

1 QDFlow: A Python package for physics simulations of quantum 2 dot devices

3 **Donovan L. Buterakos^{1,2*}, Sandesh S. Kalantre^{1,2,4}, Joshua Ziegler²,**
4 **Jacob M Taylor^{1,2,3,5} and Justyna P. Zwolak^{1,2,3†}**

5 **1** Joint Center for Quantum Information and Computer Science, University of Maryland,
6 College Park, MD 20742, USA

7 **2** National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

8 **3** Department of Physics, University of Maryland, College Park, MD 20742, USA

9 **4** Department of Physics, Stanford University, Stanford, CA 94305, USA

10 **5** Axiomatic AI, Inc., Cambridge, MA 02139, USA

11 * dbuterak@umd.edu, † jpzwolak@nist.gov

12 Abstract

13 Recent advances in machine learning (ML) have accelerated progress in calibrating and
14 operating quantum dot (QD) devices. However, most ML approaches rely on access to
15 large, representative datasets designed to capture the full spectrum of data quality en-
16 countered in practice, with both high- and low-quality data for training, benchmarking,
17 and validation, with labels capturing key features of the device state. Collating such
18 datasets experimentally is challenging due to limited data availability, slow measure-
19 ment bandwidths, and the labor-intensive nature of labeling. QDFlow is an open-source
20 physics simulator for multi-QD arrays that generates realistic synthetic data with ground-
21 truth labels. QDFlow combines a self-consistent Thomas-Fermi solver, a dynamic capac-
22 itance model, and flexible noise modules to simulate charge stability diagrams and ray-
23 based data closely resembling experiments. With an extensive set of parameters that
24 can be varied and customizable noise models, QDFlow supports the creation of large,
25 diverse datasets for ML development, benchmarking, and quantum device research.

Copyright attribution to authors.

This work is a submission to SciPost Physics Codebases.

26 License information to appear upon publication.

Publication information to appear upon publication.

Received Date

Accepted Date

Published Date

28 Contents

30	1 Introduction	2
31	2 Physics simulation	4
32	2.1 Nanowire model	4
33	2.2 Thomas-Fermi solver	6
34	2.3 Capacitance model	7
35	3 Data generation	9
36	4 Noise	11
37	5 Benchmarking and limitations	14

38	6 Conclusion	16
39	References	17
40		
41		

42 1 Introduction

43 Among the various quantum computing platforms, quantum dots (QDs) stand out for their
44 scalability potential, compact size, and long coherence times [1]. Operating QD devices, how-
45 ever, remains a formidable challenge, with complexity growing rapidly—often exponentially—
46 as the number of qubits increases. Recent advances in integrating machine learning (ML) with
47 quantum device operation have begun to mitigate these difficulties, offering promising auto-
48 mated control and calibration strategies. For example, ML algorithms have been developed for
49 the fabrication [2, 3], characterization [4, 5], tuning [6–13], and gate virtualization [14, 15]
50 of QD devices.

51 Developing robust ML models requires access to large and diverse datasets representative
52 of the multi-dimensional parameter space typical of QD devices. Crucially, for supervised ML
53 applications, these datasets must also include metadata that identifies key features, such as
54 the global state (i.e., the number of QDs formed), the charge configuration, and the type of
55 transition lines present. Unfortunately, large volumes of high-quality experimental data can be
56 challenging to obtain as companies and research groups often keep such data proprietary [16].
57 Limited measurement bandwidth in real-world experiments also constrains the efficient explo-
58 ration of the entire high-dimensional parameter space in a reasonable time. Generating accu-
59 rate feature labels for publicly available data is a labor-intensive and time-consuming task
60 that can produce subjective and potentially erroneous labels. Physics-based simulations offer
61 a practical solution: they enable the generation of arbitrarily large datasets while providing di-
62 rect access to the ground-truth charge states, thereby simplifying the labeling process needed
63 for ML training.

64 Here, we introduce QDFlow, an open-source Python package for simulating QD systems
65 and generating synthetic data tailored for ML training and applications. The core physics
66 engine in QDFlow employs the Thomas-Fermi approximation to numerically solve for the semi-
67 classical charge density $n(x)$ along a one-dimensional (1D) nanowire. While the current state-
68 of-the-art devices are typically realized by confining charges (electrons or holes) within a two-
69 dimensional (2D) heterostructure, the QDs are ultimately formed within quasi-1D channels
70 within those heterostructures, motivating our choice of a 1D model. In practice, the simulated
71 data produced by QDFlow closely resembles that of linear QD arrays in 2D heterostructures.
72 ML models trained on QDFlow-generated data have been shown to generalize effectively to
73 larger 2D QD arrays [14].

74 There are several open-source QD device simulators that rely on the constant capacitance
75 model, treating the array of QDs and their associated electrostatic gates as nodes in a network
76 of fixed capacitors [17–19]. Additionally, Ref. [19] allows the capacitances to vary with re-
77 spect to the number of charges n by introducing an n -dependent correction to the capacitance
78 matrix. In contrast, in QDFlow the capacitance parameters are physics-informed, obtained di-
79 rectly from the self-consistent Thomas-Fermi solution rather than imposed heuristically. All key
80 physical observables—such as current, charge states, and sensor readouts—are derived from a
81 capacitance model constructed based on the computed charge density $n(x)$. This ensures that
82 the capacitances evolve dynamically with gate voltages, yielding a more realistic description of
83 device behavior. Furthermore, QDFlow allows for modeling regions with low barriers between




Figure 1: Diagram illustrating the QDFlow library organizational structure. Each of QDFlow’s four modules is listed, along with the most important classes or functions within those modules.

84 dots, leading the dots on either side to combine into a single centralized dot. Finally, QDFlow
 85 incorporates a flexible noise module, enabling the addition of experimentally relevant effects
 86 such as thermal broadening, charge offset drift, and voltage fluctuations. These features make
 87 the simulated data qualitatively comparable to experimental measurements while maintain-
 88 ing full access to the ground truth labels required for ML applications. Building on QFlow—a
 89 legacy implementation of the QD simulator that applied the Thomas-Fermi approximation to
 90 model charge densities and stability diagrams [20, 21]—QDFlow extends these methods into
 91 a flexible, open-source framework that integrates physics-informed capacitance modeling with
 92 realistic noise processes tailored for ML applications.

93 QDFlow has already demonstrated its utility in advancing ML-driven QD research, con-
 94 firming that ML models trained exclusively on simulated data can be successfully deployed
 95 in experimental settings, on samples fabricated in an academic cleanroom as well as on an
 96 industrial 300-mm process line, and on 1D and 2D QD arrays. The legacy version of the
 97 simulator was used to generate the QFlow-lite dataset [21, 22], which enabled the train-
 98 ing of several ML models for global state recognition. The utility of these ML models was
 99 demonstrated both offline, by navigating the voltage space within pre-measured experimen-
 100 tal datasets [13, 20], and in closed-loop experiments [9, 12, 23]. The dataset also supported
 101 the development of a novel classification framework for simple high-dimensional geometri-
 102 cal structures, known as the *ray-based classification* (RBC) framework [24]. The expanded
 103 dataset, QFlow 2.0: Quantum dot data for machine learning [22], generated
 104 using the Thomas-Fermi solver with integrated realistic noise processes, further advanced ML-
 105 based approaches to QD tuning. In particular, models trained with data from the QFlow 2.0
 106 dataset have been successfully applied to tasks such as data quality assessment [25], physics-
 107 informed RBC and ray-based navigation in 1D QD arrays [13], and the development of a full
 108 virtualization stack for 1D and 2D QD arrays in Ge/SiGe [14] and in Si/SiGe. These successes
 109 confirm the compatibility of QFlow 2.0-generated with real-world experiments. They also
 110 highlight the value of physics-informed synthetic datasets for accelerating the development of
 111 automated control tools for QD systems.

