

# Efficiently measuring a quantum device using machine learning

1 D. T. Lennon<sup>†1</sup>, H. Moon<sup>†1</sup>, L. C. Camenzind<sup>2</sup>, Liuqi Yu<sup>2</sup>, D. M. Zumbühl<sup>2</sup>, G. A. D. Briggs<sup>1</sup>, M.  
2 A. Osborne<sup>3</sup>, E. A. Laird<sup>4</sup> & N. Ares<sup>1</sup>

3 <sup>1</sup>*Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom*

4 <sup>2</sup>*Department of Physics, University of Basel, 4056 Basel, Switzerland*

5 <sup>3</sup>*Department of Engineering, University of Oxford, Parks Road, Oxford OX1 3PJ, United Kingdom*

6 <sup>4</sup>*Department of Physics, Lancaster University, Lancaster, LA1 4YB, United Kingdom*

7 **Scalable quantum technologies such as quantum computers will require very large num-**  
8 **bers of quantum devices to be characterised and tuned. As the number of devices on chip**  
9 **increases, this task becomes ever more time-consuming, and will be intractable on a large**  
10 **scale without efficient automation. We present measurements on a quantum dot device per-**  
11 **formed by a machine learning algorithm in real time. The algorithm selects the most infor-**  
12 **mative measurements to perform next by combining information theory with a probabilis-**  
13 **tic deep-generative model that can generate full-resolution reconstructions from scattered**  
14 **partial measurements. We demonstrate, for two different current map configurations that**  
15 **the algorithm outperforms standard grid scan techniques, reducing the number of measure-**  
16 **ments required by up to 4 times and the measurement time by 3.7 times. Our contribution**  
17 **goes beyond the use of machine learning for data search and analysis, and instead demon-**  
18 **strates the use of algorithms to automate measurements. This work lays the foundation for**

<sup>†</sup>Both authors contributed equally and are displayed in alphabetical order

20 **Introduction**

21 Semiconductor quantum devices hold great promise for scalable quantum computation. In par-  
22 ticular, individual electron spins in quantum dot devices have already shown long spin coherence  
23 times with respect to typical gate operation times, high fidelities, all-electrical control, and good  
24 prospects for scalability and integration with classical electronics <sup>1</sup>.

25 A crucial challenge of scaling spin qubits in quantum dots is that electrostatic confinement  
26 potentials vary strongly between devices and even in time, due to randomly fluctuating charge traps  
27 in the host material. Characterising such devices, which requires measurements of current or con-  
28 ductance at different applied biases and gate voltages, can be very time consuming. It is normally  
29 carried out following simple scripts such as grid scans, which are sequential measurements taken  
30 from a 2D grid for a pair of voltages. We call a set of voltages that defines the state of a quantum  
31 dot a configuration. Measurement of some configurations is more informative for characterising a  
32 quantum dot than the other configurations; measuring uncertain signals is more informative than  
33 measuring predictable signals. However, grid scans do not prioritise measurement of informative  
34 signals, instead just acquiring measurements according to simple rules (e.g. following a raster pat-  
35 tern). Current efforts in the field of automation of quantum dot are focused on tuning <sup>2-9</sup>, a large  
36 portion of these relying on grid scanning techniques for measurement. An optimised measurement  
37 method that can prioritise and select important configurations is thus key for fast characterisation  
38 and automatic tuning. Our method is thus complementary to automating tuning of quantum devices

39 and holds the potential to increase the efficiency of these approaches when combined.

40 In this paper, we present an algorithm that performs efficient real-time data acquisition for  
41 a quantum dot device (Fig. 1a). It starts from a low-resolution uniform grid of measurements,  
42 creates a set of full-resolution reconstructions, calculates the predicted information gain (i.e. the  
43 acquisition map), selects the most informative measurements to perform next, and repeats this  
44 process until the information gain from new measurements is marginal.

45 In order to select measurements based on information theory, we require a corresponding  
46 uncertainty measure (of random variables)<sup>10-12</sup>, and hence a probabilistic model of unobserved  
47 variables. One typical approach is to use a Gaussian process<sup>13</sup>. Here, we use a conditional vari-  
48 ational auto-encoder (CVAE)<sup>14</sup>, which is capable of generating high-resolution reconstructions  
49 given partial information and is fast enough for real-time decisions. Deep generative models such  
50 as adversarial networks (GAN)<sup>15</sup>, the variational auto-encoder (VAE)<sup>16</sup> and its extensions, such  
51 as CVAE, have shown great success in multi-modal distributions and complex non-stationary pat-  
52 terns of data<sup>17,18</sup>, similar to those of observed in quantum device measurements. These are the  
53 main advantages of CVAE over a basic Gaussian process. Also, CVAE is more computationally  
54 efficient at generating multiple full-resolution reconstructions. Although progress has been made  
55 addressing the limitations of Gaussian processes, deep generative models are overall a better fit to  
56 the requirements for efficient quantum device measurements. Deep generative models have been  
57 used for: speech synthesis<sup>19</sup>; generating images of digits and human faces<sup>20,21</sup>; transferring image  
58 style<sup>22,23</sup>; and inpainting missing regions of images<sup>24</sup>. Recently, VAE models have been used in

59 scientific research to optimise molecular structures<sup>25–28</sup>. In spite of their suitability, these models  
60 have not previously been applied to efficient data acquisition. An advantage of deep generative  
61 models over simple interpolation techniques, such as nearest-neighbour and bilinear interpolation,  
62 is that deep generative models can learn likely patterns from training data and incorporate them  
63 into its reconstructions. Our method, as it is data-driven, it is generalizable to different transport  
64 regimes, measurement configurations, and more complex device architectures if an appropriate  
65 training set is available.

## 66 **Results**

67 **The device** Our device is a laterally defined quantum dot fabricated by patterning Ti/Au gates  
68 over a GaAs/AlGaAs heterostructure containing a two-dimensional electron gas (Fig. 1b). In this  
69 device, electrons are subject to the confinement potential created electrostatically by gate voltages.  
70 Gate voltages  $V_1$  to  $V_4$  tune the tunneling rates while  $V_G$  mainly shifts the electrical potential inside  
71 the quantum dot. The current through the device is determined both by these gate voltages and by  
72 the bias voltage  $V_{\text{bias}}$ . Measurements were performed at 30 mK.

73 The quantum dot is characterised by acquiring maps of the electrical current as a function of a  
74 pair of varied voltages, which we call a current map configuration. We first focus on varying  $V_G$  and  
75  $V_{\text{bias}}$  for fixed values of  $V_1$  to  $V_4$ . Figure 1c) shows a typical example. Diamond-shaped regions or  
76 ‘Coulomb diamonds’ indicate Coulomb blockade, where electron tunnelling is suppressed<sup>29</sup>. Most  
77 current maps have large areas in which the current is almost constant, and consequently measure-

78 ments in these regions slow down informative data acquisition. Our algorithm must therefore give  
79 measurement priority to the informative regions of the map. An overview of an algorithm-assisted  
80 measurement of a current map is shown in Fig. 1d.

81 **Training the reconstruction model** The role of the reconstruction model is to characterise likely  
82 patterns in a training data set, given by a mixture of measured and simulated current maps. We can  
83 utilise these likely patterns to predict the unmeasured signals from partial measurements.

