Random telegraph signal analysis with a recurrent neural network

N. J. Lambert,[1,2,*] A. A. Esmail,[3] M. Edwards,[3] A. J. Ferguson,[3] and H. G. L. Schwefel[1,2]
[1]*Department of Physics, University of Otago, Dunedin 9016, New Zealand*
[2]*The Dodd-Walls Centre for Photonic and Quantum Technologies, University of Otago, Dunedin 9016, New Zealand*
[3]*Microelectronics Group, Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, United Kingdom*

(Received 13 February 2020; accepted 30 June 2020; published 21 July 2020)

We use an artificial neural network to analyze asymmetric noisy random telegraph signals, and extract underlying transition rates. We demonstrate that a long short-term memory neural network can outperform other methods, particularly for noisy signals and measurements with limited bandwidths. Our technique gives reliable results as the signal-to-noise ratio approaches one, and over a wide range of underlying transition rates. We apply our method to random telegraph signals generated by quasiparticle poisoning in a superconducting double dot, allowing us to extend our measurement of quasiparticle dynamics to new temperature regimes.

DOI: 10.1103/PhysRevE.102.012312

I. INTRODUCTION

An asymmetric random telegraph signal (RTS) is a signal which exhibits stochastic switches between two levels $y = y_1$ and $y = y_2$, with the asymmetry coming from unequal dwell times at the two levels. They are a common result of measurements on a wide variety of physical systems, including ion channels in cells [1], semiconductor devices such as transistors [2,3], quantum dots [4,5] and optoelectronic devices [6], high-T_c superconductors [7], and single-Cooper-pair boxes [8], as well as being the building block of $1/f$ noise [9]. The transition rates from 1 (2) to 2 (1), $\Gamma_{1(2)}$, where in general $\Gamma_1 \neq \Gamma_2$, are the accessible parameters describing the dynamics of the underlying system, and it is often desirable to extract them from the measured time sequence.

The most straightforward way to do this is to sample the time domain signal at some rate f_s, divide it into periods in each of states 1 and 2 [Fig. 1(a)], histogram the dwell times $\tau_{1(2)}$, and fit $k\, e^{-\Gamma_{1(2)}\tau_{1(2)}}$ to the resulting distribution. However, the presence of noise and a finite measurement bandwidth can result in the measured statistics not representing the underlying system accurately [10]. The problem is twofold: Noise whilst in one state can result in a *false* time period in the other state being detected [Fig. 1(b)], and a limited bandwidth means that *genuine* short-period excursions to the other state are not seen [Fig. 1(c)]. This later effect also joins together the two time periods on either side of the missed period, resulting in a *false* long period.

A variety of solutions to this problem have been proposed [11]. Some focus on optimizing the threshold at which the signal is divided into states 1 and 2 [12]. Naaman and Aumentado modeled the detector as a separate process [10], and derived corrections to the measured rates. Other techniques include wavelet edge detection [13], autocorrelation methods [14], cross correlation methods [15], hidden Markov models [16], measurement of the factorial cumulants of the signal [17], and analysis of the probability density function of the signal [18–20].

In this paper, we demonstrate that a recurrent neural network can be used to extract the underlying rates from noisy, bandwidth limited random telegraph signals. Artificial neural networks (ANNs) comprise an input layer holding the data to be analyzed, an output layer giving the result of the analysis, and one or more "hidden" layers of any number of nodes [21]. The connections between nodes within the layer, and between the subsequent layers, can be arranged in a wide variety of topologies, with different topologies having varying performance characteristics for different tasks. The weights of the internode connections are tuned using gradient descent methods to achieve the optimal mapping between the possible input vectors and the desired output vectors, a process referred to as "training." Neural networks have proven to be a versatile tool, capable of tackling diverse problems including image recognition, financial fraud detection, and natural language processing. They have also proven useful for the physical sciences, and have been applied to astrophysical images [22,23], meta-analysis of the scientific literature [24], the generation of quantum error correction algorithms [25], and the characterization of quantum dots [26] and dopants [27] in semiconductors.

