

PAPER

# Digital-analog quantum learning on Rydberg atom arrays

To cite this article: Jonathan Z Lu *et al* 2025 *Quantum Sci. Technol.* **10** 015038

View the [article online](#) for updates and enhancements.

## You may also like

- [Pseudomode treatment of strong-coupling quantum thermodynamics](#)

Francesco Albarelli, Bassano Vacchini and Andrea Smirne

- [Quantum state tomography based on infidelity estimation](#)

Yong Wang, Lijun Liu, Tong Dou *et al.*

- [Driving superconducting qubits into chaos](#)

Jorge Chávez-Carlos, Miguel A Prado Reynoso, Rodrigo G Cortiñas *et al.*

# Quantum Science and Technology



## PAPER

### Digital-analog quantum learning on Rydberg atom arrays

RECEIVED

11 April 2024

REVISED

20 October 2024

ACCEPTED FOR PUBLICATION

12 November 2024

PUBLISHED

27 November 2024

Jonathan Z Lu<sup>1,2,3,\*</sup> , Lucy Jiao<sup>1</sup> , Kristina Wolinski<sup>1,4</sup>, Milan Kornjača<sup>2</sup> , Hong-Ye Hu<sup>1</sup> , Sergio Cantu<sup>2</sup>, Fangli Liu<sup>2</sup>, Susanne F Yelin<sup>1</sup> and Sheng-Tao Wang<sup>2</sup>

<sup>1</sup> Department of Physics, Harvard University, Cambridge, MA 02138, United States of America

<sup>2</sup> QuEra Computing Inc., 1284 Soldiers Field Road, Boston, MA 02135, United States of America

<sup>3</sup> Presently at Department of Mathematics, MIT.

<sup>4</sup> Presently at Department of Physics, Princeton.

\* Author to whom any correspondence should be addressed.

E-mail: [lujz@mit.edu](mailto:lujz@mit.edu), [fliu@quera.com](mailto:fliu@quera.com) and [swang@quera.com](mailto:swang@quera.com)

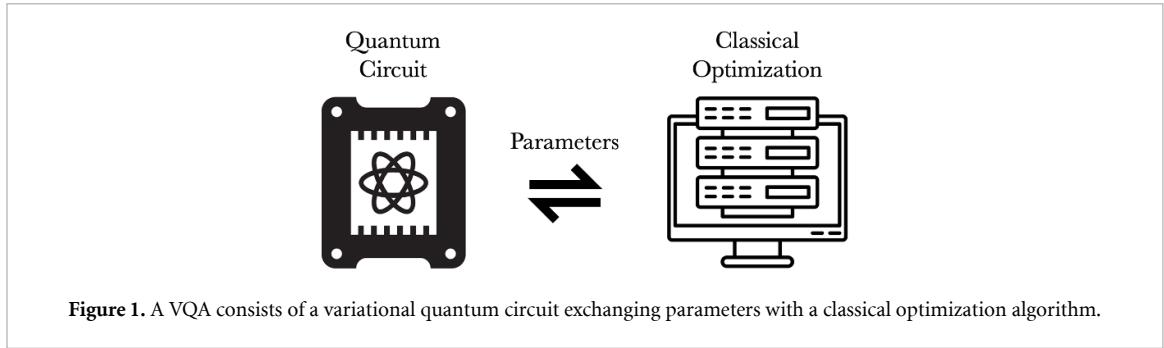
**Keywords:** digital-analog (DA) algorithms, quantum algorithms, quantum machine learning, Rydberg atoms, variational quantum algorithm

#### Abstract

We propose hybrid digital-analog (DA) learning algorithms on Rydberg atom arrays, combining the potentially practical utility and near-term realizability of quantum learning with the rapidly scaling architectures of neutral atoms. Our construction requires only single-qubit operations in the digital setting and global driving according to the Rydberg Hamiltonian in the analog setting. We perform a comprehensive numerical study of our algorithm on both classical and quantum data, given respectively by handwritten digit classification and unsupervised quantum phase boundary learning. We show in the two representative problems that DA learning is not only feasible in the near term, but also requires shorter circuit depths and is more robust to realistic error models as compared to digital learning schemes. Our results suggest that DA learning opens a promising path towards improved variational quantum learning experiments in the near term.

#### 1. Introduction

Neutral atoms trapped by optical tweezers hold great potential for a scalable quantum information processing [1–5]. Programmable quantum simulators comprised of an array of Rydberg atoms have already scaled to several hundred qubits, enabling recent experiments that have demonstrated intriguing emergent many-body quantum phenomena including quantum phase transitions [6–11], many-body quantum scars [12–14], and topological phases of matter [15–19]. Rydberg atom simulators have even been shown to naturally encode hard computational problems such as the Maximum Independent Set problem, leading to new connections between quantum physics and NP-complete computational problems [20–25]. The rapid progress towards larger scale Rydberg quantum processors motivates us to examine their applications for near-term quantum algorithms. In the case of Rydberg system using Rubidium atoms, a simplified picture of its quantum simulator is a system of atoms trapped in a desired arrangement, individually in their ground state (denoted  $|0\rangle$ ), that is globally driven for some finite time  $t$  under the Rydberg many-body Hamiltonian [26]. This capability alone is enough to realize many exotic physical properties of Rydberg arrays discussed above, and present-day simulators may be able to simulate other complex physical models, such as those in high-energy physics [27, 28]. Beyond global driving, another high-fidelity control that can be added to a Rydberg atom array system is single-qubit rotations. These gates operate in a space of hyperfine states rather than the Rydberg two-level system, which we will describe in more detail in the following sections. This ability to locally change the basis of each qubit enables a new host of applications, including quantum machine learning (QML). Recent theoretical studies have indicated that QML, in some cases, can provide practical benefits for certain problems [29–36]. The possibility of executing quantum computations with practical utility motivates us to study explicit constructions of such learning algorithms. In this paper, we explore quantum learning algorithms from near-term Rydberg atom simulators—their construction and potential benefits over digital variational quantum circuit models [37, 38].



Variational quantum algorithm (VQA) is one of the most widely considered QML models. Experimentally, the VQA recipe yields algorithms generally compatible (under reasonable circuit depth) with near-term quantum devices [39]. A remarkably wide class of VQA and VQA-variants have been proposed in recent years for applications that span a range from purely computational and physical problems to drug discovery [39–46].

We begin by briefly reviewing the general structure of a VQA so that we may connect it directly to our Rydberg construction. The VQA formulation follows two guiding principles: divide the tasks of machine learning between quantum and classical computers, and allocate to each resource the problems it is best at solving [39]. Depicted schematically in figure 1, a VQA, therefore, consists of a quantum circuit and a classical computer interacting with each other. The quantum circuit contains a set of parameters that may be varied to learn an underlying problem.

In its simplest form, a learning task involves a training and a testing component. The training portion involves optimizing the parameters of the quantum circuit with respect to a loss function  $\mathcal{L}$  derived from the circuit measurements. This optimization is generally done classically with, e.g. simplex search and gradient descent [47–50]. Such techniques require that the classical optimizer iteratively query the quantum circuit for both its parameters and its gradient; the optimizer then updates the parameter based on the gradient value. A classical analogy to this process might be a neural network optimized by gradient descent via the same iterative query process, though a VQA does not have any classical neural networks in its construction. The testing portion consists of querying the optimized quantum circuit with new input data and analyzing the output.

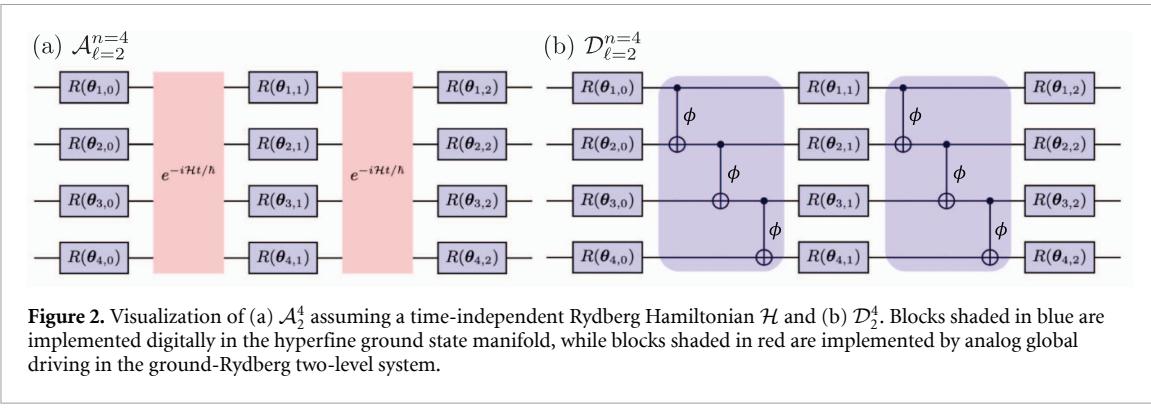
### 1.1. VQAs with Rydberg atom system

One important task to be addressed on the quantum circuit side is the choice of parameterized circuit; that is, choosing an Ansatz that is most suitable for learning. A trainable variational Ansatz typically is comprised of (a) gates corresponding to a collection of optimizable parameters and (b) gates that generate entanglement. These gates need not be disjoint generally. Correspondingly, in our model of learning on globally-driven Rydberg atom arrays, we have two resources: (a) single-qubit rotation gates  $R(\theta)$  with Euler angles  $\theta$  and (b) global time evolution  $e^{-i\mathcal{H}t/\hbar}$  under the Rydberg Hamiltonian  $\mathcal{H}$  (see equation (2)). Gates of the former type (as well as any discrete multi-qubit gates) are referred to as *digital* gates while global driving operations are referred to as *analog* gates. Analog gates are derived from the native Rydberg Hamiltonian dynamics, and while they do not correspond to an arbitrary unitary on the qubits they are applied on, we explore their effectiveness as a QML entangling layer. A general family of Ansätze (parameterized over system size  $n$  and circuit depth) assuming a time-independent Hamiltonian is given by a recursive definition over the number of layers  $\ell$ :

$$\mathcal{A}_\ell^n \left( \{\theta_{i,j}\}_{i=1,j=0}^{n,\ell} \right) = \mathcal{A}_{\ell-1}^n \left( \{\theta_{i,j}\}_{i=1,j=0}^{n,\ell-1} \right) \cdot e^{-i\mathcal{H}t/\hbar} \cdot \bigotimes_{i=1}^n R_i(\theta_{i,\ell}) \quad (1)$$

with a base case  $\mathcal{A}_0^n \left( \{\theta_{i,0}\}_{i=1}^n \right) = \bigotimes_{i=1}^n R_i(\theta_{i,0})$ . The number of layers  $\ell$  is equivalently the number of time evolution gates, such that the depth of a  $\mathcal{A}_\ell^n$  is  $d = 2\ell + 1$ . An example for  $n = 4, \ell = 2$  is given in figure 2(a). All circuits of the form  $\mathcal{A}_\ell^n$  are referred to as hybrid digital–analog (DA) learning circuits.

In this paper, we conduct a thorough numerical study of the performance and limits of practical quantum learning on DA circuitry that can be feasibly implemented in Rydberg atom arrays. To analyze the capabilities of DA learning under realistic conditions, we compare DA circuits with a natural digital counterpart  $\mathcal{D}_\ell^n$ , shown in figure 2(b). In  $\mathcal{D}_\ell^n$ , each time evolution operator is replaced with a layer of  $n - 1$



adjacent generalized controlled-NOT gates between nearest neighbors. These generalized gates are endowed with a global parameter  $\phi$  such that the gate is precisely CNOT—mapping basis elements  $|b_1, b_2\rangle \mapsto |b_1, b_1 \oplus b_2\rangle$ —when  $\phi = \pi/4$ , but can be more general. The extra parameter gives additional power to the digital scheme, which in turn strengthens the performance gap between digital and DA schemes overall.

In order to conduct a comprehensive comparative numerical study, we choose two representative learning problems on which we apply DA learning. The first is a supervised learning problem where the input is classical information encoded into a product state: recognition of handwritten digits from the MNIST dataset [51]. The second is an unsupervised learning problem where the input is quantum information: learning phase boundaries of various Hamiltonians based on the ground state as input. Both of our algorithms can be directly implemented on a Rydberg atom array that admits global driving/time evolution and single-qubit operations. We will show that in both cases, DA methods have benefits over purely digital techniques in performance and in robustness to error, poising them as a promising technique to execute VQAs in the near term.

In section 2, we describe the time evolution generated by the Rydberg Hamiltonian and its relation to experimental constraints, and discuss realistic noise models for both digital and DA circuitry. In section 3, we discuss our results from learning both classical and quantum information. In section 4, we summarize our work and discuss an outlook on DA learning.

## 1.2. Related work, contributions, and limitations

The idea of DA circuitry has been proposed in the literature. Hybridizing single-qubit digital gates and analog evolution has been studied for quantum simulation purposes under simple models, such as the Ising Hamiltonian for the simulation of spin models [52, 53]. Numerical simulations of quantum Fourier transforms and phase estimation algorithms implemented with DA architectures have given evidence of their comparative robustness to noise [54, 55]. It is even possible that DA quantum computation may open an avenue to a realizable practical quantum advantage [56].

Our primary contributions are twofold. First, we provide a comprehensive comparative study of DA architectures on representative learning problems spanning classical and quantum data. It is natural to study performance on both forms of data; such a juxtaposition has been considered experimentally as well [57]. Second, we co-design our constructions with Rydberg devices, requiring only single-qubit digital gates and global driving operations. We show explicitly that physically astute choices of analog pulse parameters lead to high-performance DA circuits. The current state of the Rydberg platform moreover suggests that our proposed learning techniques can be implemented within the next few years on Rydberg devices at the order of hundreds of qubits or more [5, 10, 12, 15, 26]. Time evolution on Rydberg atom arrays has been extensively studied, and high-fidelity single-qubit gates can be implemented at scale via hyperfine state manipulation [58–61]. Consequently, we propose these architectures for the purpose of near-term realization on Rydberg devices. At the same time, however, the robustness of our results suggests that DA architectures may be beneficial more generally for QML tasks, and should be viewed as a generic technique to boost learning schemes rather than an attempt to deal with the handicap of near-term device limitations.

Although both schemes have hyperparameters in the entangling layer, one might argue that choosing all of the Hamiltonian parameters and the evolution time in the analog scheme are more challenging computationally than the choice of  $\phi$  in the digital scheme. For example, a simple choice of  $\phi$  is  $\pi/4$ ; that is to use CNOT gates. In general, if there are many hyperparameters in a model, they may be practically challenging to optimize by brute-force computation. However, for the case of Rydberg Hamiltonians, we are able to circumvent this problem by arguing for the optimality of specific hyperparameters guided by the

physics of Rydberg atoms. These arguments, which we present in the following sections, are physically motivated and yield excellent results. Consequently, this naïve difficulty when guided by physics reveals itself to actually be a significant *advantage* of DA learning, allowing for improved learning without blindly optimizing additional parameters through the leveraging of physical reasoning. Moreover, we will show that the simple choice of  $\phi = \pi/4$  is substantially worse than the optimal choice.

