search for pulses that maximize gate fidelity for an ensemble of simulated noise sampled from an assumed distribution function. The control is typically (although not necessarily⁸) split into piecewise-constant segments^{9–15}. Although one may strive to approach a smooth pulse by making the pieces sufficiently small and penalizing sudden jumps, here we use efficient adaptive solvers as opposed to a

large number of discrete segments with fixed timestep solvers. For broadband noise, a smooth pulse approach to improve gate fidelity by suppressing the leading order effects of an assumed noise power spectral density (PSD) was reported in Ref. 16. A common theme in all these approaches is that the optimization is performed over the combined effect of the noise and control on fidelity, which requires assumptions regarding the spectrum or the statistics of the noise and averaging over noise realizations.

Here, we propose a new approach that addresses all these issues simultaneously and allows us to find smooth pulses for implementing robust quantum gates in any qubit platform. The generic approach described here is not limited to two-level systems and can suppress the effects of *both* leakage and quasistatic noise, regardless of the underlying statistics of the noise. This is achieved by suppressing the *sensitivity* to arbitrary quasistatic noise (as opposed to maximizing fidelity for a given noise sampling) by taking into account the exact functional relation between noise and its effect on the resulting quantum gate; this is enabled by recently introduced physics-informed neural network (PINN) frameworks¹⁷ which implement deep neural networks (DNNs) that can be used to minimize a cost function while at the same time respecting a given set of differential equations (see also^{18–20}).

From a computational point of view, sampling-based approaches which focus on maximizing the gate fidelity for a given stochastic perturbation require additional optimization cycles with different noise realizations whose values are sampled from a given distribution function, which comes at a significant computational cost. In contrast, our approach eliminates the necessity for averaging over noise traces by focusing on the noise sensitivity. Furthermore, operating over smooth functions rather than a large number of piecewise-continuous pulses with fixed-width segments allows additional speed up when using adaptive ordinary differential equation (ODE) solvers. Another advantage of our functional approach is that one can impose exact boundary conditions and functional constraints (such as explicit bandwidth constraints, maximum drive amplitude, boundary conditions for the drive, or even functional constraints such as DRAG⁶) from the outset. This is in contrast with the usual approach of

penalizing any deviation of such constraints by including them in the cost function, which makes it more difficult to find solutions by making the optimization landscape more complicated, leading to additional false minima and stiffness. Furthermore, as we demonstrate below, it is also feasible to suppress higher order effects of noise within our approach, which are important for strong errors and long pulse durations. Finally, in terms of methodology, within the context of DNN based quantum control, our approach is a new direction beyond sampling based learning protocols for finding shaped pulses, both robust and nonrobust.

II. METHOD

We start from a generic definition of a quantum system whose Hamiltonian is sum of a control Hamiltonian and error Hamiltonian, $H(t, \mathbf{p}) = H_c(t, \mathbf{p}) + H_\epsilon(t, \mathbf{p})$ with

$$H_c(t, \mathbf{p}) = \sum_i h_i(t, \mathbf{p}) \Lambda_i, \quad H_\epsilon(t, \mathbf{p}) = \sum_i \epsilon_i \chi_i(t, \mathbf{p}) \Lambda_i, \quad (1)$$

where $\mathbf{p} = \mathbf{p}(t)$ is a time-dependent control parameter characterizing the driving fields h_i , ϵ_i are the quasistatic stochastic noise strengths, Λ_i are traceless generators of the Lie algebra $\mathfrak{su}(n)$ that obey $\text{tr}(\Lambda_i \Lambda_j) = n \delta_{ij}$, and $\chi(t, \mathbf{p})$ represents any dependence of the noise Hamiltonian H_ϵ on the control fields. The latter becomes relevant when a term in the Hamiltonian is a nontrivial function of a noisy parameter, e.g., single-qubit microwave driving $\Omega(t)$ with multiplicative amplitude error $\delta\Omega = \epsilon\Omega(t)$ or tunable qubit-qubit coupling $J(V)$ susceptible to fluctuations δV in the voltage or flux tuning parameter as $\delta J = \delta V \partial_V J(V)$.

