


Using Deep Learning to Understand and Mitigate the Qubit Noise Environment

David F. Wise,^{1,2,*} John J.L. Morton^{1,2,3} and Siddharth Dhomkar^{1,†}

¹London Centre for Nanotechnology, University College London, London WC1H 0AH, United Kingdom

²Quantum Motion Technologies, Windsor House, Cornwall Road, Harrogate, England, HG1 2PW

³Department of Electronic & Electrical Engineering, University College London, London WC1E 7JE, United Kingdom

 (Received 3 May 2020; revised 11 December 2020; accepted 7 January 2021; published 28 January 2021)

Understanding the spectrum of noise acting on a qubit can yield valuable information about its environment, and crucially underpins the optimization of dynamical decoupling protocols that can mitigate such noise. However, extracting accurate noise spectra from typical time-dynamics measurements on qubits is intractable using standard methods. Here, we propose to address this challenge using deep-learning algorithms, leveraging the remarkable progress made in the field of image recognition, natural language processing, and more recently, structured data. We demonstrate a neural-network-based methodology that allows for extraction of the noise spectrum associated with any qubit surrounded by an arbitrary bath, with significantly greater accuracy than the current methods of choice. The technique requires only a two-pulse echo decay curve as input data and can further be extended either for constructing customized optimal dynamical decoupling protocols or for obtaining critical qubit attributes such as its proximity to the sample surface. Our results can be applied to a wide range of qubit platforms, and provide a framework for improving qubit performance with applications not only in quantum computing and nanoscale sensing but also in material characterization techniques such as magnetic resonance.

DOI: [10.1103/PRXQuantum.2.010316](https://doi.org/10.1103/PRXQuantum.2.010316)

Robust isolation of a qubit from unwanted noise in its environment is a key factor in technologies such as quantum computers and sensors. There is a long history in magnetic resonance of developing so-called dynamical decoupling protocols, including Carr-Purcell-Meiboom-Gill (CPMG) [1–3], periodic dynamical decoupling (PDD) [4], concatenated dynamical decoupling (CDD) [5], and Uhrig dynamical decoupling (UDD) [6], as general methods to preserve spin coherence in the presence of noise. In practice, the experimental noise spectrum varies significantly between different qubits in ways which are non-trivial to either predict, or indeed to accurately extract from most common measurements [7–9]. As a result, it is difficult to predict *a priori* which of the several possible dynamical decoupling protocols would provide optimal suppression of decoherence. Indeed, one could envisage

constructing a decoupling protocol customized for a particular qubit, but this is impossible without knowing the actual qubit noise spectrum with sufficient accuracy. Additionally, it is worth emphasizing that the knowledge of obscured qubit environment is in itself a valuable outcome of a precise noise spectroscopy.

Significant advances have been made in machine-learning and specifically deep-learning techniques, for example in the fields of computer vision and natural language processing, and more recently they have been applied to problems in physics and quantum engineering [10]. Deep feed-forward neural networks have been used to enhance extraction of material parameters in scanning probe microscopy [11], for the derivation of the parameters associated with a $1/f$ noise spectrum from randomized benchmarking experiments [12], and for processing of magnetic resonance spectroscopy data, associated with nuclear magnetic resonance (NMR), electron paramagnetic resonance (EPR), and double electron-electron resonance (DEER) experiments [13–17]. In quantum control, neural networks have been used to enhance dynamical decoupling protocols [18], to suppress dephasing [19], and to improve the robustness of quantum control operations [20]. In addition, deep-reinforcement-learning techniques have been applied to quantum metrology both as an efficient experiment design heuristic and to increase sensor

*d.wise.15@ucl.ac.uk

†s.dhomkar@ucl.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

sensitivity by more than an order of magnitude over comparable approaches [12,21,22].

Here, we propose that the challenges of accurately obtaining qubit noise spectra can be efficiently handled by employing deep-learning algorithms. We show how a deep neural network can be trained to extract the noise spectrum from simple and widely used time-dynamics measurements on qubits, such as the two-pulse ‘‘Hahn’’-echo curves, and compare the accuracy of a deep-learning approach with that of standard techniques. Finally, we examine a neural-network-based technique of processing noisy experimental data and discuss potential uses of an accurate noise spectrum for quantum control. In addition to the possibility of optimizing dynamical decoupling to extend qubit coherence, we explore how useful information about the qubit environment such as its proximity to particular noise sources can be deduced from its noise spectrum [23].

I. NOISE SPECTROSCOPY USING DYNAMICAL DECOUPLING

The noise spectrum of a qubit is an effective proxy to probe its surroundings and provides valuable information for qubit characterization. Furthermore, once the environmental noise spectrum is known, it is possible, in principle, to run an optimization protocol that minimizes ‘‘decoherence’’ [which can be thought of as proportional to the overlap of the filter function $F(\omega t)$ associated with a given dynamical decoupling pulse sequence and the actual noise spectrum $S(\omega)$] by varying parameters in the sequence such as separation between pulses [7], or their amplitude, phase, and duration. Alternatively, dynamical decoupling protocols can be tailored under real-time experimental feedback [8], however, such methods can be experimentally expensive to perform, limiting their applicability. Therefore, we focus here on using the *minimum* experimental data from a given qubit, such as a single

coherence-decay curve under a particular pulse sequence, and using purely computational methods to identify the noise spectrum and optimized decoupling sequence.

The time dependence of qubit coherence under applied pulse sequences can be used for extracting useful information about the noise sources in the environment [9,24–26]. Such sequences, with overall duration of $t = n\tau$, can typically be described as a set of n π pulses, each with duration τ_π applied at time t_k , and possessing a characteristic filter function [27]:

$$F(\omega t) = \left| 1 + (-1)^{n+1} e^{i\omega n\tau} + 2 \sum_{k=0}^n (-1)^k e^{i\omega t_k} \cos\left(\frac{\omega\tau_\pi}{2}\right) \right|^2. \quad (1)$$

Provided that the noise spectrum, $S(\omega)$, acting on the qubit is stationary, Gaussian, and couples only along the qubit’s z axis it can be extracted from the measured time-dependent coherence-decay curve, $C(t)$, by solving the following integral equation:

$$\chi(t) = -\ln C(t) = \int_0^\infty \frac{d\omega}{2\pi} S(\omega) \frac{F(\omega t)}{\omega^2}, \quad (2)$$

where $\chi(t)$ is known as the decoherence functional [28]. However, in practice accurately solving such an integral equation is nontrivial, and without a sufficiently faithful noise spectrum, the technique cannot yield an optimized protocol to significantly suppress decoherence.

It is possible to simplify Eq. (2) by assuming the filter function at a given delay time to be a Dirac δ function localized at a desired frequency ω_0 (which for CPMG sequences corresponds to $n\pi/t$) [25]. This assumption permits a simple mapping of the coherence-decay curve onto a corresponding noise spectrum:

$$S(\omega_0) = \frac{-\pi \ln C(t)}{t}. \quad (3)$$

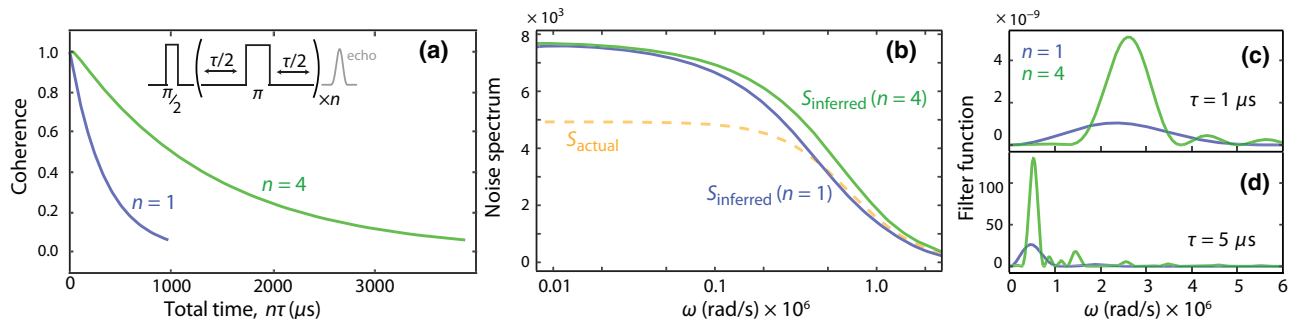


FIG. 1. The challenge of extracting noise spectra from coherence-decay measurements. (a) Decoherence curves under dynamical decoupling sequences (see inset) as a function of total sequence duration, $n\tau$, simulated using Eq. (2) from some underlying noise spectrum [dashed curve in (b)]. Solid curves in (b) show noise spectra extracted from the coherence decays using the δ -function approximation to the filter function in Eq. (3). (c),(d) Actual filter functions corresponding to sequences comprising one and four π pulses at two τ values. Simulations assume a π -pulse duration of 100 ns.

Figure 1 shows an application of this simplified approach to obtain noise spectra corresponding to decoherence curves that are simulated from some underlying noise spectrum, using the integral equation in Eq. (2). However, as demonstrated in Fig. 1(c) the δ -function approximation is valid only where the spacing between pulses is much longer than the pulse duration. Increasing the number of π pulses produces a narrower peak in the filter function but with the expense of increased harmonics. As a result, the noise spectra inferred from the decoherence curves using Eq. (3) are a poor fit to the actual spectrum used to generate them.