84 Deep generative models represent this pattern characterisation in a low-dimensional real-  
85 valued latent vector  $\mathbf{z}$ , which can be decoded to produce a full-resolution reconstruction. The  
86 latent space representation and the decoding are learned during training. Our CVAE consists of  
87 two convolutional neural networks, an encoder and a decoder. The encoder is trained to map full-  
88 resolution training examples of current maps  $Y$  to the latent space representation  $\mathbf{z}$ . The encoder  
89 also enforces that the distribution  $p(\mathbf{z})$  of training examples in latent space is Gaussian.

90 The decoder is trained to reconstruct  $Y$ , from the representation  $\mathbf{z}$  combined with an  $8 \times 8$   
91 subsample of  $Y$ . As a result,  $\mathbf{z}$  attempts to represent all the information that is missing from the  
92 subsampled data. In a plain VAE, the input of a decoder is only  $\mathbf{z}$ . If a decoder takes additional  
93 input except  $\mathbf{z}$ , then it is called CVAE, and we found that CVAE generates better reconstructions  
94 than VAE for the considered measurements. The chosen loss function, which the CVAE tries to  
95 minimise, is a measure of the difference between the training data and the corresponding recon-  
96 struction. To avoid blurry reconstructions, we define a contextual loss function that incorporates  
97 both pixel-by-pixel and higher-order differences like edges, corners, and shapes. Detailed descrip-

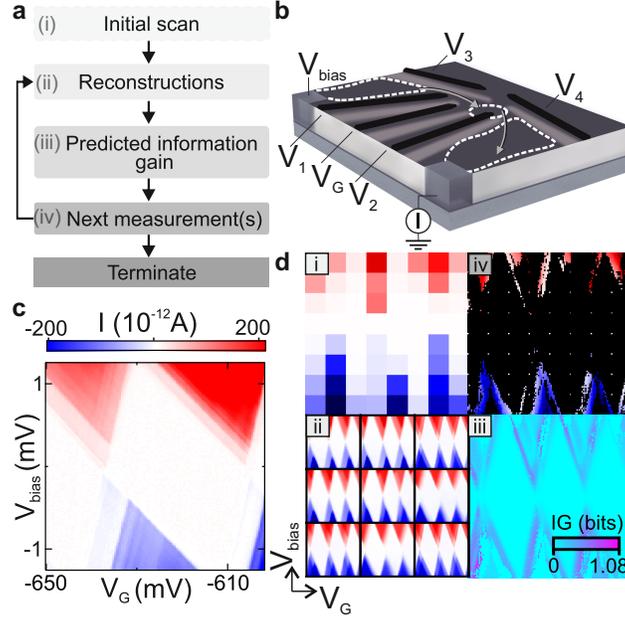


Figure 1: **Overview of the algorithm and the quantum dot device.** **a**, Schematic of the algorithm's operation. Low-resolution measurements (i) are used to produce reconstructions (ii), which are used to infer the predicted information gain acquisition map (iii). Based on this map, the algorithm chooses the location of the next measurement (iv). The process is repeated until a stopping criterion is met. **b**, Schematic of the device. A bias voltage  $V_{\text{bias}}$  is applied between ohmic contacts to the two-dimensional electron gas. We apply gate voltages labelled  $V_1$  to  $V_4$  and  $V_G$ . **c**, A measured current map as a function of  $V_{\text{bias}}$  and  $V_G$ . The Coulomb diamonds are the white regions where electron transport is suppressed, and most of the information necessary to characterise a device is contained just outside these diamonds. **d**, Sequential decision algorithm in **a** illustrated with an example of a specific current map. In panel (iv), unmeasured pixels are plotted in black; however, initial measurements (i) are represented so as to fill the entire panel (that is, the sparse grid of measurements is represented as a low-resolution image).

98 tion of these networks and their training can be found in the Supplementary sections Training and  
99 loss function, and Network specification.

100 The model is trained using both simulated and measured current maps. We choose to work  
101 with current maps of resolution  $128 \times 128$ . The simulation is based on a constant-interaction model  
102 (see Methods). To measure the current maps for training, we set the bias and gate voltage ranges  
103 randomly from a uniform distribution. The training data set consists of 25,000 simulations and  
104 25,000 real examples generated by randomly cropping 750 measured current maps. The current  
105 maps were subjected to random offsets, rescaling, and added noise to increase the variability of the  
106 training set.

107 **Generating reconstructions from partial data** After training, only the trained decoder network  
108 is used in the algorithm of Fig. 1a to reconstruct full-resolution current maps from partial data.  
109 At each stage, the known partial current map is denoted  $Y_n$ , where  $n \leq 128^2 = 16,384$  is the  
110 number of pixels to be measured. To generate a reconstruction, the decoder takes as input the  
111 initial  $8 \times 8$  grid scan  $Y_{64}$ , together with a latent vector  $\mathbf{z}$  sampled from the posterior distribution  
112  $p(\mathbf{z}|Y_n)$  (see Methods for detail equations and Fig. S1 for the decoder diagram). Note that the pos-  
113 terior density is calculated by the prior density  $p(\mathbf{z})$  and a likelihood function, which is comparing  
114 reconstructions and the partial data. Multiple posterior samples are drawn from  $p(\mathbf{z}|Y_n)$  by the  
115 Metropolis-Hastings (MH) method to approximate  $p(\mathbf{z}|Y_n)$ . From these multiple samples  $\mathbf{z}_m$ , cor-  
116 responding reconstructions are then generated, denoted  $\hat{Y}_m$ . In this paper we set  $m = 1, \dots, 100$ .  
117 The continuous posterior  $p(\mathbf{z}|Y_n)$  is then approximated by a discrete posterior of samples  $P_n(m)$ ,

118 which denotes how probable  $\hat{Y}_m$  is. We refer to  $P_n(m)$  as the posterior distribution of reconstruc-  
119 tions.

120 **Making measurement decisions** With each iteration of the decision algorithm, an acquisition  
121 map is computed from the accumulated partial measurements and the resulting reconstructions.  
122 This acquisition map assigns to each potential measurement location (i.e. to each pixel location  
123 in the current map) an information value for the posterior distribution of reconstructions (Fig. 2).  
124 The  $(n + 1)$ th measurement, whose result is  $y_{n+1}$ , is one pixel taken from the true current map and  
125 changes our posterior distribution from  $P_n(m)$  to  $P_{n+1}(m)$ , rendering different reconstructions  
126 more or less probable.

127 The acquisition map is the expected information gain  $\text{IG}(x)$  at each potential measurement  
128 location  $x$ . Our algorithm calculates it by a weighted sum over reconstructions:

$$\text{IG}(x) \equiv \sum_m P_n(m) \times \text{IG}_m(x) \quad , \quad (1)$$

129 where  $\text{IG}_m(x)$  is the Kullback-Leibler divergence between the distributions  $P_n$  and  $P_{n+1}$ , calcu-  
130 lated such that  $y_{n+1}$  at location  $x$  is taken from reconstruction  $\hat{Y}_m$ . The most informative point is  
131  $x_{n+1}^* \equiv \text{argmax}_x \text{IG}(x)$ . This criterion is equivalent both to minimising the expected information  
132 entropy of the posterior distribution and to Bayesian active learning by disagreement (BALD<sup>10</sup>,  
133 see Methods). The difference of the proposed method and BALD is that the proposed method uses  
134 random reconstructions of data, which can be multi-modal, whereas BALD assumes that data is  
135 normally distributed.