112 QDFlow is available for download from the [Python Package Index](#), with the source code
 113 released under the GNU General Public License in the QDFlow GitHub repository [26]. Com-
 114 prehensive API documentation is provided via docstrings embedded in the source code and as
 115 HTML pages hosted on [Read the Docs](#). The library includes type hints for all classes and func-
 116 tions to support clarity, maintainability, and extensibility. Unit and benchmark tests are also
 117 distributed with the QDFlow repository to facilitate validation and performance evaluation.

118 2 Physics simulation

119 QDFlow has three main modules: `simulation`, `noise`, and `generate`, and one utility
120 module, `distribution`, as depicted in Fig. 1. The `simulation` and `noise` modules are
121 part of the `physics` package. The core physics-based engine of the simulator is contained
122 in the `simulation` module. It uses a Thomas-Fermi solver to find the stable charge config-
123 uration and sensor output of a particular QD device defined by a set of physical parameters.
124 The QDFlow Thomas-Fermi solver was first introduced in Ref. 21, but has since been refined
125 and extended within QDFlow to improve flexibility, physical relevance, and integration with
126 downstream ML workflows. The `PhysicsParameters` dataclass, which is used to initialize
127 the simulation, specifies over twenty parameters governing the properties of the QD device.
128 These parameters include both material characteristics and device-specific features such as
129 gate geometry and positioning. Importantly, the gate voltages—experimentally relevant con-
130 trol knobs—are explicitly included among the simulation inputs. By sweeping these voltages,
131 QDFlow produces the final outputs: 2D CSDs and 1D rays, directly mirroring the tuning pro-
132 cedures used in real QD experiments.

133 The `noise` module is responsible for adding noise to the final datasets, as well as for apply-
134 ing certain post-processing to the data. The `generate` module contains high-level functions
135 to assist in generating datasets. It is the module that the user would most often interact with.
136 Finally, the `distribution` module, contained within the `util` package, contains classes
137 defining random variable distributions.

138 To generate data with QDFlow, the user first chooses whether to run the default configu-
139 ration or adjust the distributions and ranges over which physics parameters are randomized.
140 Next, they create one or more sets of randomized device parameters, and for each device,
141 generate a CSD using the functions in the `generate` module. Once the physics parameters
142 have been specified, an instance of the `ThomasFermi` class is instantiated. This class serves
143 two main purposes: first, it solves for the charge density function $n(x)$; and second, it uses
144 $n(x)$ to construct a capacitance model and compute physical quantities such as the device's
145 charge state and sensor response. QDFlow then runs the physics simulation for every pixel in
146 each diagram and compiles the results into a `CSDOutput` dataclass, which is returned to the
147 user. By repeating this process over a range of gate voltages, QDFlow generates data that can
148 be assembled into CSDs or rays, depending on the application. The output, stored as NumPy
149 arrays, can be directly analyzed and plotted, or optionally augmented with noise to emulate
150 experimental data.

151 In the following sections, we provide a more detailed account of the nanowire model
152 physics underlying the simulation. We then explain how the Thomas-Fermi approximation
153 is applied to construct the capacitance model that drives the CSD simulation.

154 2.1 Nanowire model

155 QDFlow employs a 1D physics model in which charges are assumed to be confined to a linear
156 nanowire that lies along the x -axis. The ends of the nanowire are connected to electron reser-
157 voirs, and a bias voltage can be applied between them. Electrostatic gates are positioned at a
158 height h below the xy -plane, and are modeled as infinite cylindrical conductors with central
159 axis parallel to the y -axis, as shown in Fig. 2(a). The arrangement makes our nanowire model
160 a hybrid between a true nanowire device and other QD device architectures. Gates biased to
161 low potential act as plunger gates, while those biased to high potential act as barrier gates (for
162 positive charge carriers, with the convention reversed for negative carriers).

163 The plunger and barrier gates define an electrostatic potential $V(x)$, where x is the distance
164 along the nanowire. Note that because we are using a 1D model, we are only concerned
165 with the potential along the x -axis. The potential at a distance r from the center of a single




Figure 2: (a) The nanowire model used in QDFlow. (b) The potential $V(x)$ created by the electrostatic gates (left) and the charge density $n(x)$ induced by the potential (right).

166 cylindrical gate (in the absence of other gates) can be expressed as the potential of a screened
 167 line charge:

$$V_{\text{gate}}(r) = V_h \frac{\mathcal{K}_0(r/\lambda)}{\mathcal{K}_0(h/\lambda)} \quad (1)$$

168 where V_h is the potential at a reference distance $r = h$ (chosen as the separation between the
 169 gate and the nanowire), λ is the screening length, and $\mathcal{K}_0(z)$ is the modified Bessel function
 170 of the second kind. Specifically, we note that V_h is not the voltage of the gate itself; rather, it
 171 is the voltage that the nanowire feels due to the gate (in the absence of other gates). This is
 172 essentially the voltage of the gate multiplied by the lever arm of the gate.

173 Because the presence of nearby gates induces additional charges on each of the gates,
 174 and consequently changes Eq. (1), we cannot obtain $V(x)$ by simply summing $V_{\text{gate}}(r)$ over
 175 all gates. Determining the exact potential in the presence of multiple gates is a challenging
 176 electrostatics problem, even in the purely classical setting. It requires solving the screened
 177 Laplace equation with boundary conditions determined by the voltages on each of the gates.
 178 While such a calculation can be performed numerically, it must be repeated whenever any gate
 179 voltage is changed.

180 To make the problem tractable, we adopt the simplifying assumption that the induced
 181 charges on each gate are rotationally symmetric about the axis of the gate. Under this ap-
 182 proximation, the induced charges act as a line charge that lies along the central axis of the
 183 gate. Because the gate is rotationally symmetric, the potential $V_{\text{gate}}(r)$ also acts as a single line
 184 charge at the center of the gate, and thus the induced charges effectively rescale $V_{\text{gate}}(r)$ by a
 185 constant factor. Let V'_i be the rescaled value of V_h for gate i after including the effects of the
 186 induced charges on gate i , and let V_i be the value of V_h necessary for $V_{\text{gate}}(r)$ to give the actual
 187 potential of gate i . Then by the superposition principle, V_i can be determined by adding the
 188 potential contributions from each of the gates:

$$V_i = \sum_j A_{ij} V'_j \quad (2)$$

189 where A_{ij} denotes the ratio between the contribution from gate j to the potential at gate i
 190 and the effective potential V'_j . In the absence of other gates, $V_i = V'_i$, so A_{ii} is simply 1. The
 191 contribution from one gate to another is given by $V_{\text{gate}}(r)$, where r is the distance between
 192 gates. Thus, A_{ij} can be expressed as follows:

$$A_{ij} = \begin{cases} 1 & \text{if } i = j \\ V_{\text{gate}}(x_j - x_i) / V_{\text{gate}}(\rho_j) & \text{otherwise,} \end{cases} \quad (3)$$

193 where x_j is the x-coordinate of the central axis of gate j and ρ_j is the radius of the gate j .
 194 Calculating and inverting the matrix \mathbf{A} allows us to determine the effective potentials V'_i from
 195 the applied gate voltages V_i . We then obtain $V(x)$ by summing $V_{\text{gate}}(\sqrt{(x - x_i)^2 + h^2})$ over
 196 all gates and using the effective potentials V'_i in place of V_h in Eq. (1).