II. ANN ARCHITECTURE

Recurrent ANNs are a class of ANN which act sequentially in a particular direction along an input data array, with an internal memory allowing correlations between data points to affect the output. They are therefore particularly suitable for time sequence data. We use a long short-term memory (LSTM) [28] architecture, a recurrent ANN designed for time sequences in which related information can have significant temporal separation, such as speech [29] and handwriting [30] recognition, and musical analysis [31].

Our ANN has an input size of 10^5 time samples, followed by an LSTM layer with internal size 128 (Table I). The subsequent fully connected hidden layers are of sizes 128, 128,

*nicholas.lambert@otago.ac.nz

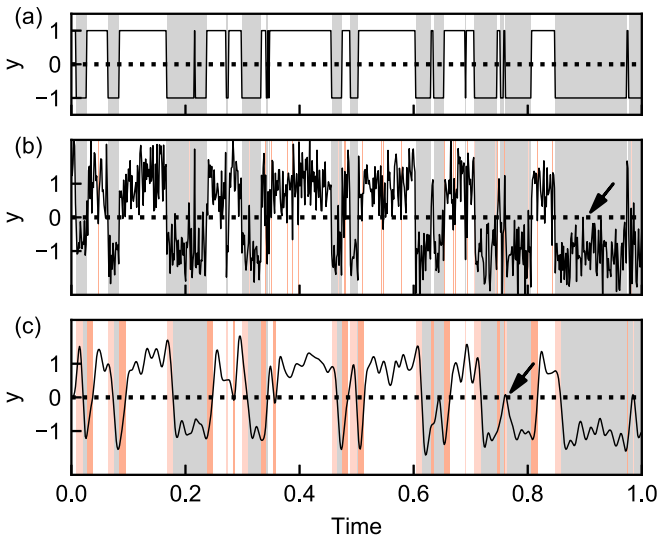

FIG. 1. Filtered noisy RTSs, with apparent periods in each state marked as gray and white. (a) The underlying signal has periods in states 1 and 2, corresponding to signal levels ± 1. (b) The same signal with added Gaussian white noise of standard deviation 0.4. The additional noise creates false excursions from one state into the other (marked in red, example arrowed). (c) The signal in (b) passed through a digital low pass filter. Now short periods in one state can be missed (marked in red, example arrowed).

128, and 10, with exponential linear unit (ELU) activation [32] functions. The output uses rectified exponential linear unit (RELU) activation to ensure a non-negative output value. Although both Γ_1 and Γ_2 are to be determined, the output of the network is chosen to be single valued, and trained to output only Γ_1 for the input $y(t)$. Γ_2 is determined by using the inverted signal, $\bar{y}(t) = -y(t)$, as input for the same ANN. This approach gives faster and more accurate training.

We implement the ANN in Python using Keras [33] (see Supplemental Material [34]). The ANN was typically trained over 250 sets (termed "epochs") of 100 RTSs, with 100 evolutionary steps per epoch, taking around 13.6 h. The ANN is trained using synthetic RTSs, generated with independent rates Γ_1 and Γ_2 uniformly distributed on the interval between $10^{-3} f_s$ and f_s. Rather than $\Gamma_{1,2}$, we train for $\log_{10}(\Gamma_{1,2})$, to compress the output space.

We find that a realistic noise model is necessary for accurate analysis of real data. Two components of additive noise are generated: One with a $1/f$ power spectrum and randomized phase, representing, for example, noise processes in a semiconductor substrate, and one with a flat power spectrum and randomized phase, representing instrumentation noise. The amplitude of the noise added to the training data can be fixed to reflect the measured experimental noise, or varied over a wider distribution. We find that, in general, the ANN cannot be successfully trained if presented with noisy data initially. Instead the training is started with noise free data, and the noise amplitude is increased every 20 training epochs until the desired level is reached. Finally, the generated signals are normalized such that they have mean 0 and standard deviation 1.