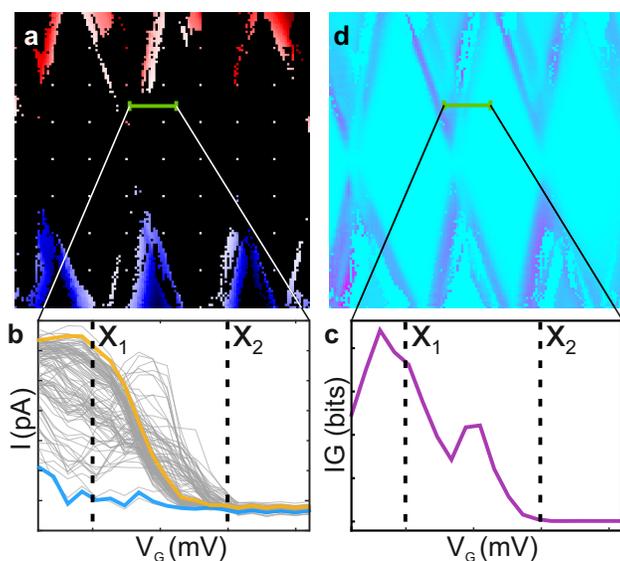


Figure 2: **Computing the acquisition map.** **a**, Partial current map. To illustrate the first step in the computation of the acquisition map, we consider a trace (green) through an unmeasured region of the map. **b**, For the unmeasured trace in **a**, reconstructions provide 100 different predictions. Blue and yellow traces highlight two of these predictions. The objective is to determine the most informative measurement location. At  $x_2$ , all predictions are similar, so measuring here will have little impact on the posterior distribution of reconstructions. At  $x_1$ , predictions are dissimilar and therefore  $x_1$  is a more informative measurement location, with a larger effect on the posterior distribution of reconstructions. **c**, Information gain computed for the unmeasured trace in **a**. **d**, Acquisition map of information gain computed from the partial measurements in **a**, and plotted over the entire current map range.

136 We devised two methods to make decisions based on the acquisition map; a pixel-wise  
137 method, and a batch method. The pixel-wise method selects the single best location in the ac-  
138 quisition map. In practice, this is often not optimal in terms of measurement time, because it  
139 does not take account of the time needed to ramp the gate voltage between measurement locations  
140 (which is limited by details of the measurement electronics and the device settling time). To take  
141 account of this limitation, we also devised a batch method, which selects multiple locations from  
142 the acquisition map, and then acquires measurements by taking a fast route between them. This  
143 reduces the measurement time compared with the pixel-wise method.

144 **Experiments** To test the algorithm, it was used to acquire a series of current maps in real time.  
145 First, the device was thermally cycled, to randomise the charge traps and therefore present the  
146 algorithm with a configuration not represented in its training data. Gate voltages  $V_1$ - $V_4$  were set  
147 to a combination of values, and the algorithm was tasked to measure the corresponding current  
148 map using both the batch and the pixel-wise methods. This step was repeated for ten different  
149 combinations of bias and gate voltages. Fig. 3 presents data acquired by the algorithm at selected  
150 acquisition stages, together with selected reconstructions. As expected, reconstructions become  
151 less diverse as more measurements are acquired. The reconstructions do not necessarily replicate  
152 the measured current map for large  $n$ . This is because reconstructions have a limited variability  
153 given by the training data. Decisions are made based on the learned patterns from the training  
154 data, which implies that this training data should contain at least general patterns which are to be  
155 characterised but does not need to include all possible features in a current map.

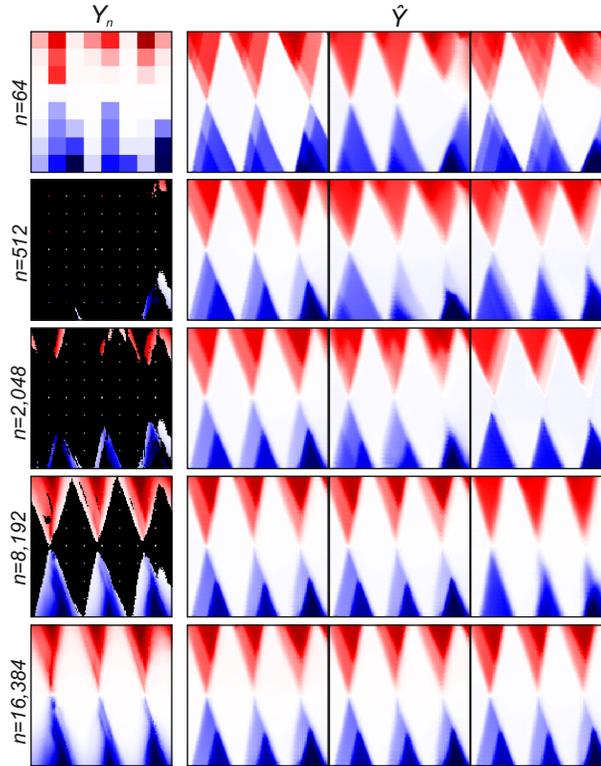


Figure 3: **Updating reconstructions using information from new measurements.** In each row, the first column shows the algorithm-assisted measurements, using the batch method, for a given  $n$ . The remaining three columns contain example reconstructions given the corresponding  $n$  measurements. As  $n$  increases, the diversity of the reconstructions is reduced and their accuracy increased. As expected, the uncertainty is almost eliminated in the last row. The residual remaining variance is because slightly different reconstructions are nearly equally consistent with the data.

156 As seen, the algorithm gives high priority to regions of the map where the current is rapidly  
157 varying, and avoids regions of nearly constant current, such as the interiors of Coulomb diamonds.  
158 This strategy is an emergent property of the algorithm and is wise; little information about the  
159 device characteristics can be found in low-current gradient regions of the current map. This pref-  
160 erence derives from the comparison between reconstructions, which exhibit the greatest disagree-  
161 ment outside Coulomb diamonds. This is also seen in Fig. 4a, which shows two representative  
162 measurement sequences using the batch method. The batch method collects grouped measure-  
163 ments while the pixel-wise method distributes measurements more uniformly, given that in this  
164 case, the acquisition map is more frequently updated to take account for recently acquired infor-  
165 mation. Results for other current maps, including for the pixel-wise method, are shown in the  
166 Supplementary Figures 2 to 6.

167 We compared the performance of the algorithm with an alternating grid scan method. This  
168 type of grid scan starts with  $8 \times 8$  measurements and alternately increases the vertical and the hori-  
169 zontal grid size by 2 (i.e.  $16 \times 8$ ,  $16 \times 16$ ,  $32 \times 16$ , etc.), without performing the same measurement  
170 twice. Over the ten different current maps, the average time for full-resolution data acquisition  
171 with the alternating grid scan method is 554 seconds. This time is limited by our bias and gate  
172 voltage ramp rate and chosen settling time. The batch method can be implemented with any batch  
173 size however for direct comparison with the alternating grid scan we selected increasing batches  
174 of  $32 \times 2^b$ , where  $b$  is the batch number starting from 1.

175 Two types of computation are required to make a measurement decision: sampling recon-

176 structions using the MH method, and constructing the acquisition map. One MH sampling iteration  
 177 takes 63 ms. For experiments, multiple sampling iterations are performed after each batch decision  
 178 and measurement while acquisition is suspended. Since sampling can be performed simultaneously  
 179 with measurement acquisition, from now on our measurement times exclude the time for sampling.  
 180 To compute a single acquisition map takes approximately 50 ms using a NVIDIA GTX 1080 Ti  
 181 graphics card and Tensorflow<sup>30</sup> implementation. The acquisition map must be computed for every  
 182 batch or every pixel measurement, except for the initial  $8 \times 8$  grid scan and the final acquisition  
 183 step (which has no choice of which pixel(s) to measure). To acquire a full resolution current map  
 184 thus requires 7 computations (350 ms) for the batch method, and 16,319 computations (816 s) for  
 185 the pixel-wise method. For the batch method, the computation time is negligible compared to the  
 186 measurement time, but for the pixel-wise method it is a limiting factor in the measurement rate.