By treating H_ϵ as the interaction Hamiltonian, the solution of the time-dependent Schrödinger equation can be expressed as $U(t) = U_\epsilon^\dagger(t) U_c(t)$ where $i\dot{U}_c(t) = H_c(t, \mathbf{p}) U_c(t)$ is the time-evolution operator as determined by the control Hamiltonian and $i\dot{U}_\epsilon(t) = U_c^\dagger(t) H_\epsilon(t, \mathbf{p}) U_c(t)$ accounts for the effects of the noise. For weak noise ($\left\| \int_0^T dt U_c^\dagger(t) H_\epsilon(t, \mathbf{p}) U_c(t) / \hbar \right\| \ll \pi$), U_ϵ can be calculated perturbatively using Magnus expansion as

$$\begin{aligned} U_\epsilon(T) &= \exp \left(-\frac{i}{\hbar} \int_0^T dt U_c^\dagger(t) H_\epsilon(t, \mathbf{p}) U_c(t) \right) + \mathcal{O}(\epsilon_i^2) \\ &= \exp \left(-\frac{i}{\hbar} \sum_i \epsilon_i \mathcal{E}_i(T) \right) + \mathcal{O}(\epsilon_i^2), \end{aligned} \quad (2)$$

where $\mathcal{E}_i(T) = \int_0^T dt U_c^\dagger(t) [\chi_i(t, \mathbf{p}) \Lambda_i] U_c(t)$. A robust quantum gate is one that is insensitive to the first order effects of ϵ_i , i.e. $\partial_{\epsilon_i} U(T)|_{\epsilon_i=0} = 0$. This condition can be achieved by choosing a control field for which $\|\mathcal{E}_i(T)\|$ vanishes or becomes negligibly small. The control field must also be chosen such that within the logical subspace

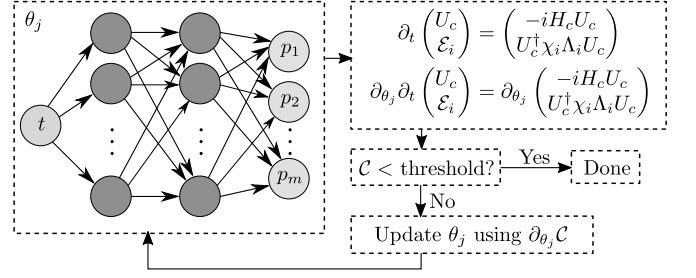


FIG. 1. (Color online) Flowchart of the physics-informed neural network for finding robust smooth pulses. Backpropagation efficiently works in conjunction with the adjoint sensitivity analysis of the coupled ODE system, and all derivatives with respect to optimization parameters θ_j (weights, biases and time-scaling parameter) are computed using automatic differentiation²¹.

$U_c(T)$ has the same result as a given target operation, U_0 . These requirements are equivalent to stating that the following cost function needs to vanish:

$$\mathcal{C} = \left(1 - \frac{|\text{tr}([PU_c(T)\mathcal{P}^\dagger]U_0^\dagger)]|^k}{D^k} \right) + \sum_i \left(\frac{w_i \epsilon_i^{\max}}{\hbar \tilde{D}} \|\mathcal{P}\mathcal{E}_i(T)\| \right)^l, \quad (3)$$

where $D \leq n$ is the dimension of the logical subspace (the part of the full Hilbert space that is relevant for quantum computation), \mathcal{P} denotes the projection operator onto the logical subspace, $\|\mathcal{P}\mathcal{E}_i(T)\|^2 = \text{tr}[\mathcal{P}\mathcal{E}_i(T)\mathcal{E}_i(T)^\dagger\mathcal{P}^\dagger]$ is the noise sensitivity of the logical subspace, $\tilde{D} = (2n - D)D$, ϵ_i^{\max} is the maximum tolerable value of the stochastic noise strength ϵ_i , $w_i \lesssim 1$ are hyperparameters for the optimization step which can be tuned or annealed with a schedule to avoid local minima. The first term in parentheses is similar to the gate infidelity in the absence of stochastic noise which ensures that the final unitary is U_0 , and the second term is the sensitivity which ensures that the implemented gate is immune to the first order effects of the stochastic noise ϵ_i ; this is different from the existing approaches where the cost function is taken to be the noisy gate infidelity which also necessitates an averaging over a particular noise distribution^{8-10,12-15}. The exponents k and l determine the relative weighting of the fidelity and the sensitivity terms in the cost function. In fact, the two terms could be wrapped inside any monotonically increasing functions, but in our results for simplicity we will simply raise to a power of either 1 or 2.

