Optimal Control of Unknown Collective Spin Systems via a Neural Network Surrogate

Yaofeng Chen[1] and Li You[1,2,3,4*]

[1]*State Key Laboratory of Low-Dimensional Quantum Physics, Department of Physics, Tsinghua University, Beijing 100084, China*
[2]*Beijing Academy of Quantum Information Sciences, Beijing 100193, China*
[3]*Frontier Science Center for Quantum Information, Beijing 100084, China*
[4]*Hefei National Laboratory, Hefei 230088, China*

(Received 13 July 2025; accepted manuscript online 15 August 2025)

Quantum optimal control (QOC) relies on accurately modeling system dynamics and is often challenged by unknown or inaccessible interactions in real systems. Taking an unknown collective spin system as an example, this work introduces a machine-learning-based, data-driven scheme to overcome the challenges encountered, with a trained neural network (NN) assuming the role of a surrogate model that captures the system's dynamics and subsequently enables QOC to be performed on the NN instead of on the real system. The trained NN surrogate proves effective for practical QOC tasks and is further demonstrated to be adaptable to different experimental conditions, remaining robust across varying system sizes and pulse durations.

DOI: 10.1088/0256-307X/42/10/100601 CSTR: 32039.14.0256-307X.42.10.100601

1. Introduction. Quantum optimal control (QOC) is a crucial technology in quantum information processing. By optimizing control pulse parameters to maximize or minimize a figure of merit (FoM), QOC enables the design of high-fidelity quantum gates[1,2] and the preparation of novel quantum states.[3–5] An important platform for QOC consists of collective spin systems, such as spinor atomic Bose–Einstein condensates,[6,7] interacting atom-cavity systems,[8] and atomic vapors. These systems are highly promising for generating spin squeezing[9] and hold significant potential for applications in quantum metrology[10] and beyond.

Traditional QOC methods are based on accurate modeling of system dynamics. However, unknown interactions often encountered in real-world systems hinder complete descriptions of these dynamics. In collective spin systems, interactions become increasingly complex with increased ensemble sizes. While well-known models such as one-axis twisting[9] and the Lipkin–Meshkov–Glick models[11–13] serve as useful approximations, they eventually fail to capture the full dynamical details; therefore, they cannot be directly applied when unknown interactions exist.

Confronted with difficulties, closed-loop optimization methods like the remote dressed chopped-random-basis optimization (RedCRAB)[14–16] and model-free reinforcement learning[1,17–19] have been developed and employed. Through iteratively optimizing control pulses based on experimental measurements, these methods bypass the need for explicit system modeling. However, two major shortcomings still remain. First, they require high-rep-rate operations and low-latency measurements, which are challenging for many experimental platforms. Second, since real systems are not mathematically differentiable, only gradient-free optimization techniques are directly applicable.

To address the aforementioned limitations, this work proposes a machine-learning-based scheme where a neural network (NN) serves as a surrogate for the actual system. As illustrated in Fig. 1, the NN is trained on experimental data to learn the system's unknown dynamics, and subsequently, optimizations are carried out with the trained NN rather than with the physical system itself. Such a scheme is highly influenced by the recent surge of efforts that demonstrate the capability of NNs in capturing quantum dynamics.[20–32] Compared to previous QOC schemes, the NN surrogate scheme we present offers several advantages: it is entirely data-driven and requires no explicit system modeling; it eliminates direct interaction between the optimizer and the physical system making the scheme more experimentally accessible; and in addition, due to the differentiable nature of NNs, it allows for gradient-based optimization techniques, broadening the range of applicable algorithms. Our scheme can also potentially lead to reduced experimental costs. As we show below, the NN surrogate can be trained with only averaged spin components, yet still yields excellent predictions for second-order correlations, thus relieving the need to measure these values in QOC tasks such as spin squeezing and quantum sensing. Moreover, once trained, the NN surrogate can be reused in different QOC tasks, eliminating repetitive optimization experiments that are encountered by closed-loop optimizations.

*Corresponding author. Email: lyou@mail.tsinghua.edu.cn
© 2025 Chinese Physical Society and IOP Publishing Ltd. All rights, including for text and data mining, AI training, and similar technologies, are reserved.

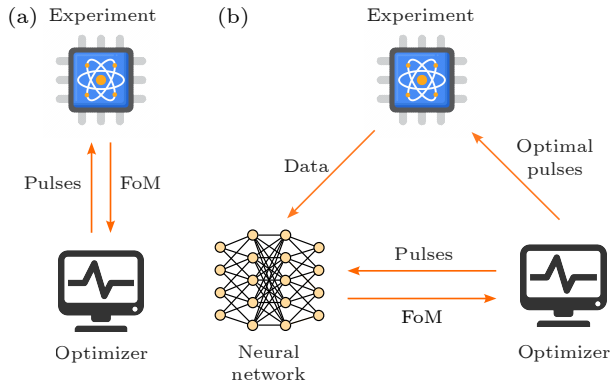

Fig. 1. Comparison of closed-loop optimization and the NN surrogate scheme. (a) In closed-loop optimization, the optimizer interacts directly with the experimental system. Each proposed pulse shape is applied in real-time, and immediate measurements are required to evaluate the FoM for iterative updates. (b) In the NN surrogate scheme, an NN first learns the system's dynamics from data. Once trained, it replaces the physical system by simulating responses to proposed pulse shapes, providing FoM predictions without needing to carry out direct experiments.