197 **2.2 Thomas-Fermi solver**

198 Having established how to compute the effective gate potentials and construct $V(x)$, we now
 199 turn to the resulting charge distribution along the nanowire. Define $n(x)$ to be the linear
 200 charge density at a point x along the nanowire, which is determined in response to the po-
 201 tential $V(x)$, as illustrated in Fig. 2(b). However, due to the Coulomb interaction between
 202 charges, the presence of an induced charge in the nanowire will create a correction to $V(x)$.
 203 This self-interaction results in the following integral equations, which must be solved self-
 204 consistently:

$$n(x) = \int_{\epsilon'(x)}^{\infty} \frac{g_0}{1 + e^{\beta(\epsilon - \mu)}} d\epsilon = \frac{g_0}{\beta} \text{sp}[\beta(\mu - \epsilon'(x))] \quad (4)$$

$$\epsilon'(x) = qV(x) + \int_{\mathbb{R}} K(x, x') n(x') dx' \quad (5)$$

205 The physical parameters μ , g_0 , and β in Eq. (4) indicate the Fermi level, the density of states
 206 in the conduction band (which is constant in 2D), and the inverse temperature, respectively,
 207 and $\text{sp}(z) = \ln(1 + e^z)$ is the softplus function. Parameter q in Eq. (5) controls the sign of the
 208 charge carriers, with -1 indicating electrons and $+1$ indicating holes, while $K(x, x')$ gives the
 209 strength of the Coulomb interaction between points x and x' , and is defined as follows:

$$K(x, x') = \frac{K_0}{\sqrt{(x - x')^2 + \sigma^2}} \quad (6)$$

210 where K_0 defines the energy scale of the interaction and σ is a softening parameter added to
 211 prevent a numerical singularity at $x = x'$, which occurs due to the 1D model breaking down at
 212 scales less than the radius of the nanowire. The value of σ can be chosen to be $3\pi r/8$, where
 213 r is the nanowire radius, to maintain consistency with the potential energy of two uniformly
 214 charged disks as the spacing between them approaches zero [27], or alternatively, a custom
 215 interaction $K(x, x')$ can be provided.

216 The Coulomb integral in Eq. (5) is formally taken over the entire nanowire. Because
 217 $K(x, x')$ scales for large x' as $1/|x'|$, this introduces concerns that the integral might diverge.
 218 At the same time, the integrand is weighted by $n(x')$, which becomes exponentially small for
 219 $V(x') - \mu \gg \beta^{-1}$. This condition is satisfied at the external barriers, see Fig. 2(b), provided that
 220 the external barrier voltages are sufficiently high. In addition, it is assumed that at distances
 221 far away from the nanowire, the system is connected to an electron reservoir where $n(x')$ is
 222 large. However, the Coulomb interaction in semiconductors tends to include a screening term
 223 that suppresses contributions past a certain range. Thus, in practice, it is sufficient to evaluate
 224 the integral between the two external barrier gates.

225 For convenience, we define a linear operator \mathbf{K} to be the result of evaluating the Coulomb
 226 integral as follows:

$$\mathbf{K}f(x) = \int_{\mathbb{R}} K(x, x') f(x') dx' \quad (7)$$

227 This allows us to combine Eqs. (4) and (5) to obtain:

$$n(x) = \frac{g_0}{\beta} \text{sp}[\beta(\mu - qV(x) - \mathbf{K}n(x))]. \quad (8)$$

228 The basic method we employ to solve Eq. (8) is successive iteration. Starting from an initial
 229 guess $n_0(x)$, the right-hand side of Eq.(8) is evaluated with $n(x) = n_0(x)$, yielding an updated
 230 function $n_1(x)$. This procedure is then repeated until $n(x)$ converges, if at all. The conver-
 231 gence tolerance and the maximum number of allowed iterations are specified through the

232 NumericsParameters dataclass, which can be provided when instantiating the Thomas-
 233 Fermi class. If the iteration does not converge to the specified tolerance within the allowed
 234 number of iterations, a ConvergenceWarning is issued.

235 The convergence can be problematic for certain parameter regimes. For example, if we
 236 define $\Delta(x)$ to be the difference between an initial guess $n_0(x)$ and the true value $n(x)$:

$$n_0(x) = n(x) + \Delta(x), \quad (9)$$

237 then evaluating the right-hand side of Eq. (8) yields:

$$n_1(x) = \frac{g_0}{\beta} \text{sp}[\beta(\mu - qV(x) - \mathbf{K}n(x) - \mathbf{K}\Delta(x))] \quad (10)$$

238 We now use the approximation $\text{sp}(z) \approx z$, which is valid for $z \gg 1$. Although this assumption
 239 does not always hold (particularly for small β), it is useful for analyzing certain convergence
 240 issues that may arise. Under this approximation, Eq. (10) simplifies to:

$$n_1(x) \approx n(x) - g_0 \mathbf{K}\Delta(x) \quad (11)$$

241 If all eigenvalues of $g_0 \mathbf{K}$ are smaller than 1, the error term $-g_0 \mathbf{K}\Delta(x)$ will be smaller in
 242 magnitude than the initial error $\Delta(x)$, and successive iterations will therefore converge to
 243 $n(x)$. Conversely, if $g_0 \mathbf{K}$ possesses eigenvalues greater than 1, the iterative scheme will gen-
 244 erally diverge. Physically, this divergence corresponds to strong coupling between charges, a
 245 regime that is well known to cause convergence difficulties in condensed matter systems [28].
 246 Fortunately, this issue can be partially mitigated by solving Eq. (9) for $\Delta(x)$, substituting the
 247 result into Eq. (11), and solving for $n(x)$, yielding the following expression:

$$n(x) \approx (\mathbf{1} + g_0 \mathbf{K})^{-1} [g_0 \mathbf{K}n_0(x) + n_1(x)] \quad (12)$$

248 If we discretize the x -axis, the operator $(\mathbf{1} + g_0 \mathbf{K})^{-1}$ can be computed through direct matrix
 249 inversion. This expression can then be incorporated into the successive iteration scheme by
 250 applying Eq. (12) after each iteration. Although there are still parameter regimes where the
 251 process diverges, this modified approach drastically enlarges the domain of convergence. In
 252 the weak-interaction limit, where the eigenvalues of $g_0 \mathbf{K} \ll 1$, the right-hand side of Eq. (12)
 253 simplifies to $n_1(x)$ to leading order in $g_0 \mathbf{K}$. Thus, in this limit, the method naturally recovers
 254 the standard successive iteration procedure.

255 2.3 Capacitance model

256 After calculating $n(x)$, QDFlow employs a capacitance model to determine the stable charge
 257 configuration and other properties. Similar techniques have been implemented in other QD
 258 simulations [17–19]. In most of those approaches, the capacitance matrix is assumed to be
 259 constant, i.e., the interdot capacitances remain fixed as the gate voltages are swept. The
 260 simulation introduced in Ref. 19 allows for variable capacitances by applying a correction
 261 to the capacitance matrix based on the particle number. In contrast, QDFlow derives the
 262 capacitance matrix directly from the charge density $n(x)$, which depends explicitly on the
 263 gate voltages. This feature enables charge-transition slopes and spacings to vary across a single
 264 CSD. Moreover, constructing the capacitance model from $n(x)$ naturally captures transitions
 265 between a double dot and a merged single dot as the interdot barrier is lowered.

266 The first step in creating the capacitance model is determining the regions of the nanowire
 267 where significant amounts of charge are induced. This is achieved by applying a threshold
 268 to $n(x)$, configurable through the NumericsParameters dataclass, and identifying con-
 269 tinuous intervals of points that lie above the threshold. This will result in a set of intervals

270 of the form $[a_i, b_i]$, which we call “charge islands.” The thresholding is also responsible for
 271 determining whether or not adjacent QDs should be handled as individual dots. Specifically,
 272 if $n(x)$ exceeds the threshold throughout the region between the two QDs, they are merged
 273 and treated as a single dot. Otherwise, they are considered to be two separate QDs with a
 274 potential barrier between them.