A neural network is a network of “neurons” (see Fig. 1), where a vertical column of neurons form a layer. The overall action of the i th layer of the neural network is to map a d_i dimensional input vector \mathbf{x} onto a d_{i+1} dimensional output vector $\mathcal{L}_i(\mathbf{x}) = \sigma_i(\mathcal{W}_i \mathbf{x} + \mathbf{b}_i)$, where \mathcal{W}_i is a matrix and \mathbf{b}_i a vector respectively containing the “weights” and “biases” of the i th layer, and σ_i is known as the “activation function.” Our DNN represents a function $t \rightarrow \mathbf{p}(t)$, where $\mathbf{p}(t)$ is a smooth curve param-

eterized by the internal degrees of freedom of the DNN, \mathbf{b}_i and \mathcal{W}_i , as $\mathbf{p}(t) = \mathcal{L}_N \circ \dots \circ \mathcal{L}_1(t)$. Intermediate layers $1 < i < N$ are commonly referred to as hidden layers. For concreteness, we will take σ_i to be the element-wise-acting tanh function, although functions which are sufficiently smooth and do not significantly affect the bandwidth requirements of the resulting pulse shape are also viable.

The goal of a DNN optimizer is to vary the weights and biases until the cost function \mathcal{C} is minimized. This is achieved by using local gradient based optimization in conjunction with the backpropagation algorithm²². This requires calculation of \mathcal{C} in addition to its parameter gradient $\partial\mathcal{C}$, i.e., the partial derivatives of \mathcal{C} with respect to each of the elements of \mathcal{W}_i and \mathbf{b}_i at each iteration, which can be computationally prohibitive. For this reason, we differentiate $\mathcal{E}_i(t)$ analytically and transform the problem of calculating \mathcal{C} into solving a coupled system of ODEs

$$\partial_t \begin{pmatrix} U_c(t) \\ \mathcal{E}_i(t) \end{pmatrix} = \begin{pmatrix} -iH_c(t, \mathbf{p})U_c(t)/\hbar \\ U_c^\dagger(t) [\chi_i(t, \mathbf{p})\Lambda_i] U_c(t) \end{pmatrix} \quad (4)$$

subject to the initial conditions $U_c(0) = \mathbb{1}$, $\mathcal{E}_i(0) = 0$. Similarly, $\partial\mathcal{C}$ is calculated by taking partial derivatives of Eq. (4) with respect to the elements of \mathcal{W}_i and \mathbf{b}_i as in Ref.². This form allows a more efficient calculation of the cost function and its parameter gradient, as required by the backpropagation algorithm, by using the adjoint sensitivity method^{23,24} on Eq. (4). In practice, we perform the straightforward but unwieldy task of calculating the parameter gradients of the coupled ODE system Eq. (4) by using reverse mode automatic differentiation^{21,25,26}.

The above analysis can be extended to higher order error correction, which is relevant when the noise is not sufficiently weak for a first order treatment. For example, second order error correction can be achieved by appending $\partial_t \mathcal{E}_{ij}^{(2)}(t) = [\partial_t \mathcal{E}_i(t) \mathcal{E}_j(t) - \mathcal{E}_j(t) \partial_t \mathcal{E}_i(t)]/2$ to Eq. (4) with the initial condition $\mathcal{E}_{ij}^{(2)}(0) = 0$, and introducing $\sum_{i,j} (w_{ij} \epsilon_i^{\max} \epsilon_j^{\max} \|\mathcal{P} \mathcal{E}_{ij}^{(2)}(T)\| / \tilde{D} \hbar^2)^l$ to the cost function, which ensures that the second order term in the Magnus expansion for $U_\epsilon(T)$ vanishes, and comes at only minor additional computational cost.²⁷