A second challenge in obtaining accurate noise spectra from coherence-decay curves is handling experimental noise in the measurement itself in an unprejudiced way. Depending on the type of qubit, the dominant noise sources can vary significantly, e.g., in the case of flux qubits and Si quantum dots the primary source of noise behaves as $1/f^\alpha$, where f is the frequency and $\alpha \sim 1$ [29,30], telegraphic noise [31] dominates for GaAs quantum dots, bulk nitrogen-vacancy (N-V) centers in diamond are affected by Lorentzian-type noise [25,32], whereas near-surface N-V centers are prone to double-Lorentzian-type noise [23]. Despite this rich variety of noise sources, decoherence curves are often fit with a stretched-exponential function with essentially two parameters—coherence lifetime, T_2 , and power of the stretched exponential, p :

$$C(t) = e^{-(t/T_2)^p}. \quad (4)$$

Whilst this equation holds for decoherence in the presence of certain simple noise forms (e.g., for simple $1/f^\alpha$ noise) [31], in practice, qubits may be surrounded by multiple noise sources with varied functional forms, which can only be approximately represented by Eq. (4), and fitting

experimental decay curves can remove or obscure valuable information regarding the true noise spectrum.

II. PROPOSED METHODOLOGY

As illustrated in the discussion above, solving Eq. (2) to extract the noise spectrum specific to the qubit under investigation is nontrivial, however, the reverse operation (evaluating the coherence decay for a given noise spectrum and decoupling sequence) is relatively simple. This asymmetry lends itself to a neural-network-based learning approach to achieve a significant improvement in accuracy, owing to the capacity of neural networks to act as universal function approximators [33,34]. To successfully implement the deep-learning technique we first develop an efficient method for generating a sufficiently large and diverse set of training data. We split this training data into train-validation-test tranches and explore a selection of possible network architectures before selecting the most effective option and training it for optimal performance on the validation set [35]. Finally, we verify the chosen network's performance on the test set. A concise flow chart of the methodology is shown in Fig. 2(a) with detailed description provided below.

A. Generation of training data

We simulate a variety of noise spectra assuming commonly applicable models described above and compute the corresponding coherence-decay curves using Eq. (2). Input and output data used in training have 151 data points each with log spacing in the time domain. Frequency range is determined by the time range of the coherence-decay curves following the aforementioned relationship $\omega = n\pi/t$ [as given by Eq. (1)]. In the first instance, we

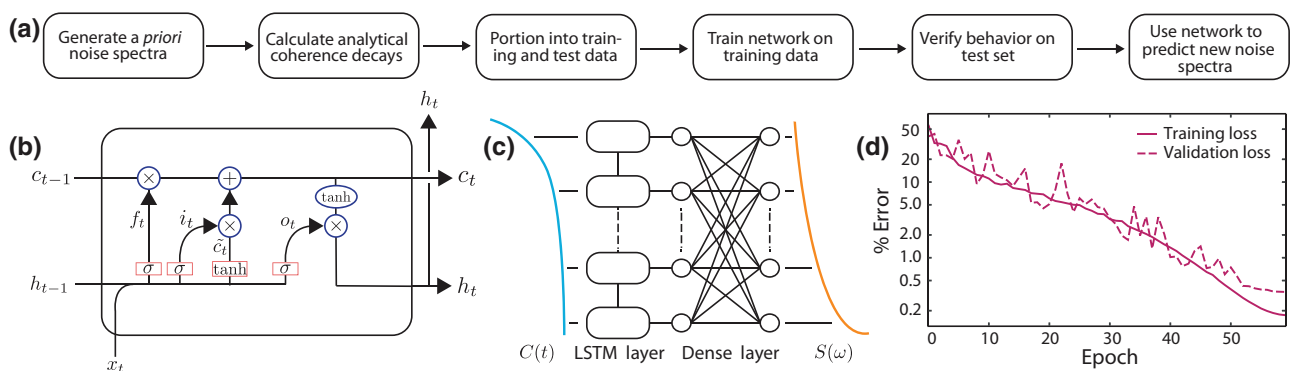


FIG. 2. Our neural-network approach. (a) Flow chart of the methodology. (b) A detailed view of a long short-term memory (LSTM) cell. c_t is the cell state vector (providing the memory of the network), x_t the cell input (one point in a coherence decay), and h_t the cell output, which is fed both to the output dense layer and the next cell in the series. Red boxes represent dense layers with either sigmoid (σ) or tanh activations. Blue circles represent pointwise operations. (c) A sketch of the whole network: points on the coherence decay are fed into the LSTM layer, the output of the LSTM layer is then input into a dense layer and the output of this dense layer is compared with a target noise spectrum. (d) An example network training graph showing training loss and validation loss (mean absolute percentage error) over training.

generate three different forms of noise spectra for training purposes.

1. A noise spectrum generated from a stretched-exponential coherence decay. The δ -function approximation shown in Eq. (3) is used to first produce an approximate noise spectrum for a given coherence decay. Equation (2) is then used to produce a new coherence decay with an exact relationship to the noise spectrum. It is this new coherence decay that is then used as an input for neural-network training.

2. $1/f$ noise, with functional form A/f^α , generated using simulated parameters chosen to be in line with expected physical constants.

3. A noise spectrum with a Lorentzian form:

$$\frac{\Delta^2 \tau_c}{\pi} \frac{1}{1 + (\omega \tau_c)^2}, \quad (5)$$

where Δ and τ_c are coupling strength and correlation time, respectively, again generated using appropriate simulated parameters.

In theory, these data can then be used to train a neural network to output the noise spectrum corresponding to a given coherence decay measured using a Hahn-echo sequence. Specifically, network training can be performed using coherence decays as inputs to the network, with noise spectra as the target outputs. Before training, the generated data must be split into training, validation, and test data sets. Training data is actively used to update the network parameters, whilst the network performance on validation data is monitored during training to avoid overfitting [see Fig. 2(d)] [35] and used for tuning the network hyperparameters (its overall structure, loss function etc.). The test data is held out for final evaluation of performance of the network on hitherto unseen data; error and loss rates quoted throughout this paper are always on the test set unless otherwise specified.

B. Choice of network architecture

Many different variants of neural network exist, the simplest being a deep feed-forward neural network [35], consisting of multiple layers each with a number of units. The first layer takes as input the training data X , with the outputs of each layer fed as input to the next. The final layer output is compared with the training data y to calculate the network loss. Each layer in a feed-forward network can be represented by multiplication by a matrix of weights, W followed by addition of a bias vector, b and finally an activation function, g . The action of the layer, h ,

on its inputs x can be written:

$$h(x) = g(x^\top W + b). \quad (6)$$

For each training step, the parameters of the neural network, ϕ are updated so that the overall function of the network f is brought closer to the desired function f^* . This is done by taking the gradient of the chosen loss function with respect to the network parameters and updating those to reduce the loss, through a process known as back propagation [35].

One problem with a feed-forward network is that all inputs are connected to all units in each layer, which has the effect of destroying local correlations. For our purposes, a network architecture that preserves these local correlations is desirable. There are two obvious choices for working with time-ordered data such as a coherence-decay curve: *convolutional* and *recurrent* (RNN) networks [36,37]. After exploration the architecture found to be most effective for the task at hand was a long short-term memory network [38], a special case of the RNN. LSTMs introduce the capacity for long-range correlations in input data to be preserved and several recent studies have found the LSTM paradigm to be effective in working with time-series data [39–41].

A simple RNN is a sequence of discrete units, each of which processes a single time step of the input data, x_t and produces an output, h_t . This output is concatenated with the input of the next layer (x_{t+1}) to allow information earlier in the sequence to influence the output of layers later in the sequence. Although useful for maintaining short-term correlations, one subtlety of the back-propagation method of training neural networks is that early layers have relatively little impact on the gradient of the loss function of the network as a whole, meaning that early information is effectively “forgotten” by the network. To combat this, LSTMs introduce the concept of a “cell state,” c_t , to which information can be added (remembered, i_t) or subtracted (forgotten, f_t), allowing the maintenance of important long-term information. A schematic of a single LSTM cell can be seen in Fig. 2(b) and its output is given by the following equations:

$$f_t = \sigma_g[W_f(x_t \oplus h_{t-1}) + b_f], \quad (7)$$

$$i_t = \sigma_g[W_i(x_t \oplus h_{t-1}) + b_i], \quad (8)$$

$$o_t = \sigma_g[W_o(x_t \oplus h_{t-1}) + b_o], \quad (9)$$

$$\tilde{c}_t = \tanh[W_c(x_t \oplus h_{t-1}) + b_c], \quad (10)$$

$$c_t = f_t \times c_{t-1} + i_t \times \tilde{c}_t, \quad (11)$$

$$h_t = o_t \times \tanh(c_t). \quad (12)$$

Here σ_g refers to the sigmoid activation function, W s refer to the weight matrices associated with each LSTM cell operation, b s refer to the bias vectors of each LSTM cell

operation and x_t and h_t refer to the input and output vectors at the relevant time step. The \oplus operation specifies a concatenation of the two vectors x_t and h_{t-1} . Many descriptions of the mathematical details and intuitions behind RNNs and LSTMs can be found [35,42,43]. To produce an output of the correct size, the output from the LSTM is connected to a dense layer matching the size of the y training data [see Fig. 2(c)]. Although the activation function for the output layer of a neural network used for regression is typically a linear function, in this case we select an exponential function. This aids the training of the neural network as weights and biases are typically initialized for relatively small output values (< 5) whilst the noise spectra used for regression in this case tend to have values $> 10^4$.