2. Optimization with an NN Surrogate. We now demonstrate the application of our NN surrogate scheme to a closed collective spin system. Quite generally, we consider situations where the form of the internal spin–spin interactions cannot be assumed. Consequently, the learning of the internal Hamiltonian essentially leads to estimating $O(d^2)$ unknown parameters, where d is the dimension of the Hilbert space. As the system size increases, this results in rapidly growing complexity for the learning task.

Fortunately, for most QOC tasks in collective spin systems, such as spin squeezing or quantum sensing, it suffices to focus on Gaussian-like states and their dynamics. In such cases, mean-field approximations become asymptotically accurate in the large-system limit, while residual quantum fluctuations can be effectively captured using the truncated Wigner approximation (TWA).[33,34] The quantum fluctuations are taken into account by sampling from the Wigner quasi-probability distribution $\mathcal{W}$ of the initial state, typically a spin coherent state. As such, TWA describes Gaussian-like dynamics with a fixed number of variables. For instance, in a spin-1/2 system, the collective spin $\boldsymbol{S} = \sum_i \boldsymbol{s}_i$ can be represented using Schwinger bosons[35]

$$S_+ = a_\uparrow^\dagger a_\downarrow, \quad S_- = a_\downarrow^\dagger a_\uparrow, \quad S_z = \frac{1}{2}(a_\uparrow^\dagger a_\uparrow - a_\downarrow^\dagger a_\downarrow). \tag{1}$$

Under this mapping, the TWA-based dynamics involve only two complex variables, $\boldsymbol{\alpha} = (\alpha_\uparrow, \alpha_\downarrow)$, obtained in the mean-field limit $a \to \sqrt{N}\boldsymbol{\alpha}$. This dimensional reduction implies that the complexity of the dynamical equations, and thus of their NN surrogate, remains insensitive to the system size, allowing for more efficient learning and optimal control for large collective spin systems.

When described by nonlinear mean-field ordinary differential equations (ODEs), TWA-based dynamics align naturally with the neural ODE (NODE) architecture,[36] where an NN parameterizes an ODE as $\frac{\mathrm{d}\boldsymbol{\alpha}}{\mathrm{d}t} = f(\boldsymbol{\alpha}, t; \boldsymbol{\theta})$, with $f(\boldsymbol{\alpha}, t; \boldsymbol{\theta})$ being an NN defined by parameters $\boldsymbol{\theta}$. Forward propagation is performed using standard ODE solvers (e.g., Runge–Kutta), while backward propagation utilizes the adjoint ODE method [see Supplementary Materials (SM)]. Although a relatively recent approach, NODE has been successfully applied to reconstruct quantum dynamics from measurements.[37,38] Compared to traditional architectures like recurrent neural networks,[39,40] it can explicitly encode the mathematical structure of the dynamics, reducing the need to learn from data. The improved training efficiency and reduced data requirements make NODE a strong candidate for modeling unknown quantum dynamics.

Building on NODE, we unify the learning and optimal control of unknown dynamics within a single framework. Consider a system with a single one-body control field. The dynamical equation under the TWA is given by

$$\mathrm{i}\frac{\mathrm{d}\boldsymbol{\alpha}}{\mathrm{d}t} = u(t)h_c\boldsymbol{\alpha} + h_0(\boldsymbol{\alpha}), \tag{2}$$

where h_c is a 2×2 Hermitian matrix representing the control field and $h_0(\boldsymbol{\alpha})$ is the nonlinear function arising from the unknown collective spin interactions. In the learning stage, we parameterize $h_0(\boldsymbol{\alpha})$ with an NN, leading to

$$\frac{\mathrm{d}\boldsymbol{\alpha}}{\mathrm{d}t} = -\mathrm{i}[u(t)h_c\boldsymbol{\alpha} + h_0(\boldsymbol{\alpha}; \boldsymbol{\theta})]. \tag{3}$$

As illustrated in Fig. 2(a), we train $h_0(\boldsymbol{\alpha}; \boldsymbol{\theta})$ using supervised learning. Specifically, we apply N_{train} randomly generated control pulses $u_k(t)$ to the system, where k indexes different pulses. The durations of $u_k(t)$ match the temporal control duration T. We measure the evolution of some physical observables $\langle O\rangle_{\text{true}}$ for the dataset. The loss function $\mathcal{L}_{\text{train}}$ is then computed by comparing $\langle O\rangle_{\text{true}}$ with the predicted evolution $\langle O\rangle_{\text{pred}}$ of the NODE. More specifically, we evaluate the training loss using the mean squared errors (MSEs), i.e.,

$$\mathcal{L}_{\text{train}} = \frac{1}{\mathcal{N}_{\text{train}}} \sum_{O,k,n} [\langle O(u_k, t_n)\rangle_{\text{pred}} - \langle O(u_k, t_n)\rangle_{\text{true}}]^2, \tag{4}$$

where $\mathcal{N}_{\text{train}} = N_O N_{\text{train}} N_t$ with N_O the number of observables, N_{train} the number of training pulses, and N_t the number of time points measured. Here, $\langle O(u_k, t_n)\rangle$ denotes $\langle O\rangle$ measured at t_n under $u_k(t)$.