We note that using a DNN to parametrize the time-dependence of the Hamiltonian eliminates the nontrivial problem of finding a suitably parametrized ansatz function to optimize over which will yield convergent solutions^{3-5,28}. We rather find that DNN optimizers are able handle two orders of magnitude larger number of optimization parameters in a comparable amount of time, making it feasible to explore a larger functional space by simply increasing the number of neurons and layers. Since both the adjoint sensitivity method and backpropagation algorithm are well established computational methods, we here focus on presenting our results on robust quantum control, and refer the reader to the literature²²⁻²⁴ for their details. For our numerical results, we use DiffEqFlux.jl²⁴, which is a Julia package implementing a PINN optimizer, with optional support

for hardware DNN accelerators. Before moving on to explicit solutions, we remark that the gate time T needs to be decided prior to optimization. For Hamiltonians that can contain non-adjustable terms (such as drift terms), a solution will generally not exist for an arbitrary T , and finding a suitable value can be laborious. This problem can be solved by introducing a new optimization parameter α , a time-independent scaling factor, to the right hand side of Eq. (4) and to the noise terms ϵ_i in Eq. (3); α can be seen as a scaling factor for time or the total Hamiltonian. To avoid long gate times approaching to inverse of the incoherent decay rates Γ_i , one can introduce a factor $e^{-\sum_i \Gamma_i T/\alpha}$ to the trace fidelity term in the cost function \mathcal{C} .

III. EXAMPLES

A. Exchange-coupled spin qubits

As our first proof-of-concept example, we consider a pair of spin qubits in a semiconductor double quantum dot, with one electron in each dot. The overlap between the electron wavefunctions is determined by the gate voltages, which provides voltage tunable exchange coupling between the spin degrees of freedom of the electrons. Single-qubit operations are realized by modulation of the magnetic field that is generated by an on-chip wire. In a rotating frame, the spin Hamiltonian of this system can be written as²⁸

$$H_c(t) = \frac{J}{4} \sigma_z \otimes \sigma_z + \frac{1}{2} g \mu_B B_x(t) \sigma_x \otimes \mathbb{1}, \quad (5)$$

whose Lie algebra is isomorphic to a two-level $\mathfrak{su}(2)$ generated by $\{\sigma_z \otimes \sigma_z, \sigma_x \otimes \mathbb{1}, \sigma_y \otimes \sigma_z\}$. The exchange coupling, J , which in some devices is essentially fixed due to bandwidth limitations^{28,29}, is susceptible to charge noise induced fluctuations, which can be modelled by the noise Hamiltonian $H_\epsilon = \epsilon_J J \sigma_z \otimes \sigma_z / 4$. An entangling CZ-equivalent gate ($e^{-i\sigma_z \otimes \sigma_z \pi/4}$) can be produced naively by setting $B_x = 0$ and waiting a time $\hbar\pi/J$. However, instead we search for a smooth pulse $B_x(t)$ which can correct both first and second order quasistatic fluctuations in J by parameterizing the magnetic field as $g\mu_B B_x(t)/2 = (J/4)A(t)p_1(t)\sin(p_2(t))$. Here, $A(t) = \coth(\kappa T)[\tanh(\kappa t/T) - \tanh(\kappa(t-T)/T)] - 1$ is a smoothed unit square pulse (κ determines the degree of the smoothing), and its purpose is to enforce the condition that the magnetic field is turned on only for the duration of the pulse, $B_x(0) = B_x(T) = 0$. The resulting smooth pulse shown in Fig. 2 produces a CZ gate with extraordinarily high gate fidelities (defined as $\mathcal{F} = |\text{tr}(U(T)U_0^\dagger)/4|^2$) above 99.99% for exchange errors as large as 15%, and is faster than the smooth pulse reported in Ref.¹. 99% of the spectral density of this pulse is contained within the bandwidth $\approx 11/T$, which corresponds to a very conservative value of $\Delta f \approx 2\text{MHz}$ for $J = 1\text{MHz}$.




FIG. 2. (Color online) (a) Pulse shape implementing a CZ gate that is robust against the first and second order effects of quasistatic errors in exchange coupling, obtained with $\kappa = 6$, $k = 2$, $l = 2$, $T = 4.8h/J$ using a neural network with two hidden layers of length 32. The detailed numerical parameters of this pulse are tabulated in Appendix A. (b) Gate infidelity as a function of the error strength ϵ_J for the robust pulse shape (corrected) and a simple, undriven implementation (naive).