TABLE I. Configuration of our neural network.

Layer	Type	Size	Activation
Input		10^5	
1	LSTM	128	Sigmoid
2	Fully connected	128	ELU
3	Fully connected	128	ELU
4	Fully connected	128	ELU
5	Fully connected	10	ELU
Output		1	RELU

III. ANN VALIDATION

Once trained, the effectiveness of the ANN can be tested by applying it to RTS data sequences. In Fig. 2 we compare the results of testing ANNs trained for different noise levels on synthetic data of length 10^5 samples to alternative methods. The ANNs are applied to RTSs with logarithmically spaced transition rates in the range $10^{-3} f_s < \Gamma_{1,2} < f_s$, signal levels $y = \pm 1$, and noise with standard deviations 0 (no additive noise), 0.2, 0.4, and 0.6 (examples in upper panels). In Figs. 2(a)–2(d) we show the results of analysis of 100 RTSs for each rate pair (Γ_1, Γ_2) using a double-sided Page-Hinkley algorithm [35] (middle panels) and using an ANN (lower panels). The red dashed grid marks the underlying rates of the synthesized signals.

The ANNs are more effective for all noise levels. Their accuracy reduces when one or both rates are low ($\Gamma_{1,2} < 3 \times 10^{-3} f_s$), and so there are fewer events to analyze, or equivalently the total power in the signal is low. Signals with rates within a factor of ~ 3 of the sample rate and noises of amplitudes $\geqslant 0.4$ are also challenging, reflecting the similarity between a short period in a particular state and a spike due to noise. Nevertheless, the ANNs still perform well for regimes in which the Page-Hinkley algorithm fails.

In Fig. 3 we study the effect of filtering the RTS. The noise amplitude for synthetic signals is now fixed at 0.4, a noise level at which the age-Hinkley algorithm fails to deal with the unfiltered signal, and they are filtered using a fifth order digital Butterworth filter with 3 dB cutoffs of $f_c = f_s/3.162$, $f_c = f_s/10$, and $f_c = f_s/31.62$, and again analyzed using appropriately trained ANNs for each cutoff frequency. The ANN is generally robust against low pass filtering when $f_c > \Gamma_1 + \Gamma_2$. If the rates are outside this regime, a significant spectral content of the signal is above the passband of the filter, and the ANN cannot extract reliable rates.

We now apply our neural network to RTSs due to Cooper pair breaking in a superconducting double dot (SCDD) [19,36,37]. The device comprises two aluminium superconducting islands coupled to each other by a Josephson junction, and to metallic leads by superconducting-insulator-normal tunnel junctions [Fig. 4(a)]. The state of the device is described by the charge on each island, and is labeled $(q_{\text{left}}/e, q_{\text{right}}/e)$. In the ground state, both q_{left}/e and q_{right}/e are even, due to the Cooper pairing of electrons in the superconducting state.

When in the ground state the anticrossing between the (0,2) and (2,0) charge states gives the device a finite quantum capacitance, but this can be removed by the destruction of