C. Network training

The neural networks used for this study are defined using the Keras module of the Tensorflow 2.0 deep-learning library distributed and maintained by Google [44,45]. We carry out training using the previously discussed train-validation-test split. We use one-cycle

learning, as proposed by Smith [46], to control the learning rate during training. Hyperparameter tuning is undertaken to determine optimum network features using Bayesian search and the Weights and Biases library [47].

III. RESULTS AND DISCUSSION

Figure 3 (and Fig. 8) compares the performance of the traditional δ -function approximation for calculating the noise spectrum against our neural-network-based approach. We show results for synthetic coherence-decay curves generated from three different sources of noise: Lorentzian, $1/f$, and stretched exponential. The recorded performance of the neural network is for a *single* network trained on the three noise sources simultaneously, applied to the three different test data sets. Detailed statistics on the performance of the network are shown in Table I.

Our results show that the neural-network approach significantly outperforms the methodology based on a δ -function approximation to the sequence filter function for deducing noise spectra from qubit coherence-decay curves. Of particular interest is that a single network is able

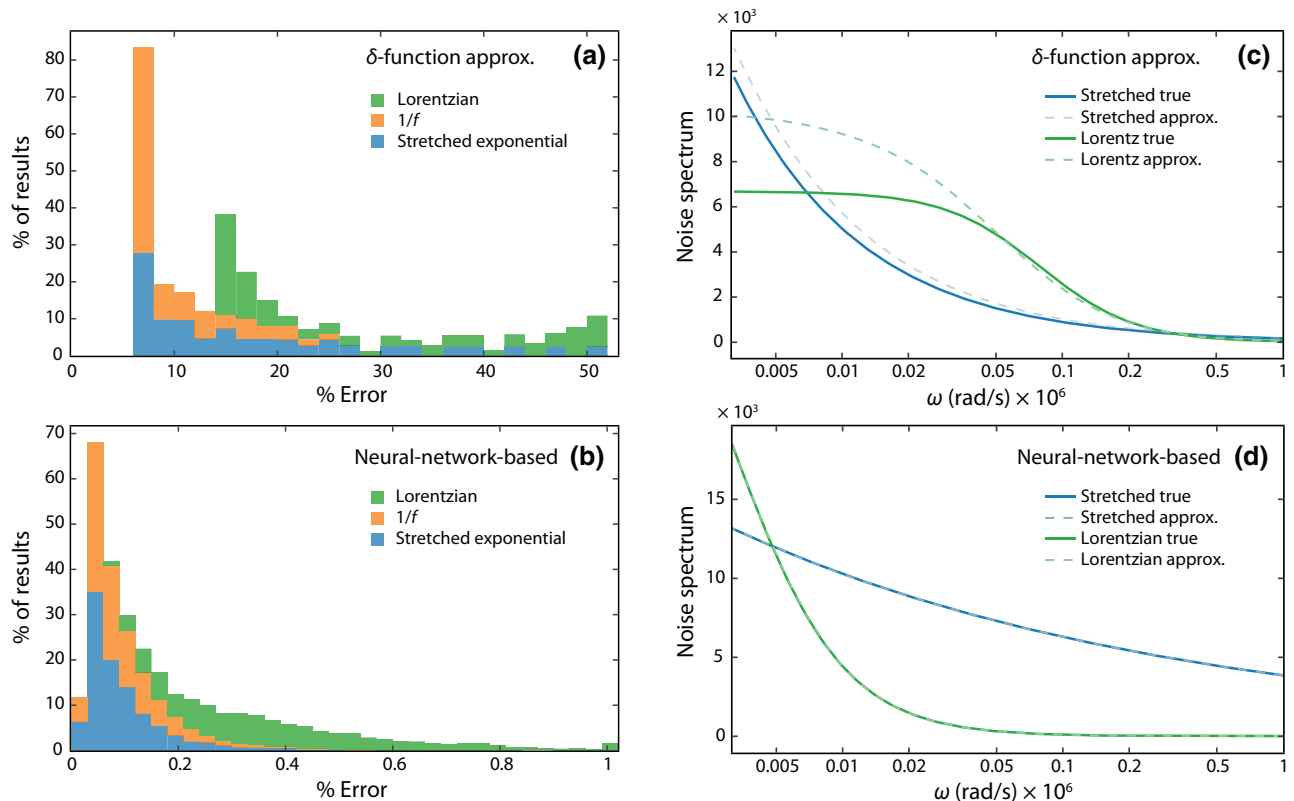


FIG. 3. Comparison of two approaches for producing a noise spectrum from a measured coherence-decay curve. Stacked histograms of errors in the predicted noise spectra from the holdout test data set are shown using (a) a δ -function approximation, and (b) our neural network. Three different forms of noise spectrum are studied, each with approximately 8500 individual noise spectra and associated decay curves. In (b), the tail of the histogram is clipped for clarity and the maximum error is 2.3%. Also shown are the true (solid curves) and predicted (dashed curves) noise spectra obtained under (c) the δ -function approximation and (d) the neural-network approach. The noise spectra used are at the 50th percentile error for each approach.

TABLE I. Comparison of error statistics in % in noise-spectrum extraction using the neural-network approach introduced here versus the δ -function approximation for the pulse-sequence filter function.

Noise model	NN approach		δ -function approx.	
	Mean	σ	Mean	σ
Stretched exp.	0.098	0.1	17.6	12.2
$1/f$	0.10	0.1	10.2	5.0
Lorentzian	0.39	0.21	28.4	13.7

to successfully reproduce noise spectra of different functional forms. A current limitation of this approach is that the network has only been trained on a set of coherence decays with a decoherence time constant T_2 s in a relatively small range between approximately 120 to 600 μ s. In

principle, this can be circumvented by training a network on the same principles but for different length coherence decays. For example, in Figs. 4(a) and 4(b) a network is trained on coherence decays with T_2 s between 10 – 140 μ s, generated using the phenomenological stretched-exponential method discussed above. This neural-network approach can therefore be applied generally, given additional training and data generation to cover the parameter range of interest. A more generalized approach would use the time-vector data for training, in addition to the signal data, and represents a promising direction for further development.

Next, we demonstrate the performance of neural-network-based approach for quantum systems experiencing multiple simultaneous noise sources. Figures 4(c) and 4(d) show the results of a network that is trained to