B. Transmon qubit

To illustrate simultaneous noise and leakage suppression, we next turn to the transmon qubit, although we remark in passing that this scenario is also relevant in the context of resonator coupled spin qubits³⁰ and encoded exchange-only spin qubits³¹. The effective Hamiltonian for the transmon can be written as³²

$$H_c(t) \approx \delta(t)a^\dagger a + \frac{\Delta}{2}a^\dagger a(a^\dagger a - \mathbb{1}) + \frac{\Omega(t)a + \Omega^*(t)a^\dagger}{2} \quad (6)$$

in a rotating frame, where Δ is the anharmonicity and $\Omega(t)$ is the complex envelope of the microwave drive, whose frequency is detuned from the qubit frequency by $\delta(t)$. The first two levels encode a logical qubit, and we consider the first four levels in our calculations, but we will revisit this point below. Thus, when calculating \mathcal{C} , we project $U(T)$ onto the qubit subspace. The single-sided projection in $||\mathcal{PE}_i(T)||$ allows us to leave out the effects of the noise solely on leakage subspace³³; that being said, it is possible to protect the leakage subspace as well by simply omitting this projection. Our goal is to find a smooth pulse for implementing a gate that can suppress both leakage and shifts in detuning $\delta(t) \rightarrow \delta(t) + \epsilon$ which can be caused by calibration errors³⁴ or stochastic phase errors³⁵.

At this point, we recall that an established method of suppressing leakage in Josephson junction based qubits is to use pulse shapes that obey a particular family of differential relations between $\Omega(t)$ and Δ , known as DRAG⁶. These relations ensure that the leakage inducing terms remain small throughout the pulse. It is possible to enforce DRAG conditions by construction, for instance by augmenting Eq. (4) with the relation $\dot{\Omega}_P(t) = p_1(t)\sin(p_2(t))$ and parameterizing the drive as $\Omega(t) = A(t)\Omega_P(t) - i\partial_t[A(t)\Omega_P(t)]/2\Delta$, $\delta(t) = 0$. However, we will proceed without doing so in order to avoid limiting the search space: DRAG is a sufficient condition for suppressing leakage, but it is not a necessary one since what matters is whether the qubit




FIG. 3. (Color online) (a) Pulse shape implementing a $X_{\pi/2}$ gate that is robust against the leading order effects of quasistatic errors in detuning and leakage errors, obtained with $\kappa = 5$, $k = 1$, $l = 2$, $T = 3h/\Delta$ using a neural network with two hidden layers of length 32. The detailed numerical parameters of this pulse are tabulated in Appendix A. (b) Gate infidelity as a function of the error strength $\epsilon/|\Delta|$.

subspace time-evolution operator is equal to the target unitary at the final time $t = T$, regardless of any leakage that may be present during intermediate times $0 < t < T$. We thus parameterize the driving field as $\Omega(t) = \Omega_x(t) + i\Omega_y(t) = 2\Delta A(t)[p_1(t)\sin(p_2(t)) + ip_3(t)\sin(p_4(t))]$, $\delta(t) = \Delta A(t)p_5(t)\sin(p_6(t))$, and target a $X_{\pi/2}$ gate. This is in particular a useful example, because when used together with error-free virtual Z rotations implemented by shifting the rotating frame³⁶, $X_{\pi/2}$ gates are sufficient to implement arbitrary single qubit rotations. The resulting pulse shape is shown in Fig. 3. For a typical anharmonicity value of $\Delta/h \sim -100\text{MHz}$ ³², the pulse duration is $T \sim 30\text{ns}$, for which the effects of relaxation can be neglected. Under this assumption, the resulting gate fidelities, defined by $\mathcal{F} = |\text{tr}([\mathcal{P}U(T)\mathcal{P}^\dagger]U_0^\dagger)/2|^2$, remain above 99.99% as long as the shift in detuning remains below 1.5% of Δ . When the effect of two additional higher leakage states are taken into account for this pulse shape, the baseline fidelity remains the same. Compared to a nonrobust pulse based on DRAG^{6,32,34}, which can take at least $T \approx 2.1h/|\Delta| = 21\text{ns}$ and reaches the same infidelity threshold at 0.03% of Δ as shown in Fig. 3, this shaped pulse significantly improves the error threshold against detuning errors. 99% of the spectral density of the pulse is contained within the bandwidth $\approx 15/T$, which corresponds to $\Delta f \approx 500\text{MHz}$ for $\Delta = -100\text{MHz}$, which is ≈ 3 times higher than that of DRAG, but still well achievable with arbitrary waveform generators such as Tektronix AWG70002A. We remark that when the robustness constraint is removed, our method easily produces faster nonrobust gates that well outperform the fastest DRAG pulses, ultimately limited by the bandwidth of the drive. More stringent experimental bandwidth constraints can be accommodated by running the search with an appropriately increased gate time T ; although it is harder to find solutions for shorter gate times, one can always find solutions at longer times.