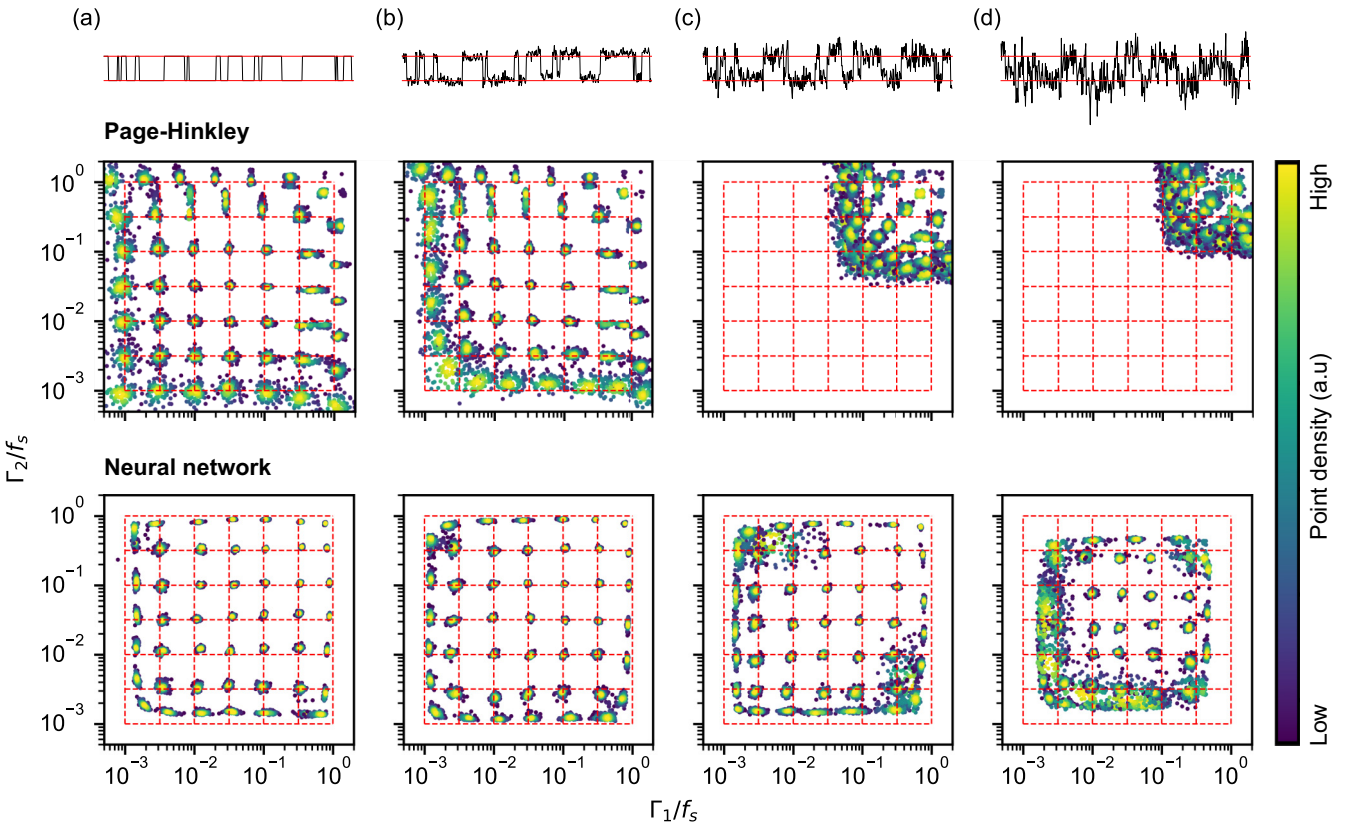

FIG. 2. Application of the trained neural network to synthetic data with varying additive noise. Top panels, example RTSs with levels ± 1 (red lines) and noise of standard deviations (a) 0, (b) 0.2, (c) 0.4, and (d) 0.6. Middle panels, analysis of RTSs with rates $10^{-3} f_s \leqslant \Gamma_{1,2} \leqslant f_s$ using the Page-Hinkley algorithm to extract dwell times, and a fit to an exponential decay. Points are colored according to density in (Γ_1, Γ_2), with lighter points corresponding to higher point densities. The red grid marks the underlying rates of the synthesized data. The addition of moderate amounts of noise causes the method to fail completely. Bottom panels show analysis of the same RTSs using the trained neural networks, showing greatly increased robustness against noise.

coherence due to the presence of unpaired electrons, a process known as quasiparticle poisoning [38–40]. By monitoring the quantum capacitance of the SCDD via radio frequency reflectometry, the loss and recovery of this coherence can be observed. Time domain measurements [Fig. 4(d), upper panel] yield an RTS with one rate determined by the vulnerability of the Cooper pairs to breaking, and the other determined by the recombination dynamics.