## 2. Theory

We consider a model of a Rydberg system that is described by two distinct modes; the device can switch between the two modes by certain laser pulse sequences. In the digital mode, an atom in the Rydberg array can be in any one of  $d$  states and is manipulated according to the standard circuit model. In this paper, we set  $d = 2$  and choose two states from the Rydberg hyperfine ground state manifold to encode the digital qubit [58, 59]. In the analog mode, each atom is described by a two-level system consisting of a ground state  $|g\rangle$  and an excited state  $|r\rangle$ . The dynamics of the Rydberg atom array are governed by the Rydberg Hamiltonian [26]

$$\begin{aligned} \frac{\mathcal{H}(t)}{\hbar} = & \sum_{j=1}^n \frac{\Omega_j(t)}{2} \left( e^{i\phi_j(t)} |g_j\rangle\langle r_j| + e^{-i\phi_j(t)} |r_j\rangle\langle g_j| \right) \\ & - \sum_{j=1}^n \Delta_j(t) \hat{n}_j + \sum_{j < k} V_{jk} \hat{n}_j \hat{n}_k \end{aligned} \quad (2)$$

where  $\Omega_j$  is the Rabi frequency of atom  $j$ ,  $\phi_j$  is a phase factor,  $\Delta_j$  is the laser detuning frequency,  $\hat{n}_j = |r_j\rangle\langle r_j|$  is the Rydberg projection operator, and

$$V_{jk} = \frac{C_6}{|\vec{r}_j - \vec{r}_k|^6} \quad (3)$$

is a van der Waals interaction potential with  $C_6 = 862690 \times 2\pi \text{ MHz} \mu\text{m}^6$  [26]. The detuning  $\Delta_j(t)$ , frequency  $\Omega_j(t)$  and atomic distance  $|\vec{r}_j - \vec{r}_k|$  are all parameters of experimental control. The time dependence of the Rydberg Hamiltonian has been used for techniques such as adiabatic evolution of ground states [20, 62]. We will discuss the role of adiabatic evolution for the preparation of ground states as quantum training data in the following section. However, to obtain the simplest operations with a tractable number of hyperparameters in our DA model, we restrict to the time-independent setting. We further simplify by enforcing site independence; that is, we let all parameters take the same value across all atoms, i.e.  $\Omega_i = \Omega_j = \Omega$ ,  $\forall i, j$  and similarly for the other parameters, and set  $\phi_j = 0$ .

As an operator,  $\mathcal{H}$  depends both on the geometry of the underlying Rydberg atom array and a dimensionless detuning parameter  $\Delta/\Omega$ . We adopt a simple geometry—a one-dimensional chain of atoms with open boundary conditions, spaced apart by a distance  $a$ . The lattice parameter  $a$  can be made dimensionless by introducing a characteristic length scale, the so-called blockade radius  $R_b$  defined by

$$\frac{C_6}{R_b^6} = \Omega \quad (4)$$

the length at which the interaction strength equals the Rabi frequency. The blockade radius physically corresponds to the length scale at which the interaction potential term in the Hamiltonian begins to dominate. For a ground state within this regime, among any cluster of Rydberg atoms contained in a disk of radius  $R_b$ , with high probability, no more than one of the atoms may be in the Rydberg state [10, 12, 14].

Consequently, the Hamiltonian is fixed by a frequency parameter  $\Delta/\Omega$  and a lattice parameter  $R_b/a$ . The evolution operation also includes a time parameter  $t$ . Collectively,  $(\Delta/\Omega, R_b/a, t)$  define the hyperparameters over the DA learning model.

For the remainder of this paper, we will rely on context to determine whether the system is in the hyperfine or ground-Rydberg state space, and refer to both two-level systems as  $|0\rangle$  and  $|1\rangle$ . In order to facilitate numerical simulations, we will assume that the transfer between the hyperfine and  $|g\rangle - |r\rangle$  manifolds is performed ideally. This is not the case in physical realizations, as the transfer between the two manifolds is performed by single qubit pulses in the presence of non-zero Rydberg interactions. These interactions implement weakly interacting short-time Hamiltonian evolution layers before and after each single qubit rotation layer that can be absorbed in a modified entangled evolution layer. As such, they do not change the DA machine learning algorithm we present, nor are they expected to significantly alter its performance. Disregarding the interacting evolution during manifold transfer allows us to treat our atoms as

an effective two-level system instead of three-level, thus leading to significant computational savings. The physical implementation of qubits, however, is different in the digital rotation and analog entanglement stages of our proposal. A brief discussion of experimental realizations of a Rydberg DA system is given in appendix A.

## 2.1. Noise models

To simulate realistic conditions on a Rydberg device, we adopt noise models that align with results from experimental measurements on *Aquila*, a 256-atom programmable Rydberg simulator [26]. In this work, we ignore the infidelity of single-qubit gates as they generally have much higher fidelity than multi-qubit gates [61, 63]. For feasibility of numerical simulation, we only consider a model of coherent errors independent among every time evolution operation that captures noise in the Rabi frequencies and detunings as well as position perturbations of each atom. These are given as

1. noisy detuning  $\tilde{\Delta} \sim \Delta + \mathcal{N}(0, 0.1 \text{ MHz})$ ,
2. noisy Rabi frequency  $\tilde{\Omega} \sim \Omega \cdot \mathcal{N}(1, 0.01)$
3. and perturbed atomic coordinates according to a Gaussian process  $\tilde{r}_j \stackrel{\text{iid}}{\sim} r_j + \mathcal{N}(0, 0.1 \mu\text{m})$

where  $\mathcal{N}(\mu, \sigma)$  is a Gaussian distribution with mean  $\mu$  and variance  $\sigma^2$ . These noise parameters are consistent with those of [26]. Although this noise model does not capture all possible sources of noise, it includes a large and important subset of experimental errors and provides a fair comparison of digital versus DA algorithms within our scope.

To define the digital noise model, we first establish the generalized form of the CNOT gate. The general recipe to obtain a gate is to write it as  $U = e^{-i\theta H_U/\hbar}$  for a corresponding Hamiltonian  $H_U$  and time  $\theta$ , and to use a physical system described by  $H_U$  to implement  $U$ . As a concrete example, one may achieve a controlled- $Z$  gate by a two-body Ising interaction with unit time evolution and interaction strength  $\pi/4$ , i.e.  $CZ_{12} = e^{-i\frac{\pi}{4}(Z_1+Z_2)+\frac{\pi}{4}Z_1Z_2}$ . Similarly, we may write

$$CX_{1 \rightarrow 2} = e^{-i\frac{\pi}{4}(I_1-Z_1)(I_2-X_2)} \quad (5)$$

where  $CX_{1 \rightarrow 2}$  is a controlled-NOT gate with the first qubit as the control. This equivalent definition of the CNOT gate is derived in appendix B. More generally, define

$$CX_{1 \rightarrow 2}(\phi) = e^{-i\phi(I_1-Z_1)(I_2-X_2)} \quad (6)$$

to be the generalized CNOT gate, with  $\phi \in [0, \pi/4]$ . As a simple but realistic digital independent noise model, we sample  $\theta$  from a Gaussian distribution with mean  $\mu = \phi$  and standard deviation  $\sigma$  chosen such that the resulting gate fidelity is 99%. As with the DA case, infidelities of the single-qubit gates are neglected. We estimated numerically the desired standard deviation of the noisy CX gate to be  $\sigma = 0.065$ . Our analysis of noisy circuits introduces noise at every instance of the circuit evaluation, including both the training and the testing phases. We do not, so as to save simulation time, re-sample noise in estimation of the gradient and instead use the same noise on the loss function evaluation as well as its gradient in each step.

## 2.2. Physical hyperparameters

In appendix B, we show by direct computation that an approximately optimal choice of  $\phi$  is  $\pi/8$  for the digital model with noise, while  $\phi = \pi/4$ —the controlled-NOT gate—is optimal in the absence of noise. We are, however, primarily interested in model comparisons in the presence of noise. Thus, henceforth, references to the noisy digital or generalized CNOT gate refer to the  $CX_{1 \rightarrow 2}(\pi/8)$  gate, and we maintain the standard CNOT gate for the digital circuit without noise.

We next consider the choice of hyperparameters associated with the DA model.  $\Omega$  fixes the scale of the Hamiltonian, so we set  $\Omega = 2\pi \times 4 \text{ MHz}$  which in turn determines  $R_b$  by equation (4). We argue that the remaining hyperparameters are approximately decoupled, can be chosen primarily through physical justifications rather than variational methods, and are universal. These hyperparameters are constant from layer to layer in the circuit.

The generic task of optimizing hyperparameters for a learning model is difficult and may be performed by brute-force grid search. However, in our case, we argue that the optimal value of each hyperparameter does not depend strongly on those of other hyperparameters. This decoupling occurs because each hyperparameter represents an underlying physical operation related to coupling strength or entanglement in a distinct fashion. Consequently, rather than conduct a single  $M$ -dimensional search for  $M$  hyperparameters, we independently conducted  $M$  one-dimensional searches in a coordinate-wise optimization. Furthermore, some of these searches can be simplified or eliminated by physical arguments, as we discuss below. Note that

although this heuristic argument may be useful, it is not necessary, as it suffices to by any means obtain a set of hyperparameters that demonstrate substantial improvement over the digital model.

First, we discuss the choice of the evolution time  $t$ . For  $t \ll 2\pi/\Omega$ , the quench is so fast that it is close to the identity and therefore does not produce sufficient entanglement. Indeed,  $t$  should be at least on the scale of  $\sim 2\pi/\Omega$  to generate sufficient entanglement. If  $t \gg 2\pi/\Omega$ , the quench will cause most states (except for quantum scar states [12, 14]) to thermalize, resulting in a loss of information. Therefore, an optimal  $t$  exists in an intermediate zone. An intuitive choice is  $t = 2\pi/\Omega$  itself because the quench passes through exactly one full Rabi oscillation, a physical timescale. Hence, we adopt this choice of quench timescale in this paper. Nonetheless, it is reasonable that any time on the scale of  $2\pi/\Omega$  will produce a trainable Ansatz.

Next, we consider the Hamiltonian parameters  $\Delta/\Omega$  and  $R_b/a$ . While we cannot deduce from physical arguments the exact optima, we are able to bound the range containing the optima, and then numerically determine the exact values. A practically realizable regime for  $\Delta/\Omega$  is  $[0, 4]$  [9]. We find similar performance for any  $\Delta/\Omega$  in this region, and we choose  $\Delta/\Omega = 0.8$ . Likewise, a reasonable regime for  $R_b/a$  is between 0 and 1: if  $R_b \leq a$ , the Rydberg atoms will be interacting so strongly that it will be impractical to transfer atoms between the hyperfine ground state manifold and Rydberg manifold. Moreover, the strong interactions in the  $R_b/a > 1$  regime reduce the effective Hilbert space dimension, limiting the space the Ansatz can explore. Within this choice, however, there is a physically optimal region by the following plausibility argument. As  $R_b/a \rightarrow 1$ , the atoms move closer together and, therefore, interact more strongly, which should produce more entanglement as a resource for learning. However, at the same time, moving atoms closer implies that the relative effect of positional fluctuations in our noise model becomes more prominent.

To approximately optimize this tradeoff, we estimate the gate fidelity  $\mathcal{F}$  of the quench  $e^{-i\mathcal{H}t/\hbar}$  (for one layer), defined by

$$\mathcal{F}(U) = \int d\tilde{U} \int d\psi \langle \psi | U^\dagger \tilde{U} | \psi \rangle \quad (7)$$

where  $d\tilde{U}$  denotes the measure of the noise model and  $d\psi$  denotes the Haar measure. The gate fidelity measures the average robustness of our entangling gate over all states. Our goal is to set  $R_b/a$  as close as possible to 1 while maintaining a high gate fidelity. We expect that the gate fidelity essentially characterizes an operator's error robustness under a given error model for quantum learning tasks. This is because  $\mathcal{F}$  measures the fidelity between the noisy and noiseless output, averaged over random states and the error distribution. In a QML training task, as the model iteratively explores the space of operators, we expect the noisy gate to encounter many typical quantum states. Thus, the gate's performance based on the average state determines its typical performance during training.

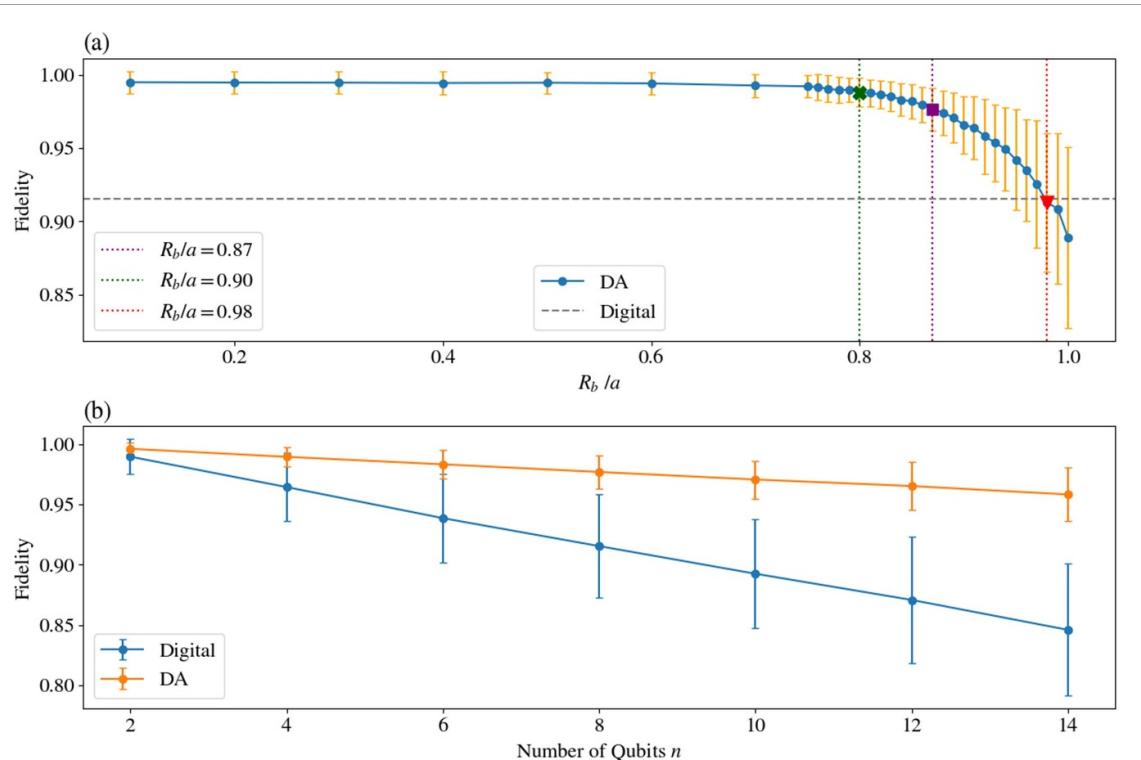
We numerically estimate the gate fidelity by discretizing the integral over  $d\tilde{U}$  over 500 samples of the noise model and show the result in figure 3. The technique is identical for DA and digital models: in each case we sample a gate from their respective noise model discussed previously and calculate the gate fidelity. Figure 3(a) demonstrates that the  $R_b/a \rightarrow 1$  limit shows a dramatic drop in  $\mathcal{F}$ , for a fixed  $n = 8$ . In particular, around  $R_b/a = 0.98$ , the fidelity is similar to that of a digital layer. The fidelity is close to 1 around  $R_b/a = 0.8$ , but we may expect—and confirm in the next section—that such a  $R_b/a$  results in interactions too weak for optimal machine learning performance without noise. A reasonable trade-off is  $R_b/a = 0.87$ , whose approximate optimality we explore further in the next section. At this tradeoff value, the gate fidelity of the DA layer is substantially larger than that of the digital layer.