275 Once the charge islands are identified, the energy E of the resulting capacitance model is
 276 defined as follows:

$$E = \sum_{i,j} E_{ij}(Q_i - Z_i)(Q_j - Z_j) \quad (13)$$

$$Z_i = \int_{a_i}^{b_i} n(x)dx \quad (14)$$

$$E_{ij} = \frac{1}{Z_i Z_j} \left[c_k \delta_{ij} \int_{a_i}^{b_i} n(x)^2 dx + \frac{1}{2} \int_{a_i}^{b_i} \int_{a_j}^{b_j} K(x, x') n(x) n(x') dx dx' \right] \quad (15)$$

277 where Z_i is the (potentially noninteger) charge induced by the gates on island i under the
 278 Thomas-Fermi approximation, and Q_i is the integer number of charges on island i under a
 279 specific charge configuration \vec{Q} . The c_k term of Eq. (15) incorporates the kinetic energy of
 280 the charges. Since, for our purposes, the energy matrix fully characterizes the system, we do
 281 not compute the capacitances explicitly and instead work directly with the energy matrix. If
 282 desired, the capacitance matrix C can be obtained from the energy matrix via the relationship
 283 $C = (2E)^{-1}$.

284 After calculating the energy matrix, the next step is to determine the charge configuration
 285 \vec{Q} that minimizes the total energy E , subject to the constraint that all Q_i must be nonneg-
 286 ative integers. This is an instance of an integer optimization problem, which in general is
 287 NP-complete. However, for a moderate number of gates, a brute-force search is sufficient to
 288 find the minimum. In particular, we first locate the minimum in the continuous space, which
 289 occurs at \vec{Z} , and then evaluate $E(\vec{Q})$ over all \vec{Q} such that for each integer Q_i , $|Q_i - Z_i| < 1$.
 290 Once a stable charge configuration is identified, the potential at each of the sensors is calcu-
 291 lated under the assumption that each island i hosts a line of charge with total charge $q Q_i$ and
 292 charge density proportional to $n(x)$. The Coulomb potential at each of the sensors arising from
 293 these charge islands is calculated, and the result is normalized by dividing by the potential of
 294 a single point charge located at a point on the nanowire closest to the sensor in question. This
 295 means that a single transition should have a height of no more than 1 after normalization.

296 Finally, QDFlow allows the current across the nanowire to be found. For this calculation,
 297 the left and right sides of the nanowire are assumed to be connected to electron baths with po-
 298 tentials V_L and V_R , respectively. The dynamics of the charges are modeled using a semi-classical
 299 approach, treating them as particles that travel at the Fermi velocity. Each time they collide
 300 with a barrier, the particles have a chance to either tunnel through it or be reflected back. The
 301 tunneling probability across each barrier is determined by the transmission coefficient, which
 302 we calculate using the WKB approximation. This allows the tunneling rates between islands
 303 and the tunneling rates to and from the external charge baths to be obtained. These tunnel
 304 rates are then used to define a Markov graph which encodes the dynamics of the transitions
 305 between charge states. The current through the nanowire is obtained by evaluating the net
 306 rate at which charges enter and leave the charge baths at the steady state of this Markov graph.

307 3 Data generation

308 The data generation is carried out within the `generate` module. A single instance of the
 309 `ThomasFermi` class calculates quantities of interest for a single point in voltage-space only
 310 based on the device configuration specified in the `physics` module. To generate a complete
 311 CSD, a new simulation instance must be created for each pixel. However, since the gate volt-
 312 ages of neighboring pixels vary only slightly, it follows that the corresponding charge density
 313 $n(x)$ will also not change significantly between adjacent pixels. To optimize QDFlow per-
 314 formance, the result of the $n(x)$ calculation at one pixel is used as an initial condition when
 315 calculating $n(x)$ at adjacent pixels. This means that $n(x)$ must only be calculated from scratch
 316 once for each diagram.

317 QDFlow contains convenience functions for generating CSDs and rays in the `generate`
 318 module. Since the primary purpose of QDFlow is to generate data for training specialized ML
 319 models, it is essential that the resulting dataset captures the full range of variability observed
 320 in contemporary QD devices. To achieve this, QDFlow includes functionality to randomize
 321 nearly all physics parameters and to control the distributions from which each parameter is
 322 drawn. This capability is implemented via the `PhysicsRandomization` dataclass, which
 323 specifies each physics parameter as either a fixed value (when no randomization is desired,
 324 e.g., to allow regeneration of the same QD device), or a `Distribution` from which to draw
 325 the random values. A code example in Listing 1 shows how to import QDFlow and how to
 326 initialize a random configuration of physical parameters, with μ drawn from a distribution
 327 provided by the user.

328 The `Distribution`, defined in the `distribution` module, is an abstract base class
 329 that encodes how a given parameter should be randomized. Several standard distributions,
 330 including `Uniform`, `Normal`, and `LogNormal`, are implemented in QDFlow as wrappers to
 331 the NumPy random generator functions of the same name. In addition, user-defined distri-
 332 butions can be easily created by extending the `Distribution` class and implementing the
 333 `draw()` method to generate random values in an arbitrary manner. The `CorrelatedDis-`
 334 `tribution` class handles cases where it is desired or necessary for multiple random variables
 335 to be related to one another in some way.

336 This randomization framework, along with the dozens of configurable physics parame-
 337 ters, enables QDFlow to generate a highly diverse set of CSDs. Figure 3 shows six examples

```
from qdflow import generate
from qdflow.util import distribution

# Create a new dataclass instance that contains the default
# randomization distributions for each physics parameter
phys_rand = generate.PhysicsRandomization.default()

# Change the range from which mu can be drawn
phys_rand.mu = distribution.Uniform(0, 1.2)

# Generate a list of 6 sets of random device parameters
n_devices = 6
phys_params = generate.random_physics(phys_rand, n_devices)
```

Listing 1: Example code to generate a list of 6 randomized sets of device parameters. First, a `PhysicsRandomization` object is created, which defines the ranges and distributions, as appropriate, from which the physics parameters should be randomized. Distributions for each parameter can then be set as desired.




Figure 3: (a)-(f) Examples of CSDs generated with randomized physics parameters. (g) Examples of ray data generated along the rays shown on the CSD in panel (f).

338 of CSDs generated with QDFlow. A snippet of code allowing to generate CSD data from a
 339 `PhysicsParameters` object is shown in Listing 2.

340 While CSDs are extremely useful for visualizing charge states, they require extensive data
 341 collection. In practice, only a relatively small subset of points—the charge transitions—are
 342 of primary importance. As the number of dots grows, so does the dimensionality of the gate
 343 voltage space, rapidly making the exploration of the complete, multidimensional voltage space
 344 infeasible. This challenge is typically handled by measuring multiple 2D CSDs, each defined
 345 by a different pair of gates.

346 In Ref. 12, an alternative method for assessing the charge state in QD devices, with 1D rays
 347 measured in multiple directions in the voltage space used in place of the 2D CSDs. This method
 348 greatly reduces the amount of data required for assessing the charge state of the device, but
 349 sacrifices some of the intuitive human interpretability provided by CSDs, necessitating the
 350 use of ML tools. To support the development of ML methods for ray-based analysis, QDFlow

```
from qdflow import generate
import numpy as np

# Create a set of physics parameters
phys = generate.default_physics(n_dots=2)

# Set ranges (in mV) and resolution of plunger gate sweeps
V_x = np.linspace(2., 16., 100)
V_y = np.linspace(2., 16., 100)

# Generate a charge stability diagram
csd = generate.calc_2d_csd(phys, V_x, V_y)

# Obtain the sensor readout in the form of a numpy array
sensor_num = 0
sensor_readout = csd.sensor[:, :, sensor_num]
```

Listing 2: Example code demonstrating how to generate a CSD from a `PhysicsParameters` object.

```

from qdflow import generate
import numpy as np
from scipy.stats import qmc

# Create a set of physics parameters
phys = generate.default_physics(n_dots=2)

# Generate quasirandom points inside a given area
v_min, v_max = 2., 16.
point_generator = qmc.Halton(d=2, scramble=False)
initial_points = qmc.scale(point_generator.random(n=50), v_min, v_max)

# Define a list of rays that will extend out from each point
ray_length = 3. # length of rays in mV
num_rays = 8
rays = ray_length * np.array([[np.cos(2*np.pi*i/num_rays),
    np.sin(2*np.pi*i/num_rays)] for i in range(num_rays)])

# Generate ray data
resolution = 100 # points per ray
ray_data = generate.calc_rays(phys, initial_points, rays, resolution)

```

Listing 3: Example code demonstrating how to generate ray data from a `PhysicsParameters` object.