Because the transitions between charge configurations are driven by the energy difference between the states, it is interesting to study the transition rates as a function of temperature and applied magnetic field, which affect

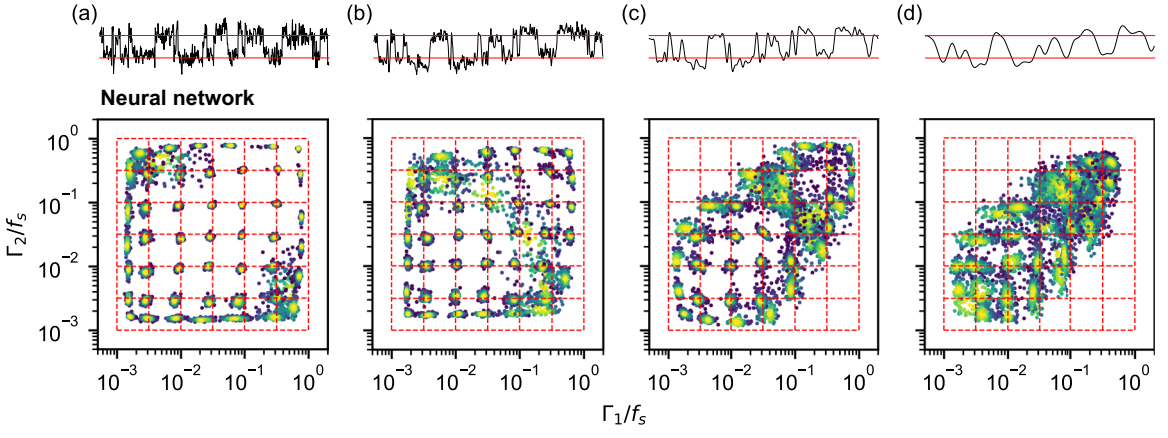

FIG. 3. Application of the trained neural network to synthetic data with noise standard deviation 0.4, passed through digital low-pass filters with cutoff frequencies (a) $f_c > f_s$ [identical to Fig. 2(c)], (b) $f_c = f_s/3.162$, (c) $f_c = f_s/10$, and (d) $f_c = f_s/31.62$. Example time signals are shown in the top panels and extracted rates in the lower panels. Points are colored as in Fig. 2.

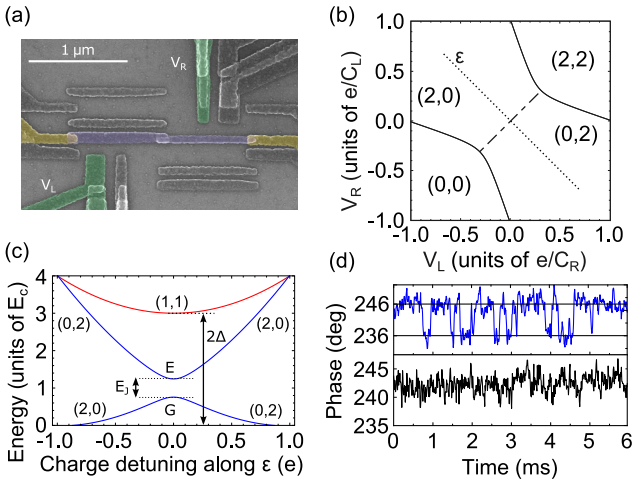

FIG. 4. (a) False color SEM of a superconducting double dot. Purple regions are the superconducting islands, yellow regions are normal-metal grounded reservoirs, and electrostatic control gates are green. Other gates are not used in these experiments. (b) The charge stability diagram, showing the ground state charge configuration as a function of gate voltages V_L and V_R. (c) Charge state energies along ϵ in panel (b). The anticrossing between the (2,0) and (0,2) charge states is mediated by the Josephson energy (E_J) of the interdot tunnel junction, while the additional energy of the (1,1) state is due to the superconducting gap Δ. (d) Top panel: rf reflected phase at $B = 0$, $T = 25$ mK, $\epsilon = 0$, showing an RTS. Bottom panel: rf reflected phase at $B = 180$ mT, $T = 100$ mK, $\epsilon = 0$ with no RTS apparent.

the poisoned state's free energy via changes in the superconducting gap [41]. For low temperatures ($T = 35$ mK) and magnetic fields ($B < 120$ mT), rates can be determined by applying thresholding methods to the measured signal. But at higher temperatures and fields, the quality of the phase signal is sufficiently degraded [Fig. 4(d), lower panel] that this technique does not work.