187 To quantify the algorithm’s performance, we have devised a measure based on the observa-  
 188 tion that the most informative regions of the current map are those where the current varies strongly  
 189 with  $V_G$  and  $V_{\text{bias}}$ . We therefore define the local gradient of the current map at each location  $x$  as

$$v(x) \equiv \|\nabla Y(x)\|_2 = \sqrt{\left(\frac{\partial I(x)}{\partial V_G}\right)^2 + \left(\frac{\partial I(x)}{\partial V_{\text{bias}}}\right)^2}, \quad (2)$$

190 where  $I(x)$  is current measurement at  $x$ , and the derivatives are calculated numerically. The error  
 191 measure  $r(n)$  of a partial current map is the fraction of the total gradient that remains uncaptured,  
 192 i.e.

$$r(n) \equiv 1 - \frac{V(n)}{V(N)} \quad (3)$$

193 where  $V(n) \equiv \sum_{i=1}^n v(x_i)$  is the total acquired gradient and  $x_i$  is the location of the  $i$ th measure-

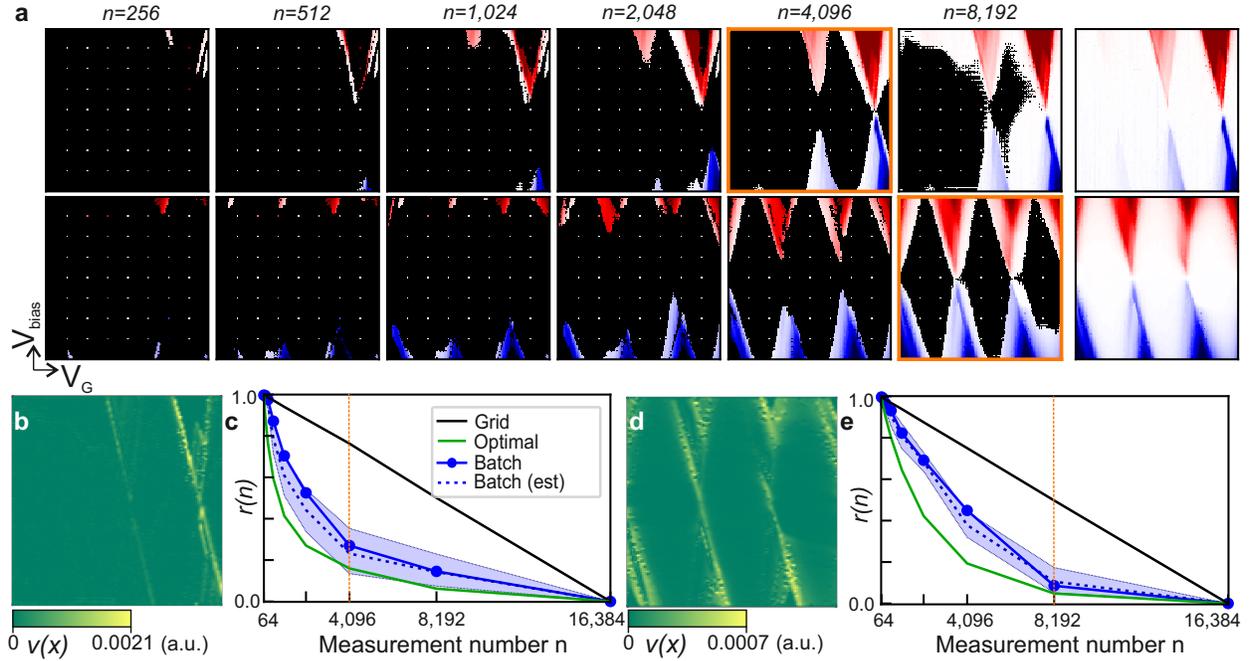


Figure 4: **Measurements of Coulomb diamonds performed by the algorithm.** **a**, Sequential batch measurement in two different experiments. Each row displays algorithm assisted measurements of the current map as a function of  $V_{\text{bias}}$  and  $V_G$  for different values of  $n$ . The last plot in each row is the full-resolution current map. **b**, **d**, Current gradient map (defined by Eq. (2)) for each example in **a**. **c**, **e**, Measure of the algorithm's performance  $r(n)$ , real-time estimate of  $r(n)$  across reconstructions (with 90% credible interval shaded), and optimal  $r(n)$  for both examples in **a**. The black line is the value of  $r(n)$  corresponding to the alternating grid scan method. The vertical orange line indicates the value of  $n$  determined by the stopping criterion. The corresponding current map in **a** is highlighted in orange.

194 ment. This error can only be calculated after all measurements have been performed. However,  
 195 we can utilise the  $m$ th reconstruction to generate an estimate  $\tilde{r}_m(n)$  in real time by replacing  
 196  $\|\nabla Y(x)\|_2$  with  $\|\nabla \hat{Y}_m(x)\|_2$ . The estimates from multiple reconstructions yield a credibility inter-  
 197 val for  $r(n)$ . For an optimal algorithm, the error would be  $\bar{r}(n) = 1.0 - \frac{V^*(n)}{V(N)}$ , where  $V^*(n)$  is the  
 198 sum of the largest  $n$  values of  $v(x)$ . This would be achieved if each measurement location were  
 199 chosen knowing the full-resolution current map, and thus the location of the the highest unmea-  
 200 sured current gradient. No decision method can exceed this bound. For the real time estimates of  
 201  $r(n)$ , we have increased the number of reconstructions to 3,000 by adding different noise patterns  
 202 that are present in typical measured current maps (see Supplementary section Noisy reconstruc-  
 203 tion). This increase in the variability of the reconstructions is needed to avoid an overconfident  
 204 estimation of  $r(n)$ .

205 Performances for two experiments are shown in Fig. 4c, e. Grid scans reduce  $r(n)$  linearly  
 206 with increasing  $n$ . The decision algorithm outperforms a simple grid scan and is nearly optimal.  
 207 When most of the current gradient is localised, the grid scan is far from optimal and even the  
 208 decision algorithm has room for improvement. In this case, the performance of the algorithm is  
 209 determined by how representative the training data is. Quantitative analysis of all 10 experiments  
 210 is in Supplementary Figures 5 and 6.

211 We propose a simple stopping criterion that uses the estimated reduction of the error  $r(n)$  to  
 212 determine when to stop measuring a given current map, in a scenario where more experiments are  
 213 waiting to be conducted. For a given current map containing  $n$  measured pixels, the error after the

214 next measurement batch is estimated for reconstruction  $m$  to be  $\tilde{r}_m(n + \Delta)$ , where  $\Delta$  is the size of  
 215 the batch. Thus the estimated rate at which the error decreases is  $\beta_m \equiv |\tilde{r}_m(n + \Delta) - \tilde{r}_m(n)|/\Delta$ .  
 216 In the worst case among the candidate reconstructions, this rate is  $\beta \equiv \min_m \beta_m$ . However, if  
 217 the algorithm begins to measure a new map, for which no reconstructions yet exist, the error of  
 218 that map will decrease at a rate of at least  $\alpha \equiv 1/N$ ; this is the slope achieved by a grid scan  
 219 and the worst case of the decision algorithm (black lines in Fig. 4c, e). Hence when  $\beta < \alpha$ , it is  
 220 beneficial to halt measurement and move onto a new current map that is awaiting measurement.  
 221 Since  $\alpha$  and  $\beta$  are the worst-case estimates for each case, the criterion is conservative. The stopping  
 222 points by this criterion are shown in Fig. 4c, e, with orange dashed lines. The total average time  
 223 (measurement time plus decision time) to reach the stopping criterion was 237 s, compared with  
 224 554 s to measure the complete current map by grid scan, reducing the time needed by a factor  
 225 between 1.84 and 3.70 across all 10 test cases. A more sophisticated stopping criterion utilising  
 226 the number of remaining unmeasured current maps and a total measurement budget is presented in  
 227 Methods.