IV. CONCLUSION

In summary, we have introduced a method for performing dynamically corrected quantum gates with practical smooth pulses that is broadly applicable to large Hilbert spaces beyond two-level systems, which is necessary for robust control of multi-qubit devices and leakage into excited states, and is practically extensible to correction of errors beyond first order. Our approach is the first robust smooth pulse shaping method in this context that does not rely on sampling, and leverages physics-informed deep neural networks for computational advantages over existing similar protocols. In addition to noise cancellation, our generic approach can also be used to reduce or eliminate the need for careful recalibration cycles during experiments. A future direction is the extension of this approach to suppress time-dependent broadband noise.

ACKNOWLEDGMENTS

This research was sponsored by the Army Research Office (ARO), and was accomplished under Grant Number W911NF-17-1-0287.

Appendix A: Parameters for shaped pulses

The pulse shapes given in the main text can be approximated using a truncated Chebyshev series in the form $\sum_{n=0}^{N-1} c_n T_n(2t/T - 1)$. Up to the first three decimals, the first few coefficients c_n for $g\mu_B B_x(t)/J$ are $\{0.056, 0.248, -0.306, 0.273, -0.844, 0.608, 0.591, -0.556, 0.293, -0.904, 0.375, 0.503, -0.134, -0.329, -0.083, 0.209, 0.131, -0.076, -0.113, -0.015, 0.056, 0.096, -0.04, -0.107, 0.026, 0.088, -0.017, -0.062, 0.014, 0.036\}$. Similarly, for $\Omega_x(t)/\Delta$, $\Omega_y(t)/\Delta$, and $\delta(t)/\Delta$, the coefficients are respectively $\{0.029, 0.027, -0.093, -0.06, 0.02, 0.016, 0.016, 0.005, 0.011, 0.004, 0.008, 0.003, 0.005, 0.002, 0.003, 0.001, 0.002\}$, $\{0.0626, -0.0458, -0.1149, 0.0547, 0.0569, 0.0019, -0.0082, -0.0065, 0.0016\}$ and $\{0.425, 0.213, 0.255, -0.076, -0.24, -0.051, -0.149, -0.031, -0.116, -0.025, -0.085, -0.017, -0.055, -0.01, -0.032, -0.005, -0.015\}$. Approximate pulse shapes reconstructed from these Chebyshev coefficients result in similar fidelity curves which reach the 10^{-4} threshold around the same value as the ones given in the main text. The code used to produce these results and the resulting internal state of the neural networks can be found in the supplemental files.