Time signals with an acquisition time of 1s and comprising 10^5 phase measurements were taken at temperatures of 75 mK, 100 mK, and 125 mK, and at fields of 150 mT $< H <$ 200 mT. The demodulated phase signal was filtered with a cutoff of 15 kHz before sampling. For each field and temperature, measurements were made at points along a line in the charge stability diagram corresponding to transfer of charge from one island to the other, labeled ϵ in Fig. 4(b). Signals were found to have a low signal-to-noise ratio, with no obvious random telegraph behavior [Fig. 4(d), lower panel]. Rates can nevertheless be extracted using our ANN. In Fig. 5(a) we plot the rates at zero detuning for increasing field at three temperatures, and observe the super-exponential behavior previously seen at lower temperatures [19,36].

To further support the validity of the rates extracted using our ANN, we compare the mean phase for each measured signal with the expected mean value for an RTS having extracted rates Γ_1 and Γ_2,

$$\bar{y} = \frac{\Gamma_1}{\Gamma_1 + \Gamma_2}. \tag{1}$$

In each case this is a measure of the excited state occupancy. In Fig. 5(b) we plot the measured mean signals (left) and

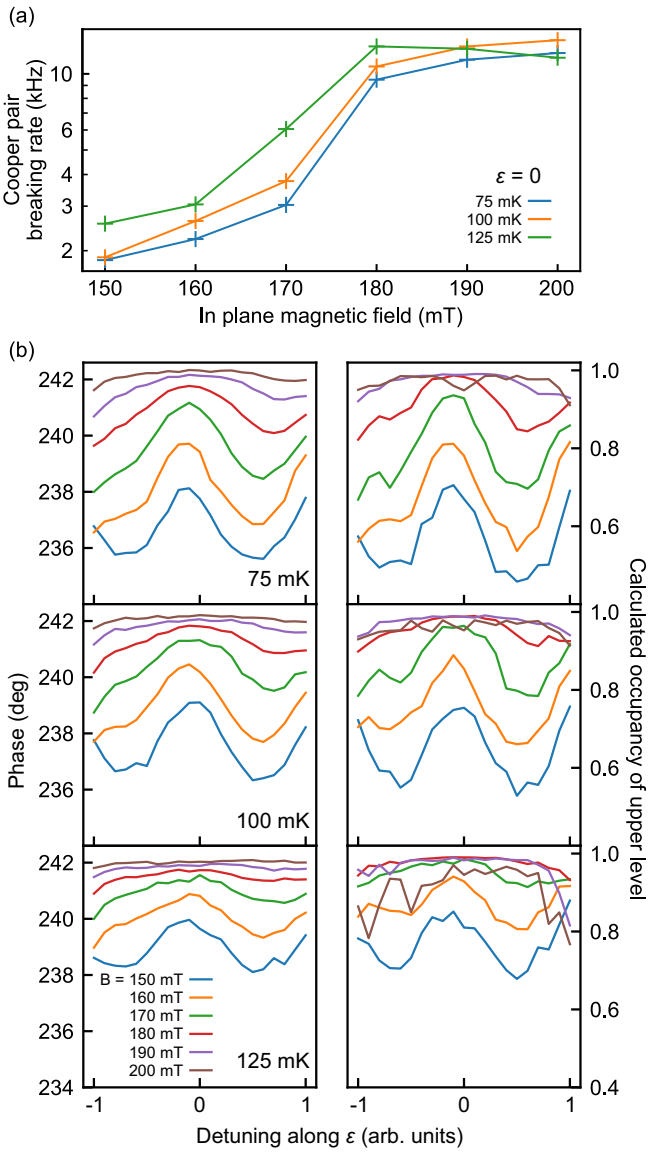

FIG. 5. Application of trained ANN to experimental data from an SCDD. (a) Cooper pair breaking rates (log scale) as a function of applied field for increasing temperatures at zero detuning from the charge state anticrossing. (b) Left panels: Time average phase as a function of ϵ detuning. The measured phase is a measure of the proportion of time spent in the excited state. By fitting the probability density function of a filtered noisy RTS to the measured probability density [18–20] we determine phases of $y_1 = (232.2 \pm 0.1)^\circ$ and $y_2 = (242.1 \pm 0.2)^\circ$ for the ground and excited states. Right panels: The occupancy of the excited state deduced from the rates Γ_1 and Γ_2 extracted by the trained ANN.

deduced excited state occupancy (right). The agreement is excellent, demonstrating the efficacy of our ANN for analysis of experimental data.

IV. CONCLUSION

In summary, we find that an LSTM recurrent neural network is a powerful tool for the determination of the transition rates underlying noisy random telegraph

signals with finite sampling rates. The network architecture presented is simple, yet versatile enough to apply to different signal parameters, and can be trained in a reasonable time on relatively modest hardware. Furthermore, there is scope for increasing the depth of the network, or adapting its architecture, for application to problems in different regimes. This technique has allowed us to analyze previously inaccessible signals, and is particularly useful for measurement of delicate quantum systems for which measurements are difficult, and the SNR likely to be low.

ACKNOWLEDGMENTS

We would like to acknowledge support from the MBIE (NZ) Endeavour Smart Ideas fund, Hitachi Cambridge Laboratory, and UK EPSRC Grant No. EP/K027018/1. A.J.F. was supported by a Hitachi Research fellowship.