After training, we fix $\boldsymbol{\theta} = \boldsymbol{\theta}_0$ and proceed to optimal control. We parameterize the pulse shape $u(t)$ using an ansatz $u(t; \lambda)$ with parameters λ, yielding the dynamical equation for the optimal control stage as

$$\frac{\mathrm{d}\boldsymbol{\alpha}}{\mathrm{d}t} = -\mathrm{i}[u(t; \lambda)h_c\boldsymbol{\alpha} + h_0(\boldsymbol{\alpha}; \boldsymbol{\theta}_0)]. \tag{5}$$

To facilitate gradient-based optimization, we treat $u(t; \lambda)$ as an NN parameterized by λ. This transforms Eq. (5) into another NODE, allowing gradients of the FoM $\mathcal{F}$ with respect to λ to be computed via NODE's backward propagation. Consequently, gradient-based QOC is reformulated as an unsupervised learning task, with loss function given by the FoM prediction [see Fig. 2(b)].

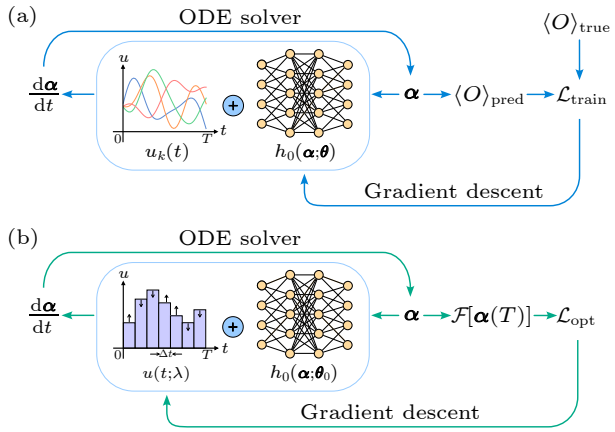

Fig. 2. Learning and optimal control with the NODE. (a) Learning the system dynamics with NODE is performed in a supervised manner. A set of pulse shapes $u_k(t)$ is randomly generated and applied to the system, producing the true evolution $\langle O\rangle_{\rm true}$. These pulses are then fed into the NODE, whose integration yields predictions $\langle O\rangle_{\rm pred}$. The loss between $\langle O\rangle_{\rm pred}$ and $\langle O\rangle_{\rm true}$ is used to train the parameters $\boldsymbol{\theta}$. (b) Gradient-based optimal control with NODE is formulated as an unsupervised learning task. The NN parameters are fixed at $\boldsymbol{\theta}_0$, while the pulse shape is parameterized as $u(t;\lambda)$. The optimization loss $\mathcal{L}_{\rm opt}$ is simply the FoM prediction $\mathcal{F}(\boldsymbol{\alpha})$, which is used to update λ via gradient descent.

3. Example. We demonstrate the NN surrogate scheme using a numerically simulated unknown system. While the scheme is general and can accommodate arbitrary interaction forms, for concreteness, we consider a system governed by quadratic spin interactions, with the internal Hamiltonian given by

$$H_0 = \sum_{\alpha,\beta\in\{x,y,z\}} \chi_{\alpha\beta} S_\alpha S_\beta. \tag{6}$$

A system consisting of $N = 100$ spins is simulated numerically. For the interaction matrix $\chi_{\alpha\beta}$, we choose the following parameters (energy scale set by χ_0):

$$\chi_{\alpha\beta} = \chi_0 \begin{bmatrix} -0.412882 & 0.4136497 & 0.3361028 \\ 0.4136497 & 0.2010825 & -0.242741 \\ 0.3361028 & -0.242741 & 0.3000709 \end{bmatrix} \times 10^{-2}, \tag{7}$$

consisting of random numbers in $[-1/2N, 1/2N]$. This ensures that no physical properties can be utilized for controlling the system, consistent with our assumption that no information about the system is known a priori. We remark that, for the same reason, we cannot reduce the learning of the dynamics to estimating the coefficients of $\chi_{\alpha\beta}$, as this exploits the knowledge of the internal interaction being quadratic.