351 includes functionality for generating ray-based datasets, as shown in Listing 3.

352 4 Noise

353 The simulations described thus far capture many essential physical features of QD devices but
 354 omit one critical ingredient: noise. In experimental data, noise strongly influences both the
 355 visibility of charge transitions and the reliability of automated analysis.

356 To more faithfully emulate experimental conditions, QDFlow includes the `noise` module,
 357 which contains functionality for adding noise to both CSDs and rays, as well as several postpro-
 358 cessing functions designed to mimic effects of experimental measurements. The module imple-
 359 ments several types of noise, including the white noise, pink ($1/f$) noise, telegraph noise, and
 360 latching effects, as well as stray transitions arising from nearby unintended dots [25,29]. Post-
 361 processing functions include adding gate-sensor coupling, adding a sech^2 x blur, and adding
 362 Coulomb peak effects. Each noise in QDFlow can be controlled individually, with its mag-
 363 nitude defined relative to the scale of the CSD data, or as a predefined mixture. Similar to
 364 the physical parameter randomization discussed in Sec. 3, QDFlow supports designating the
 365 distributions from which each noise parameter is drawn via the `NoiseRandomization` dat-
 366 aclass.

367 Figure 4(a) shows an example of a noiseless CSD simulated with QDFlow. CSDs with
 368 noise implementations adapted from QFlow—white, pink, and telegraph noise and Coulomb
 369 peak—are depicted in panels (b), (c), (d), and (e), respectively. CSDs with latching, sech^2
 370 blur, unintended QD, and sensor-gate coupling—new to QDFlow—are presented in panels
 371 (f), (g), (h), and (i), respectively.

372 The simplest, white noise, is implemented by adding to each pixel a value drawn from a
 373 normal distribution with standard deviation given by the magnitude of the white noise. Pink
 374 noise is implemented by generating white noise in Fourier space with a uniform random phase




Figure 4: Examples of noise added to a CSD. (a) The original CSD data, (b) white noise, (c) pink noise, (d) telegraph noise, (e) Coulomb peak, (f) latching, (g) Sech blur, (h) unintended QD, and (i) sensor-gate coupling.

375 and magnitude proportional to $1/\sqrt{f_x^2 + f_y^2}$, where f_x and f_y are the components of each point
 376 in Fourier space, and then applying an inverse Fourier transform. This configuration allows for
 377 greater variability compared to telegraph noise. Alternatively, QDFlow also provides the op-
 378 tion to add pink noise correlated along only the primary measurement axis, which corresponds
 379 to the more physical model of pink noise correlated in time.

380 Telegraph noise is applied along an axis corresponding to the primary measurement direc-
 381 tion in experimental data. It consists of adding a value (drawn from a normal distribution with
 382 nonzero mean) to a line of several contiguous pixels. The length of the added line is randomly
 383 drawn from a geometric distribution. This allows the distribution of lengths of the telegraph
 384 noise to follow an exponential, as expected for two-level systems with finite excited-state life-
 385 times. This process is then repeated across the CSD, alternating the sign of the value added
 386 each time.

387 Transitions from spurious QDs are emulated by adding functions of the form $\tanh((\vec{x} - \vec{x}_0) \cdot \vec{a})$,
 388 where \vec{x} gives the coordinates of each pixel, \vec{x}_0 is the location on the CSD of the transition,
 389 and \vec{a} determines how strongly each of the gates plotted on the x - and y -axes are coupled
 390 to the unintended dot. Values of \vec{x}_0 and \vec{a} are randomized, but a single \vec{a} is used if multiple
 391 unintended transitions appear on a single diagram.

392 The latching noise implemented in QDFlow can be controlled using one of two meth-
 393 ods. The first and simpler legacy method to simulate latching is to shift each row of pixels
 394 by a random number of pixels drawn from a geometric distribution. This process produces
 395 a latching-like effect along the charge transitions; however, it is somewhat unrealistic since




Figure 5: (a)-(c) Examples of CSDs with all noise types combined. (d) Ray data without noise (dotted line) and with noise added (solid line).

396 all transitions on the same row are shifted by the same amount, and because pixels far away
 397 from transitions are also displaced. A more physically realistic method relies on the nanowire
 398 simulation to calculate the sensor readout for both an excited charge state and a stable state
 399 at each pixel, similar to what is implemented in Ref. 18. The excited state chosen corresponds
 400 to the charge configuration most recently occupied prior to the most recent transition when
 401 sweeping the gate voltages. When latching noise is added, the sensor readout from the ex-
 402 cited state replaces the stable-state readout for the first few pixels after each transition, with
 403 the number of pixels randomized each instance. In general, the second method is preferable
 404 as it more accurately reflects the experimental conditions; however, the first, legacy method
 405 is provided as a fallback when excited-state data are unavailable or computationally inconve-
 406 nient, or impossible to obtain.

407 In addition to noise, several postprocessing functions can be applied to the data. A sensor-
 408 gate coupling in the form of a linear gradient along a random direction can be added. Con-
 409 volving with a $\text{sech}^2 x$ kernel along the measurement axis introduces smoothing of the sharp
 410 transitions.

411 Finally, it is important to note that the physics simulation returns the value of the potential
 412 at each of the sensors; however, experimentally, the conductance of the sensor is the quantity
 413 that is measured. Therefore, we convert from potential to conductance by using a Coulomb
 414 peak lineshape of the form $G \propto \text{sech}^2[A(V - V_0)]$, where G is the conductance of the sensor, A
 415 is a parameter that determines the width of the Coulomb peak, V is the potential at the sensor
 416 (the simulation output), and V_0 is the peak center [30].

417 CSDs with a mixture of noises optimized for compatibility with experimental data are
 418 presented in Fig. 5(a)-(c). Figure 5(d) shows two rays with added noise. The exact amounts
 419 of each type of noise are randomized for each diagram. A code example showing how to
 420 generate noisy CSD data for a previously simulated sample CSD is shown in Listing 4.

421 Given the computational complexity of Thomas-Fermi calculations, the `noise` module is
 422 configured to assume that a complete noiseless CSD (or a ray-based data) has already been
 423 generated using the `generate` module. This approach gives us several advantages. First, it
 424 significantly reduces the computational overhead since multiple noise realizations with differ-
 425 ent relative noise strengths can be generated from a single noiseless CSD. Secondly, the mod-
 426 ular approach adopted in QDFlow provides flexibility that cannot be achieved experimentally,
 427 where it is not possible to calibrate individual noise sources to the desired level. Additionally,
 428 it allows us to specify the magnitudes of each of the noise types relative to the local scale of
 429 the surrounding data points, which is important since the scale of the data points can vary
 430 across large CSDs. Finally, it allows for systematic benchmarking of ML algorithms since the
 431 noise level and type can be modified independently of the underlying physical configuration.
 432 For example, one can generate a single noiseless dataset and then apply different realizations
 433 of noise to study the robustness of a given algorithm under varying experimental conditions,

```

from qdflow.physics import noise
from qdflow.util import distribution
import numpy as np

# Use data from previous example
data = np.load("sensor_readout.npy")

# Create a new dataclass instance that contains the default
# randomization distributions for each noise type
noise_rand = noise.NoiseRandomization.default()

# How much noise to add, relative to the transition height upper bound
noise_amount = 0.15

# Use a CorrelatedDistribution to randomize the white, pink, and
# telegraph noise so that the total always equals noise_amount
num_dists = 3
dists = distribution.SphericallyCorrelated(num_dists,
                                             noise_amount).dependent_distributions()
noise_rand.white_noise_magnitude = dists[0].abs()
noise_rand.pink_noise_magnitude = dists[1].abs()
noise_rand.telegraph_magnitude = dists[2].abs()

# Generate a random set of noise parameters
noise_params = noise.random_noise_params(noise_rand)

# Add noise to the data
noisy_data = noise.NoiseGenerator(noise_params).calc_noisy_map(data)

```

Listing 4: Example code to add noise to a charge stability diagram. First, a `NoiseRandomization` object is created, which defines the ranges and distributions, as appropriate, from which the noise parameters should be randomized. A `CorrelatedDistribution` is used to randomize the different noise types while ensuring that the total noise amount is constant.