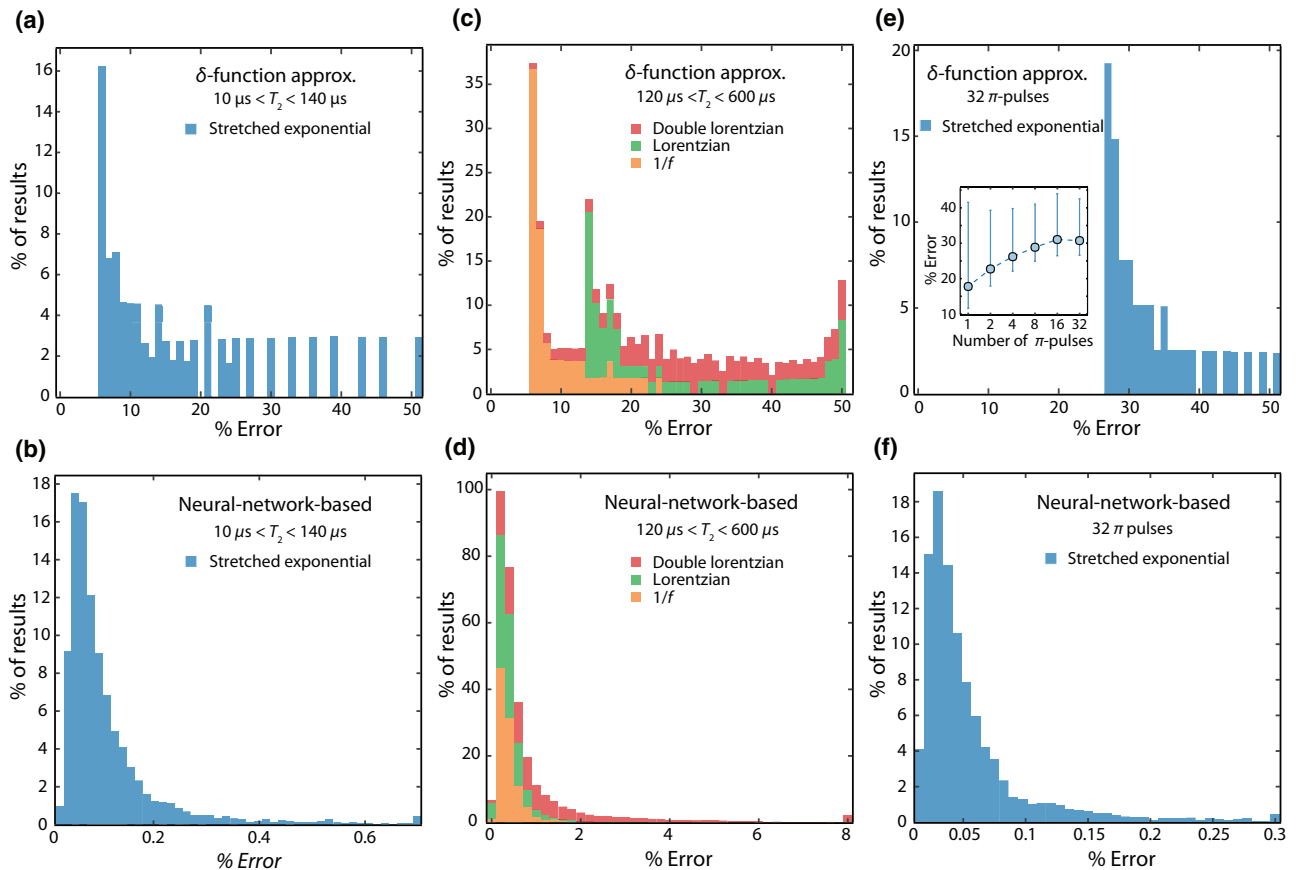


FIG. 4. Comparison between estimated error in noise-spectra predictions for three different neural networks and those obtained using a δ -function approximation. Histograms of error rates for predicting noise spectra generated using Eq. (3) from coherence decays with a coherence-decay time constant T_2 between 10 – 140 μ s obtained by means of (a) the δ -function approach and (b) a trained neural network. Stacked histograms of error rates for predicting noise spectra having three different functional forms: $1/f$, Lorentzian, and double Lorentzian obtained by means of (c) the δ -function approach and (d) a third trained neural network, the tail of the neural-network distribution is clipped for clarity with maximum error in this case 200.1% with high errors caused by noise spectra with large portions very close to 0 where small deviations in predicted noise spectrum magnify mean percentage error. Histograms of error rates for predicting stretched-exponential noise spectra generated from coherence decays constructed by implementing dynamical decoupling protocol comprising 32 π pulses extracted by means of (e) the δ -function approach and (f) a second trained neural network; inset in (e) shows mean percent errors along with the maximum deviation as a function of the number of π pulses for the conventional technique.

TABLE II. Error statistics in % for neural network and δ -function noise-spectrum approach, for a network trained on $1/f$, Lorentzian, and double Lorentzian noise spectra.

Noise model	NN approach		δ -function approx.	
	Mean	σ	Mean	σ
$1/f$	0.41	0.46	10.2	5.0
Lorentzian	0.42	0.40	28.4	13.7
Double Lorentzian	1.72	3.63	30.8	11.9

predict noise spectra arising from three different models: $1/f$, Lorentzian, and double Lorentzian [sum of two Lorentzians with different Δ , and τ_c following Eq. (5)]. Detailed statistics are shown in Table II. Double-Lorentzian noise spectra are characteristic of quantum systems experiencing noise from two uncorrelated sources, as observed in near-surface spins such as N-V centers in diamond [23]. Our results show how a single trained network can successfully predict cases where a coherence decay is a result of a system experiencing multiple noise sources, whilst also differentiating multiple single-noise sources. Furthermore, training data comprising varied models allows the network to learn complex correlations and subsequently output robust predictions in the context of previously unseen data (see Appendix B, Fig. 9). This proof-of-principle demonstration suggests that expanding the range of noise spectra that the network is able to identify, particularly where a noise spectrum is made up of multiple different noise sources, is a fruitful direction for future research.

Having shown the robustness of the neural-network approach to a range of coherence-decay times and a combination of noise spectra, in Figs. 4(e) and 4(f) we show the performance of a network trained to reproduce noise spectra from coherence decays generated using the phenomenological method described above but in this case using a filter function associated with a 32 π -pulse CPMG sequence. The similarity between the δ function and the actual filter function of a CPMG sequence with a sufficiently large number of π pulses is generally used to justify the approach of employing Eq. (3) to extract noise spectra. Surprisingly, we find the δ -function approach to be less accurate in this case than when a Hahn-echo filter function is used as shown in Fig. 4. The neural network appears unaffected by the change in pulse sequence and remains able to successfully reproduce the noise spectra with a low error rate. We speculate that the reason for the poor performance of the δ -function approach in this case is that, whilst the central peak of the filter function becomes narrower at higher pulse numbers, the contributions of higher frequency harmonics become more significant [48] [also see Figs. 1(e) and 1(f)]. Therefore, for monotonic noise spectra, such as those considered in Fig. 4, the advantages of using multiple π -pulse measurements together with δ -function approximation are not evident. Nevertheless,

the implementation of multi- π -pulse protocols becomes indispensable for nonmonotonic feature extraction owing to the narrowness of the filter function that they produce. Hence, researchers have developed elaborate strategies for mitigating the errors introduced by the higher harmonics of the CPMG filter function.

One of these state-of-the-art approaches, developed by Alvarez and Suter [9], employs sequences involving multiple π pulses to infer the underlying noise spectrum. The approach relies on the fact that if the sequence maintains a constant τ , and measures the qubit decay as a function of the number of π pulses then the decay rate, R , of the qubit during dynamical decoupling becomes time independent and its relationship with the noise spectrum can be defined as follows:

$$R = \sum_{k=1}^{\infty} A_k^2 S(k\omega_0), \quad (13)$$

where $\omega_0 = \pi/\tau$ and A_k s are related to the sensitivity amplitudes of the filter function at the probing frequencies [see Appendix C, Fig. 10(a)]. Therefore, by measuring signal at carefully selected values of τ and repeating the experiment several times by employing multiple π -pulse decoupling sequences, accurate estimation of bath noise spectrum can be achieved (see Fig. 10 for more details on the actual methodology).

To benchmark the neural-network-based approach against more elaborate experimental methods, we train a network with a set of noise spectra that have an underlying $1/f$ functional form overlaid with an additional Lorentzian-shaped feature with randomly selected amplitude, peak position, and width. Figure 5 demonstrates that the neural-network-based methodology performs at least as well as Alvarez and Suter, despite the former using only a simple Hahn-echo experiment. Both approaches are significantly superior to the simple δ -function approach. Notably, the results establish that the network is capable of extracting more complex nonmonotonic noise spectra with high accuracy, circumventing the need of lengthy and technically challenging measurements. Further work will be required to upgrade the network, both in terms of the architecture and the diversification of the training data, for extending its applicability to varied qubit environments.

There exist more intricate strategies, for example, discrete prolate spheroidal sequences (DPSSs) [49] have a filter function that is devoid of additional lobes and higher harmonics resulting in suppression of spectral leakage, while a Gaussian enveloped dynamic sensitivity control sequence (gDYSCO) [50] allows generation of a filter function with minimal spectral leakage at the expense of reduced sensitivity and broader main peak. Though more effective than the basic CPMG approach, these sequences can be challenging to implement in practice and require repeated measurements [49,50], while in contrast

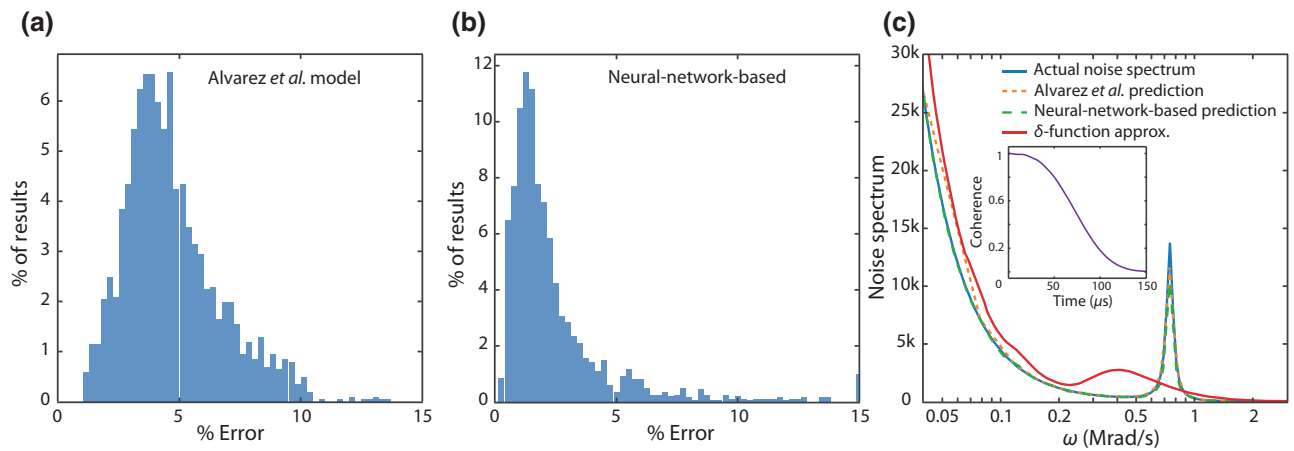


FIG. 5. Comparison of the neural-network approach with an established approach. (a) Histograms of errors in noise extraction when employing Alvarez *et al.* methodology [9], (b) histograms of errors in predicting noise via the trained neural network, (c) a plot comparing predictions of different methodologies for a representative noise spectrum. The inset shows the decoherence curve corresponding to the underlying noise spectrum. Note: all the artificial noise spectra are generated using a $1/f$ model along with a Lorentzian-shaped feature whose parameters are chosen randomly.

our neural-network-based approach is shown to accurately extract underlying noise spectra using only a single measurement of a Hahn-echo decay curve. Nevertheless, in principle, elements of both techniques could be combined in the pursuit of even better noise-spectrum extraction, while minimizing the increase in experimental complexity.