228 **Generalising the algorithm** The algorithm described here does not require assumptions about  
 229 the physics of the acquired data, such as requiring that it show Coulomb diamonds. Provided that  
 230 training data are available, it should also work for other kinds of measurements. To test this, we  
 231 applied it to a different current map configuration also encountered in quantum dot tuning. In this  
 232 case the current flowing through the device is measured as a function of two gate voltages ( $V_1$   
 233 and  $V_2$ ), while keeping other voltages fixed ( $V_G$ ,  $V_{\text{bias}}$ ,  $V_3$  and  $V_4$ ). In these current maps, Coulomb  
 234 blockade leads to large areas where the current scarcely changes, with diagonal features of allowed

235 current. For the training set, we measured 382 current maps with a resolution of  $251 \times 251$  which  
236 we randomly cropped to a resolution of  $128 \times 128$  and subjected to simple image augmentation  
237 techniques (as for the previous training set).

238 We tested the performance of the algorithm in this new scenario by taking two different  
239 combinations of  $V_G$ ,  $V_{\text{bias}}$ ,  $V_3$  and  $V_4$  and measuring the corresponding current maps in real time  
240 (Fig. 5). The device was thermally cycled after the training set was acquired and also between the  
241 acquisition of the two current maps in Fig. 5. The algorithm focuses on measuring regions of high  
242 current gradient, the corner edges and, in particular, the Coulomb peaks close to these.

243 In the top row of Fig. 5a,  $n = 4,096$  was chosen by the stopping criterion. In the bottom row,  
244 the corners edges extended further in the current map and the stopping criterion chose  $n = 8,192$ .  
245 This reduced the time needed to measure the current maps by 3.36 and 1.50, respectively, for the  
246 two test cases when compared with the alternating grid scans.

## 247 **Discussion**

248 The proposed measurement algorithm makes real-time informed decisions on which measurements  
249 to perform next on a single quantum dot device. Decisions are based on the disagreement of  
250 competing reconstructions built from sparse measurements. The algorithm outperforms grid scan  
251 in all cases, and in the majority of cases shows nearly optimum performance. The algorithm  
252 reduced the time required to observe the regions of finite current gradient by factors ranging from  
253 1.5 to 3.7 times. Optimisation of batch sizes or a variable scan resolution might reduce this time

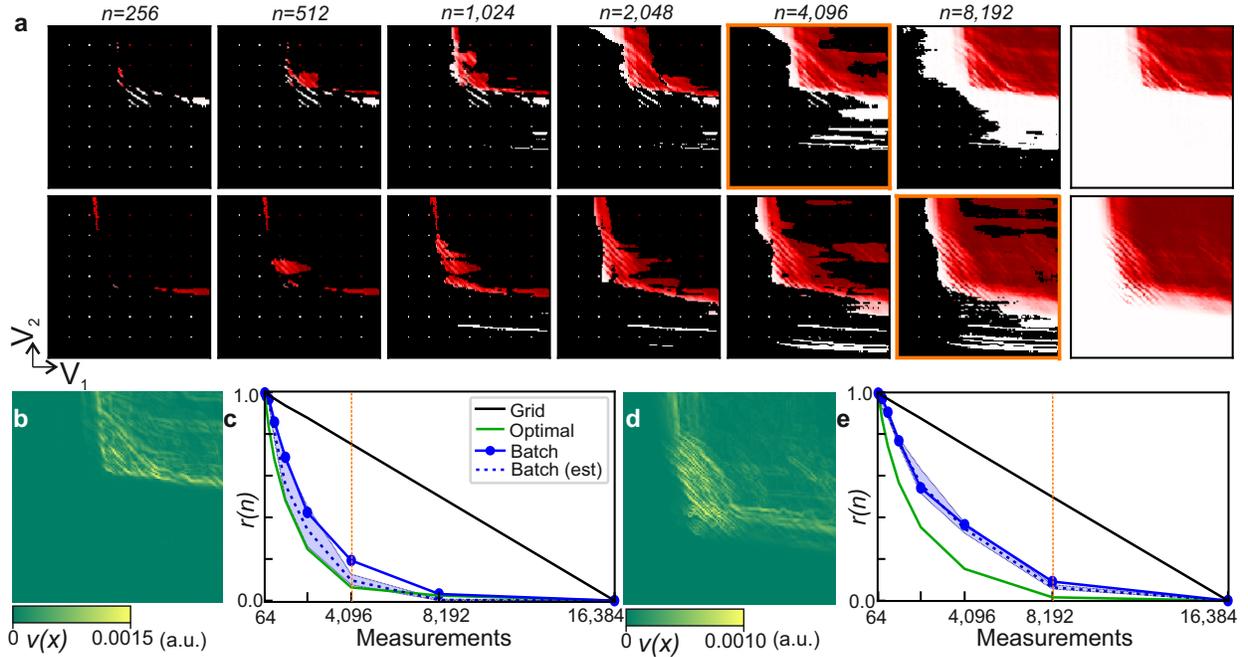


Figure 5: **Measuring a different current map.** **a**, Sequential batch measurement. Each row displays the algorithm-assisted measurements of a current map as a function of  $V_1$  and  $V_2$  for different values of  $n$ . The last plot in each row is the full-resolution current map. **b**, **d**, Current gradient map for both examples in **a**. **c**, **e**, Measure of the algorithm's performance  $r(n)$ , average real-time estimate of  $r(n)$  with 90% credible interval, and optimal  $r(n)$  for both current maps in **a**. The black line is the value of  $r(n)$  corresponding to the alternating grid scan method. The dashed orange line indicates the value of  $n$  determined by the stopping criterion. The corresponding current map in **a** is highlighted in orange. The alternating grid scan took 2,267 s and 2,333 s to acquire all measurements in the two cases. The batch method took 673 s and 1,552 s to reach the stopping criterion.

254 further, however, the performance gain is limited by the spread of the information gain over the  
255 scan range. This is evidenced in both Figs 4c,e and Fig 5c,e, where we show that even an optimal  
256 algorithm does not significantly outperform the algorithm.

257 Our algorithm with no modifications can be re-trained to measure different current maps. It  
258 simply requires a diverse data set of training examples from which to learn. The decision algorithm  
259 performed well even when trained on a small data set of only 382 current maps (at a resolution of  
260  $251 \times 251$ ), implying that it is robust to limited training data sets. Our algorithm focused on observ-  
261 ing all informative regions present in the current map, making it generalisable to different types  
262 of measurements and devices. The acquisition function can still be specifically designed to focus  
263 on specific transport features such as Coulomb peaks or Coulomb diamond edges. In additional  
264 experiments we demonstrate how this can be achieved by applying additional transformations to  
265 the reconstructions (see Supplementary Section Context-aware decision for stability diagrams).