434 as was done in Ref. 13.

435 5 Benchmarking and limitations

436 To benchmark QDFlow, we find the mean runtime of the physics simulation for 10,000 ran-
 437 domly generated sets of physics parameters. Each simulation is run with at most 1,000 it-
 438 erations and a relative convergence tolerance of 10^{-3} . We are interested in comparing the
 439 contribution to the overall runtime from calculating $n(x)$ with the remaining portion of the
 440 simulation as the number of QDs increases. For each number of dots N , we obtain two run-
 441 times: (i) the time required to compute the charge density profile $n(x)$, and (ii) the remaining
 442 time to construct the capacitance model, minimize Eq. (13), and evaluate the sensor output.
 443 The results for small- and mid-sized QD arrays are shown in Fig. 6(a) and Fig. 6(b), respec-
 444 tively.

445 For a small number of dots ($N \leq 20$), the main bottleneck is the initial part of the simulation
 446 where $n(x)$ is calculated. However, as the number of dots N is increased, the minimization of
 447 Eq. (13) over integer charge configurations $\vec{Q} \in \mathbb{Z}^N$ for which the runtime scales exponentially
 448 in N begins to dominate. In practice, we expect most use cases to involve no more than ten




Figure 6: Average runtime versus number of dots N for (a) Thomas-Fermi solver and (b) capacitance model solver. Benchmarks were performed on a 2.8 GHz AMD Ryzen 5 7520U processor.

449 QDs, as long-range interactions tend to be negligible due to screening from the material and
 450 the nearby gates; thus, we focus attention on the cost of computing $n(x)$ in this regime.

451 Because QDFlow calculates the capacitance model at each point in the CSD by explicitly
 452 solving for $n(x)$, it requires several orders of magnitude more runtime to generate CSDs
 453 than simulations based on constant capacitance models. Thus, for applications that require
 454 real-time data, faster simulations such as QDsim [17] or QArray [18] may be more suitable.
 455 On the other hand, for applications such as generating datasets for training and validation
 456 of ML models for autonomous control of QD systems, the higher-fidelity data provided by
 457 QDFlow—including effects such as QDs merging across low barriers—often justifies the additional
 458 computational cost.

459 In addition to its relatively long runtime, QDFlow has several modeling limitations. While
 460 QDFlow is based on an underlying physical nanowire model, it is designed only to produce
 461 qualitatively realistic behavior. For example, gates are modeled as infinite cylinders arranged
 462 in a simple, idealized geometry, whereas in real devices, the shapes and arrangements of gates
 463 can be considerably more complex. Furthermore, the electron density $n(x)$ is calculated only
 464 in 1D, an approximation that can introduce errors for systems with non-negligible transverse
 465 extent. Despite being a 1D simulation, data generated with QDFlow is qualitatively similar to
 466 that from 2D QD arrays, provided that crosstalk is not too large. Influence from other dots can
 467 be included either during the simulation phase by extending the nanowire to include additional
 468 dots on the left and right sides, or during the noise-adding phase as unintended dot noise.
 469 However, the ability to model crosstalk from other dots remains somewhat limited, as QDFlow
 470 cannot reproduce data features from 2D architectures with tight couplings between nearby
 471 dots. Finally, the simulation uses a semiclassical approach and therefore does not include
 472 certain quantum effects, such as finite tunnel coupling effects and charge-state hybridization,
 473 which would be captured by more complex models such as the Hubbard model. Taken together,
 474 these approximations prevent QDFlow from generating quantitatively accurate data.

475 QDFlow also lacks certain capabilities present in other simulators. For example, it cannot
 476 currently model closed systems with a fixed number of charges isolated from the source
 477 and drain contacts—a functionality available in QArray [18]. Additionally, QDFlow does
 478 not currently model quantum effects due to finite tunnel coupling, which can cause broad-
 479 ening of the interdot transitions and rounding of sharp corners around the triple-points due
 480 to hybridization of charge states. For applications involving virtualization of gates, quantum
 481 effects are generally small enough that a semiclassical approximation is sufficient. However,
 482 for applications such as readout, where fine control of gates is needed, quantum effects are
 483 more important. If these features are relevant, QDarts contains a simulation of the finite
 484 tunnel coupling effect [19]. Finally, QDFlow is configured to model the state that minimizes

485 the energy of the capacitance model, implicitly assuming that the system can transition freely
486 between charge configurations. However, this assumption can break down for large arrays of
487 QDs. For example, if a charge must transition between non-neighboring QDs passing through
488 an intermediate QD that forms a large barrier, it may get “stuck” on one side even if it is ener-
489 getically favorable to move to the other side. This will result in much more extreme latching
490 effects than QDFlow can model.

491 QDFlow has been tested both modularly and holistically, through extensive unit testing and
492 by comparing its output to experimental data. The QDFlow repository includes unit tests for
493 every function in the package, ensuring coverage of all lines of code. The validity of QDFlow
494 has been confirmed by performance of multiple specialized ML models trained on QDFlow-
495 generated data and deployed in real-world applications, as we discuss in Sec. 1 and in Sec. 6.

496 6 Conclusion

497 Progress toward scalable quantum information technologies based on QD systems depends
498 critically on overcoming the complexity of device operation and calibration with increasing
499 number of QDs. Novel ML-based methods have emerged as powerful tools to address these
500 challenges, but their effectiveness relies on access to large, diverse, and accurately labeled
501 datasets. QDFlow was developed precisely to meet this need.

502 QDFlow differs from existing QD simulations in that it fully simulates the charge density
503 function $n(x)$. By integrating a self-consistent Thomas-Fermi solver with a dynamic capac-
504 itance model, QDFlow provides a physics-informed simulation framework that goes beyond
505 constant-capacitance approximations. This enables the generation of CSDs and ray-based data
506 with features that evolve naturally with gate voltages, mimicking experimental behavior such
507 as dot merging and transition slope variation.

508 The modular data generation tools allow for extensive randomization over physical param-
509 eters, yielding highly diverse synthetic datasets suitable for ML applications, while the noise
510 module introduces experimentally relevant effects—including thermal broadening, telegraph
511 noise, latching, and unintended transitions—in a controllable fashion. Together, these fea-
512 tures make QDFlow uniquely positioned to support both the development and benchmarking
513 of ML algorithms implemented in a wide range of tuning procedures, device architectures,
514 and material platforms. Early use cases, such as the QFlow-lite and QFlow 2.0 datasets,
515 have already demonstrated QDFlow’s ability to accelerate the training of ML models for global
516 state recognition [9, 20], ray-based navigation and charge tuning [12, 13], data quality assess-
517 ment [25], detection of spurious QDs [31], and virtualization of QD arrays [13, 14]. As an
518 open-source, extensible platform with comprehensive documentation, QDFlow is designed to
519 serve as both a research tool and a community resource.