We show in the following section that a larger gate fidelity does indeed give rise to practical performance improvements. *A priori*, we may already expect an advantage in the DA fidelity scaling over  $n$  as compared to the digital. In general, the digital gate fidelity drops quickly because we require  $n - 1$  generalized CNOT gates per layer. Since each gate is subject to an independent source of noise, the total noise should multiply, thus decreasing the error exponentially as  $\sim 0.99^{n-1}$ . By contrast, the analog gate, being a single global gate with a single source of error, does not have the same obvious exponential downfall and thus opens up the possibility of improved error robustness. Indeed, figure 3(b) demonstrates that this advantage exists for a range of system sizes  $n$ . Although we cannot determine the exact scaling of the analog gate from numerical analysis, we can see that even if it is exponential, it has a noticeably larger base than that of the digital scheme.

Figure 3 also shows the standard deviation over the fidelities of the sampled set. With lower fidelities, the variance of the output between each instance of noise also increases.

Being representative of a physical operation, each of the hyperparameters discussed above should remain optimal between various choices of the underlying computational task. We observed this independence between the two tasks we study in this work.

A final hyperparameter is the number of layers  $\ell$ . However,  $\ell$  lacks a physical interpretation and will generically be task-dependent. Specifically,  $\ell$  encodes an expressivity-trainability tradeoff. If  $\ell$  is too small,



**Figure 3.** (a) Estimated fidelity of the time evolution gate on  $n = 8$  qubits across varying  $R_b/a$ . Errors bars denote the standard deviation across instances of the noisy gates, which increase as the fidelity decreases. Three points of note are  $R_b/a = 0.80$  (cross),  $R_b/a = 0.87$  (square), and  $R_b/a = 0.98$  (triangle). The fidelity for an 8-qubit digital gate is also shown as the dotted horizontal line. Notice that the time evolution gate where  $R_b/a = 0.98$  has almost the same fidelity as the digital gate. (b) Estimated gate fidelity, fixed at  $R_b/a = 0.87$  for digital-analog circuits, compared across system sizes with that of digital circuits.

the model will have an insufficient number of parameters and thus lack the expressivity to perform the learning task. However, if  $\ell$  is too large, then the number of errors may accumulate and result in an insufficient amount of signal to train the circuit. A large  $\ell$  also suffers from higher-order effects such as potential for overfitting and an intractably large parameter dimension. In general, the optimal  $\ell$  may be found by linear or binary search for a given task. We perform the former in the following section.

Our code used for this work<sup>5</sup> is freely available for usage [64].

### 3. Results

We explore the effectiveness of Rydberg DA learning with two different problems. In the first, we train a quantum binary classifier on the MNIST dataset of handwritten digits [51]. In the second problem, we use techniques of anomaly detection to identify the boundaries of quantum phase diagrams. The input of the former problem is classical—but encoded into a product state—while the input of the latter problem is the ground state of a quantum Hamiltonian.

Note that in studying these problems we aim to provide evidence for the generic benefits of DA circuitry in VQA's. Hence, we choose the given problems not necessarily for their usefulness in themselves (though phase learning can be of useful physical interest, which we discuss below), but for their representation of practical classical and quantum learning problems, respectively. Indeed, the construction of a problem for which a VQA is indisputably superior to the best classical alternative remains an open problem and is out of the scope of our present work. We remark that while our work also does not attempt to explicitly mitigate the barren plateau problem of VQA's, our loss functions consist only of local observables so that non-barren plateaus remain a possibility [65].

#### 3.1. Digit classification

The MNIST dataset contains monochromatic images of handwritten single-digit numbers and is widely used as a machine learning benchmarking dataset. To downsample images of the MNIST dataset into a suitable

<sup>5</sup> The numerical simulations we develop in the following section were built on `Bloqade.jl` and `Yao.jl` [82, 83] with automatic differentiation techniques that can be replaced experimentally by parameter shift rules [84].

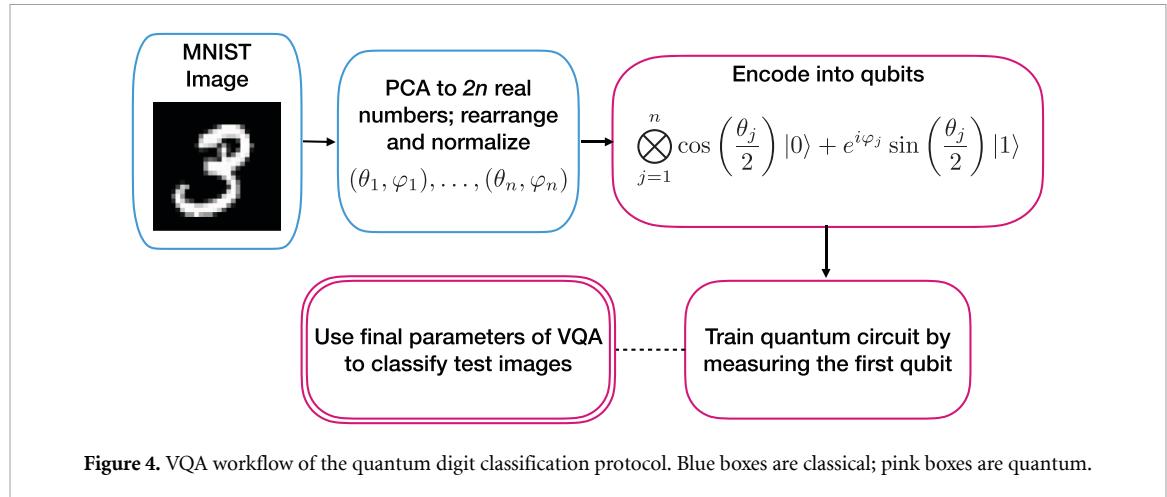


Figure 4. VQA workflow of the quantum digit classification protocol. Blue boxes are classical; pink boxes are quantum.

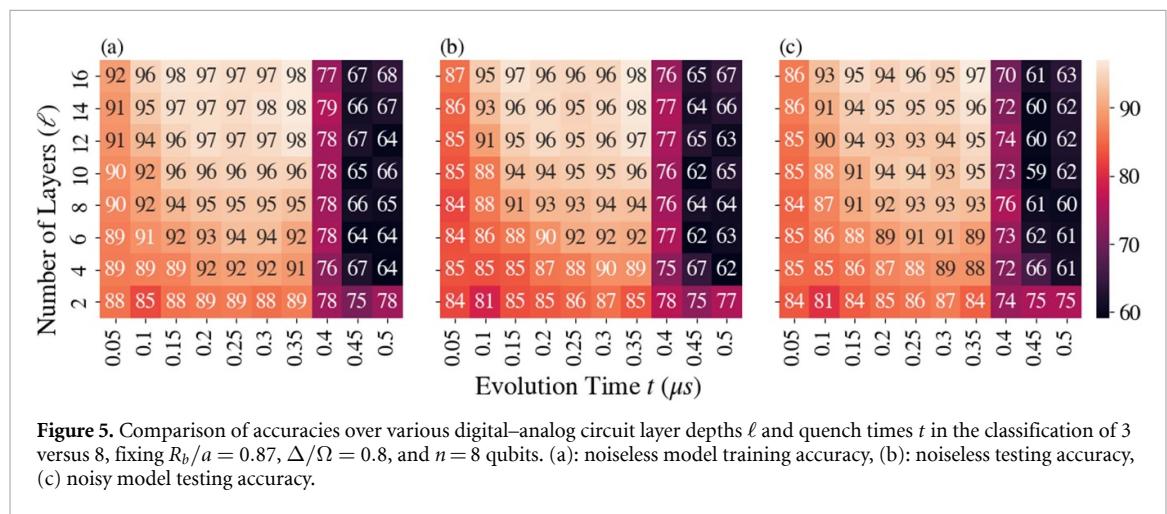


Figure 5. Comparison of accuracies over various digital-analog circuit layer depths  $\ell$  and quench times  $t$  in the classification of 3 versus 8, fixing  $R_b/a = 0.87$ ,  $\Delta/\Omega = 0.8$ , and  $n = 8$  qubits. (a): noiseless model training accuracy, (b): noiseless testing accuracy, (c) noisy model testing accuracy.

form for quantum circuitry, we transform the data via principle component analysis (PCA) into a set of  $2n$  numbers, where  $n$  is the number of qubits. Our workflow is outlined in figure 4. The PCA output values are normalized so that they represent rotation angles, which are then encoded into product states of the form

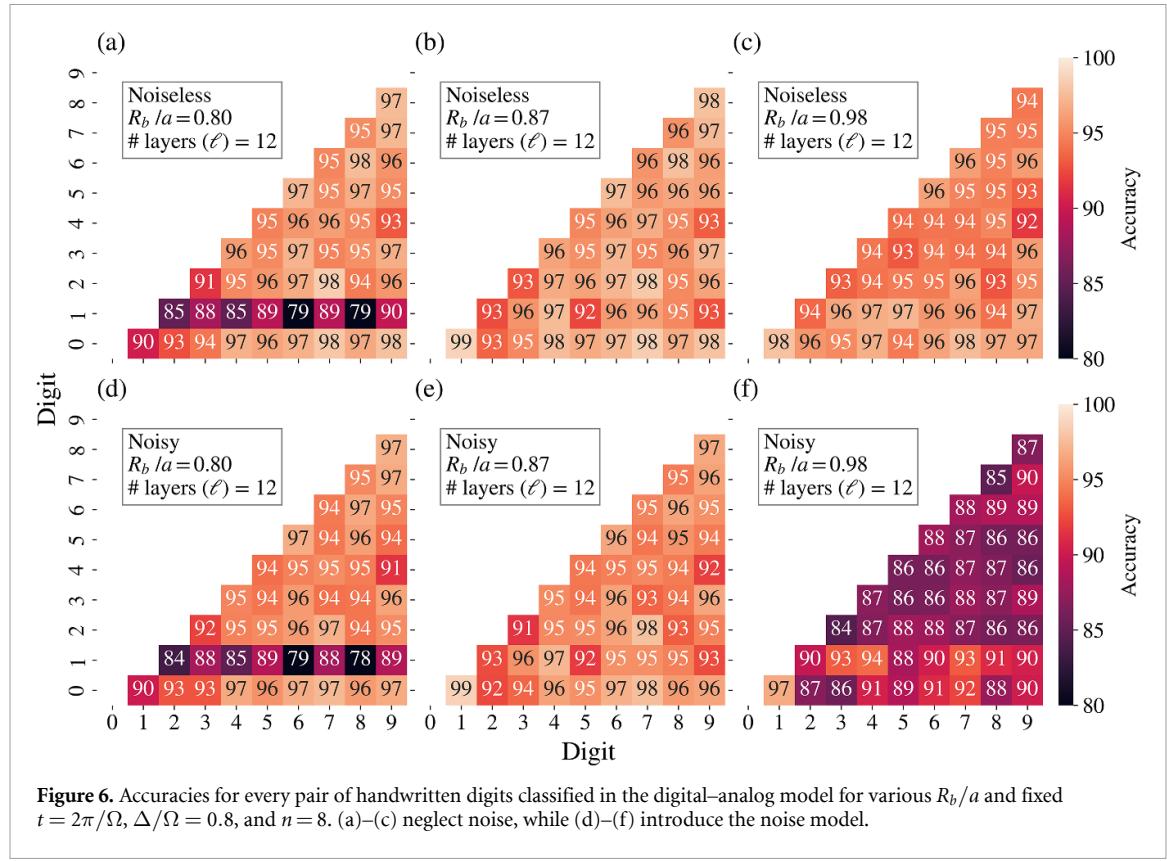
$$|\psi_{\text{in}}\rangle = \bigotimes_{j=1}^n \cos\left(\frac{\theta_j}{2}\right) |0\rangle + e^{i\varphi_j} \sin\left(\frac{\theta_j}{2}\right) |1\rangle. \quad (8)$$

Note that the principle components are ordered by importance, that is, the variance of the dataset projected along the principle component's corresponding eigenvector. We assign the first  $n$  components to the  $\theta$  angles and the last  $n$  to the  $\varphi$  angles. We then train the variational circuit by gradient descent methods on the measurement output of the first qubit, using a local cross-entropy loss function

$$\mathcal{L}(\mathbf{y}, \mathbf{q}) = -\frac{1}{m} \sum_{i=1}^m y_i \log q(y_i), \quad (9)$$

where  $y_i \in \{0, 1\}$  labels the  $i$ th digit in the training batch of size  $m$  and  $q(y) \in [0, 1]$  is the probability that the first qubit in the circuit output is 1, given  $y$  as input. To estimate  $\mathcal{L}$ , we measure the first qubit repeatedly for each fixed  $y_i$ , thereby obtaining an empirical estimate of  $q(y_i)$ . To output a classification among the two possible digits  $a$  and  $b$ , we say that the circuit classifies the digit as  $b$  if  $q(y) \geq 0.5$  and  $a$  otherwise.