266 We believe that our algorithm represents a significant first step in automating what to measure  
267 next in quantum devices. For a single quantum dot it provides a means of accelerating what can  
268 currently be achieved by human experimenters and other automation methods. When provided  
269 with an appropriate training data set our algorithm can be applied to a large variety of experiments.  
270 In particular, in any conventional qubit tuning method for which time-consuming grid scans are  
271 performed, our algorithm would allow for an improvement in measurement efficiency. It will not  
272 be long before this kind of approach enables experiments to be performed, and technology to be  
273 developed, that would not be feasible otherwise.

274 **Methods**

275 **Distribution of reconstructions and sampling** Since it is known that deep generative models  
 276 work well when the data range is from -1 to 1, all measurements are rescaled so that the maximum  
 277 value of the absolute value of the initial measurement is 1. Let  $Y$  be a random vector containing  
 278 all pixel values. Observation  $Y_n$ , where  $n \geq 1$ , is the set of pairs of location  $x_j$  and measurement  
 279  $y_j$ :  $Y_n = \{(x_j, y_j) \mid j = 1, \dots, n\}$ . Also, a subset of measurements is defined:  $Y_{n:n'} = \{(x_j, y_j) \mid$   
 280  $j = n, \dots, n'\}$ . The likelihood of observations given  $Y$  is defined by

$$p(Y_n \mid Y) \propto \exp(-\lambda \sum_{(x,y) \in Y_n} |y - Y(x)|), \quad (4)$$

281 where  $Y(x)$  is the pixel value of  $Y$  at  $x$ , and  $\lambda$  is a free parameter that determines the sensitivity  
 282 to the distance metric and is set to 1.0 for all experiments in this paper. The posterior probability  
 283 distribution is defined by Bayes' rule:

$$p(Y \mid Y_n) \propto p(Y_n \mid Y) p(Y). \quad (5)$$

284 Likewise, we can find the posterior distribution of  $\mathbf{z}$  given measurements instead of  $Y$ . Let  $\mathbf{z}'$   
 285 denote another input of the decoder, which is set to  $Y_{64}$  in the experiments. Then the posterior  
 286 distribution of  $\mathbf{z}$  can be expressed with  $\mathbf{z}'$  when  $n \geq 64$ :

$$\begin{aligned} p(\mathbf{z} \mid Y_n, \mathbf{z}') &\propto p(\mathbf{z} \mid \mathbf{z}') p(Y_n \mid p(\mathbf{z}, \mathbf{z}')) \\ &\propto p(\mathbf{z}) \int_Y p(Y_n \mid Y) p(Y \mid p(\mathbf{z}, \mathbf{z}')) dY \\ &\propto p(\mathbf{z}) p(Y_n \mid Y = \hat{Y}_{\mathbf{z}}), \end{aligned}$$

287 where  $\hat{Y}_{\mathbf{z}}$  is the reconstruction produced by the decoder given  $\mathbf{z}$  and  $\mathbf{z}'$ . Since all inputs of the  
 288 decoder are given,  $p(Y \mid \mathbf{z}, \mathbf{z}')$  is the Dirac delta function centered at  $\hat{Y}_{\mathbf{z}}$ . Also,  $p(\mathbf{z} \mid \mathbf{z}') = p(\mathbf{z})$

289 as  $\mathbf{z}$  and  $\mathbf{z}'$  are assumed independent. Proposal distribution for MH is set to a multivariate normal  
 290 distribution having centered mean and a covariance matrix equal to one quarter of the identity  
 291 matrix. For the experiments in this paper, 400 iterations of MCMC steps are conducted when  
 292  $n = 32 \times 2^b$ , where  $b$  is any integer larger than or equal to 1. We found that 400 iterations result  
 293 in good posterior samples. If  $(x_{n+1}, y_{n+1})$  is newly observed, then the posterior can be updated  
 294 incrementally:

$$\begin{aligned} p(\mathbf{z} | Y_{n+1}, \mathbf{z}') &= \frac{p(x_{n+1}, y_{n+1} | \mathbf{z}, \mathbf{z}')}{p(x_{n+1}, y_{n+1} | Y_n, \mathbf{z}')} p(\mathbf{z} | Y_n, \mathbf{z}') \\ &= \frac{p(x_{n+1}, y_{n+1} | \hat{Y}_{\mathbf{z}})}{p(x_{n+1}, y_{n+1} | Y_n, \mathbf{z}')} p(\mathbf{z} | Y_n, \mathbf{z}'), \end{aligned}$$

295 because each term in (4) can be separated.

296 **Decision algorithm** In this section, we derive a computationally simple form of the information  
 297 gain and the fact that maximising the information gain is equal to minimising the entropy. Let  
 298  $p_n(\cdot) = p(\cdot | Y_n, \mathbf{z}')$ , and any probabilistic quantity of  $y_{n+1}$  has the condition  $x_{n+1}$ , but omitted for  
 299 brevity.

The continuous version of the information gain equation is

$$\begin{aligned} &\mathbb{E}_{y_{n+1}} \left[ \text{KL}(p_n(\mathbf{z} | y_{n+1}) \| p_n(\mathbf{z})) \right] \\ &= \int_{y_{n+1}} p_n(y_{n+1}) \text{KL}(p_n(\mathbf{z} | y_{n+1}) \| p_n(\mathbf{z})) dy_{n+1} \\ &= \int_{y_{n+1}} p_n(y_{n+1}) \int_{\mathbf{z}} p_n(\mathbf{z} | y_{n+1}) \log \frac{p_n(\mathbf{z} | y_{n+1})}{p_n(\mathbf{z})} d\mathbf{z} dy_{n+1} \tag{6} \\ &= \int_{y_{n+1}} \int_{\mathbf{z}} p_n(\mathbf{z}, y_{n+1}) \log \frac{p_n(\mathbf{z}, y_{n+1})}{p_n(\mathbf{z}) p_n(y_{n+1})} d\mathbf{z} dy_{n+1} \\ &= I(\mathbf{z} | Y_n; y_{n+1} | Y_n), \end{aligned}$$

300 where KL is Kullback-Leibler divergence,  $I(\cdot; \cdot)$  is mutual information. Since  $I(\mathbf{z} | Y_n; y_{n+1} |$   
 301  $Y_n) = H(\mathbf{z} | Y_n) - H(\mathbf{z} | Y_n, y_{n+1})$ , maximising the expected KL divergence is equivalent to  
 302 minimising  $H(\mathbf{z} | Y_n, y_{n+1})$ , which is the entropy of  $\mathbf{z}$  after observing  $y_{n+1}$ .

303 Since this integral is hard to compute, we approximate probability density functions (PDFs)  
 304 with samples and substitute them into (6). Let  $n_s$  denote the number of measurements that are used  
 305 for sampling reconstructions  $\hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_M$  (the samples are converted to  $\hat{Y}_1, \dots, \hat{Y}_M$ ). Then  $p_{n_s}(\mathbf{z}) \approx$   
 306  $\frac{1}{M} \sum_m \delta_{\hat{\mathbf{z}}_m}(\mathbf{z})$ , or with the sample index  $m$ ,  $P_{n_s}(m) = 1/M$ . For any  $n \geq n_s$ , the probability is  
 307 updated with the new measurements after  $n_s$ :  $P_n(m; n_s) = \frac{p(Y_{n_s+1:n} | \hat{Y}_m)}{\sum_m p(Y_{n_s+1:n} | \hat{Y}_m)}$ , which can be derived  
 308 from importance sampling. For brevity, the sampling distribution information  $n_s$  is omitted for  
 309 the remaining section. Likewise,  $p_n(y_{n+1}) = \int_{\mathbf{z}} p_n(y_{n+1} | \mathbf{z}) p_n(\mathbf{z}) \approx \sum_m P_n(m) p_n(y_{n+1} | \mathbf{z}_m)$ .  
 310 Lastly, we use the value of  $\hat{Y}_m$  at  $x_{n+1}$  for a sample of  $p_n(y_{n+1} | \mathbf{z}_m)$  for simple and efficient  
 311 computation. As a result, the information gain is approximated by:

$$\begin{aligned} & \mathbb{E}_{y_{n+1}} \left[ \text{KL}(p_n(\mathbf{z} | y_{n+1}) \| p_n(\mathbf{z})) \right] \\ & \approx \sum_m P_n(m) \text{KL}(P_{n+1} \| P_n). \end{aligned}$$

312 **Simulator for Training data** To aid the training of the model simulated training data was used to  
 313 prevent over-fitting. Simulated data produced via a simple implementation of the constant inter-  
 314 action model<sup>29</sup> was used along with basic data augmentation techniques. These techniques were  
 315 not intended to be physically accurate but instead to produce quickly a diverse set of examples that  
 316 contain features that mimic real data.