A. Handling experimental data

The utility of our proposed approach requires a robust method for ingesting experimental coherence curves, including unavoidable experimental measurement noise. A traditional approach would be to apply a least-squares fit of a stretched-exponential function [given in Eq. (4)] to the noisy data to gain a best approximation to the underlying decay. However, such an approach necessarily makes an

assumption regarding the particular form of the noise, restricting the possible noise spectra that can be inferred to those that produce a stretched-exponential coherence decay.

We propose instead a function-agnostic approach to removing experimental noise. Once again, we look to LSTM-based neural networks to perform this task—an approach that has already been successfully used for denoising of electrocardiogram (ECG) data [41]. We train the network by applying random noise of $\pm 5\%$ to the coherence curves generated from (a) the three different noise models described above, or (b) the phenomenological stretched-exponential curves. This is used as the input X data, whilst the original noiseless curves are used as the target y . To evaluate the success of this neural-network technique we compare its ability to reconstruct a given

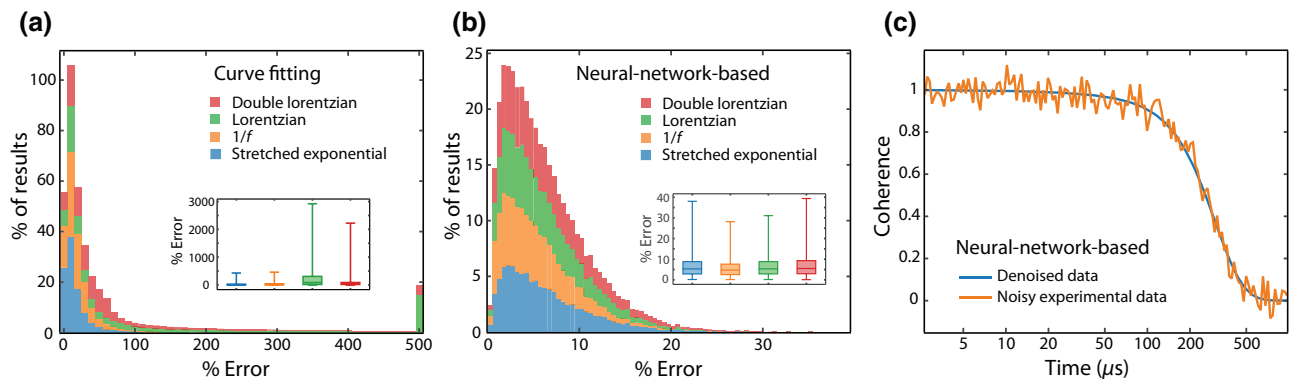


FIG. 6. Denoising the experimental data using curve fitting or neural networks. Stacked histograms of errors in smoothing synthetic experimental decoherence curves from four classes of noise model are shown using two denoising approaches: (a) curve fitting to Eq. (4), or (b) the neural-network approach. In both cases, insets show mean errors and the maximum deviation. The percent errors are calculated using logarithmic values of decoherence curve. (c) A synthetic noisy decoherence curve (from a $1/f$ noise model) along with a denoised curve obtained from the neural network.

TABLE III. Error statistics in % of $\log C(t)$ for neural-network and stretched-exponential-fitting denoising approaches.

Noise model	NN approach		Curve fitting	
	Mean	σ	Mean	σ
Stretched exp.	6.49	4.84	16.79	20.53
$1/f$	5.52	3.82	21.21	23.06
Lorentzian	6.35	4.42	233.67	341.21
Double Lorentzian	6.72	4.91	96.72	156.02

coherence decay against the approach of fitting a stretched exponential to the same noisy data.

To infer the impact of the errors on reconstruction of the noise spectrum, we compare the mean absolute percentage errors in $\log C(t)$ of ground truth versus the noisy data, as the noise spectrum depends explicitly on this as seen in Eq. (2)—the results are shown in Fig. 6, with detailed

statistics in Table III. Performance for the neural network is comparable to that of fitting for stretched exponential and $1/f$ coherence decays, although it does show some improvement. However, the advantage of the network becomes clear when we examine errors for coherence decays derived from Lorentzian and double-Lorentzian noise spectra: whilst the neural network performs similarly on all decay types, the error in curve fitting increases drastically for these cases. The results demonstrate that a functional-form agnostic approach to denoising—as permitted by the neural network—is essential for an accurate and general extraction of the noise spectrum from experimental data.

IV. OUTLOOK

Accurately deducing the noise environment of a quantum system has many potential applications: for qubits,

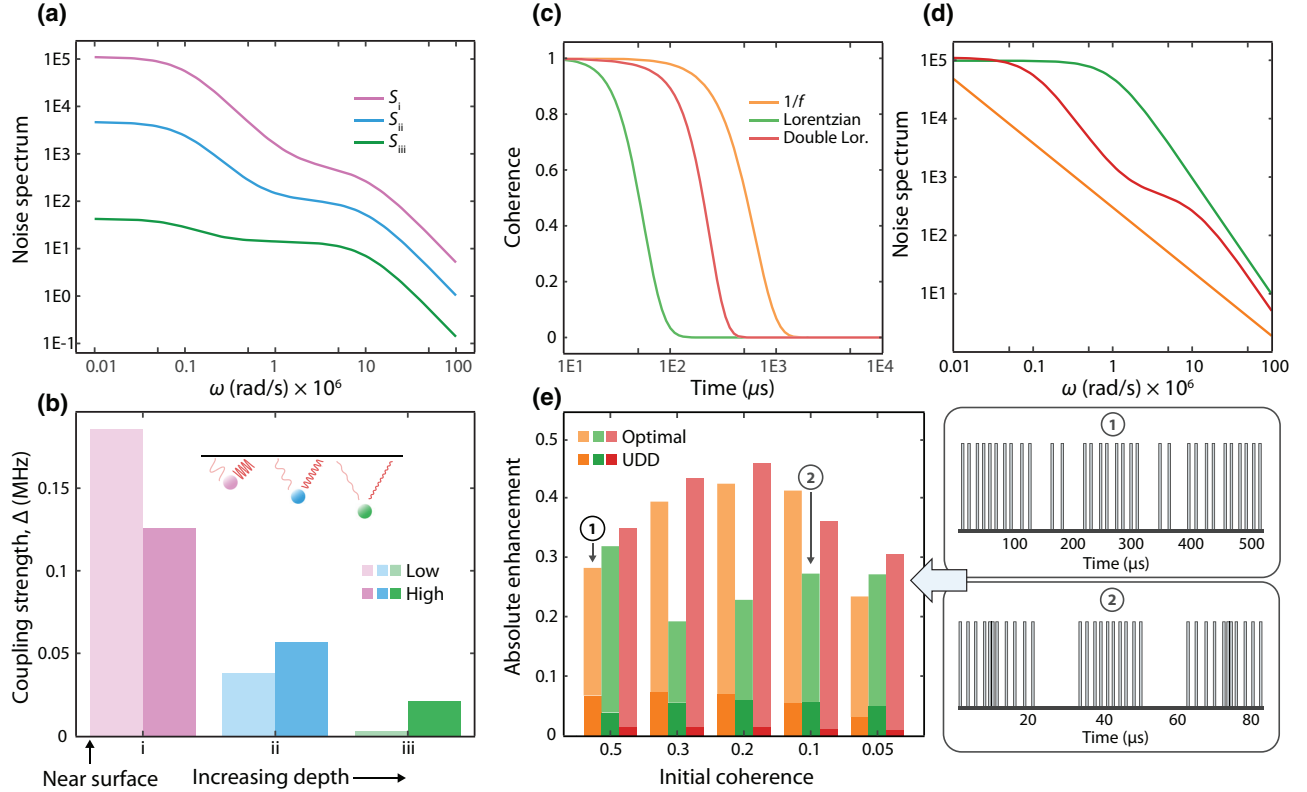


FIG. 7. Potential applications of the methodology include defect depth prediction [23] and the generation of customized dynamical decoupling sequences. (a) Simulated double-Lorentzian-type noise-spectra characteristic of near-surface spin defects, for different increasing depths (i–iii). (b) Defect-bath coupling strengths for low- and high-frequency noise corresponding to noise spectra shown in (a); the variations in the coupling strengths are indicative of changes in the defect position with respect to the surface. Generation of optimized dynamical decoupling sequences: (c) specimen decoherence curves derived by applying 32 π -pulse CPMG to the noise spectra shown in (d). (e) Histograms demonstrating coherence improvement after implementing the SLSQP algorithm to optimize π -pulse positions; enhancements obtained by employing UDD protocol [6] are also included for comparison. Extracted optimal protocols corresponding to two histograms denoted by ① and ② are shown alongside the bar plot, illustrating an increasing divergence from the equivalent CPMG sequence as coherence decreases. Initial coherence refers to the value of coherence for a 32-pulse CPMG sequence, absolute enhancement refers to the increased coherence at the same point in time achieved with a UDD and optimized sequence. If the qubit had an initial coherence of 0.5 at 100 μ s with a CPMG-32 sequence and had a coherence of 0.8 at 100 μ s with an optimized sequence, then the absolute enhancement is 0.3.

knowledge of the noise spectrum can be used to extend coherence times through bespoke dynamical decoupling sequences [8,28], while for spin defects in solid-state systems, the noise spectrum has been shown to provide information on parameters such as defect depth [23]. Figure 7 illustrates both such applications.