We begin by verifying our physical arguments in the previous section about the optimality of the quench time  $t$ . As an illustrative example, consider a classification between 3 and 8, intuitively a particularly difficult instance because one number is—in writing—approximately contained in the other. Figure 5 depicts the landscape of classification accuracies over  $t$  as well as the layer parameter  $\ell$ , fixing  $R_b/a$  and  $\Delta/\Omega$ . We observe that indeed near  $t = 2\pi/\Omega = 0.25 \mu\text{s}$ , the accuracy is optimal for all  $\ell$  in the absence of noise both on the training dataset shown in figure 5(a) and the test dataset shown in figure 5(b). This optimality is nearly preserved up to 1% differences in the accuracy on the test dataset in the presence of noise, as shown in



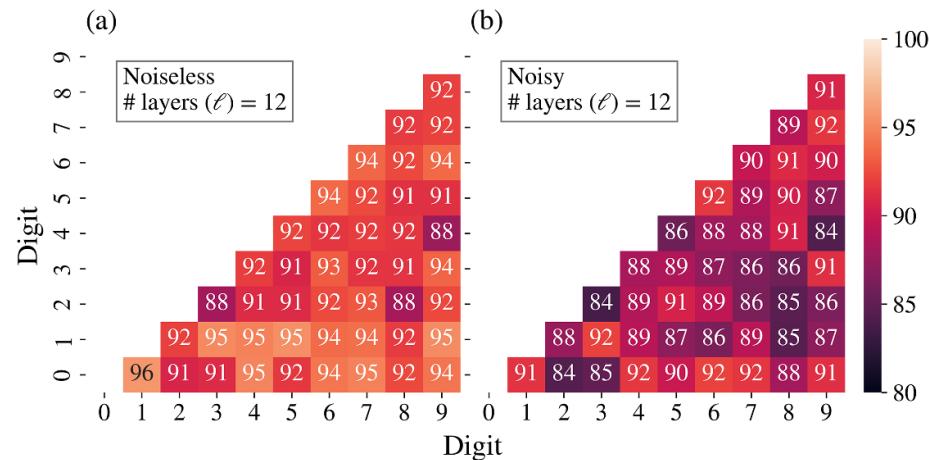
**Figure 6.** Accuracies for every pair of handwritten digits classified in the digital-analog model for various  $R_b/a$  and fixed  $t = 2\pi/\Omega$ ,  $\Delta/\Omega = 0.8$ , and  $n = 8$ . (a)–(c) neglect noise, while (d)–(f) introduce the noise model.

figure 5(c). Up to the numbers of layers we considered, increasing  $\ell$  weakly improved accuracies, but at some point the introduction of additional parameters will serve only to slow training and overfit data rather than improve performance. The entire region of  $t = 0.15$ – $0.35\mu\text{s}$  provides roughly equivalent performance. We proceed by fixing  $t = 2\pi/\Omega$  and  $\ell = 12$  for DA analysis. This optimality is not unique to 3 and 8; for example, we show an additional comparison between 1 and 9 in figure 18 in the appendix.

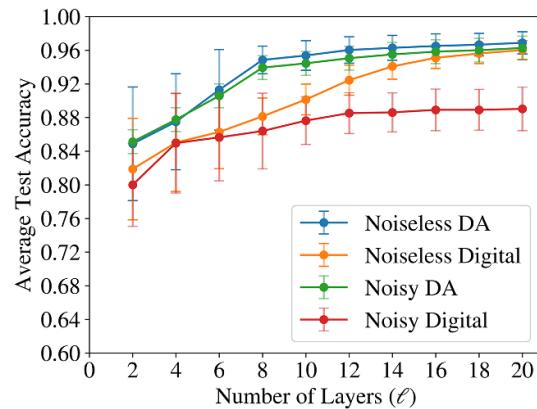
Next, we verify the approximate optimality of  $R_b/a = 0.87$ . Figure 6 shows choices of  $R_b/a$  in the three regimes discussed in the previous section, 0.8, 0.87, and 0.98, for both noiseless (a–c) and noisy (d–f) instances. While  $R_b/a = 0.8$  results in little ( $\sim 1\%$ ) loss of accuracy due to noise, the performance is quite low even without noise. On the other side of the spectrum,  $R_b/a = 0.98$  fares very well in the absence of noise, but suffers heavy losses in accuracy when the noise model is introduced. The choice of  $R_b/a = 0.87$  provides the predicted tradeoff between these effects, with high accuracy that is not lost in the presence of noise. While not shown here, there is a neighborhood of  $R_b/a = 0.87$  that provides similar accuracies and hence small deviations around  $R_b/a = 0.87$  produce similar results. The importance of this optimality is, rather, that it exists in an intermediate area between  $R_b/a = 0.8$  and  $R_b/a = 0.98$ , but in general the choice of  $R_b/a = 0.87$  suffices across different problems, such as the phase boundary problem discussed later.

Fixing the physically optimal DA gate hyperparameters now, we proceed to compare the performance of the DA model to that of the digital model. Figure 7 exhibits the accuracies across every pair of digits under the digital model, both with and without noise. Noiselessly, we found (see appendix for an example on 1 and 9) that the digital model performs better with increasing depths up to at least  $\ell = 20$ , but does not improve substantially beyond  $\ell = 12$ . As we show later in figure 8, in the presence of noise, the digital model strictly performs best at  $\ell = 12$  due to the noise overpowering the benefit of additional layers beyond 12. We therefore fix  $\ell = 12$  for both cases, which place them on equally optimal footing with the DA analysis. Juxtaposing figures 6(b) and 7(a) demonstrates that the DA and digital-models in their optimal hyperparameter regimes produce similar accuracies in the absence of noise. However, juxtaposing figures 6(e) and 7(b) shows that in the presence of noise, the digital model suffers considerably in accuracy, while the DA model has no significant loss.

This robustness to noise persists across a large range of  $\ell$ . As shown in figure 8, the DA model maintains a consistent advantage in noise robustness over the digital model for  $\ell$  up to 20. Moreover, the robustness to noise of the digital model begins to accelerate negatively after  $\ell = 12$ , due to the effect of the noise on  $\Theta(\ell n)$  gates overpowering the expressiveness of additional layers. Meanwhile, the DA model improves, albeit slightly, with increasing  $\ell$ , suggesting that it requires a substantially larger  $\ell$  for noise to overpower learning.



**Figure 7.** Accuracies for every pair of handwritten digits classified with the digital model, both (a) without and (b) with noise. The number of layers  $\ell = 12$  and qubits  $n = 8$  are fixed for both models.



**Figure 8.** Comparison of classification accuracy, averaged over all pairs of digits, between  $R_b/a = 0.87$ ,  $t = 2\pi/\Omega$ ,  $\Delta/\Omega = 0.8$ , and  $n = 8$  digital-analog (DA) and  $n = 8$  digital models with their respective noise models. Standard deviation, shown in the error bars, is measured across digit pairs.

This increased depth capability implies that DA models can learn tasks that require more expressivity (i.e. layers) due to their complex nature, extending the near-term utility of VQA's.

In summary, with optimal hyperparameters, the DA system demonstrated a substantial improvement in noise robustness over the digital model. In turn, it allows for improved accuracy in the near term and potential for learning tasks that require larger numbers of layers. A simple explanation for this improvement is the gate fidelity itself. Even for  $n = 8$  qubits, the fidelity of a digital layer is  $\sim 0.918$ , while the DA layer has a fidelity of  $\sim 0.971$ . Importantly, the gate fidelity of the DA layer appears to decrease substantially slower in  $n$  than that of the digital layer. Further, the DA model seems to require fewer layers than digital models to achieve similar performance, though the difference in this classification problem is not significant enough to definitively come to such a conclusion.

### 3.2. Boundaries of quantum phases

We proceed to study a DA learning scheme on quantum data. In this scheme, an unsupervised DA VQA classifies ground states of a chosen Hamiltonian according to their quantum phases. Quantum phases are of fundamental importance in many-body quantum systems. They are of particular relevance in Rydberg atom arrays, in which a plethora of exotic quantum phases have been predicted and discovered [6, 8–10, 66, 67]. To classify phases, we apply an enhanced and appropriately quantized technique common in classical machine learning algorithms known as anomaly detection. In the classical machine learning setting, anomaly detection is commonly used to identify statistical outliers in datasets, with a broad range of applications such as detecting fraudulent bank activity [68]. Recently, Kottmann *et al* [69] first considered a quantized version of anomaly detection for the learning of quantum phases. We consider an analogous version built on a Rydberg DA VQA.

Classically, anomaly detection circuits (learning agents) are trained to recognize a specific region of the data space that we consider to be ‘normal’. This training might be done, for example, by optimizing a neural network to output 0 on training data sampled from the normal subspace. Effectively, the training serves to overfit a learning agent to a desired region in space. In the testing phase, the learning agent recognizes ‘anomalous’ data (those outside the normal subspace) by outputting values that are far from zero, thus performing a classification of normality versus anomaly [70]. If multiple types of anomalies are close in the appropriate metric on the data space, then an anomaly detector can directly distinguish between the anomalies by labeling each as a cluster in data space. In this sense, anomaly detection is a clustering algorithm trained on a single cluster.

Such ideas can be extended to learn quantum information by the following procedure. First, replace the neural network with a Rydberg DA circuit  $\mathcal{A}_d^n$ . Next, to ensure the circuit outputs zero on training data, we use the average Rydberg density as the training loss function, given by

$$\begin{aligned}\mathcal{L}[\mathcal{A}(\{\theta_{ij}\})] &= \frac{\langle 0^{\otimes n} | \mathcal{A}^\dagger(\{\theta_{ij}\}) \mathcal{P}_1 \mathcal{A}(\{\theta_{ij}\}) | 0^{\otimes n} \rangle}{n} \\ &= \frac{1}{n} \sum_{k=1}^n \|\langle 1_k | \psi_{\text{out}} \rangle\|^2\end{aligned}\quad (10)$$

where  $|\psi_{\text{out}}\rangle = \mathcal{A}(\{\theta_{ij}\})|0^{\otimes n}\rangle$  is the circuit output and  $\mathcal{P}_1 = \sum_{k=1}^n |1_k\rangle\langle 1_k|$  is the total one-qubit projector onto the Rydberg state. Minimizing equation (10) on the training data ensures that the DA circuit outputs a number very close to 0 in the normal subspace. Moreover, experimental estimation of the average Rydberg density can be done in a number of computational basis measurements independent of  $n$ , rendering equation (10) a scalable loss function.

Quantum anomaly detection can be directly extended to quantum phase boundary learning by choosing one phase arbitrarily as the ‘normal’ phase and letting the rest be anomalies of a different type. Naively, one might expect that each phase would have its own distinct cluster and thus have its own range of circuit outputs. A subtlety, however, is that the circuit may not be sufficiently expressive to have a different value for each phase. Nonetheless, it may still take a dramatically different value when *crossing the boundary* between two phases. As a result, anomaly detection generically can only hope to learn the boundaries of phases rather than quantitatively cluster them. In practice, however, one can visually cluster the phases by qualitatively examining the learned phase diagram.

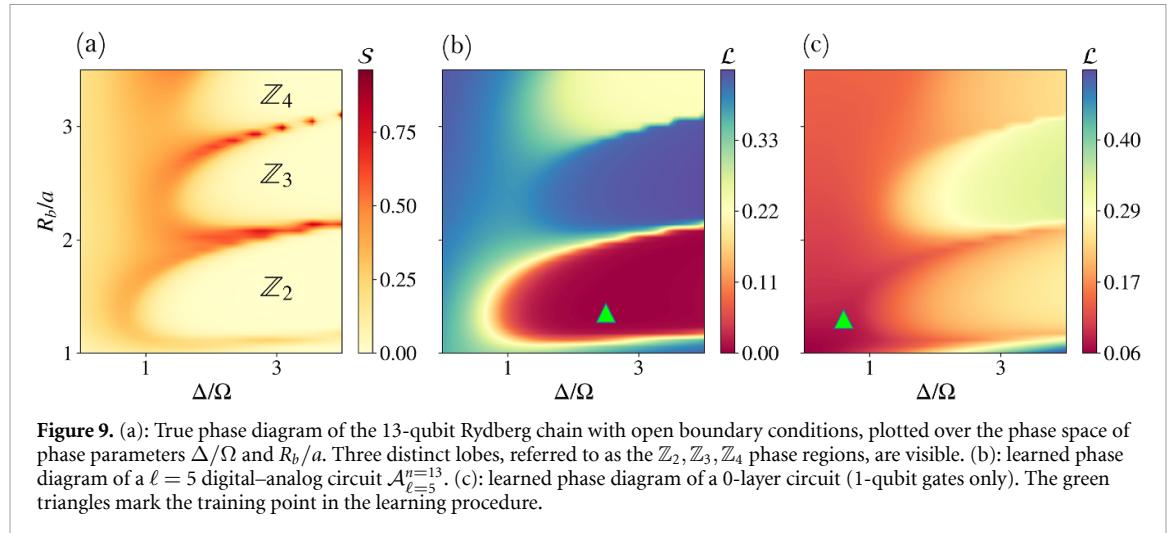
This picture is still slightly oversimplified because the second-order phase transitions are not perfectly discontinuous jumps but rather continuous transitions with a large gradient (with respect to the Hamiltonian parameters) in the corresponding order parameter. This effect serves to smoothen out the outputs of the DA learning circuit, in particular for finite systems, but does not hinder its efficacy at learning phases; rather, the gradient landscape itself indicates the location of the boundaries.

The quantum data we associate with each point in phase space is the ground state of a chosen Hamiltonian  $\mathcal{H}_{\text{learn}}(\phi)$  (generally not the Rydberg Hamiltonian  $\mathcal{H}$ ) with Hamiltonian parameters  $\phi$ . We demonstrate that it suffices to carve out the entire phase diagram with a single ground state of a chosen phase. This training dataset size of 1 is particularly appealing for near-term realization because training with large datasets becomes very resource-intensive especially on noisy devices.

We remark that DA circuitry enables an extended picture of phase clustering—in which the quantum data and the learning are performed consecutively in one complete DA circuit—by chaining the DA learning circuit with a step of adiabatic evolution. Adiabatic evolution enables the preparation of approximate ground states at a given phase starting from a ground state of an easy-to-prepare phase (e.g.  $|0\rangle^{\otimes n}$ ) by time-dependent quench dynamics

$$|g(\phi)\rangle = e^{-\frac{i}{\hbar} \int_0^1 ds \mathcal{H}_{\text{learn}}[s]} |0\rangle^{\otimes n} \quad (11)$$

where  $s$  parameterizes a smooth path in phase space—in which every point has a gapped Hamiltonian—such that  $\mathcal{H}_{\text{learn}}[s=0]$  is the Hamiltonian in the easy-to-prepare phase and  $\mathcal{H}_{\text{learn}}[s=1] = \mathcal{H}_{\text{learn}}(\phi)$ . Data preparation can, therefore, be incorporated into the DA circuit as a preliminary analog step before the first 1-qubit layer in figure 2(a). Note, however, that this analog step is driven according to a Hamiltonian  $\mathcal{H}_{\text{learn}}$  that is generally different the Rydberg Hamiltonian,  $\mathcal{H}$ . Adiabatic evolution has been explored extensively not only in Rydberg systems [10], but also much more generally as a technique for encoding computations into ground states of local Hamiltonians [71–74]. Adiabatic evolution produces approximate ground states that lead to a deformed version of the phase diagram that becomes the ground truth for learning [18, 19, 75]. This unified picture of data preparation and learning further demonstrates the utility of the DA paradigm.