520 QDFlow represents a paradigm shift among QD simulators. Whereas other simulations
521 typically rely on constant-capacitance models that impose static couplings regardless of de-
522 vice state, QDFlow ties these parameters directly to the underlying physics through its self-
523 consistent charge density, producing capacitances and observables that evolve dynamically
524 with gate voltages. This distinction not only improves the connection with the experiment but
525 also allows for the capture of nontrivial behaviors—such as QDs merging, fluctuating slopes,
526 and disorder-induced effects—that are inaccessible to static-capacitance approaches.

527 As QD systems advance toward larger arrays and integration into functional quantum pro-
528 cessors, the need for such realism will only grow. QDFlow package is currently being used
529 to generate datasets for multiple ongoing research projects and we intend to support it for
530 the foreseeable future. Although the current release is stable, we anticipate adding further
531 modules and functionality to the package. Possible future extensions of QDFlow, including

532 multi-dimensional modeling, adding additional quantum effects, hybridization with exper-
533 imental feedback loops, and systematic studies of robustness under different noise and disor-
534 der regimes, could establish it as a cornerstone for bridging theory, experiment, and ML in
535 the quest for scalable quantum technologies. We have designed the core class structure to
536 allow for seamless addition of new physical parameters and functions, and plan to update the
537 generate module to provide support for parallelization. By releasing QDFlow as an open-
538 source package, we aim to foster a shared foundation for accelerating progress in quantum dot
539 technologies. We anticipate that the package will not only continue to advance automated QD
540 control but also provide a flexible testbed for exploring broader questions at the intersection
541 of condensed matter physics, quantum information, and machine learning.

542 Acknowledgements

543 **Funding information** D.B. was supported in part by an ARO grant no. W911NF-24-2-0043.
544 S.S.K. acknowledges financial support from the S.N. Bose Fellowship during this project. This
545 research was performed in part while J.Z. held an NRC Research Associateship award at NIST.
546 The views and conclusions contained in this paper are those of the authors and should not be
547 interpreted as representing the official policies, either expressed or implied, of the U.S. Govern-
548 ment. The U.S. Government is authorized to reproduce and distribute reprints for Government
549 purposes, notwithstanding any copyright noted herein. Any mention of commercial products
550 is for information only; it does not imply recommendation or endorsement by NIST.

551 **Code availability** QDFlow is available on the [Python Package Index](https://pypi.org/project/QDFlow/), with the source code re-
552 leased under the GNU General Public License, and can be installed using `pip install QD-`
553 `Flow` command. The associated GitHub repository is <https://github.com/QDFlow/QDFlow->
554 `solver`. Documentation is available at <https://qdflow-sim.readthedocs.io/>. Any discovered
555 bugs should be reported using GitHub [issues](#). If you find this package useful, please star the
556 repository and cite this paper.

557 References

- 558 [1] G. Burkard, T. D. Ladd, A. Pan, J. M. Nichol and J. R. Petta, *Semiconductor spin qubits*,
559 *Rev. Mod. Phys.* **95**(2), 025003 (2023), doi:[10.1103/RevModPhys.95.025003](https://doi.org/10.1103/RevModPhys.95.025003).
- 560 [2] A. B. Mei, I. Milosavljevic, A. L. Simpson, V. A. Smetanka, C. P. Feeney, S. M. Seguin, S. D.
561 Ha, W. Ha and M. D. Reed, *Optimization of quantum-dot qubit fabrication via machine*
562 *learning*, *Appl. Phys. Lett.* **118**(20), 204001 (2021), doi:[10.1063/5.0040967](https://doi.org/10.1063/5.0040967).
- 563 [3] C. Shen, W. Zhan, K. Xin, M. Li, Z. Sun, H. Cong, C. Xu, J. Tang, Z. Wu, B. Xu, Z. Wei,
564 C. Xue *et al.*, *Machine-learning-assisted and real-time-feedback-controlled growth of inas/-*
565 *gaas quantum dots*, *Nat. Commun.* **15**(1), 2724 (2024), doi:[10.1038/s41467-024-47087-w](https://doi.org/10.1038/s41467-024-47087-w).
- 566 [4] D. Schug, T. J. Kovach, M. A. Wolfe, J. Benson, S. Park, J. P. Dodson, J. Corrigan,
567 M. A. Eriksson and J. P. Zwolak, *Automation of quantum dot measurement analysis*
568 *via explainable machine learning*, *Mach. Learn.: Sci. Technol.* **6**(1), 015006 (2025),
569 doi:[10.1088/2632-2153/ada087](https://doi.org/10.1088/2632-2153/ada087).
- 570 [5] E. Corcione, F. Jakob, L. Wagner, R. Joos, A. Bisquerra, M. Schmidt, A. D. Wieck, A. Lud-
571 *wig*, M. Jetter, S. L. Portalupi, P. Michler and C. Tarín, *Machine learning enhanced evalua-*

573 *tion of semiconductor quantum dots*, Sci. Rep. **14**(1), 4154 (2024), doi:[10.1038/s41598-024-54615-7](https://doi.org/10.1038/s41598-024-54615-7).

574

575 [6] J. D. Teske, S. S. Humpohl, R. Otten, P. Bethke, P. Cerfontaine, J. Dedden, A. Ludwig, A. D. Wieck and H. Bluhm, *A machine learning approach for automated fine-tuning of semiconductor spin qubits*, Appl. Phys. Lett. **114**(13), 133102 (2019), doi:[10.1063/1.5088412](https://doi.org/10.1063/1.5088412).

576

577

578 [7] R. Durrer, B. Kratochwil, J. Koski, A. Landig, C. Reichl, W. Wegscheider, T. Ihn and E. Gre-
579 plova, *Automated tuning of double quantum dots into specific charge states using neural net-
580 works*, Phys. Rev. Appl. **13**(5), 054019 (2020), doi:[10.1103/PhysRevApplied.13.054019](https://doi.org/10.1103/PhysRevApplied.13.054019).

581

582 [8] J. Darulová, S. J. Pauka, N. Wiebe, K. W. Chan, G. C. Gardener, M. J. Manfra, M. C. Cassidy
583 and M. Troyer, *Autonomous tuning and charge-state detection of gate-defined quantum dots*,
Phys. Rev. Appl. **13**(5), 054005 (2020), doi:[10.1103/PhysRevApplied.13.054005](https://doi.org/10.1103/PhysRevApplied.13.054005).

584

585 [9] J. P. Zwolak, T. McJunkin, S. S. Kalantre, J. Dodson, E. R. MacQuarrie, D. Savage,
586 M. Lagally, S. Coppersmith, M. A. Eriksson and J. M. Taylor, *Autotuning of double-
587 dot devices in situ with machine learning*, Phys. Rev. Appl. **13**(3), 034075 (2020),
doi:[10.1103/PhysRevApplied.13.034075](https://doi.org/10.1103/PhysRevApplied.13.034075).

588

589 [10] N. M. van Esbroeck, D. T. Lennon, H. Moon, V. Nguyen, F. Vigneau, L. C. Camenzind,
590 L. Yu, D. M. Zumbühl, G. A. D. Briggs, D. Sejdinovic and N. Ares, *Quantum device
591 fine-tuning using unsupervised embedding learning*, New J. Phys. **22**(9), 095003 (2020),
doi:[10.1088/1367-2630/abb64c](https://doi.org/10.1088/1367-2630/abb64c).

592

593 [11] H. Moon, D. T. Lennon, J. Kirkpatrick, N. M. van Esbroeck, L. C. Camenzind, L. Yu,
594 F. Vigneau, D. M. Zumbühl, G. A. D. Briggs, M. A. Osborne, D. Sejdinovic, E. A. Laird
595 *et al.*, *Machine learning enables completely automatic tuning of a quantum device faster
596 than human experts*, Nat. Commun. **11**(1), 4161 (2020), doi:[10.1038/s41467-020-17835-9](https://doi.org/10.1038/s41467-020-17835-9).