317 The constant interaction model makes the assumptions that all interactions felt by a confined

318 electrons within the dot can be captured by a simple constant capacitance  $C_\Sigma$  which is given by  
 319  $C_\Sigma = C_S + C_D + C_G$  where  $C_S$ ,  $C_D$  and  $C_G$  are capacitances to the source, drain and gate re-  
 320 spectively. Making this assumption the total energy of the dot  $U(N)$  where  $N$  is the number of  
 321 electrons occupying the dot, is  $U(N) = \frac{(-|e|(N-N_0)+C_S V_S+C_D V_D+C_G V_G)^2}{2C_\Sigma} + \sum_{n=1}^N E_n$  where  $N_0$  compen-  
 322 sates for the background charge and  $E_n$  is a term that represents occupied single electron energy  
 323 levels that is characterised by the confinement potential.

324 Using this we derive the electrochemical potential  $\mu(N) = U(N) - U(N - 1) = \frac{e^2}{C_\Sigma}(N -$   
 325  $N_0 - \frac{1}{2}) - \frac{|e|}{C_\Sigma}(V_S C_S + V_D C_D + V_G C_G) + E_n$ .

326 To produce a training example random values are generated for  $C_S$ ,  $C_D$  and  $C_G$ . The energy  
 327 levels within a randomly generated gate voltage window and source drain bias window are then  
 328 counted. To aid generalisation to real data we randomly generated energy level transitions (which  
 329 are also counted) as well as slightly linearly scaled  $C_\Sigma$ ,  $C_S$ ,  $C_D$ , and  $C_G$  with  $N$ . This linear scaling  
 330 was also randomly generated and results in produced diamonds that vary in size with respect to  $V_G$ .  
 331 Examples of the training data produced by this simulator can be seen in Supplementary Figure 1.

332 **Stopping criterion** Utility, denoted by  $u$ , is the ratio of total measured gradient to the total gra-  
 333 dient of a stability diagram:  $u(n) = 1.0 - r(n)$ . Here, we assume that we have  $K$  more stability  
 334 diagrams to be measured. The location of each diagram is defined by a different voltage range, and  
 335  $k = 0, \dots, K$  is the index of the diagrams, where  $k = 0$  is the index of the diagram that we are  
 336 currently measuring.

337 Let  $T$  denote the total measurement budget for the current and remaining stability diagram.

338 In this paper we assume that a unit budget for measuring one pixel is 1.0. The total utility is

$$\begin{aligned}
 u_{\text{tot}} &= \sum_{k=0}^K u_k(t_k) \\
 &= u_0(t_0) + u_{\text{next}}(T - t_0),
 \end{aligned}$$

339 where  $u_k(\cdot)$  is the utility from measuring  $k$ th diagram,  $t_k$  is the planned budget for  $k$ th diagram  
 340 satisfying  $\sum_{k=0}^K t_k = T$ , and  $u_{\text{next}}(T - t_0) = \sum_{k=1}^K u_k(t_k)$ .

341 Let  $t$  denote the already spent budget on the current diagram,  $t \leq t_0$ . If we stop the measure-  
 342 ment then  $t_0 = t$ , or  $t_0 = t + \Delta$  if we decide to continue the measurement, where  $\Delta$  is a predefined  
 343 batch size. For the decision, the utilities of two cases are compared: when  $t_0 = t$ ,

$$u_{\text{tot}} = u_0(t) + u_{\text{next}}(T - t). \quad (7)$$

344 Otherwise,  $t_0 = t + \Delta$  and

$$u_{\text{tot}} = u_0(t + \Delta) + u_{\text{next}}(T - (t + \Delta)) \quad . \quad (8)$$

345 If (8) < (7), it is better to stop and move to the next voltage range. Rearranging the inequality leads  
 346 to

$$u_0(t + \Delta) - u_0(t) < u_{\text{next}}(T - t) - u_{\text{next}}(T - (t + \Delta)). \quad (9)$$

347 The left-hand-side (lhs) of (9) means the difference of utility if we invest  $\Delta$  budget more on the  
 348 current diagram, and the right-hand-side the difference when  $\Delta$  more budget is used for remaining  
 349 diagrams. As we discussed in Results section, we can calculate multiple slope estimates  $\beta_m$  for  
 350 spending  $\Delta$  to the current diagram:  $u_0(t + \Delta) - u_0(t) \approx \beta_m \Delta$ .

351 The right-hand-side (rhs) of (9) can be approximated by  $\alpha\Delta$  if  $K = \infty$ , where  $\alpha = 1/16,384$   
 352 is the slope of grid scan measuring a new stability diagram. Note that  $\alpha$  can be considered as the  
 353 empirical worst case performance of the decision algorithm measuring a new diagram as it holds  
 354 for all the experiments we have conducted. If  $\Delta = N$ , this approximation is the exact quantity for  
 355 any algorithms as all algorithms satisfy  $r(0) = 1.0$  and  $r(N) = 0.0$ . Since  $\alpha$  can be interpreted as  
 356 the worst case estimate, we also approximate lhs of (9) with the worst case estimate  $\beta = \min_m \beta_m$ .

357 If  $K < \infty$ , and the remaining budget  $T - t$  is more than the budget to measure all of remain-  
 358 ing diagrams, there is no utility after all measurements are finished. Hence, the approximation is  
 359 capped:

$$u_{\text{next}}(T - t) = \alpha \min(T - t, N \times K), \quad (10)$$

360 where  $K$  is the number of remaining diagrams to be measured.

361 As a result, the stopping criterion when  $K = \infty$  is

$$\beta < \alpha. \quad (11)$$

362 The stopping criterion when  $K < \infty$  is

$$\beta < \frac{\alpha(\min(T - t, N \times K) - \min(T - (t + \Delta), N \times K))}{\Delta}. \quad (12)$$

363 The rhs of (12) is always less than or equal to  $\alpha$ , and more total budget  $T$  makes it low, which leads  
 364 to late stopping or no stopping.

365 **Code Availability** A documented implementation of the algorithm in a github repository is avail-  
 366 able at <https://doi.org/10.5281/zenodo.2537934>.

367 **Data Availability** The data sets used for the training of the model are available from the corre-  
368 sponding author upon reasonable request.