Using an approach based on the δ -function approximation, Romach *et al.* find evidence for a double Lorentzian form for $S(\omega)$ for shallow N-V centers in diamond, attributed to the presence of two distinct noise sources with differing correlation times. The faster correlation time is attributed to surface-modified phonons and the slower to spin-spin coupling between a bath of surface spins. The coupling strength, Δ , to each of these noise sources decreases as defect depth increases. Our technique provides a potential method for accurately estimating defect depth through simple coherence-decay measurements, at a much lower experimental cost than previous approaches. Although, the presented noise spectra in Fig. 7(a) are artificial, the objective here is to show that the precise extraction of parameters corresponding to the underlying noise spectrum such as coupling strengths [in line with Eq. (5)] provides an excellent means to detect qubit’s surface proximity as sketched in Fig. 7(b).

Figures 7(c)–7(e) show the application of dynamical decoupling optimization to three different classes of noise spectra. Using the sequential least-squares programming (SLSQP) minimization technique [51], provided by the SciPy library [52], we significantly enhance residual coherence over what can be achieved using a 32-pulse CPMG or UDD sequence. The SLSQP algorithm allows for the constrained optimization of π -pulse position, starting from 32-pulse CPMG positions, for a given total sequence time to minimize the value of $\chi(t)$, which in turn maximizes the value of $C(t)$. The optimization process

is constrained to ensure pulse ordering is preserved and to prevent overlaps between pulses. Underlying noise spectra are used to synthesize coherence-decay curves under 32-pulse CPMG, and to generate bespoke pulse sequences (two examples of which are illustrated) that enhance the coherence at specific points in time. While UDD offers some limited enhancement in coherence over CPMG, substantial increases are seen using the optimized sequences, up to values of 4–8 \times .

Whilst the results presented here are promising, the potential avenues for improvement are equally clear. The technique currently requires a specific input dimension and time scale, though an encoder-decoder architecture, as seen in sequence-to-sequence models and transformers, provides a possible solution [53,54]. In addition, whilst the optimization techniques described above appear successful, the application of deep-learning techniques to optimizing qubit coherence times is a promising route for exploration. Furthermore, recent results have shown the successful application of deep reinforcement learning to quantum-sensing problems, suggesting they may well have application in this field [21,22]. It would also be instructive to explore the impact of pulse errors on the noise-spectrum coherence-decay relationship. Pulse errors are unavoidable in experiments and being able to account for their effects would add further applicability to the techniques developed here. We anticipate that the performing experiments under the influence of artificial noise sources and known pulse errors will allow greater insight into these effects. In addition, there has been recent theoretical work to explore the generation of coherence decays from noise spectra when pulse errors are present [55]. Incorporating these techniques into the training data generation process detailed above would allow the training of networks where such pulse errors are accounted for.

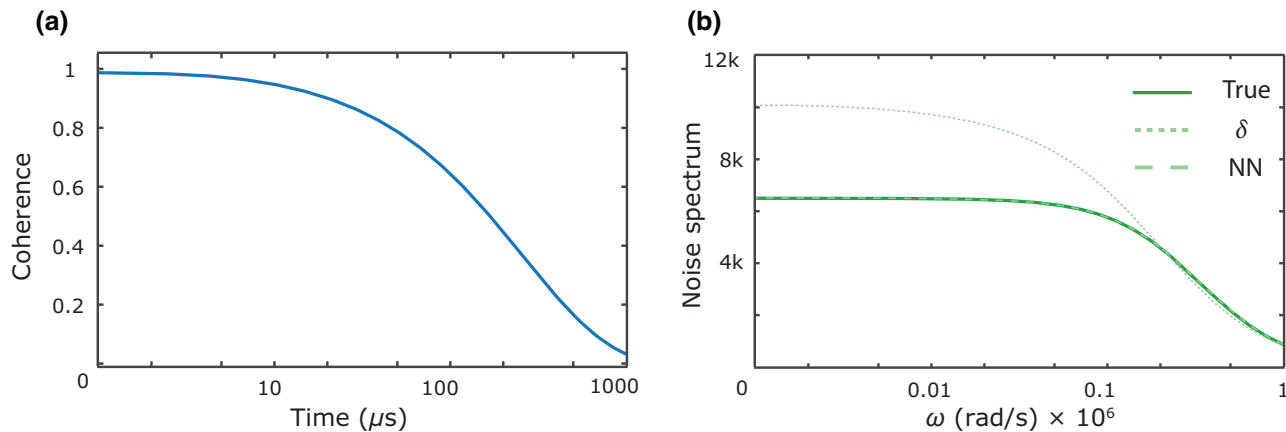


FIG. 8. Neural network versus δ -function approximation performance for inferring the Lorentzian noise spectrum shown in (b) from the coherence-decay curve shown in (a). The true (target) noise spectrum is shown in solid green in (b), the δ -function approximation result is shown by a dotted pale-green line, the neural-network result is shown by a dashed pale-green line.

In summary, we demonstrate a multifaceted toolbox that uses neural networks to accurately deduce the environmental noise spectrum of a qubit. We show that an LSTM network, properly trained with diverse training data, is capable of predicting the precise noise spectrum from a coherence decay curve recorded using dynamical decoupling protocols comprising variable numbers of π pulses. To treat experimental data with measurement noise, we again employ an LSTM network to perform effective denoising that preserves the original functional form of the decoherence curve. This unprejudiced way of smoothing noisy experimental data avoids the use of a predetermined fitting function, enabling more accurate reconstruction of the qubit environment. Using these

deep-learning techniques one can even extract information about multiple noise sources, which can then be used to infer salient properties of the qubit under investigation. Finally, we show how the extracted noise spectrum can be used to generate a customized dynamical decoupling sequence that enhances coherence time significantly beyond what can be achieved with standard protocols. We survey a variety of possible avenues related to the application of deep learning for problems related to qubit decoherence and establish that approach has several potential applications in different quantum technologies. Despite this promise, these concepts remain in their infancy and there is much scope for further improvement and development.