Although adiabatic evolution is necessary for experimental ground state preparation, the manner of ground state preparation is not significant for purposes of algorithm analysis. Consequently, we prepare learning data directly via exact diagonalization in numerical simulation.

The most obvious learning example would be to set  $\mathcal{H}_{\text{learn}} = \mathcal{H}$ , the Rydberg Hamiltonian itself. We first map out a phase diagram classically by using the entanglement entropy

$$S[\rho_{AB}] = -\text{Tr}[\rho_A \log \rho_A], \quad (12)$$

where  $\rho_{AB}$  is the density matrix of the ground state of the Hamiltonian, partitioned into roughly equal-sized subsystems, and  $\rho_A$  and  $\rho_B$  are the reduced density matrices. Figure 9(a) displays the phase diagram on a 13-qubit chain, comprised of the well-known 3 lobes of ordered phases separated by high-entanglement boundaries, surrounded by the disordered phase [76, 77]. Each ordered phase  $\mathbb{Z}_n$  is close in fidelity to the state  $|10\cdots010\cdots01\cdots\rangle$  where each string of 0's is of length  $n - 1$ .

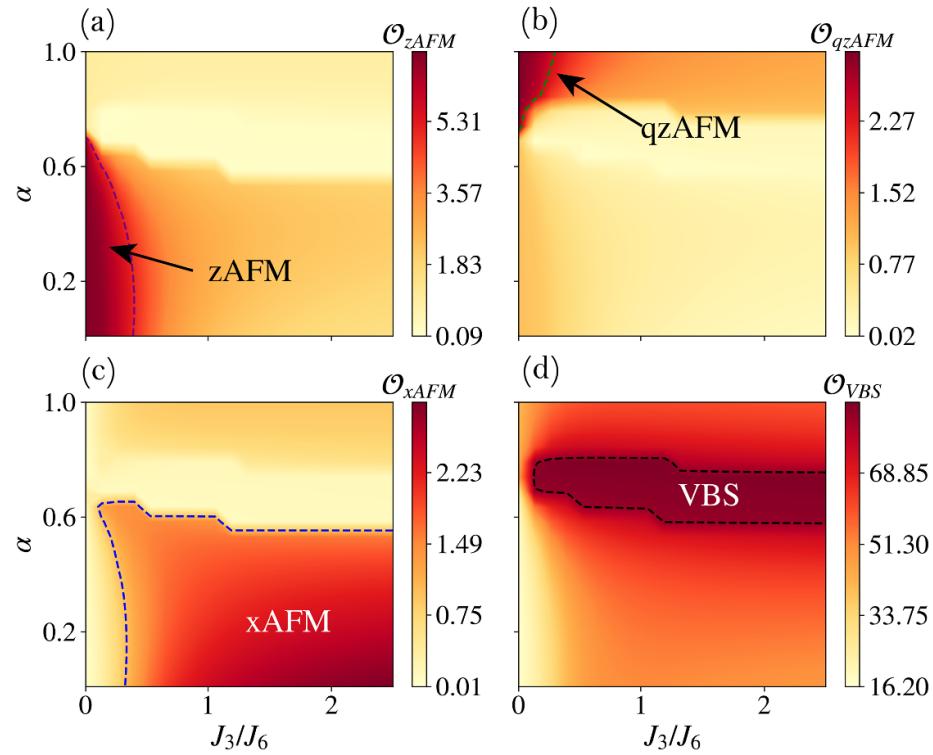
We apply the anomaly detection VQA on a DA circuit of depth  $\ell = 5$ , using the same quench hyperparameters as in the digit classification VQA. Figure 9(b) illustrates the output of the DA learning circuit trained on a point in the  $\mathbb{Z}_2$  lobe, demonstrating the learnability of the phase diagram. However, this Rydberg Hamiltonian, at least for a small system size, is not a good instance upon which we can analyze the DA VQA because all but the disordered phase are adiabatically connected to their respective computational-basis product states. Even the disordered phase has a relatively low entanglement entropy, being adiabatically connected to the  $|+\rangle^{\otimes N}$  state. As such, one might expect that no entanglement in the learning circuit is even necessary to reproduce the phase diagram. This intuition turns out to be correct, as shown in figure 9(c). Here, a  $\ell = 0$  circuit which has only a single rotation gate per qubit and no entanglement, successfully carves out the phase diagram of the Rydberg chain.

Instead, we keep the one-dimensional spin chain structure of the atoms but consider a 2-body XXZ Hamiltonian with periodic boundary conditions, governed by antiferromagnetic nearest-neighbor (NN) and next-nearest-neighbor (NNN) interactions [78, 79]

$$\begin{aligned} \mathcal{H}_{\text{learn}} = & J_3 \sum_r (X_r X_{r+1} + Y_r Y_{r+1}) \\ & + \alpha (X_r X_{r+2} + Y_r Y_{r+2}) \\ & + J_6 \sum_r [Z_r Z_{r+1} + \alpha^2 Z_r Z_{r+2}] \end{aligned} \quad (13)$$

where,  $X, Y, Z$  are the Pauli matrices,  $J_3$  and  $J_6$  give the strengths of spin-exchange and Ising interactions for NN pairs, respectively. The strength parameter  $\alpha$  is defined as  $(d_1/d_2)^3$ , where  $d_1$  is NN distance and  $d_2$  is NNN distance. This Hamiltonian cannot be learned by a non-entangling circuit, as shown in appendix C.4 figure 20.

Analysis of the XXZ Hamiltonian in the ground state has revealed the existence of four distinct phases, shown in figure 10 [78, 79]. When  $J_3 = 0$ , the system is strictly classical. In this region, two interactions compete: the NN Ising interactions, which induce antiferromagnetic (zAFM) order in the  $z$ -direction, and the NNN Ising interactions, which induce Quadrupled antiferromagnetic (qzAFM) order. The effect of each



**Figure 10.** Order parameters of the ground state for the 8-qubit XXZ Hamiltonian with periodic boundary conditions. (a)  $\mathcal{O}_{zAFM}$ , the zAFM order parameter (equation (15)). (b)  $\mathcal{O}_{qzAFM}$ , the qzAFM order parameter (equation (16)). (c)  $\mathcal{O}_{xAFM}$ , the xAFM order parameter (equation (17)), a proxy for the XY-QLRO phase region. (d)  $\mathcal{O}_{VBS}$ , the VBS order parameter (equation (20)). In each plot, the dashed lines represent the respective phase boundaries, defined for (a)–(d) respectively by contouring around the largest 18%, 3%, 50%, and 9% values.

interaction is determined by  $\alpha$ . The zAFM (figure 10(a)) and qzAFM (figure 10(b)) phases are identified by the two-point  $Z$  correlator  $C_z(r)$ , defined as

$$C_z(r) = \langle \psi | Z_0 Z_r | \psi \rangle \quad (14)$$

where  $r$  indexes the qubit,  $r=0$  represents the chosen ‘central’ site in the chain due to periodic boundary conditions (here,  $r=2$ ), and  $\psi$  is the ground state of the Hamiltonian.

The zAFM and qzAFM order parameters are defined from this correlator as

$$\mathcal{O}_{zAFM} = \sum_{r=1}^n e^{i\pi r} C_z(r), \quad (15)$$

$$\mathcal{O}_{qzAFM} = \sum_{r=1}^n e^{i\pi r/2} C_z(r), \quad (16)$$

where  $n$  the number of sites in the chain.

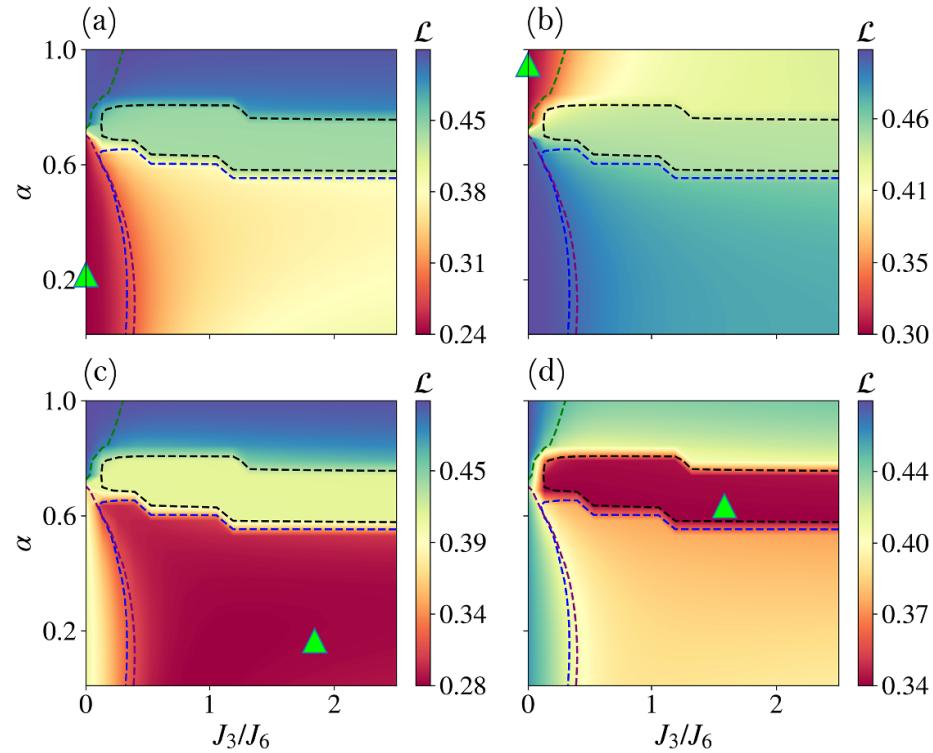
For  $\alpha \sim 0$ , the zAFM order dominates when  $J_3 < J_6$ , while a symmetric gapless XY phase with quasi-long-range order (XY-QLRO) prevails when  $J_3 > J_6$  (figure 10(c)). This quantum, highly entangled phase does not have a local order parameter, but can be qualitatively inferred by the xAFM order parameter that decays algebraically in the phase, and is given by

$$\mathcal{O}_{xAFM} = \sum_{r=1}^n e^{i\pi r} C_x(r) \quad (17)$$

where

$$C_x(r) = \langle \psi | X_0 X_r | \psi \rangle \quad (18)$$

is the two-point  $X$  correlator.



**Figure 11.** Four learned phase diagrams for the 8-qubit XXZ Hamiltonian with parameters  $\alpha$  and  $J_3/J_6$ , achieved on a digital-analog circuit  $\mathcal{A}_{\ell=2}^{n=8}$  without noise. The green triangles indicate the point of training. The dashed contours outlining each phase are the same as in figure 10, demonstrating remarkably that the digital-analog circuit reproduces all of the boundaries specified by the four order parameters at once for each training area.

The final phase region occurs roughly where  $\alpha \approx 0.7$  and  $J_3/J_6 \rightarrow \infty$ . This phase (figure 10(d)) is referred to as valence bond solid (VBS), which consists of dimerized patterns of spin singlets. The VBS phase is characterized by the two-point correlator

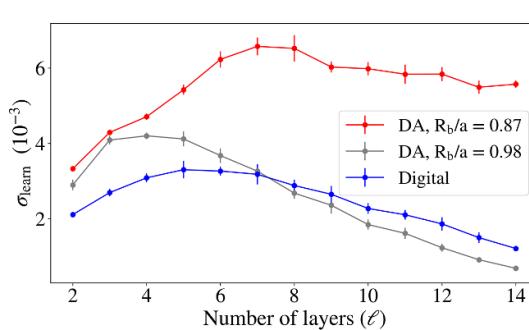
$$C_{\text{VBS}}(r) = \langle \psi | (\vec{\sigma}_1 \cdot \vec{\sigma}_0 - \vec{\sigma}_0 \cdot \vec{\sigma}_{-1}) \times (\vec{\sigma}_{r+1} \cdot \vec{\sigma}_r - \vec{\sigma}_r \cdot \vec{\sigma}_{r-1}) | \psi \rangle \quad (19)$$

where  $\vec{\sigma}_i = (X_i, Y_i, Z_i)$  is the Pauli spin matrix vector. Likewise, the VBS order parameter is defined as

$$\mathcal{O}_{\text{VBS}} = \sum_{r=1}^n e^{i\pi r} C_{\text{VBS}}(r). \quad (20)$$

These order parameters are by symmetry real. In contrast to Rydberg density wave phases, VBS and QLRO cannot be adiabatically connected to product states and are entangled. The same holds for zAFM and qzAFM in a perfect system with periodic boundaries. For example, under such conditions the ground state of the zAFM is a cat state  $\propto |0101\dots\rangle + |1010\dots\rangle$ .

Figures 11(a)–(d) depicts the output of DA anomaly detection on the 8-qubit XXZ chain, with four different training points chosen each in one phase region. Although the circuits were simulated noiselessly here, we observe already two noteworthy results. First, generically to DA learning, the successful demarcation of phases alone is a nontrivial result, because we used the same physically justified hyperparameters as in the noiseless digit classification algorithm. This result provides stronger evidence that the quench dynamics parameters can be optimally chosen without tailoring them to a specific problem, and support the validity of our physical justifications. It moreover provides an explicit demonstration of the versatility of the analog layer as a tool for entanglement generation. Second, regardless of which phase acted as the training point, each learning outcome produced demarcations of the entire phase diagram, which if relying on order parameters alone would have required examining multiple order parameters simultaneously to deduce. In this sense, anomaly detection may have additional utility for the study of Hamiltonians whose phases are not visible by a single global order parameter.



**Figure 12.** Comparison of  $\sigma_{\text{learn}}$  estimated numerically as a function of  $\ell$  on digital-analog  $\mathcal{A}_\ell^{n=8}$  and digital  $\mathcal{D}_\ell^{n=8}$  circuits for the XXZ Hamiltonian. Each circuit is trained on a zAFM point. Error bars denote 1 standard deviation over randomness in the noise. The advantage in sharpness is visible for a good hyperparameter  $R_b/a = 0.87$ , but does not persist when the choice of hyperparameter  $R_b/a = 0.98$  is bad.