To demonstrate our scheme and showcase its efficacy, we consider two prototypical QOC tasks relevant to collective spin systems and quantum metrology. The first task concerns enhanced spin squeezing. Squeezed spin states[9] are quantum states whose quantum fluctuations along a specific direction are suppressed. As a result, they can be used to reach measurement sensitivity beyond the standard quantum limit along the squeezed direction. Enhanced spin squeezing thus leads to improved quantum measurement precision and represents an important QOC task in collective spin systems.[4] The degree of spin squeezing is quantified by the squeezing parameter $\xi^2 = N(\Delta \boldsymbol{S})^2_{\rm min}/|\langle \boldsymbol{S}\rangle|^2$,[41] where $\langle \boldsymbol{S}\rangle$ is the mean collective spin and $(\Delta \boldsymbol{S})^2_{\rm min}$ represents the quantum fluctuation along the squeezed direction in the tangent plane perpendicular to $\langle \boldsymbol{S}\rangle$. Thus, ξ^2 serves as the FoM in the first task.

The second task focuses on improving the quantum Fisher information (QFI) $F_{\rm Q}$ for a given measurement direction. It complements the first task by enabling one to improve the measurement precision along the desired direction directly, in comparison to the nominal approach of creating squeezing first and subsequently rotating the squeezed spin state. When one measures a phase ϕ encoded in a probe state $|\psi\rangle$ via $\mathrm{e}^{-\mathrm{i}\boldsymbol{S}\cdot\hat{\boldsymbol{n}}\phi}$, where $\hat{\boldsymbol{n}}$ denotes the pre-selected measurement direction, the precision is limited by the quantum Cramér–Rao bound $(\Delta\phi)_{\rm QCR} = 1/\sqrt{F_{\rm Q}}$. Therefore, the improved $F_{\rm Q}$ gives the plausible precision limit along $\hat{\boldsymbol{n}}$ and is a key property for quantum sensing.[42] For a closed system in a pure state, $F_{\rm Q}$ is given by the variance $F_{\rm Q} = 4(\Delta S_{\hat{\boldsymbol{n}}})^2$. Here, we take $\hat{\boldsymbol{n}}$ along the y-axis as an example, which gives the FoM $F_{\rm Q} = 4(\Delta S_y)^2$.

The remaining setup is as follows. Since the initial state choice is not critical in the above two tasks, we take the polar state $|S_z = N/2\rangle$ for simplicity. The control field is chosen as the effective magnetic field along the x-axis, and the pulse duration for optimal control is assumed to be $T = 5/\chi_0$. Thus, the total Hamiltonian is given by $H = u(t)S_x + H_0$, leading to $h_{\rm c} = \sigma_x/2$ in Eqs. (3) and (5). The Wigner quasi-probability distribution of the initial conditions $\boldsymbol{\alpha}(0)$ is given by

$$\mathcal{W}(\boldsymbol{\alpha}) = \left(\frac{2N}{\pi}\right)^2 \exp[-2N(|\alpha_\uparrow - 1|^2 + |\alpha_\downarrow|^2)]. \tag{8}$$

Inspired by Ref. [37], we construct the training dataset using $N_{\rm train} = 20$ pulses of the form $u_k(t) = \chi_0 \sum_{m=1}^{10} \sin(2\pi\omega_{k,m}t)$, where $\omega_{k,m}$ is sampled uniformly from $[-1/2, 1/2]$. We use the normalized spin components $\langle \boldsymbol{s}\rangle = \langle \boldsymbol{S}\rangle/N$ as observables in the datasets, which are easily accessible in experiments and have been shown to be sufficient for learning the dynamics[37] (the normalization over N is for training stability). We exclude $\langle s_y\rangle$ from the dataset, as both $\langle s_y\rangle$ and $\langle s_z\rangle$ describe transverse responses with respect to the control axis (x-axis), and thus contribute redundant information. The resulting datasets consist of $\langle s_x\rangle$ and $\langle s_z\rangle$, measured at $N_t = 50$ evenly spaced time points over $[0, T]$.

In practice, training data are obtained from experimental measurements, as opposed to the simulated ones used here from numerical evolutions based on exact diagonalization (ED). Realistically, the system size N may vary between different experimental realizations. To incorporate such variation, for each data point, we average over 10 exact simulations with $\Delta N/N \in [-10\%, 10\%]$; see the SM for details.

Table 1. Mean squared errors (MSEs) obtained with the learned NODE under training, validation, and test pulses.