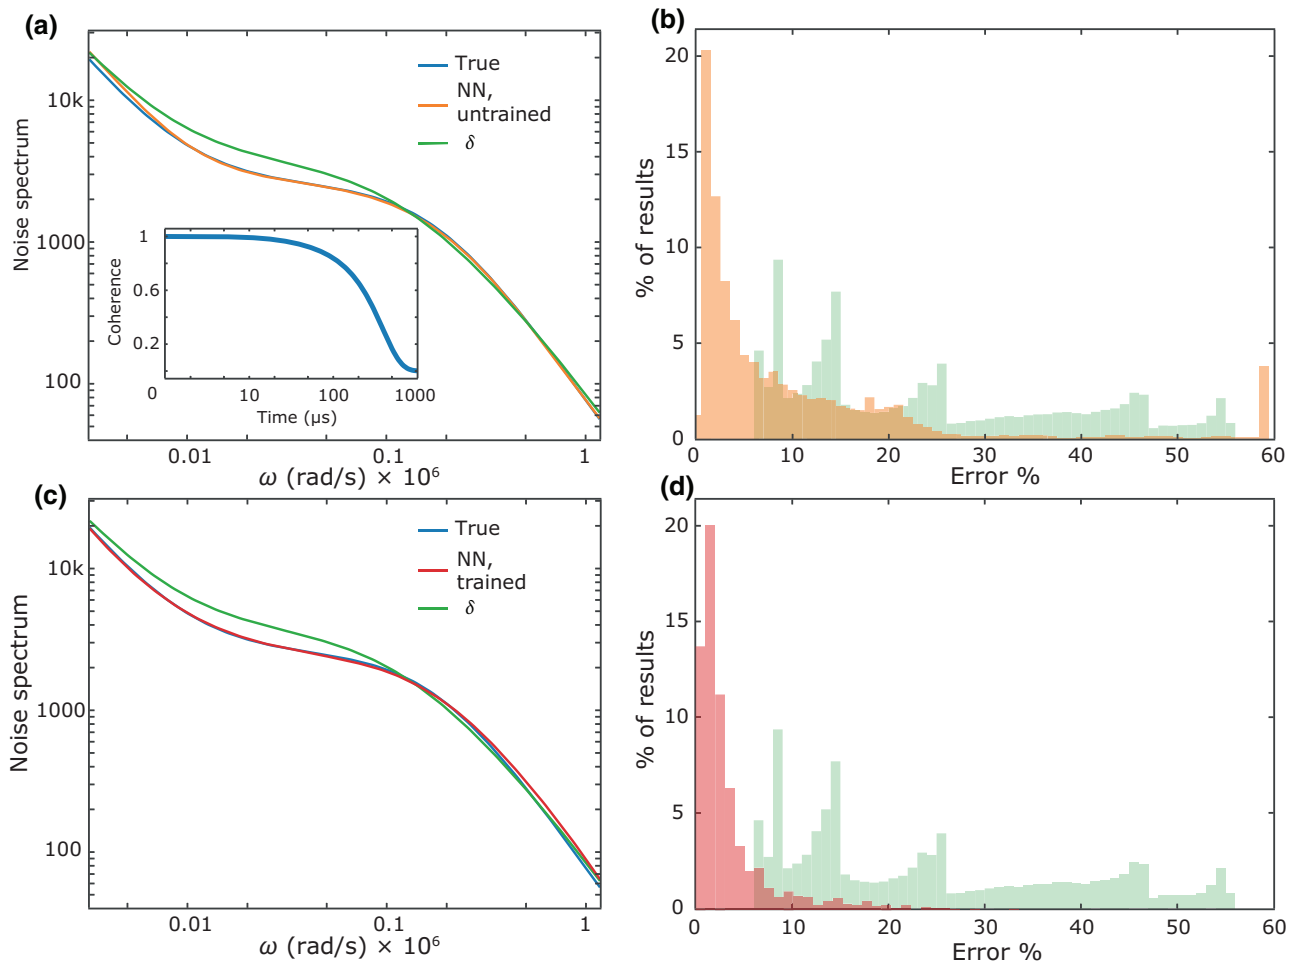


FIG. 9. Neural-network performance for inferring noise spectra from coherence-decay curves derived from $1/f +$ Lorentzian noise, having been trained only on $1/f$, Lorentzian, and double Lorentzian data. (a) example target noise spectrum in blue, resulting in the coherence-decay curve in the inset, with the approximate noise spectra derived from the δ -function approximation and the untrained neural network. (b) The distribution of errors across a set of data of this functional form. The mean error for the untrained neural network is 11.9% whilst the mean error for the δ -function approximation is 24.4%. The distribution of the untrained neural network has a long tail due to noise spectra with large portions close to 0, where mean percentage error is magnified. (c),(d) The same data, but with the neural network now fine tuned on a small $1/f +$ Lorentzian dataset (entirely separate from the data upon which it is evaluated to maintain train-test split). In this case the mean error drops to 3.5%.

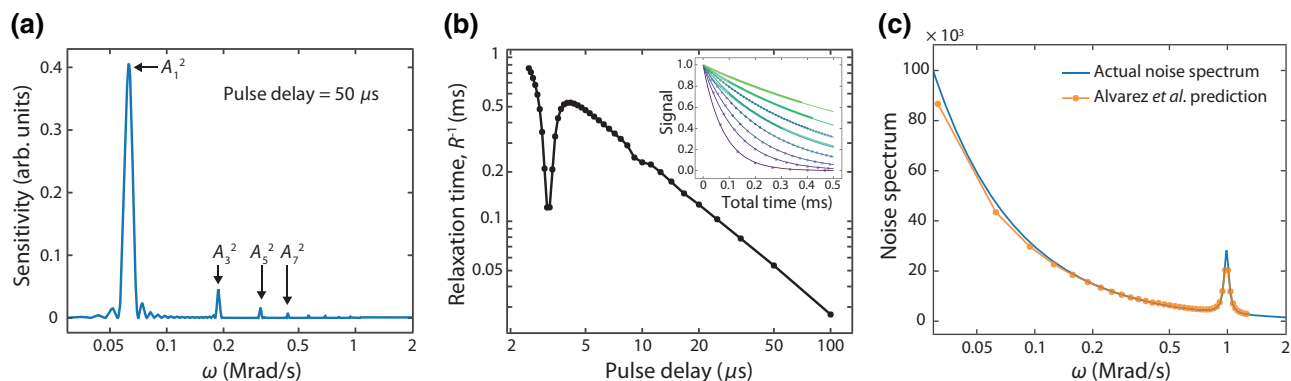


FIG. 10. Application of Alvarez-Suter methodology to the simulated data to predict the underlying noise spectrum. (a) Plot demonstrating how A_k s are associated with the sensitivities of the filter function corresponding to a given pulse delay. (b) Relaxation times extracted via fitting the signal at various pulse delays. The inset shows the actual signal for select pulse delays. Each data point here corresponds to the value of coherence recorded for the given pulse delay and the total duration of the sequence; the total duration can be varied by changing the number of π pulses present in the protocol. (c) Comparison between the actual noise spectrum and the predicted noise spectrum.

ACKNOWLEDGMENTS

We thank Prasanna Vaidya, Gyde, India, for the fruitful discussions on neural networks. This research has received funding from the Engineering and Physical Sciences Research Council (EPSRC) via the Centre for Doctoral Training in Delivering Quantum Technologies (EP/L015242/1), as well as from the European Research Council (ERC) via the LOQO-MOTIONS grant (H2020-EU.1.1., Grant No. 771493).

APPENDIX A: ALTERNATIVE COMPARISON BETWEEN δ -FUNCTION-BASED VERSUS NEURAL-NETWORK-DRIVEN APPROACHES

Figure 3 in the main text compares the performance of the δ -function approximation and the neural-network approach at the 50th percentile in (c) and (d). For added transparency, we include here the performance of both approaches for the reconstruction of the same Lorentzian noise spectrum in Fig. 8.

APPENDIX B: NETWORK PERFORMANCE ON UNSEEN NOISE FORMS

To evaluate the generality of the neural networks used here and their potential for application to potentially novel qubit platforms with unknown noise forms, we take the neural network used to evaluate noise spectra of forms $1/f$, Lorentzian, and double Lorentzian in the main text and apply it to coherence decays derived from noise spectra with $1/f + \text{Lorentzian}$ form. This is performed with no further training of the network. The results are shown below in Fig. 9. In addition, the figure shows the performance of the same network after being trained on a set of data with the target $1/f + \text{Lorentzian}$ form.

Notably the untrained neural network performs significantly better than the δ -function-approximation approach, with a mean error of 11.9% vs 24.4%. The neural-network error distribution has a long tail due largely to noise spectra with significant portions valued close to 0, where mean percentage error is magnified.

To evaluate the network performance further, a small amount of data is generated and the network trained further (fine tuned) and its performance evaluated once more. After this extra training, mean error for the network is reduced to 3.5% and the tail of the distribution becomes significantly shorter. Generation of this extra data takes approximately 10 min, whilst additional training requires 15 min. This extra step is taken to highlight the adaptability and practicality of the approach described in the paper. If a new noise form is identified that the network is not trained to reconstruct, then only a small overhead is required to fine tune the network for this task, with the expectation of greatly increased accuracy of noise-spectrum reconstruction.

APPENDIX C: ALVAREZ-SUTER METHODOLOGY

We provide explicit illustration of the approach proposed by Alvarez *et al.* [9] in Fig. 10. Note that an additional improvement in the noise-spectrum prediction can be obtained if an approximate functional form of the noise spectrum can be derived. However, this approach is not readily generalizable and thus Eq. (13) in the main paper is used here for performing the calculations.

- [1] E. Hahn, Spin echoes, *Phys. Rev.* **80**, 580 (1950).
- [2] H. Y. Carr and E. M. Purcell, Effects of diffusion on free precession in nuclear magnetic resonance experiments, *Phys. Rev.* **94**, 630 (1954).