369 **Acknowledgements** We acknowledge discussions with J.A. Mol and S.C. Benjamin. This work was sup-  
370 ported by the EPSRC National Quantum Technology Hub in Networked Quantum Information Technology  
371 (EP/M013243/1), Quantum Technology Capital (EP/N014995/1), Nokia, Lockheed Martin, the Swiss NSF  
372 Project 179024, the Swiss Nanoscience Institute and the NCCR QSIT. This publication was also made pos-  
373 sible through support from Templeton World Charity Foundation and John Templeton Foundation. The  
374 opinions expressed in this publication are those of the authors and do not necessarily reflect the views of  
375 the Templeton Foundations. We acknowledge J. Zimmerman and A. C. Gossard for the growth of the Al-  
376 GaAs/GaAs heterostructure.

377 **Author Contributions** D.T.L. and N.A. and the machine performed the experiment. H.M. developed the  
378 algorithm under the supervision of M.A.O. The sample was fabricated by L.C.C., L.Y., and D.M.Z. The  
379 project was conceived by G.A.D.B., E.A.L., M.A.O., and N.A. All authors contributed to the manuscript  
380 and commented and discussed results.

381 **Competing Interests** The authors declare that there are no competing interests.

382 **Correspondence** Correspondence and requests for materials should be addressed to Natalia Ares (email:  
383 natalia.ares@materials.ox.ac.uk).

- 385 1. Vandersypen, L. M. K. *et al.* Interfacing spin qubits in quantum dots and donors-hot, dense,  
386 and coherent. *npj Quantum Information* **3**, 34 (2017).
- 387 2. Baart, T. A., Eendebak, P. T., Reichl, C., Wegscheider, W. & Vandersypen, L. M. K. Computer-  
388 automated tuning of semiconductor double quantum dots into the single-electron regime. *Ap-  
389 plied Physics Letters* **108**, 213104 (2016).
- 390 3. Ito, T. *et al.* Detection and control of charge states in a quintuple quantum dot. *Scientific  
391 Reports* **6**, 4–6 (2016).
- 392 4. Van Diepen, C. J. *et al.* Automated tuning of inter-dot tunnel coupling in double quantum  
393 dots. *Applied Physics Letters* **113** (2018).
- 394 5. Volk, C. *et al.* Loading a quantum-dot based "Qubyte" register 1–12 (2019). Preprint at  
395 <http://arxiv.org/abs/1901.00426>.
- 396 6. Stehlik, J. *et al.* Fast charge sensing of a cavity-coupled double quantum dot using a Josephson  
397 parametric amplifier. *Physical Review Applied* **4**, 1–10 (2015).
- 398 7. Kalantre, S. S. *et al.* Machine Learning techniques for state recognition and auto-tuning in  
399 quantum dots 1–15 (2017). Preprint at <https://arxiv.org/abs/1712.04914>.
- 400 8. Frees, A. *et al.* Compressed optimization of device architectures for semiconductor quantum  
401 devices. *Physical Review Applied* **11**, 1 (2019).
- 402 9. Teske, J. D. *et al.* A machine learning approach for automated fine-tuning of semiconductor  
403 spin qubits. *Applied Physics Letters* **114**, 1–8 (2019).

- 404 10. Houthby, N., Huszár, F., Ghahramani, Z. & Lengyel, M. Bayesian Active Learning for Classi-  
405 fication and Preference Learning (2011). Preprint at <https://arxiv.org/abs/1112.5745>.
- 406 11. Ankenman, B., Nelson, B. L. & Staum, J. Stochastic Kriging for Simulation Metamodeling.  
407 *Operations Research* **58**, 371–382 (2010).
- 408 12. Sacks, J., Welch, W. J., Mitchell, T. J. & Wynn, H. P. Design and Analysis of Computer  
409 Experiments. *Statistical Science* **4**, 409–423 (1989).
- 410 13. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning (Adaptive*  
411 *Computation and Machine Learning)* (The MIT Press, 2005).
- 412 14. Sohn, K., Lee, H. & Yan, X. Learning structured output representation us-  
413 ing deep conditional generative models. In *Advances in Neural Informa-*  
414 *tion Processing Systems* (2015). URL [http://papers.nips.cc/paper/](http://papers.nips.cc/paper/5775-learning-structured-output-representation-using-deep-conditional-g)  
415 [5775-learning-structured-output-representation-using-deep-conditional-g](http://papers.nips.cc/paper/5775-learning-structured-output-representation-using-deep-conditional-g)  
416 [pdf](http://papers.nips.cc/paper/5775-learning-structured-output-representation-using-deep-conditional-g).
- 417 15. Goodfellow, I. *et al.* Generative adversarial nets. In *Advances in Neural In-*  
418 *formation Processing Systems* (2014). URL [http://papers.nips.cc/paper/](http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf)  
419 [5423-generative-adversarial-nets.pdf](http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf).
- 420 16. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. In *International Conference*  
421 *on Learning Representations* (2014).

- 422 17. Mescheder, L., Nowozin, S. & Geiger, A. Adversarial variational bayes: Unify-  
423 ing variational autoencoders and generative adversarial networks (2017). Preprint at  
424 <http://arxiv.org/abs/1701.04722>.
- 425 18. Srivastava, A., Valkov, L., Russell, C., Gutmann, M. & Sutton, C. VEEGAN: Reducing mode  
426 collapse in GANs using implicit variational learning. In *Advances in Neural Information*  
427 *Processing Systems* (2017).
- 428 19. van den Oord, A. *et al.* WaveNet: A generative model for raw audio (2016). Preprint at  
429 <http://arxiv.org/abs/1609.03499>.
- 430 20. Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I. & Frey, B. Adversarial autoencoders  
431 (2015). Preprint at <http://arxiv.org/abs/1511.05644>.
- 432 21. Huang, Huaibo, H. R. S. Z. T. T., Li. Zhihang. IntroVAE: Introspective Vari-  
433 ational Autoencoders for Photographic Image Synthesis 1–20 (2018). Preprint at  
434 <http://arxiv.org/abs/1807.06358>.
- 435 22. Zhu, J.-Y., Park, T., Isola, P. & Efros, A. A. Unpaired image-to-image translation using cycle-  
436 consistent adversarial networks. In *International Conference on Computer Vision* (2017).
- 437 23. Taigman, Y., Polyak, A. & Wolf, L. Unsupervised cross-domain image generation. In *Inter-*  
438 *national Conference on Learning Representations* (2017).
- 439 24. Iizuka, S., Simo-Serra, E. & Ishikawa, H. Globally and locally consistent image completion.  
440 *ACM Transactions on Graphics (Proc. of SIGGRAPH 2017)* **36**, 107:1–107:14 (2017).

- 441 25. Sanchez-Lengeling, B. & Aspuru-Guzik, A. Inverse molecular design using machine learning:  
442 Generative models for matter engineering. *Science* **361**, 360–365 (2018).
- 443 26. Gómez-Bombarelli, R. *et al.* Automatic chemical design using a data-driven continuous rep-  
444 resentation of molecules. *ACS Central Science* **4**, 268–276 (2018).
- 445 27. Kusner, M. J., Paige, B. & Hernández-Lobato, J. M. Grammar variational autoencoder. In  
446 *International Conference on Machine Learning* (2017).
- 447 28. Dai, H., Tian, Y., Dai, B., Skiena, S. & Song, L. Syntax-directed variational autoencoder for  
448 structured data. In *International Conference on Learning Representations* (2018).
- 449 29. Hanson, R., Kouwenhoven, L. P., Petta, J. R., Tarucha, S. & Vandersypen, L. M. Spins in  
450 few-electron quantum dots. *Reviews of Modern Physics* **79**, 1217–1265 (2007).
- 451 30. Abadi, M. *et al.* Tensorflow: A system for large-scale machine learning. In *USENIX Sympo-*  
452 *sium on Operating Systems Design and Implementation* (2016).