[1] B. Hille, *Ion Channels of Excitable Membranes*, 3rd ed. (Sinauer Associates, Sunderland, MA, 2001).
[2] K. Kandiah, IEEE Trans. Electron Devices **41**, 2006 (1994).
[3] K. Kandiah, M. O. Deighton, and F. B. Whiting, J. Appl. Phys. **66**, 937 (1989).
[4] A. L. Efros and M. Rosen, Phys. Rev. Lett. **78**, 1110 (1997).
[5] W. Lu, Z. Ji, L. Pfeiffer, K. W. West, and A. J. Rimberg, Nature (London) **423**, 422 (2003).
[6] X. Wang, P. R. Rao, A. Mierop, and A. J. Theuwissen, in *2006 International Electron Devices Meeting, San Francisco, CA* (IEEE, 2006), pp. 1–4.
[7] G. Jung and B. Savo, J. Appl. Phys. **80**, 2939 (1996).
[8] M. D. Shaw, R. M. Lutchyn, P. Delsing, and P. M. Echternach, Phys. Rev. B **78**, 024503 (2008).
[9] M. J. Kirton and M. J. Uren, Adv. Phys. **38**, 367 (1989).
[10] O. Naaman and J. Aumentado, Phys. Rev. Lett. **96**, 100201 (2006).
[11] D. H. Slichter, Ph.D. thesis, University of California, Berkeley, 2011.
[12] Y. Yuzhelevski, M. Yuzhelevski, and G. Jung, Rev. Sci. Instrum. **71**, 1681 (2000).
[13] J. R. Prance, B. J. V. Bael, C. B. Simmons, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, and M. A. Eriksson, Nanotechnology **26**, 215201 (2015).
[14] J. Martin-Martinez, J. Diaz, R. Rodriguez, M. Nafria, and X. Aymerich, IEEE Electron Device Lett. **35**, 479 (2014).
[15] B. Küng, O. Pfäffli, S. Gustavsson, T. Ihn, K. Ensslin, M. Reinwald, and W. Wegscheider, Phys. Rev. B **79**, 035314 (2009).
[16] M. G. House, H. W. Jiang, and X. C. Zhang, Phys. Rev. B **80**, 113308 (2009).
[17] A. Kurzmann, P. Stegmann, J. Kerski, R. Schott, A. Ludwig, A. D. Wieck, J. König, A. Lorke, and M. Geller, Phys. Rev. Lett. **122**, 247403 (2019).
[18] R. Fitzhugh, Math. Biosci. **64**, 75 (1983).
[19] N. J. Lambert, A. A. Esmail, F. A. Pollock, M. Edwards, B. W. Lovett, and A. J. Ferguson, Phys. Rev. B **95**, 235413 (2017).
[20] S. Singh, E. T. Mannila, D. S. Golubev, J. T. Peltonen, and J. P. Pekola, Appl. Phys. Lett. **112**, 243101 (2018).
[21] J. Schmidhuber, Neural Networks **61**, 85 (2015).
[22] M. Banerji, O. Lahav, C. J. Lintott, F. B. Abdalla, K. Schawinski, S. P. Bamford, D. Andreescu, P. Murray, M. J. Raddick, A. Slosar, A. Szalay, D. Thomas, and J. Vandenberg, Mon. Not. R. Astron. Soc. **406**, 342 (2010).