Pulses	$\langle s_x\rangle$	$\langle s_y\rangle$	$\langle s_z\rangle$	$\langle s_x^2\rangle$	$\langle s_y^2\rangle$	$\langle s_z^2\rangle$	$\langle s_x s_y\rangle_{\rm sym}$	$\langle s_y s_z\rangle_{\rm sym}$	$\langle s_z s_x\rangle_{\rm sym}$
Training	2.11×10^{-6}	8.34×10^{-6}	6.54×10^{-6}	2.59×10^{-6}	4.00×10^{-6}	4.39×10^{-6}	1.80×10^{-6}	2.21×10^{-6}	2.01×10^{-6}
Validation	5.62×10^{-6}	1.15×10^{-5}	1.22×10^{-5}	1.54×10^{-6}	5.11×10^{-6}	6.36×10^{-6}	1.82×10^{-6}	3.50×10^{-6}	1.83×10^{-6}
Test	2.08×10^{-5}	3.52×10^{-5}	3.93×10^{-5}	5.28×10^{-6}	8.87×10^{-6}	1.00×10^{-5}	3.76×10^{-6}	7.01×10^{-6}	4.98×10^{-6}

To avoid overfitting, we construct a separate validation set by generating an additional set of $N_{\rm val}$ pulses. Data in the validation set are obtained in the same manner as those in the training set and are used to compute the validation loss $\mathcal{L}_{\rm val}$ (evaluated using MSE as well), which serves as an estimate of the NODE's performance on previously unseen data. We train the NODE for a fixed number of epochs, $N_{\rm epoch} = 2\times10^4$, and select the set of network parameters corresponding to the lowest observed $\mathcal{L}_{\rm val}$ during training as our final result.

To learn $h_0(\boldsymbol{\alpha})$, we employ a modified multi-layer fully connected network that incorporates physical knowledge into its architecture. Specifically, since the total spin is conserved in a closed collective spin system, the norm $|\boldsymbol{\alpha}|$ along each sampled TWA trajectory needs to be conserved. To respect this conservation, we decompose $h_0(\boldsymbol{\alpha}) = [\sum_{i=0}^{3} c_i(\boldsymbol{\alpha})\sigma_i]\boldsymbol{\alpha}$, where $c_i(\boldsymbol{\alpha})$ are real-valued functions of $\boldsymbol{\alpha}$, $\sigma_0 = I$ is the identity matrix, and $(\sigma_1,\sigma_2,\sigma_3) = (\sigma_x,\sigma_y,\sigma_z)$ are the Pauli matrices. The four real-valued functions c_i can thus be jointly parameterized by a single multi-layer fully connected network with four outputs, with the activation function chosen to be the rectified linear unit (ReLU). We further introduce an extra learnable parameter γ that multiplies c_i, which is initialized to $\gamma_0 = 0$ to stabilize the initial stage of training.

The hyperparameters of such a network structure include the number of hidden layers $n_{\rm layer}$ and the hidden dimension $d_{\rm h}$ of each layer. To determine their values, we train models for different combinations of $n_{\rm layer} \in \{1,2\}$ and $d_{\rm h} \in \{16,32,64,128\}$. Their training and validation loss curves are shown in the SM. Based on the results, we proceed with $n_{\rm layer} = 2$ and $d_{\rm h} = 64$, which give the lowest $\mathcal{L}_{\rm val}$ among all combinations. Nevertheless, as shown in the SM, the performance of the learned NODE is stable for other sets of hyperparameters as well, indicating that our scheme does not require sophisticated hyperparameter tuning.

The prediction errors of the learned NODE are listed in Table 1. We report the MSEs for $\langle s_i\rangle$ and $\langle s_i s_j\rangle$, since they are the most relevant quantities to the QOC tasks under study. It is shown that, while learning is based on $\langle s_x\rangle$ and $\langle s_z\rangle$ only, the learned NODE achieves consistently low errors ($\lesssim10^{-5}$) across all nine observables, indicating excellent agreement with the ED results. To ensure that the learned NODE is not overfitting, we evaluate the NODE on an independent test set generated by $N_{\rm test} = 100$ pulses. The resulting MSEs, shown in Table 1, remain at $\sim10^{-5}$, confirming that the accuracy indeed generalizes beyond the training and validation pulses and that the learned NODE effectively captures the system dynamics. This verifies that the learned NODE does not overfit the training pulses and can be applied to optimal control tasks.

With the learned NODE as the NN surrogate, we proceed towards the two aforementioned QOC tasks. Since the system is assumed to be unknown, there is no preferred initial guess pulse based on physical intuitions. We start the optimization from the control-free case of $u_{\rm init}(t) = 0$. Other details for the optimal control stage can be found in the SM.

Results are shown in Fig. 3. In both tasks, smooth optimal pulses are obtained. For these pulses, the predicted FoMs closely match their values obtained by ED, demonstrating that the NODE provides reliable gradients throughout the optimization. Consequently, significant improvements in FoMs are achieved. For the initial polar state, which constructs the baseline of the optimal control, the FoMs are found to be $\xi^2_{\rm base} = 1.0$ (0 dB) and $F_{\rm Q,base}/N^2 = 0.01$. For spin squeezing, the FoM is reduced to $\xi^2_{\rm opt} = 0.048$, equivalently a 13.2 dB gain over the initial polar state. As for the QFI, it is improved to $F_{\rm Q,opt}/N^2 = 0.576$, significantly higher than the baseline value as well. We also plot the FoMs obtained by the initial pulse in Fig. 3, which illustrates that FoMs are not improved by the initial pulse alone. These results confirm the effectiveness of using the NN surrogate for optimal control.