597

598 [12] J. P. Zwolak, T. McJunkin, S. S. Kalantre, S. F. Neyens, E. R. MacQuarrie, M. A. Eriksson
599 and J. M. Taylor, *Ray-based framework for state identification in quantum dot devices*, PRX
Quantum **2**(2), 020335 (2021), doi:[10.1103/PRXQuantum.2.020335](https://doi.org/10.1103/PRXQuantum.2.020335).

600

601 [13] J. Ziegler, F. Luthi, M. Ramsey, F. Borjans, G. Zheng and J. P. Zwolak, *Tuning arrays
602 with rays: Physics-informed tuning of quantum dot charge states*, Phys. Rev. Appl. **20**(3),
034067 (2023), doi:[10.1103/PhysRevApplied.20.034067](https://doi.org/10.1103/PhysRevApplied.20.034067).

603

604 [14] A. S. Rao, D. Buterakos, B. van Straaten, V. John, C. X. Yu, S. D. Oosterhout, L. Stehouwer,
605 G. Scappucci, M. Veldhorst, F. Borsoi and J. P. Zwolak, *Modular Autonomous Virtualiza-
606 tion System for Two-Dimensional Semiconductor Quantum Dot Array*, Phys. Rev. X **15**(2),
021034 (2025), doi:[10.1103/PhysRevX.15.021034](https://doi.org/10.1103/PhysRevX.15.021034).

607

608 [15] G. A. Oakes, J. Duan, J. J. L. Morton, A. Lee, C. G. Smith and M. F. G. Zalba, *Auto-
609 matic virtual voltage extraction of a 2x2 array of quantum dots with machine learning*,
doi:[10.48550/arXiv.2012.03685](https://arxiv.org/abs/2012.03685) (2024).

610

611 [16] J. P. Zwolak, J. M. Taylor, R. W. Andrews, J. Benson, G. W. Bryant, D. Buterakos, A. Chat-
612 terjee, S. Das Sarma, M. Eriksson, E. Greplova, M. J. Gullans, F. Hader *et al.*, *Data needs
613 and challenges for quantum dot devices automation*, npj Quantum Inf. **10**(1), 105 (2024),
doi:[10.1038/s41534-024-00878-x](https://doi.org/10.1038/s41534-024-00878-x).

614

615 [17] V. Gualtieri, C. Renshaw-Whitman, V. Hernandes and E. Greplova, *Qdsim: A user-friendly
616 toolbox for simulating large-scale quantum dot devices*, SciPost Phys. Codebases p. 46
(2025), doi:[10.21468/SciPostPhysCodeb.46](https://doi.org/10.21468/SciPostPhysCodeb.46).

617 [18] B. van Straaten, J. Hickie, L. Schorling, J. Schuff, F. Fedele and N. Ares, *Qarray: A gpu-*

618 accelerated constant capacitance model simulator for large quantum dot arrays, SciPost

619 Phys. Codebases p. 35 (2024), doi:[10.21468/SciPostPhysCodeb.35](https://doi.org/10.21468/SciPostPhysCodeb.35).

620 [19] J. A. Krzywda, W. Liu, E. van Nieuwenburg and O. Krause, *QDarts: A quantum dot array*

621 *transition simulator for finding charge transitions in the presence of finite tunnel couplings,*

622 *non-constant charging energies and sensor dots*, SciPost Phys. Codebases p. 43 (2025),

623 doi:[10.21468/SciPostPhysCodeb.43](https://doi.org/10.21468/SciPostPhysCodeb.43).

624 [20] S. S. Kalantre, J. P. Zwolak, S. Ragole, X. Wu, N. M. Zimmerman, M. D. Stewart and J. M.

625 Taylor, *Machine learning techniques for state recognition and auto-tuning in quantum dots*,

626 npj Quantum Inf. **5**(1), 6 (2019), doi:[10.1038/s41534-018-0118-7](https://doi.org/10.1038/s41534-018-0118-7).

627 [21] J. P. Zwolak, S. S. Kalantre, X. Wu, S. Ragole and J. M. Taylor, *QFlow lite dataset: A*

628 *machine-learning approach to the charge states in quantum dot experiments*, PLoS ONE

629 **13**(10), e0205844 (2018), doi:[10.1371/journal.pone.0205844](https://doi.org/10.1371/journal.pone.0205844).

630 [22] National Institute of Standards and Technology, *Qflow 2.0: Quantum dot data for machine*

631 *learning*, Database: data.nist.gov, <https://doi.org/10.18434/T4/1423788> (2022).

632 [23] A. Zubchenko, D. Middlebrooks, T. Rasmussen, L. Lausen, F. Kuemmeth, A. Chatterjee

633 and J. P. Zwolak, *Autonomous bootstrapping of quantum dot devices*, Phys. Rev. Appl.

634 **23**(1), 014072 (2025), doi:[10.1103/PhysRevApplied.23.014072](https://doi.org/10.1103/PhysRevApplied.23.014072).

635 [24] J. P. Zwolak, S. S. Kalantre, T. McJunkin, B. J. Weber and J. M. Taylor, *Ray-based classifi-*

636 *cation framework for high-dimensional data*, In *Third Workshop on Machine Learning and*

637 *the Physical Sciences (NeurIPS 2020)*, pp. 1–7. Vancouver, Canada, ArXiv:2010.00500

638 (2020).

639 [25] J. Ziegler, T. McJunkin, E. S. Joseph, S. S. Kalantre, B. Harpt, D. E. Savage,

640 M. G. Lagally, M. A. Eriksson, J. M. Taylor and J. P. Zwolak, *Toward robust au-*

641 *totuning of noisy quantum dot devices*, Phys. Rev. Appl. **17**(2), 024069 (2022),

642 doi:[10.1103/PhysRevApplied.17.024069](https://doi.org/10.1103/PhysRevApplied.17.024069).

643 [26] D. L. Buterakos, S. S. Kalantre, J. Ziegler, J. M. Taylor and J. P. Zwolak, *QDFlow: A*

644 *Python package for physics simulations of quantum dot devices*, GitHub (2025).

645 [27] O. Ciftja, *Electrostatic interaction energy between two coaxial parallel uniformly charged*

646 *disks*, Results Phys. **15**, 102684 (2019), doi:[10.1016/j.rinp.2019.102684](https://doi.org/10.1016/j.rinp.2019.102684).

647 [28] A. Alexandradinata, N. P. Armitage, A. Baydin, W. Bi, Y. Cao, H. J. Changlani, E. Chertkov,

648 E. H. da Silva Neto, L. Delacretaz, I. E. Baggari, G. M. Ferguson, W. J. Gannon *et al.*,

649 *The future of the correlated electron problem*, SciPost Phys. Comm. Rep. p. 8 (2025),

650 doi:[10.21468/SciPostPhysCommRep.8](https://doi.org/10.21468/SciPostPhysCommRep.8).

651 [29] J. Darulová, M. Troyer and M. C. Cassidy, *Evaluation of synthetic and experimental train-*

652 *ing data in supervised machine learning applied to charge-state detection of quantum dots*,

653 Mach. Learn.: Sci. Technol. **2**(4), 045023 (2021), doi:[10.1088/2632-2153/ac104c](https://doi.org/10.1088/2632-2153/ac104c).

654 [30] C. W. J. Beenakker, *Theory of coulomb-blockade oscillations in the conductance of a quan-*

655 *tum dot*, Phys. Rev. B **44**, 1646 (1991), doi:[10.1103/PhysRevB.44.1646](https://doi.org/10.1103/PhysRevB.44.1646).

656 [31] J. Ziegler, F. Luthi, M. Ramsey, F. Borjans, G. Zheng and J. P. Zwolak, *Automated extraction*

657 *of capacitive coupling for quantum dot systems*, Phys. Rev. Appl. **19**(5), 054077 (2023),

658 doi:[10.1103/PhysRevApplied.19.054077](https://doi.org/10.1103/PhysRevApplied.19.054077).