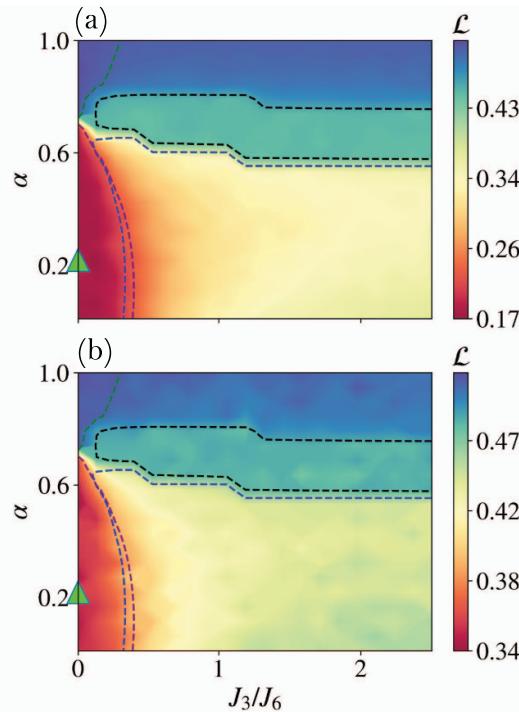
In the absence of noise, there does not appear to be a substantial advantage in the number of layers  $\ell$  required for successful learning between DA and digital circuitry. This lack of advantage is not evidence against a general advantage for DA circuitry, however, because phase boundaries of the XXZ Hamiltonian simply appear to be learnable with relatively low entanglement in the circuit (see figure 21 in appendix C.4). Ground states of local one-dimensional Hamiltonians generally have either constant entanglement entropy for gapped phases or  $\log N$  for gapless phases, which limits the number of layers needed to distinguish phases. It is likely that with even more complex Hamiltonians, higher dimensions, or dynamical phase diagrams with non-area law states, a layer advantage may emerge.

We next apply the noise models discussed in section 2 for purposes of a DA and digital circuit comparison. To obtain a quantitative metric of performance akin to accuracy in the digit classification problem, we define the sharpness of a learned phase diagram as

$$\sigma_{\text{learn}} (\{\theta_{ij}\}) = \text{SD}_\phi (\|\nabla_\phi \mathcal{L} [\mathcal{A}_\ell^n (\{\theta_{ij}\})]\|^2) \quad (21)$$

where  $\text{SD}_\phi$  is the standard deviation over the uniform distribution of the two-dimensional parameter space ( $J_3/J_6, \alpha$ ). Intuitively, the sharpness of a learned phase diagram measures how distinctly the boundaries can be identified. A well-trained circuit that outputs nearly gradient-free interiors and rapid-transitioning boundaries produces very large  $\sigma_{\text{learn}}$ . By contrast, a poorly trained circuit with less distinct, smoother transitions results in lower sharpness and less informative learning results. In practice,  $\sigma_{\text{learn}}$  is estimated by calculating numerical gradients and sample standard deviations over a discretized uniform mesh grid of the parameter space. We note that circuit noise impacts sharpness both by reducing the ability of the circuit to distinguish phases, thereby decreasing  $\sigma_{\text{learn}}$ , and introducing small local fluctuations that increase  $\sigma_{\text{learn}}$ . However, the latter effect is by definition substantially smaller than the first and can be thought of as negligible, so that  $\sigma_{\text{learn}}$  remains a well-defined measure of noise robustness in phase boundary learning. Moreover, noise-based effects on  $\sigma_{\text{learn}}$  are mitigated by averaging over multiple trials and, therefore, primarily affect the variance of  $\sigma_{\text{learn}}$ .

In figure 12, we calculate  $\sigma_{\text{learn}}$  numerically on learned phased diagrams averaged over 20 independently trained learning circuits, again with  $n = 8$  but with the training point fixed in the zAFM phase for consistency. We repeat this procedure five times to provide a simple error bound, shown in the error bars. We observe that when the hyperparameter choice is good ( $R_b/a = 0.87$ )—in terms of the analysis in section 2.2—the sharpness improves at larger depths for DA circuits, whereas the noise quickly reduces the sharpness of the digital circuits. We are primarily interested in the large- $\ell$  behavior as it indicates the noise robustness for learning problems that require greater circuit depths, e.g. 2D ground states or dynamical phases (see next section). Consequently, DA circuits appear also to be more robust than digital circuits in learning phase boundaries, in agreement with findings from digit classification. We observe, however, that this advantage is lost if the choice of hyperparameter with respect to noise is poor, as shown by the grey line ( $R_b/a = 0.98$ ). Hence, we can also explicitly see in phase boundary learning that the choice of hyperparameter makes a substantial impact on the learning advantage of DA circuits. We observe visual differences in sharpness at  $\ell = 14$  between DA and digital circuitry in figure 13, namely blurrier phase boundaries and more artifacts in



**Figure 13.** Learned phase diagrams of the XXZ Hamiltonian with respect to parameters  $\alpha$  and  $J_3/J_6$  for (a) digital-analog and (b) digital circuits with  $n = 8, \ell = 14$  in the presence of their respective noise models. Dotted contours overlay the true phase boundaries.

the digital case. Notably, the performance of the digital circuit with  $\phi = \pi/4$  is *qualitatively* worse. In figure 17 of appendix B, we show that a digital  $\phi = \pi/4$  circuit is unable to identify the VBS phase at all.

#### 4. Conclusion & outlook

We have observed numerically that both in digit classification and phase boundary learning, DA circuits achieve substantial practical advantages over their digital counterparts as the quantum agents in the VQA. In both cases, the DA circuit demonstrated a substantially greater robustness to noise. The comparatively larger gate fidelity of the DA entangling layer provided a simple explanation of this improvement. A subtlety of the error advantage was that it depended on the choice of  $R_b/a$ , based on the physical tradeoff of large separations ( $R_b/a \ll 1$ ) leading to relatively smaller position noise effects and small separations ( $R_b/a \sim 1$ ) leading to greater interaction strength. Properly trading off these effects by choice of  $R_b/a = 0.87$  enabled the noise advantage. Similarly, choices of  $t = 2\pi/\Omega$  led to optimal performance of the DA circuits. While seemingly complicated, the hyperparameter choices have physical justifications which allow them to be more easily understood, providing an advantage over digital gates which do not have any such parameters that can be tuned to physical optimality. Moreover, our work provides evidence that these choices are indeed roughly independent of the learning problem, since the same choices led to comparatively advantageous performance of DA circuits in both problems we investigated. This universality is expected, since the parameters are inherently related to the physics of the circuit and not the problem it learns. Consequently, no re-derivation of these hyperparameters is necessary from one learning problem to the next.

We showed also that for a problem in which circuits of a few layers  $\ell$  do not perform well on, DA circuitry enables learning to comparable accuracies in substantially smaller  $\ell$ . We reiterate that the fact that DA circuits can learn at all is itself notable. Practically, even without the DA advantages we have shown, it is much easier to realize analog learning because it requires much fewer local control on qubits and acts as one global operation rather than  $\Theta(n)$  local, non-commuting operations. Theoretically, our results suggest that certain quench times and Hamiltonian parameters are admissible to learning while the rest are not, raising questions that we discuss further below about the ranges of good hyperparameters.

A useful future direction would involve putting the physical justifications in this work on more rigorous theoretical grounds. Specifically, it is an interesting question as to the general robustness of the optimal

evolution time  $t = 2\pi/\Omega$ . We justified this choice by arguing that (a) sufficiently small times would not offer enough entanglement, (b) sufficiently large times are known to generically lead to thermalization, and (c)  $2\pi/\Omega$  is the natural time scale of the system. Yet this justification accounts only for the small- and large-time limits, and does not enable us to understand what range of intermediate times are generally useful. A similar optimality question concerns an analytical solution and more rigorous physical understanding of the tradeoff effect between interaction strength and gate fidelity when tuning  $R_b/a$ . Similarly, with either greater computational resources or more efficient noise simulation algorithms, a related direction of extension consists of expanding the noise model to include incoherent sources.

Another potential direction involves deforming part of the phase boundary learning problem to examine its versatility. One possibility is to study adiabatically prepared states for their own utility. While we discussed the use of adiabatic evolution to generate approximate ground states that reproduce our phase diagrams, it is in general infeasible to do so in polynomial time because the preparation of ground states is computationally hard. A classic result in quantum complexity, for example, states that even determining whether the ground state energy of a local Hamiltonian is below a certain threshold is **QMA**-hard [80], where **QMA** is a quantum analog of the classical complexity class **NP**. Any adiabatically prepared ensemble of states produces quantum state classification task which we can learn, though it may not correspond to a phase diagram. It is an interesting question as to what physics we may learn from this class of quantum classification problems.

Another option for learning the ground-state phase diagram in particularly hard instances of the Hamiltonian may be to turn to quench dynamics themselves. Consider  $|0\rangle$  quenched for some intermediate time long enough to leave a ‘signature’ of the phase behind but not long enough for thermalization and quantum chaos to ensue. Whether or not such intermediate-time signatures even always exist is an open question, but recent studies on spin chains have demonstrated promising results [81]. Assuming the presence of intermediate-time quench phase signatures, an intriguing extension of quantum phase learning would be whether and how anomaly detectors could learn the signatures without supervision. Even without signatures, the use of intermediate-time quenched states may be useful for QML predictions of dynamical phase transitions, although this may require input from states quenched with many different quench times.

## Data availability statement

The data that support the findings of this study will be openly available following an embargo at the following URL/DOI: <https://github.com/jz-lu/daqml>. Data will be available from 30 December 2024.

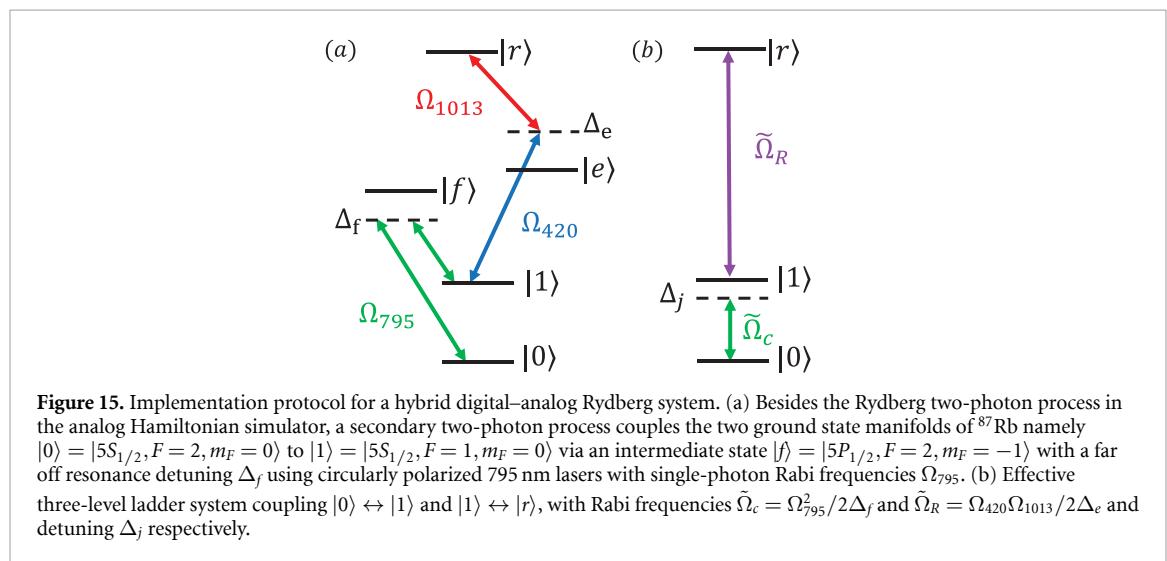
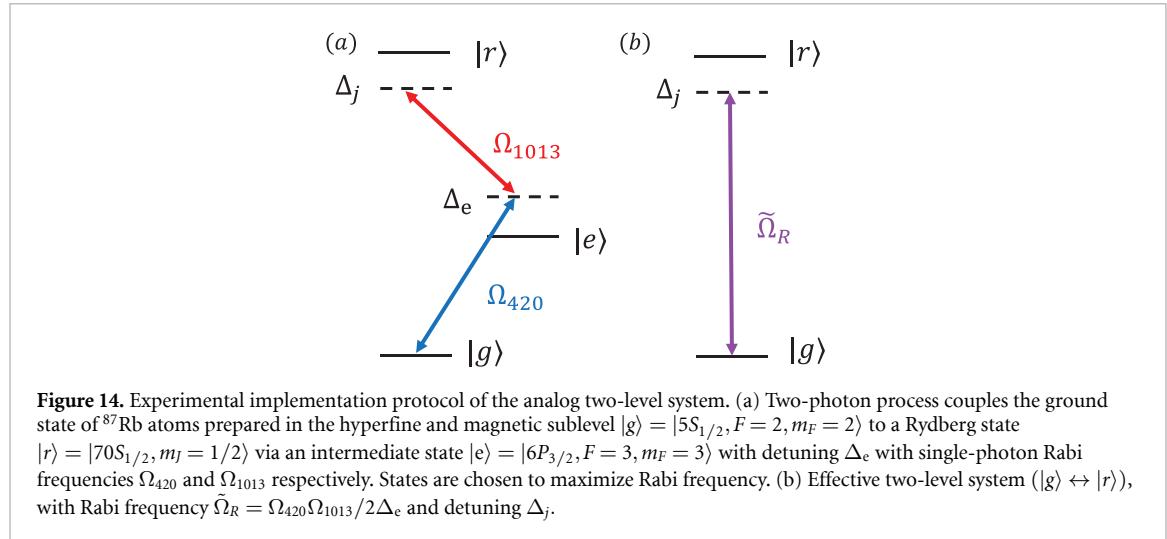
## Acknowledgments

The authors thank Jonathan Wurtz for an insightful discussion about practical error models on Rydberg systems. Authors affiliated with QuEra Computing Inc. acknowledge the support by the DARPA-STTR award (Award No. 140D0422C0035). S F Y thanks the NSF through the CUA PFC and the Qu-IDEAS HDRI for funding. H Y H is grateful for the support from the Harvard Quantum Initiative. The computations in this paper were run on the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

## Appendix A. Experimental considerations

In this appendix, we briefly discuss an experimental technique for the implementation of a DA system on Rydberg atom arrays. One such realization relies on two-photon coupling processes. A complete discussion of these experimental considerations can be found in Bluvstein *et al* [63].

Figures 14 and 15 highlight the key levels in the atomic structure in  $^{87}\text{Rb}$ . In the analog-only configuration, atoms are prepared in a specific ground state of the  $5S_{1/2}$  manifold defined by the hyperfine state and magnetic sublevel  $|g\rangle = |F = 2, m_F = 2\rangle$ . Atoms interact by coupling them to the  $70S_{1/2}$  Rydberg state of the magnetic sublevel  $|r\rangle = |m_J = 1/2\rangle$  via a two-photon scheme with 420 nm  $\sigma^+$ -polarized light and 1013 nm  $\sigma^-$ -polarized light. These states are chosen to maximize Rabi frequency at fixed laser power. In the digital implementation, qubit states are defined by the clock states in the ground state manifold  $|0\rangle = |F = 1, m_F = 0\rangle$  and  $|1\rangle = |F = 2, m_F = 0\rangle$  to minimize magnetic field sensitivity. Qubit rotations are performed by a two-photon Raman process using 795 nm  $\sigma^+$  light that couples state  $|0\rangle$  and  $|1\rangle$ . Interactions between qubits are enabled by a two-photon transition that couples state  $|1\rangle$  to the Rydberg state  $|r\rangle$ . Therefore, it is possible to program a series of pulses to enable DA control over the qubits.