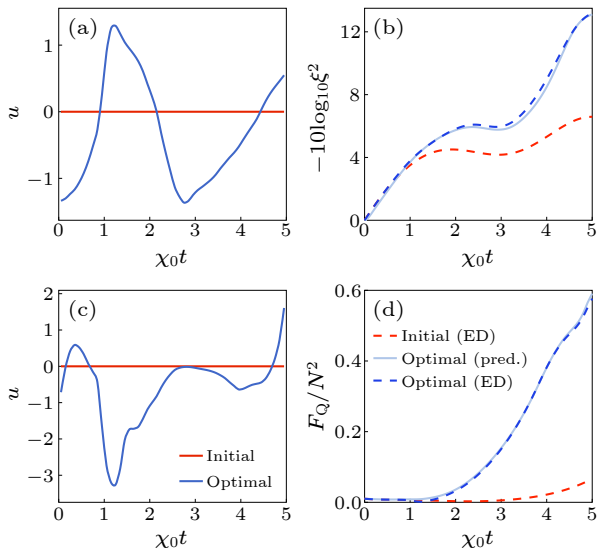

Fig. 3. Optimal control results using the NN surrogate. The first row shows results for enhanced spin squeezing, and the second row shows results for improved QFI. In the left column, the initial (red) and optimal (blue) pulse shapes are presented, while the right column depicts their corresponding evolutions of ξ^2 and $F_{\rm Q}$. The dashed lines are obtained by ED, while the faint blue solid lines indicate the NODE predictions.

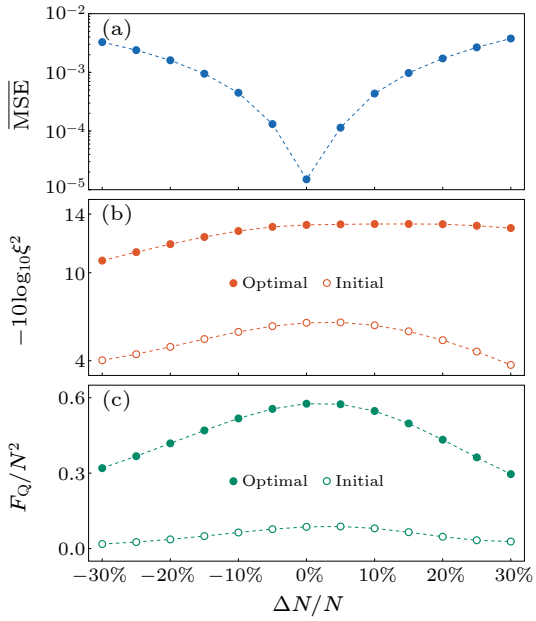

Fig. 4. Applications of the NN surrogate to varying N. (a) The prediction errors $\overline{\mathrm{MSE}}$ under the test pulses, averaged over all nine observables. (b) and (c) Compare the FoMs of optimal pulses with initial pulses for enhanced spin squeezing and improved QFI, respectively. The lines serve as guides to the eye.

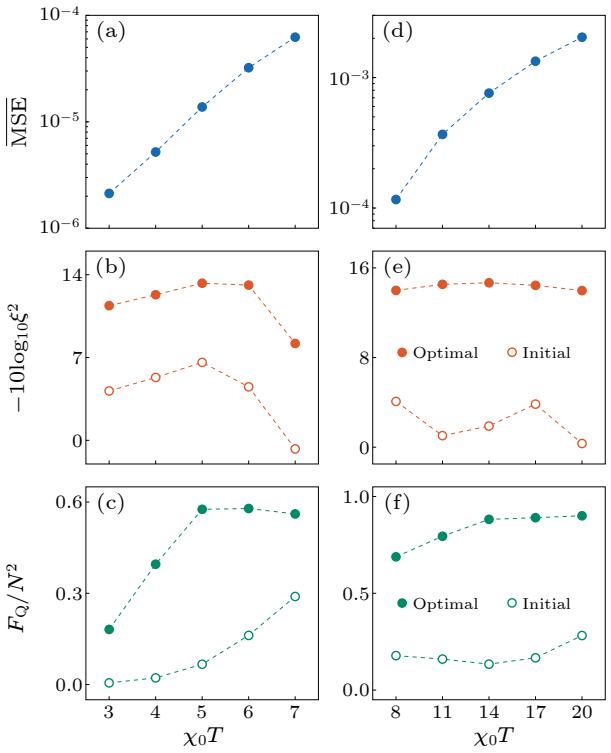

Fig. 5. Applications of the NN surrogate to different T. The left column shows results for $T \in [3,7]/\chi_0$ with $u_{\mathrm{init}} = 0$, while the right column extends to $T \in [8,20]/\chi_0$ using $u_{\mathrm{init}} = 2\chi_0$. The first row plots $\overline{\mathrm{MSE}}$ under test pulses, while the second and third rows compare the FoMs of optimal pulses with initial pulses for enhanced spin squeezing and improved QFI, respectively. The lines serve as guides to the eye.