[23] Y. D. Hezaveh, L. P. Levasseur, and P. J. Marshall, Nature (London) **548**, 555 (2017).
[24] V. Tshitoyan, J. Dagdelen, L. Weston, A. Dunn, Z. Rong, O. Kononova, K. A. Persson, G. Ceder, and A. Jain, Nature (London) **571**, 95 (2019).
[25] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, Phys. Rev. X **8**, 031084 (2018).
[26] S. S. Kalantre, J. P. Zwolak, S. Ragole, X. Wu, N. M. Zimmerman, M. D. Stewart, and J. M. Taylor, npj Quantum Inf. **5**, 6 (2019).
[27] M. Usman, Y. Z. Wong, C. D. Hill, and L. C. L. Hollenberg, npj Comput. Mater. **6**, 19 (2020).
[28] S. Hochreiter and J. Schmidhuber, Neural Comput. **9**, 1735 (1997).
[29] A. Graves, A.-r. Mohamed, and G. Hinton, in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (IEEE, New York, 2013), pp. 6645–6649.
[30] A. Graves, M. Liwicki, H. Bunke, J. Schmidhuber, and S. Fernández, in *Advances in Neural Information Processing Systems 20*, edited by J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis (Curran Associates, Inc., San Jose, CA, 2008), pp. 577–584.
[31] D. Eck and J. Schmidhuber, in *Artificial Neural Networks—ICANN 2002*, edited by G. Goos, J. Hartmanis, J. van Leeuwen, and J. R. Dorronsoro (Springer Berlin Heidelberg, Berlin, Heidelberg, 2002), Vol. 2415, pp. 284–289.
[32] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, arXiv:1511.07289.
[33] F. Chollet *et al.*, https://keras.io.
[34] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevE.102.012312 for description of implementation code.
[35] D. V. Hinkley, Biometrika **58**, 509 (1971).
[36] N. J. Lambert, M. Edwards, A. A. Esmail, F. A. Pollock, S. D. Barrett, B. W. Lovett, and A. J. Ferguson, Phys. Rev. B **90**, 140503(R) (2014).
[37] A. A. Esmail, A. J. Ferguson, and N. J. Lambert, Appl. Phys. Lett. **111**, 252602 (2017).
[38] O. Naaman and J. Aumentado, Phys. Rev. B **73**, 172504 (2006).
[39] A. J. Ferguson, N. A. Court, F. E. Hudson, and R. G. Clark, Phys. Rev. Lett. **97**, 106603 (2006).
[40] R. Lutchyn, L. Glazman, and A. Larkin, Phys. Rev. B **72**, 014517 (2005).
[41] N. J. Lambert, A. A. Esmail, M. Edwards, F. A. Pollock, B. W. Lovett, and A. J. Ferguson, Appl. Phys. Lett. **109**, 112603 (2016).