## Appendix B. The generalized CNOT gate

In the main text, we described a controlled-NOT gate by a  $\pi/4$ -time evolution under a two-qubit Hamiltonian  $e^{-i\frac{\pi}{4}(I_1-Z_1)(I_2-X_2)}$ . We here derive this expression.  $(I_1 - Z_1)(I_2 - X_2) = I - Z_1 - X_2 + Z_1X_2$ . Since all these terms commute, we may factor the exponential as

$$e^{-i\frac{\pi}{4}(I_1-Z_1)(I_2-X_2)} = e^{-i\frac{\pi}{4}} e^{i\frac{\pi}{4}Z_1} e^{i\frac{\pi}{4}X_2} e^{-i\frac{\pi}{4}Z_1X_2}. \quad (\text{B1})$$

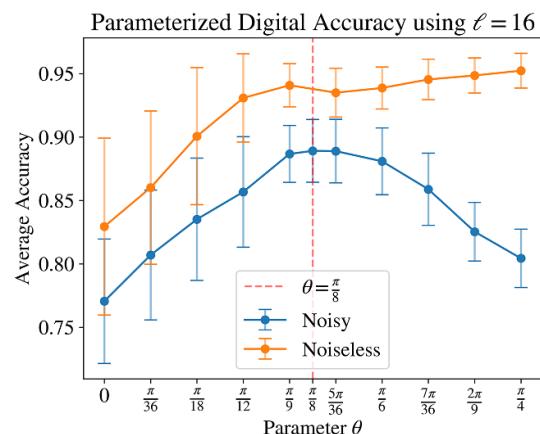
Note that if  $A^2 = I$ , then  $e^{i\theta A} = \cos(\theta)A + i \sin(\theta)A$ , and thus the above resolves to

$$e^{-i\frac{\pi}{4}} \frac{1}{\sqrt{2}} (I + iZ_1) \frac{1}{\sqrt{2}} (I + iX_2) \frac{1}{\sqrt{2}} (I + iZ_1X_2). \quad (\text{B2})$$

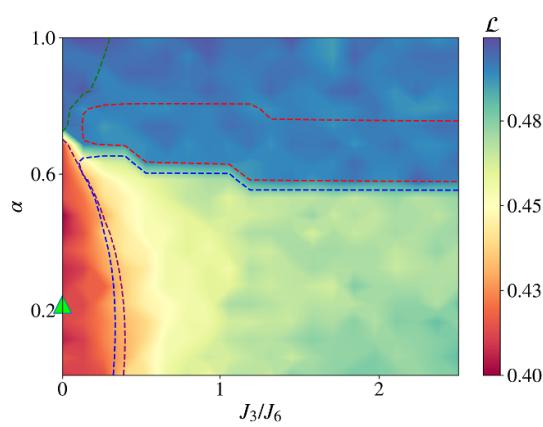
We can readily check that the above matrix is the CNOT gate

$$\text{CX}_{1 \rightarrow 2} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (\text{B3})$$

The generalized CNOT gate is given by  $\text{CX}_{1 \rightarrow 2}(\phi) = e^{-i\phi(I_1-Z_1)(I_2-X_2)}$ . Using the methods in the main text of our paper, we vary  $\phi \in [0, \pi/4]$  and compute the accuracy in classifying digits. Figure 16 demonstrates that the optimal performance occurs approximately halfway in the range, giving an optimal choice of  $\phi = \pi/8$ . We observe that the performance difference relative to the simple choice of  $\pi/4$  is substantial, with



**Figure 16.** Average accuracies over digit comparisons of the parameterized digital circuits on the MNIST dataset with varying the parameter  $\theta$ , while keeping the number of layers constant at  $\ell = 16$ . The vertical dashed line denotes the optimal point for the noisy circuit at  $\theta = \pi/8$ .



**Figure 17.** Learned phase diagram of the XXZ Hamiltonian with respect to parameters  $\alpha$  and  $J_3/J_6$  for a digital ( $\phi = \pi/4$ ) circuit with  $n = 8, \ell = 14$  in the presence of noise. Dotted contours overlay the true phase boundaries. Notably, the VBS phase (red contour) is not learned at all.

an accuracy gap of 8.5%. This gap holds true even more significantly for phase boundary learning. Figure 17 demonstrates as an example, a stark difference in the learned phase diagrams of the XXZ Hamiltonian between  $\pi/4$  and  $\pi/8$  (shown in figure 13 of the main text). In the  $\pi/4$  case, the VBS phase is not identified at all. These differences highlight that a choice of  $\pi/4$  over  $\pi/8$  causes significant, even qualitative, losses in QML performance. This gap strengthens the separation between digital and DA schemes, as in some devices the choice of  $\pi/8$  might need to be built out of more native gates, which could amplify noise and worsen performance. Yet if  $\pi/4$  is chosen, the performance will worsen anyway. By contrast, the DA scheme avoids these implementation subtleties due to its native implementation on neutral atom devices.

## Appendix C. Further discussion

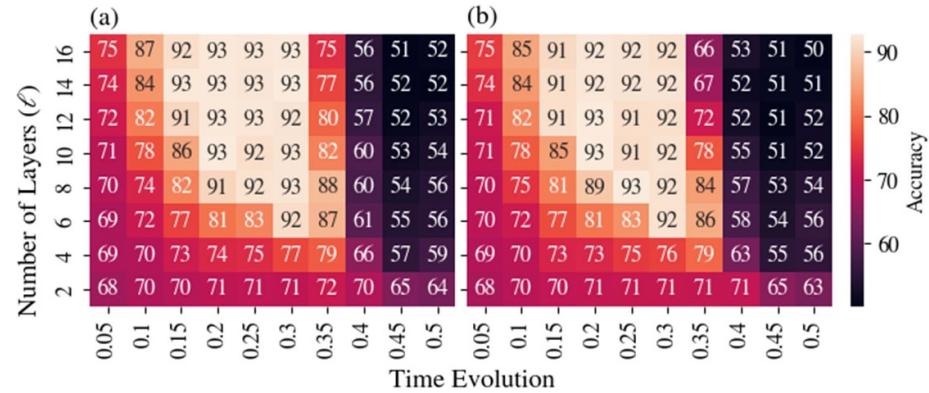
### C.1. Binary classification

#### C.1.1. Additional hyperparameter search

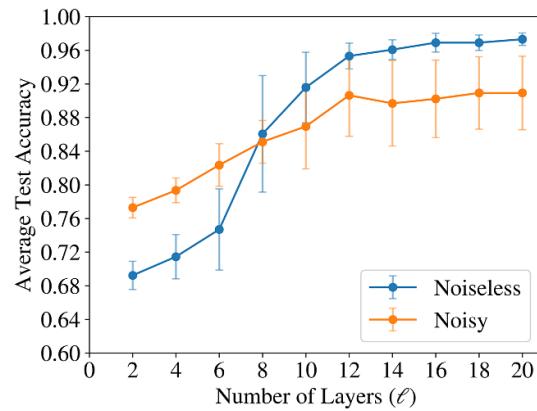
The choice for the time evolution constant may seem suboptimal when looking solely at the 3 versus 8 digit comparison, but bringing our attention to the 1 versus 9 digit comparison, we can see that the optimal region is much more constrained.

In figure 18, we see that the accuracy drops after 0.3 for the time evolution constant. This indicates we should pick our time evolution parameter to be more deeply in the optimal region, especially for robustness to noise.

Additionally, we see in figure 19 that in the noiseless 1 versus 9 comparison case, we do not see much improvement in the test accuracy after 16 layers, and coincidentally see that this is also optimal for the noisy



**Figure 18.** Hyperparameter grid search of the time evolution constant and the number of layers in the circuit on the comparison between 1 versus 9 for (a) the noiseless model and (b) the noisy model. All models are using  $R_b/a = 0.87$  and  $n = 8$ .



**Figure 19.** Comparison of the average test accuracy when using a digital circuit for the digit comparison of 1 versus 9.

model. To note the general case across digit comparisons, in figure 8, we see accuracy plateauing around  $\ell = 16$  as well. However,  $\ell = 16$  does not seem to be optimal for the noisy case in general as seen in figure 8.

### C.2. Computational details

All instances of learning used the AdaGrad optimizer from Flux with a learning rate of 0.1 for the classical optimization. During each training epoch, a randomized batch was used for training and to calculate the gradient. There are 70 total training epochs where one randomized sample was taken each time. Then for average accuracies taken in figure 5, 50 trials were taken over random starting parameters.

Also, when making any measurements during the training or test phase, a new noisy circuit is generated such that the noise is independent for any measurements. No data from the test set is used to inform the training process during gradient descent.

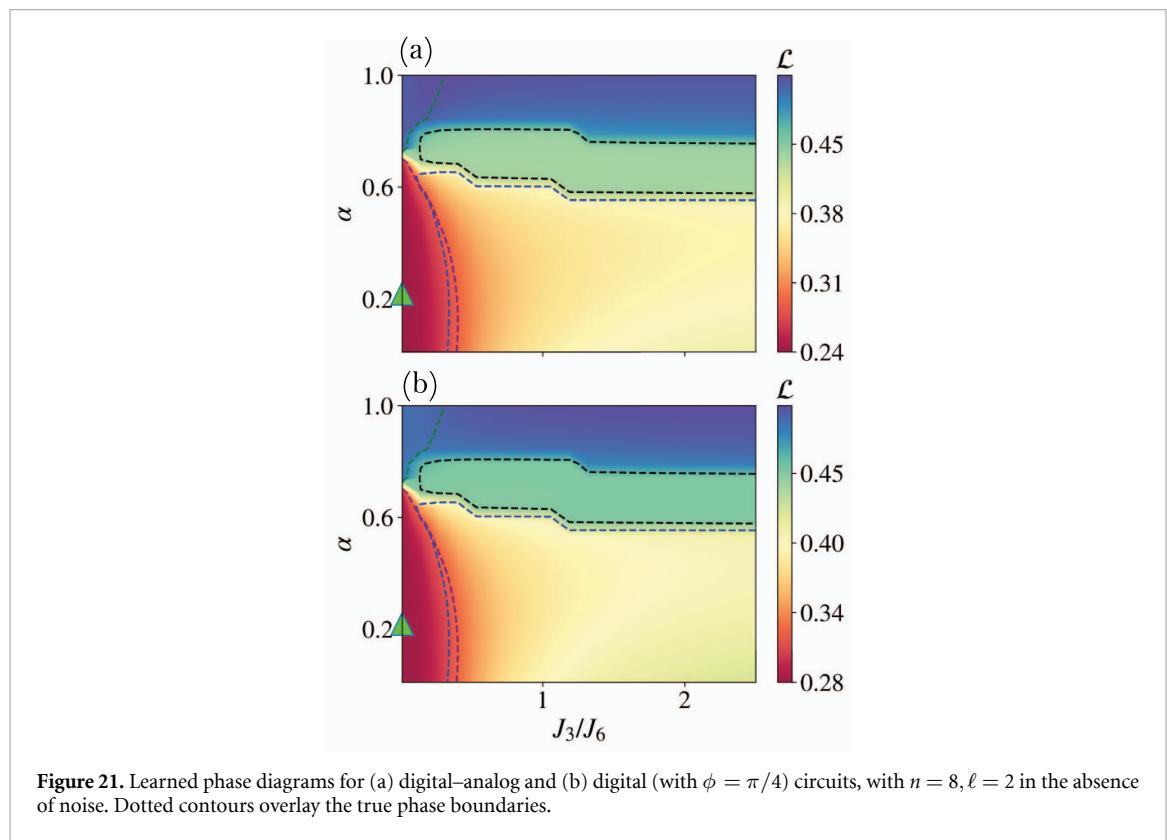
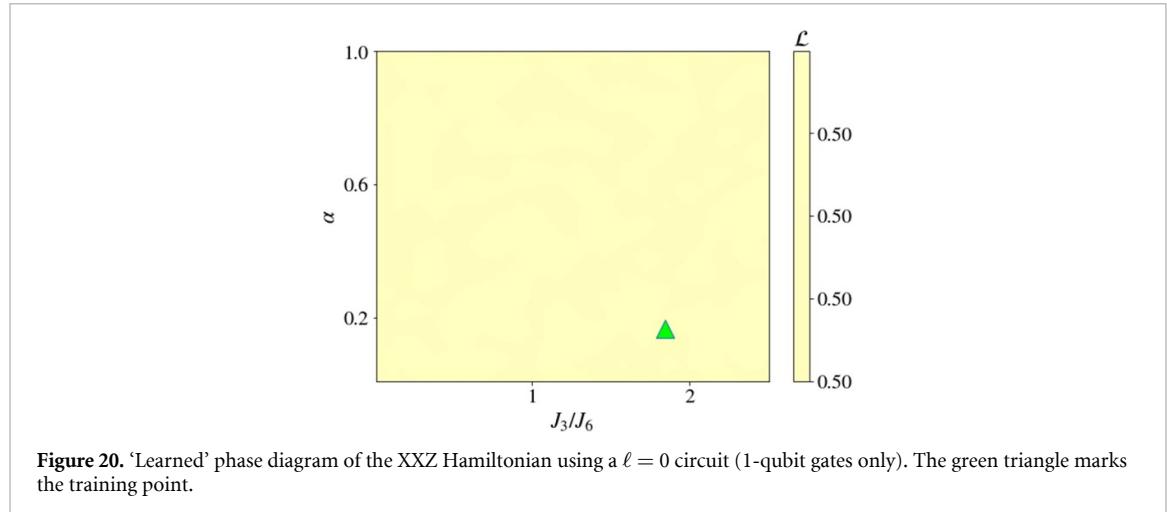
The details of numerical convergence in training are deferred to figure 22(a) the final section of the appendix.

### C.3. Phase detection

#### C.3.1. Phase diagram discussion

Two supplementary plots offer additional context for the phase learning previously described. First, the phase diagram of the XXZ Hamiltonian cannot be learned using a circuit with no entangling resource. Figure 20 depicts the result when such learning is attempted: uniform loss of 1/2 throughout the entire phase diagram. This is in contrast to the Rydberg phase diagram, which could be learned using such a circuit, as seen in figure 9(c).