To assess the flexibility of the NN surrogate, we investigate its performance when system parameters differ between training and optimal control. We first examine variations in the system size N, which have been accounted for in the training datasets with $\Delta N/N \in [-10\%, 10\%]$. Here, we evaluate the performance of the NN surrogate over an extended range of $\Delta N/N \in [-30\%, 30\%]$. To adapt the samples of $\boldsymbol{\alpha}(0)$ to $W(\boldsymbol{\alpha})$ in Eq. (8) for different N, we rescale those originally drawn with $N_0 = 100$ by $\sqrt{N_0/N}$. Using these samples, the prediction errors of the NN surrogate remain at a moderate level ($\lesssim 10^{-3}$; see Fig. 4), and optimal control using the NN surrogate consistently improves the FoMs for all $\Delta N/N \in [-30\%, 30\%]$. This confirms the robustness of the NN surrogate against fluctuations in N.

We further investigate the effect of varying the temporal control duration T, which may differ according to the specific requirements of different QOC tasks. The left column of Fig. 5 presents results for $T \in [3,7]/\chi_0$. Low prediction errors are obtained, and FoMs are improved for all cases after optimization. In the right column of Fig. 5, we further consider $T \in [8,20]/\chi_0$, extending T up to 4 times the training interval. For these T values, taking $u_{\mathrm{init}}(t) = 0$ becomes unsuitable, as it leads to extremely over-squeezed states when $T \gtrsim 7/\chi_0$, challenging the intended optimization. We thus adopt an alternative initial pulse $u_{\mathrm{init}}(t) = 2\chi_0$. As shown in Fig. 5, this choice effectively enhances the FoMs for all $T \in [8,20]/\chi_0$. These results demonstrate that the NN surrogate remains effective even when T exceeds the training interval, confirming its strong flexibility against variations in T.

4. Discussion. This work introduces a machine-learning-based scheme for performing QOC tasks in unknown collective spin systems. By using a NODE to learn the system's internal dynamics, we efficiently complete two QOC tasks. The NODE is trained with experimentally accessible data, which in our case are simulated expectations of collective spin components. Once trained, the NODE can accurately reproduce the evolution of spin components and their second-order correlations across different control pulse shapes. The learned NODE is subsequently applied as a surrogate for optimizing spin squeezing and QFI, yielding effective results in both cases. Additionally, we demonstrate the NODE's performance across varying system sizes and pulse durations, confirming its robustness and flexibility when the experimental parameters in the optimal control stage differ from those in the learning stage.

Our scheme can be further extended in several ways. It can be integrated with more advanced optimization techniques or generalized to other types of unknown systems. A promising direction is adapting it to spin lattice systems, which could involve replacing the TWA with the discrete TWA [43] or the Gutzwiller variational approach. [44–46] Furthermore, it would be exciting to explore how this scheme can be extended to open quantum systems, where dynamics at the TWA level are governed by stochastic differential equations, possibly with the neural stochastic differential equation [47] as a promising candidate for the surrogate NN architecture.

Acknowledgements. We thank Dr. Feng Chen for drawing our attention to the NODE. We are grateful to Dr. Yuqi Wang for critical advice and Ang Li for insightful discussions. This work was supported by the Innovation Program for Quantum Science and Technology (Grant No. 2021ZD0302100) and the National Natural Science Foundation of China (Grant Nos. 12361131576, 92265205, and 92476205). The implementation of the NODE is based on torchdiffeq.[48] The simulated data and the ED results are produced with JAX.[49]