- [3] S. Meiboom and D. Gill, Modified spin-echo method for measuring nuclear relaxation times, *Rev. Sci.* **29**, 688 (1958).
- [4] L. Viola and S. Lloyd, Dynamical suppression of decoherence in two-state quantum systems, *Phys. Rev. A* **58**, 2733 (1998).
- [5] K. Khodjasteh and D. A. Lidar, Fault-Tolerant Quantum Dynamical Decoupling, *Phys. Rev. Lett.* **95**, 180501 (2005).
- [6] G. S. Uhrig, Keeping a Quantum Bit Alive by Optimized π -Pulse Sequences, *Phys. Rev. Lett.* **98**, 100504 (2007).
- [7] F. Poggiali, P. Cappellaro, and N. Fabbri, Optimal Control for One-Qubit Quantum Sensing, *Phys. Rev. X* **8**, 21059 (2018), [arXiv:1712.08256](https://arxiv.org/abs/1712.08256).
- [8] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, Optimized dynamical decoupling in a model quantum memory, *Nature* **458**, 996 (2009).
- [9] G. A. Álvarez and D. Suter, Measuring the Spectrum of Colored Noise by Dynamical Decoupling, *Phys. Rev. Lett.* **107**, 230501 (2011).
- [10] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 45002 (2019).
- [11] N. Borodinov, S. Neumayer, S. V. Kalinin, O. S. Ovchinnikova, R. K. Vasudevan, and S. Jesse, Deep neural networks for understanding noisy data applied to physical property extraction in scanning probe microscopy, *Npj Comput. Mater.* **5**, 1 (2019).
- [12] X. M. Zhang, Z. Wei, R. Asad, X. C. Yang, and X. Wang, When does reinforcement learning stand out in quantum control? A comparative study on state preparation, *Npj Quantum Inf.* **5**, 1 (2019).
- [13] X. Qu, Y. Huang, H. Lu, T. Qiu, D. Guo, T. Agback, V. Orekhov, and Z. Chen, Accelerated nuclear magnetic resonance spectroscopy with deep learning, *Angew. Chem. Int. Ed.* **59**, 1 (2020).
- [14] C. Zhang and X. Wang, Spin-qubit noise spectroscopy from randomized benchmarking by supervised learning, *Phys. Rev. A* **99**, 042316 (2019).
- [15] A. T. Taguchi, E. D. Evans, S. A. Dikanov, and R. G. Griffin, Convolutional neural network analysis of two-dimensional hyperfine sublevel correlation electron paramagnetic resonance spectra, *J. Phys. Chem. L.* **10**, 1115 (2019).
- [16] S. G. Worswick, J. A. Spencer, G. Jeschke, and I. Kuprov, Deep neural network processing of DEER data, *Sci. Adv.* **4**, 1 (2018).
- [17] A. Youssry, G. A. Paz-Silva, and C. Ferrie, Beyond Quantum Noise Spectroscopy: modelling and mitigating noise with quantum feature engineering, [arXiv:2003.06827](https://arxiv.org/abs/2003.06827) (2020).
- [18] M. August and X. Ni, Using recurrent neural networks to optimize dynamical decoupling for quantum memory, *Phys. Rev. A* **95**, 012335 (2017).
- [19] S. Mavadia, V. Frey, J. Sastrawan, S. Dona, and M. J. Biercuk, Prediction and real-time compensation of qubit decoherence via machine learning, *Nat. Commun.* **8**, 1 (2017).
- [20] R.-B. Wu, H. Ding, D. Dong, and X. Wang, Learning robust and high-precision quantum controls, *Phys. Rev. A* **99**, 042327 (2019).
- [21] L. J. Fiderer, J. Schuff, and D. Braun, Neural-Network Heuristics for Adaptive Bayesian Quantum Estimation, [arXiv:2003.02183](https://arxiv.org/abs/2003.02183).
- [22] J. Schuff, L. J. Fiderer, and D. Braun, Improving the dynamics of quantum sensors with reinforcement learning, *New J. Phys.* **22**, 035001 (2020).
- [23] Y. Romach, C. Müller, T. Unden, L. J. Rogers, T. Isoda, K. M. Itoh, M. Markham, A. Stacey, J. Meijer, S. Pez-zagna, B. Naydenov, L. P. McGuinness, N. Bar-Gill, and F. Jelezko, Spectroscopy of Surface-Induced Noise Using Shallow Spins in Diamond, *Phys. Rev. Lett.* **114**, 017601 (2015).
- [24] T. Yuge, S. Sasaki, and Y. Hirayama, Measurement of the Noise Spectrum Using a Multiple-Pulse Sequence, *Phys. Rev. Lett.* **107**, 170504 (2011).
- [25] N. Bar-Gill, L. M. Pham, C. Belthangady, D. Le Sage, P. Cappellaro, J. R. Maze, M. D. Lukin, A. Yacoby, and R. Walsworth, Suppression of spin-bath dynamics for improved coherence of multi-spin-qubit systems, *Nat. Commun.* **3**, 858 (2012).
- [26] S. Hernández-Gómez, F. Poggiali, P. Cappellaro, and N. Fabbri, Noise spectroscopy of a quantum-classical environment with a diamond qubit, *Phys. Rev. B* **98**, 214307 (2018).
- [27] H. Uys, M. J. Biercuk, and J. J. Bollinger, Optimized Noise Filtration through Dynamical Decoupling, *Phys. Rev. Lett.* **103**, 040501 (2009).
- [28] Ł. Cywiński, R. M. Lutchyn, C. P. Nave, and S. Das Sarma, How to enhance dephasing time in superconducting qubits, *Phys. Rev. B* **77**, 174509 (2008).
- [29] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J. S. Tsai, and W. D. Oliver, Noise spectroscopy through dynamical decoupling with a superconducting flux qubit, *Nat. Phys.* **7**, 565 (2011).
- [30] J. Yoneda, K. Takeda, T. Otsuka, T. Nakajima, M. R. Delbecq, G. Allison, T. Honda, T. Kodera, S. Oda, Y. Hoshi, N. Usami, K. M. Itoh, and S. Tarucha, A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%, *Nat. Nanotechnol.* **13**, 102 (2018).
- [31] J. Medford, Ł. Cywiński, C. Barthel, C. M. Marcus, M. P. Hanson, and A. C. Gossard, Scaling of Dynamical Decoupling for Spin Qubits, *Phys. Rev. Lett.* **108**, 086802 (2012).
- [32] J. R. Klauder and P. W. Anderson, Spectral diffusion decay in spin resonance experiments, *Phys. Rev.* **125**, 912 (1962).
- [33] K. Hornik, Some new results on neural network approximation, *Neural Netw.* **6**, 1069 (1993).
- [34] G. B. Huang, L. Chen, and C.-k. Siew, Universal approximation using incremental constructive feedforward networks with random hidden nodes, *IEEE Trans. Neural Netw. Learn. Syst.* **17**, 879 (2006).
- [35] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016), <http://www.deeplearningbook.org>.
- [36] K. Fukushima, Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition

- unaffected by shift in position, *Biol. Cybern.* **36**, 193 (1980).
- [37] B. A. Pearlmutter, Learning state space trajectories in recurrent neural networks, *IJCNN Int Jt Conf Neural Network*, 365 (1989).
- [38] S. Hochreiter and J. Schmidhuber, Long short-term memory, *Neural Comput.* **9**, 1735 (1997).
- [39] C. T. C. Arsene, R. Hankins, and H. Yin, in *2019 27th European Signal Processing Conference (EUSIPCO)* (IEEE, New York, NY, 2019), p. 1.
- [40] J. Chen, G. Q. Zeng, W. Zhou, W. Du, and K. D. Lu, Wind speed forecasting using nonlinear-learning ensemble of deep learning time series prediction and extremal optimization, *Energy Convers. Manag.* **165**, 681 (2018).
- [41] K. Antczak, Deep Recurrent Neural Networks for ECG Signal Denoising, [arXiv:1807.11551](https://arxiv.org/abs/1807.11551).
- [42] C. Olah, Understanding LSTM Networks (2015).
- [43] W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K.-R. Muller, Explainable AI: Interpreting, explaining and visualizing deep learning, *Lect. Notes Comput. Sci.* 11700, 435 (2019), [arXiv:1606.06565](https://arxiv.org/abs/1606.06565).
- [44] F. Chollet, Keras (2015), <https://keras.io>
- [45] M. Abadi *et al.*, TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems, [arXiv:1603.04467](https://arxiv.org/abs/1603.04467) (2016).
- [46] L. N. Smith, A disciplined approach to neural network hyper-parameters: Part 1 – learning rate, batch size, momentum, and weight decay, [arXiv:1803.09820](https://arxiv.org/abs/1803.09820).
- [47] L. Biewald, *Experiment Tracking with Weights and Biases*, Tech. Rep. (2020).
- [48] M. Loretz, J. M. Boss, T. Roskopf, H. J. Mamin, D. Rugar, and C. L. Degen, Spurious Harmonic Response of Multipulse Quantum Sensing Sequences, *Phys. Rev. X* **5**, 021009 (2015), [arXiv:1412.5768](https://arxiv.org/abs/1412.5768).
- [49] L. M. Norris, D. Lucarelli, V. M. Frey, S. Mavadia, M. J. Biercuk, and L. Viola, Optimally band-limited spectroscopy of control noise using a qubit sensor, *Phys. Rev. A* **98**, 032315 (2018).
- [50] Y. Romach, A. Lazarev, I. Avrahami, F. Kleißler, S. Arroyo-Camejo, and N. Bar-Gill, Measuring Environmental Quantum Noise Exhibiting a Nonmonotonic Spectral Shape, *Phys. Rev. Appl.* **11**, 014064 (2019).
- [51] D. Kraft, Algorithm 733: TOMP—fortran modules for optimal control calculations, *ACM Trans. Math. Softw.* **20**, 262 (1994).
- [52] P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in python, *Nat. Methods* **17**, 261 (2020).
- [53] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, in *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., Cambridge, MA, 2017), p. 5998.
- [54] N. Wu, B. Green, X. Ben, and S. O'Banion, Deep Transformer Models for Time Series Forecasting: The Influenza Prevalence Case, [arXiv:2001.08317](https://arxiv.org/abs/2001.08317).
- [55] L. Z. He, M. C. Zhang, C. W. Wu, Y. Xie, W. Wu, and P. X. Chen, Effects of imperfect pulses on dynamical decoupling using quantum trajectory method, *Chin. Phys. B* **27**, 120303 (2018).