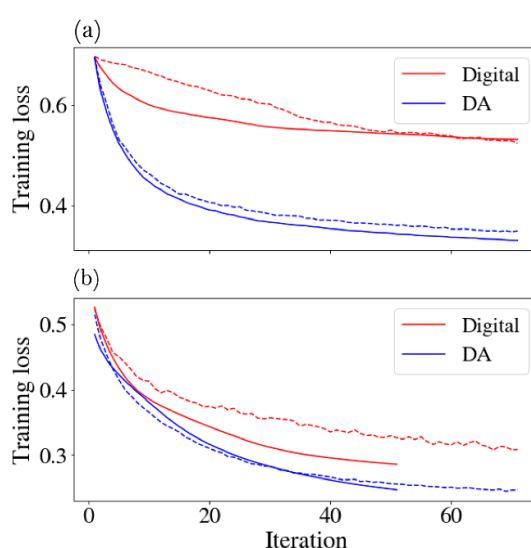
In addition, in the noiseless scheme we have asserted that learnability is not an interesting problem for comparison of digital versus DA circuitry. This is because even at low depths, both circuits are able to carve out all four phases of the XXZ Hamiltonian: figure 21(a) demonstrates the DA-learned phase diagram with circuit depth 2, while figure 21(b) shows the digitally-learned diagram at the same depth.



#### C.4. Computational details

Training was performed using a looped gradient descent technique, which at each step calculated the gradient of the loss function with respect to the circuit parameters. This was inputted into Julia’s Flux Optimizer [85, 86], which updated the circuit parameters at the end of each iteration. Training in the Rydberg phases used 50 training iterations, whereas training in the XXZ Hamiltonian phases used 70 iterations for both noisy digital and noisy DA circuits, and 50 iterations for noiseless digital and DA circuits. For the Rydberg Hamiltonian, figure 9(b) was trained at the  $\mathbb{Z}_2$  phase point (2.5, 1.3538) and figure 9(c) was trained at the disordered phase point (0.6, 1.3). For the XXZ Hamiltonian, figure 11(a) was trained on the zAFM phase point (0.01, 0.2184), figure 11(b) on the qzAFM phase point 0.01, 0.9479), figure 11(c) the XY-QLRO phase point (1.8447, 0.1663), and figure 11(d) the VBS phase point (1.5826, 0.6353). The details of numerical convergence are shown in figure 22(b) in the next section.

Testing was performed using an  $n_x \times n_y$  discretized parameter mesh, where  $n_x = n_y = 45$  for the Rydberg testing, and  $n_x = n_y = 20$  for the XXZ Hamiltonian testing. Training points were chosen directly from these meshes. The phase diagrams presented in this work represent an average of the losses achieved over 20 separate training and testing runs.



**Figure 22.** (a) 3 versus 8 digit classification training convergence for digital–analog (blue) versus digital (red) circuits in the noiseless (solid line) and noisy (dashed line) regimes. The number of qubits is  $n = 9$ , the number of layers is  $\ell = 12$  and  $R_b/a = 0.87$ . (b) XXZ Hamiltonian training convergence in the zAFM phase for digital–analog (blue) and digital (red) circuits with  $n = 8, \ell = 2$  in the noiseless (solid line) and noisy (dashed line) regimes. Each step represents the average loss over 20 runs.

One loss evaluation was performed in each iteration of the training loop (for DA, this was 70 iterations  $\times$  1 loss evaluation = 70 loss evaluations) and once per testing iteration ( $20 \times 20 = 400$  additional loss evaluations) for a total of 470 loss evaluations per run. To achieve 1% accuracy in the loss function that appears to be sufficient for reaching the performance here reported, a total of  $10\,000 \times 470 \sim 5 \times 10^6$  samples were necessary.

## Appendix D. Numerical convergence

Figure 22 compares the training loss function at each iteration between digital (with  $\phi = \pi/4$ ) and DA (red/blue), noiseless and noisy (dashed/solid), and on digit classification and phase learning (a/b). We obtained each curve by averaging over 20 random runs. In addition to the other potential improvements DA learning has shown in the main text, figure 22 also demonstrates that DA circuits may plausibly require fewer iterations to reach a desired threshold, though a more in-depth analysis would be necessary to more thoroughly investigate this possibility.

## ORCID iDs

Jonathan Z Lu <https://orcid.org/0000-0002-2806-8248>  
 Lucy Jiao <https://orcid.org/0009-0007-7937-3967>  
 Milan Kornjača <https://orcid.org/0000-0003-3828-6377>  
 Hong-Ye Hu <https://orcid.org/0000-0001-5841-831X>

## References

- [1] Saffman M, Walker T G and Molmer K 2010 *Rev. Mod. Phys.* **82** 2313
- [2] Endres M, Bernien H, Keesling A, Levine H, Anschuetz E R, Krajenbrink A, Senko C, Vuletic V, Greiner M and Lukin M D 2016 *Science* **354** 1024
- [3] Barredo D, de Léséleuc S, Lienhard V, Lahaye T and Browaeys A 2016 *Science* **354** 1021
- [4] Morgado M and Whitlock S 2021 *AVS Quantum Sci.* **3** 023501
- [5] Kornjača M *et al* 2024 Large-scale quantum reservoir learning with an analog quantum computer (arXiv: 2407.02553 [quant-ph])
- [6] Keesling A *et al* 2019 *Nature* **568** 207
- [7] Scholl P *et al* 2021 *Nature* **595** 233
- [8] Maceira I A, Chepiga N and Mila F 2022 *Phys. Rev. Res.* **4** 043102
- [9] Samajdar R, Ho W W, Pichler H, Lukin M D and Sachdev S 2020 *Phys. Rev. Lett.* **124** 103601
- [10] Ebadi S *et al* 2021 *Nature* **595** 227
- [11] Chen C, Bornet G, Bintz M, Emperaiger G, Leclerc L, Liu V S, Scholl P, Barredo D, Hauschild J, Chatterjee S, Schuler M 2023 *Nature* **616** 691
- [12] Bernien H *et al* 2017 *Nature* **551** 579
- [13] Bluvstein D *et al* 2021 *Science* **371** 1355
- [14] Turner C J, Michailidis A A, Abanin D A, Serbyn M and Papić Z 2018 *Nat. Phys.* **14** 745

[15] Semeghini G *et al* 2021 *Science* **374** 1242

[16] Verresen R, Lukin M D and Vishwanath A 2021 *Phys. Rev. X* **11** 031005

[17] De Léséleuc S, Lienhard V, Scholl P, Barredo D, Weber S, Lang N, Büchler H P, Lahaye T and Browaeys A 2019 *Science* **365** 775

[18] Kornjača M, Samajdar R, Macrì T, Gemelke N, Wang S-T and Liu F 2023 *Commun. Phys.* **6** 358

[19] Sahay R, Vishwanath A and Verresen R 2022 arXiv: 2211.01381

[20] Ebadi S *et al* 2022 *Science* **376** eab06587

[21] Kim M, Kim K, Hwang J, Moon E-G and Ahn J 2022 *Nat. Phys.* **18** 755

[22] Nguyen M-T, Liu J-G, Wurtz J, Lukin M D, Wang S-T and Pichler H 2023 *PRX Quantum* **4** 010316

[23] Pichler H, Wang S-T, Zhou L, Choi S and Lukin M D 2018 arXiv: 1809.04954

[24] Wurtz J, Lopes P L, Gemelke N, Keesling A and Wang S 2022 arXiv: 2205.08500

[25] Byun A, Kim M and Ahn J 2022 *PRX Quantum* **3** 030305

[26] Wurtz J *et al* 2023 arXiv: 2306.11727

[27] Surace F M, Mazza P P, Giudici G, Lerose A, Gambassi A and Dalmonte M 2020 *Phys. Rev. X* **10** 021041

[28] Surace F M and Lerose A 2021 *New J. Phys.* **23** 062001

[29] Huang H-Y, Kueng R and Preskill J 2021 *Phys. Rev. Lett.* **126** 190505

[30] Ristè D, Da Silva M P, Ryan C A, Cross A W, Córcoles A D, Smolin J A, Gambetta J M, Chow J M and Johnson B R 2017 *npj Quantum Inf.* **3** 16

[31] Anschuetz E R, Hu H-Y, Huang J-L and Gao X 2023 *PRX Quantum* **4** 020338

[32] Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N and Lloyd S 2017 *Nature* **549** 195

[33] Liu Y, Arunachalam S and Temme K 2021 *Nat. Phys.* **17** 1013

[34] Huang H-Y, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H and McClean J R 2021 *Nat. Commun.* **12** 2631

[35] Gyurik C, Cade C and Dunjko V 2022 *Quantum* **6** 855

[36] Li W and Deng D-L 2022 *Sci. China Phys. Mech. Astron.* **65** 220301

[37] Quek Y, D, Khatri S, Meyer J J and Eisert J 2024 *Nat. Phys.* **20** 1648–58

[38] Schuster T, Yin C, Gao X and Yao N Y 2024 A polynomial-time classical algorithm for noisy quantum circuits (arXiv: 2407.12768 [quant-ph])

[39] Cerezo M *et al* 2021 *Nat. Rev. Phys.* **3** 625

[40] Benedetti M, Fiorentini M and Lubasch M 2021 *Phys. Rev. Res.* **3** 033083

[41] Du Y, Tu Z, Yuan X and Tao D 2022 *Phys. Rev. Lett.* **128** 080506

[42] Uvarov A and Biamonte J D 2021 *J. Phys. A: Math. Theor.* **54** 245301

[43] Self C N, Khosla K E, Smith A W, Sauvage F, Haynes P D, Knolle J, Mintert F and Kim M 2021 *npj Quantum Inf.* **7** 116

[44] Jones T, Endo S, McArdele S, Yuan X and Benjamin S C 2019 *Phys. Rev. A* **99** 062304

[45] Robert A, Barkoutsos P K, Woerner S and Tavernelli I 2021 *npj Quantum Inf.* **7** 38

[46] Liu H-L, Wu Y-S, Wan L-C, Pan S-J, Qin S-J, Gao F and Wen Q-Y 2021 *Phys. Rev. A* **104** 022418

[47] Spall J C 1998 *Johns Hopkins APL Tech. Dig.* **19** 482

[48] Nakanishi K M, Fujii K and Todo S 2020 *Phys. Rev. Res.* **2** 043158

[49] Kingma D P and Ba J 2014 arXiv: 1412.6980

[50] Nelder J A and Mead R 1965 *Comput. J.* **7** 308

[51] Deng L 2012 *IEEE Signal Process. Mag.* **29** 141

[52] Parra-Rodriguez A, Lougovski P, Lamata L, Solano E and Sanz M 2020 *Phys. Rev. A* **101** 022305

[53] Garcia-de Andoin M, Saiz A, Pérez-Fernández P, Lamata L, Oregi I and Sanz M 2023 arXiv: 2307.00966

[54] Martin A, Lamata L, Solano E and Sanz M 2020 *Phys. Rev. Res.* **2** 013012

[55] García-Molina P, Martin A, de Andoin M G and Sanz M 2021 arXiv: 2107.12969

[56] Daley A J, Bloch I, Kokail C, Flannigan S, Pearson N, Troyer M and Zoller P 2022 *Nature* **607** 667

[57] Ren W *et al* 2022 *Nat. Comput. Sci.* **2** 711

[58] Levine H *et al* 2019 *Phys. Rev. Lett.* **123** 170503

[59] Jaksch D, Cirac J I, Zoller P, Rolston S L, Côté R and Lukin M D 2000 *Phys. Rev. Lett.* **85** 2208

[60] Wilk T, Gaëtan A, Evellin C, Wolters J, Miroshnychenko Y, Grangier P and Browaeys A 2010 *Phys. Rev. Lett.* **104** 010502

[61] Evered S J *et al* 2023 *Nature* **622** 268

[62] Wild D S, Sels D, Pichler H, Zanoci C and Lukin M D 2021 *Phys. Rev. Lett.* **127** 100504

[63] Bluvstein D *et al* 2022 *Nature* **604** 451

[64] Lu J Z, Jiao L, Wolinski K, Kornjaca M, Liu F, Hu H, Wang S and Yelin S Digital-analog quantum learning on Rydberg atoms (available at: <https://github.com/jz-lu/daqml>)

[65] Cerezo M, Sone A, Volkoff T, Cincio L and Coles P J 2021 *Nat. Commun.* **12** 1791

[66] Samajdar R, Choi S, Pichler H, Lukin M D and Sachdev S 2018 *Phys. Rev. A* **98** 023614

[67] Kalinowski M, Samajdar R, Melko R G, Lukin M D, Sachdev S and Choi S 2022 *Phys. Rev. B* **105** 174417

[68] Hilal W, Gadsden S A and Yawney J 2022 *Expert Syst. Appl.* **193** 116429

[69] Kottmann K, Metz F, Fraxanet J and Baldelli N 2021 *Phys. Rev. Res.* **3** 043184

[70] Nassif A B, Talib M A, Nasir Q and Dakalbab F M 2021 *IEEE Access* **9** 78658

[71] Farhi E, Goldstone J, Gutmann S and Sipser M 2000 arXiv: quant-ph/0001106

[72] Farhi E, Goldstone J and Gutmann S 2002 arXiv: quant-ph/0208135

[73] Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A and Predu D 2001 *Science* **292** 472

[74] Roland J and Cerf N J 2002 *Phys. Rev. A* **65** 042308

[75] Lu J Z, Bravo R A, Hou K, Dagnew G A, Yelin S F and Najafi K 2022 arXiv: 2206.11970

[76] Fendley P, Sengupta K and Sachdev S 2004 *Phys. Rev. B* **69** 075106

[77] Schachenmayer J, Lesanovsky I, Micheli A and Daley A J 2010 *New J. Phys.* **12** 103044

[78] Lee J Y, Ramette J, Metlitski M A, Vuletić V, Ho W W and Choi S 2023 *Phys. Rev. Lett.* **131** 083601

[79] Mudry C, Furusaki A, Morimoto T and Hikihara T 2019 *Phys. Rev. B* **99** 205153

[80] Schuch N and Verstraete F 2009 *Nat. Phys.* **5** 732

[81] Haldar A, Mallayya K, Heyl M, Pollmann F, Rigol M and Das A 2021 *Phys. Rev. X* **11** 031062

- [82] 2023 Bloqade.jl: package for the quantum computation and quantum simulation based on the neutral-atom architecture (available at: <https://github.com/QuEraComputing/Bloqade.jl/>)
- [83] Luo X-Z, Liu J-G, Zhang P and Wang L 2020 *Quantum* **4** 341
- [84] Wierichs D, Izaac J, Wang C and Lin C Y-Y 2022 *Quantum* **6** 677
- [85] Innes M, Saba E, Fischer K, Gandhi D, Rudilosso M C, Joy N M, Karmali T, Pal A and Shah V 2018 *CoRR* (arXiv: [1811.01457](https://arxiv.org/abs/1811.01457))
- [86] Innes M 2018 *J. Open Source Softw.* **3** 602