References

[1] Baum Y, Amico M, Howell S, Hush M, Liuzzi M, Mundada P, Merkh T, Carvalho A R, and Biercuk M J 2021 *PRX Quantum* **2** 040324
[2] Jandura S and Pupillo G 2022 *Quantum* **6** 172
[3] Omran A, Levine H, Keesling A, Semeghini G, Wang T. T, Ebadi S, Bernien H, Zibrov A. S, Pichler H, Choi S, Cui J, Rossignolo M, Rembold P, Montangero S, Calarco T, Endres M, Greiner M, Vuletić V, and Lukin M D 2019 *Science* **365** 570
[4] Chen F, Chen J J, Wu L N, Liu Y C, and You L 2019 *Phys. Rev. A* **100** 041801
[5] Li X X, Li D X, and Shao X Q 2024 *Phys. Rev. A* **109** 042604
[6] Kawaguchi Y and Ueda M 2012 *Phys. Rep.* **520** 253
[7] Stamper-Kurn D M and Ueda M 2013 *Rev. Mod. Phys.* **85** 1191
[8] Leroux I D, Schleier-Smith M H, and Vuletić V 2010 *Phys. Rev. Lett.* **104** 073602
[9] Kitagawa M and Ueda M 1993 *Phys. Rev. A* **47** 5138
[10] Pezzè L, Smerzi A, Oberthaler M K, Schmied R, and Treutlein P 2018 *Rev. Mod. Phys.* **90** 035005
[11] Lipkin H, Meshkov N, and Glick A 1965 *Nucl. Phys.* **62** 188
[12] Meshkov N, Glick A, and Lipkin H 1965 *Nucl. Phys.* **62** 199
[13] Glick A, Lipkin H, and Meshkov N 1965 *Nucl. Phys.* **62** 211
[14] Caneva T, Calarco T, and Montangero S 2011 *Phys. Rev. A* **84** 022326
[15] Rach N, Müller M M, Calarco T, and Montangero S 2015 *Phys. Rev. A* **92** 062343
[16] Müller M M, Said R S, Jelezko F, Calarco T, and Montangero S 2022 *Rep. Prog. Phys.* **85** 076001
[17] Bukov M, Day A G R, Sels D, Weinberg P, Polkovnikov A, and Mehta P 2018 *Phys. Rev. X* **8** 031086
[18] Fösel T, Tighineanu P, Weiss T, and Marquardt F 2018 *Phys. Rev. X* **8** 031084
[19] Sivak V V, Eickbusch A, Liu H, Royer B, Tsioutsios I, and Devoret M H 2022 *Phys. Rev. X* **12** 011059
[20] Flurin E, Martin L S, Hacohen-Gourgy S, and Siddiqi I 2020 *Phys. Rev. X* **10** 011006
[21] Mohseni N, Fösel T, Guo L Z, Navarrete-Benlloch C, and Marquardt F 2022 *Quantum* **6** 714
[22] Choi M, Flam-Shepherd D, Kyaw T H, and Aspuru-Guzik A 2022 *Phys. Rev. A* **105** 042403
[23] Mazza P P, Zietlow D, Carollo F, Andergassen S, Martius G, and Lesanovsky I 2021 *Phys. Rev. Res.* **3** 023084
[24] Banchi L, Grant E, Rocchetto A, and Severini S 2018 *New J. Phys.* **20** 123030
[25] Carleo G and Troyer M 2017 *Science* **355** 602
[26] Schmitt M and Heyl M 2020 *Phys. Rev. Lett.* **125** 100503
[27] Gutiérrez I L and Mendl C B 2022 *Quantum* **6** 627
[28] Sinibaldi A, Giuliani C, Carleo G, and Vicentini F 2023 *Quantum* **7** 1131
[29] Hartmann M J and Carleo G 2019 *Phys. Rev. Lett.* **122** 250502
[30] Reh M, Schmitt M, and Gärttner M 2021 *Phys. Rev. Lett.* **127** 230501
[31] Luo D, Chen Z, Carrasquilla J, and Clark B K 2022 *Phys. Rev. Lett.* **128** 090501
[32] Norambuena A, Mattheakis M, González F J, and Coto R 2024 *Phys. Rev. Lett.* **132** 010801
[33] Wigner E 1932 *Phys. Rev.* **40** 749
[34] Sinatra A, Lobo C, and Castin Y 2002 *J. Phys. B: At. Mol. Opt. Phys.* **35** 3599
[35] Sakurai J J and Napolitano J 2020 *Modern Quantum Mechanics* 3rd ed. (Cambridge University Press)
[36] Chen R T Q, Rubanova Y, Bettencourt J, and Duvenaud D K 2018 *Proceedings of the 32nd International Conference on Neural Information Processing Systems* p. 6572
[37] Chen L and Wu Y 2022 *Phys. Rev. A* **106** 022201
[38] Tucker K, Rege A K, Smith C, Monteleoni C, and Albash T 2024 *Phys. Rev. Appl.* **22** 044080
[39] Hochreiter S and Schmidhuber J 1997 *Neural Comput.* **9** 1735
[40] Cho K, van Merriënboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, and Bengio Y 2014 *Conference on Empirical Methods in Natural Language Processing (EMNLP)* p. 1724
[41] Wineland D J, Bollinger J J, Itano W M, Moore F L, and Heinzen D J 1992 *Phys. Rev. A* **46** R6797
[42] Xiao T, Fan J, and Zeng G 2022 *npj Quantum Inf.* **8** 1
[43] Schachenmayer J, Pikovski A, and Rey A M 2015 *Phys. Rev. X* **5** 011022
[44] Gutzwiller M C 1963 *Phys. Rev. Lett.* **10** 159
[45] Gutzwiller M C 1964 *Phys. Rev.* **134** A923
[46] Gutzwiller M C 1965 *Phys. Rev.* **137** A1726
[47] Li X, Wong T K L, Chen R T Q, and Duvenaud D 2020 *PMLR* **108** 3870
[48] Chen R T Q 2018 *torchdiffeq: Differentiable ODE solvers with full GPU support and O(1)-memory backpropagation*
[49] Bradbury J, Frostig R, Hawkins P, Johnson M. J, Leary C, Maclaurin D, Necula G, Paszke A, Vander-Plas J, Wanderman-Milne S, and Zhang Q 2018 *JAX: Composable transformations of Python+NumPy programs*