-
- * Current address: Laboratory for Physical Sciences, College Park, Maryland 20740, USA; utkan@lps.umd.edu
- ¹ E. Barnes and S. Das Sarma, Phys. Rev. Lett. **109**, 060401 (2012).
 - ² S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, Phys. Rev. Lett. **120**, 150401 (2018).
 - ³ E. Barnes, X. Wang, and S. Das Sarma, Sci. Rep. **5**, 12685 (2015).
 - ⁴ J. Zeng, C. H. Yang, A. S. Dzurak, and E. Barnes, Phys. Rev. A **99**, 052321 (2019).
 - ⁵ U. Güngördü and J. P. Kestner, Phys. Rev. A **100**, 062310 (2019).
 - ⁶ F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Phys. Rev. Lett. **103**, 110501 (2009).
 - ⁷ A. Warren, E. Barnes, and S. E. Economou, Phys. Rev. B **100**, 161303 (2019).
 - ⁸ T. Nöbauer, A. Angerer, B. Bartels, M. Trupke, S. Rotter, J. Schmiedmayer, F. Mintert, and J. Majer, Phys. Rev. Lett. **115**, 190801 (2015).
 - ⁹ D. Dong, C. Wu, C. Chen, B. Qi, I. R. Petersen, and F. Nori, Sci. Rep. **6**, 36090 (2016).
 - ¹⁰ E. Zahedinejad, J. Ghosh, and B. C. Sanders, Phys. Rev. Appl. **6**, 054005 (2016).
 - ¹¹ P. Palittapongarnpim, P. Wittek, E. Zahedinejad, S. Vedaie, and B. C. Sanders, Neurocomputing **268**, 116 (2017).
 - ¹² M. Bukov, A. G. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Phys. Rev. X **8**, 031086 (2018).
 - ¹³ R.-B. Wu, H. Ding, D. Dong, and X. Wang, Phys. Rev. A **99**, 042327 (2019).
 - ¹⁴ M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, npj Quantum Inf. **5**, 33 (2019).
 - ¹⁵ X. Ge, H. Ding, H. Rabitz, and R. B. Wu, Phys. Rev. A **101**, 52317 (2020).
 - ¹⁶ C.-H. Huang and H.-S. Goan, Phys. Rev. A **95**, 062325 (2017).
 - ¹⁷ M. Raissi, P. Perdikaris, and G. Karniadakis, J. Comput. Phys. **378**, 686 (2019).
 - ¹⁸ D. C. Psychogios and L. H. Ungar, AIChE J. **38**, 1499 (1992).
 - ¹⁹ I. Lagaris, A. Likas, and D. Fotiadis, Comput. Phys. Commun. **104**, 1 (1997).
 - ²⁰ I. Lagaris, A. Likas, and D. Fotiadis, IEEE Trans. Neural Networks **9**, 987 (1998).
 - ²¹ M. Innes, A. Edelman, K. Fischer, C. Rackauckas, E. Saba, V. B. Shah, and W. Tebbutt, arXiv:1812.01892 (2019).
 - ²² R. Rojas, in *Neural networks* (Springer, 1996) pp. 149–182.
 - ²³ R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, arXiv:1806.07366 (2018).
 - ²⁴ C. Rackauckas, M. Innes, Y. Ma, J. Bettencourt, L. White, and V. Dixit, arXiv:1902.02376 (2019).
 - ²⁵ N. Leung, M. Abdelhafez, J. Koch, and D. Schuster, Phys. Rev. A **95**, 042318 (2017).
 - ²⁶ M. Abdelhafez, D. I. Schuster, and J. Koch, Phys. Rev. A **99**, 52327 (2019).
 - ²⁷ Similarly, third order error correction can be enforced using $\partial_t \mathcal{E}_{ijk}^{(3)}(t) = [\mathcal{E}_{jk}^{(2)}(t) \partial_t \mathcal{E}_i(t) - \mathcal{E}_j(t) \partial_t \mathcal{E}_i(t) \mathcal{E}_k(t) - \mathcal{E}_k(t) \partial_t \mathcal{E}_i(t) \mathcal{E}_j(t) + \partial_t \mathcal{E}_i(t) \mathcal{E}_{kj}^{(2)}(t)]/3$, and so on.
 - ²⁸ U. Güngördü and J. P. Kestner, Phys. Rev. B **101**, 155301 (2020).

- ²⁹ W. Huang, C. H. Yang, K. W. Chan, T. Tanttu, B. Hensen, R. C. C. Leon, M. A. Fogarty, J. C. C. Hwang, F. E. Hudson, K. M. Itoh, A. Morello, A. Laucht, and A. S. Dzurak, *Nature* **569**, 532 (2019).
- ³⁰ A. Warren, E. Barnes, and S. E. Economou, *Phys. Rev. B* **100**, 161303 (2019).
- ³¹ M. Russ and G. Burkard, *J. Phys. Condens. Matter* **29**, 393001 (2017).
- ³² J. M. Gambetta, F. Motzoi, S. T. Merkel, and F. K. Wilhelm, *Phys. Rev. A* **83**, 012308 (2011).
- ³³ If the input states are guaranteed to be fully within the logical subspace, the sensitivity cost can be simplified to $||\mathcal{P}\mathcal{E}_i(T)\mathcal{P}^\dagger||$.
- ³⁴ E. Lucero, J. Kelly, R. C. Bialczak, M. Lenander, M. Mariantoni, M. Neeley, A. D. O’Connell, D. Sank, H. Wang, M. Weides, J. Wenner, T. Yamamoto, A. N. Cleland, and J. M. Martinis, *Phys. Rev. A* **82**, 042339 (2010).
- ³⁵ H. Ball, W. D. Oliver, and M. J. Biercuk, *npj Quantum Inf.* **2**, 16033 (2016).
- ³⁶ D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, *Phys. Rev. A* **96**, 022330 (2017).