AIMS Mathematics, 11(11): 38963–38993.
DOI: 10.3934/math.20261502
Received: 07 September 2026
Revised: 07 November 2026
Accepted: 17 November 2026
Published: 23 November 2026

https://www.aimspress.com/journal/Math

Research article

Oracle-free goal-oriented reliability for charge and gate capacitance in reduced Schrödinger–Poisson models via a compressed dual operator

Ahmed Alqurashi*

Department of Electrical Engineering, College of Engineering and Architecture, Umm Al-Qura University, Makkah, Saudi Arabia

* **Correspondence:** Email: akqurashi@uqu.edu.sa.

Abstract: We presented an online goal-oriented estimator of reduction errors in silicon sheet charge and its common-gate bias derivative for a self-consistent Schrödinger–Poisson model. A charge-trained compressed augmented adjoint uses one dual operator for both outputs. The method combines established tangent and dual-weighted-residual constructions with an empirical tolerance controller. We selected the controller from 12,100 candidates and validated it independently on 1024 physical tuples (8192 algorithmic states). The frozen primary rule produced 1 false acceptance and 6 false rejections; the historical comparator produced 0 false acceptances and 58 false rejections. In the independent 0.5–2% capacitance-error band (737 states), effectivity ranged from 0.533 to 1.363. Rank-nine derivative-adjoint projection error reached 7.56% on 25 diagnostic cases. Measured median paired latency reductions were 5.84%–13.80% on three cases. A separate interval procedure using full-model information verified selected output statements, including capacitance-dominated failures just above the 1% tolerance. All numerical evidence used explicitly spin-degenerate occupations and rebuilt offline assets. These certificates and sampled decisions did not establish uniform controller reliability or continuum accuracy.

Keywords: goal-oriented error estimation; reduced-order modeling; derivative quantity of interest; dual-weighted residual; verification and validation; Schrödinger–Poisson
Mathematics Subject Classification: 65N15, 65N25, 65N30, 35Q81, 82D37

1. Introduction

Projection-based model reduction is widely used for many-query parametrized systems. Certified reduced-basis methods combine offline–online decomposition with computable a posteriori error estimates, allowing reduced outputs to support design and decision-making as well as fast evaluation [1–3]. Reliability must nevertheless be assessed for the quantity used in an engineering decision. This is particularly important when that quantity is a parameter derivative rather than a

state-like functional.

Reduced models of parametrized nonlinear systems increasingly support decisions that depend on parameter *derivatives* of model outputs. Derivative accuracy depends on both the reduced state and its parametric sensitivity, so accurate states and primary outputs can still yield inaccurate derivatives. Goal-oriented (adjoint-weighted) error estimation addresses the reliability of reduced outputs, and finite-element verification has treated goal functionals defined on sensitivity fields. The online task considered here is to provide an inexpensive, oracle-free estimate of derivative error alongside the reduced output, and to use that estimate for an accept/reject decision against an engineering tolerance. We develop and stress-evaluate this approach for nonlinear, self-consistent, eigenvalue-coupled systems, using quantum-confined semiconductor electrostatics, where the derivative is an engineering deliverable. Reduced Schrödinger–Poisson (SP) models have existed for over fifteen years [4, 5]. Their gate capacitance $C_{\mathrm{gg}} = dQ/dV_g$ sets switching charge, delay, and the C–V signatures used to calibrate compact device models [6, 7]. We ask whether a user can assess, online and without recourse to the full model, whether the returned capacitance is within tolerance.

The method combines a tangent-based derivative output with an augmented adjoint in a charge-trained compressed space and a controller for reduction-error budgets of 0.5% on Q and 1% on C_{gg}. The block-triangular augmented adjoint allows both goals to share the dual operator and its factorization. Online evaluation requires no full-order output; full-order information is used in offline training, controller development, and post hoc validation. We select the primary rule systematically on development data and then freeze it before validation on separately drawn physical tuples (Section 6.1). The decisions are empirical: acceptance does not guarantee that every unseen state satisfies the budgets.

The contribution is an application-specific integration of established components. Schrödinger–Poisson reduced-basis error estimation [4] and dual-weighted-residual constructions [8–10] supply part of the foundation. Applications to reduced models [11, 12], analysis of nonlinear goal functionals [13], sensitivity-related output estimation [14, 15], and repeated-Jacobian tangent systems provide further components. The present realization combines a charge-trained dual space, direct derivative-output error estimation, shared operator construction, and a tolerance decision for a nonlinear eigenvalue-coupled device model. The conditional local compression argument explains a possible mechanism; it is neither a new general certification theorem nor a proof that rank nine suffices uniformly.

The revised calculations include an explicit spin degeneracy of two in addition to the silicon conduction-band valley multiplicities. The earlier occupation prefactor omitted the spin factor; we have corrected it consistently in the reference, reduced, baseline, and verification models. All reported training snapshots, reduced and dual bases, regression coefficients, numerical campaigns, and interval certificates have been regenerated. We identify previously examined physical designs as reused designs and exclude a newly fixed 1024-tuple validation set from training and controller selection. Checks of reference convergence, physical profiles, adjoint projections, estimator sensitivity, and complete online timing make the scope of the results explicit.

Section 2 reviews related work. Section 3 specifies the discrete model, reduced approximation, estimator, controller, and optional local verification. Sections 4–5 describe the computational protocol and reference sensitivity. Section 6 reports controller validation, effectivity, coverage, physical and algebraic diagnostics, timing, and selected-point certificates. Sections 7–8 discuss the implications and conclude the paper.

2. Related work

Reduced-order quantum device electrostatics: Pau [4] coupled the reduced-basis method with subband decomposition for self-consistent Schrödinger–Poisson nanodevice simulation, using a posteriori error estimation and greedy sampling to construct the basis. Li and Li [5] developed fixed- and free-interface component-mode-synthesis reductions of the Schrödinger–Poisson system for nanoscale electrostatic analysis. These studies established prior Schrödinger–Poisson reduction and error estimation. Here we focus on a direct bias-derivative goal, a charge-trained compressed augmented adjoint, and a shared-operator tolerance workflow.

Quantum-confined charge and capacitance modeling: Explicit physics-based models of quantum-confined charge and gate capacitance are well-developed for gate-all-around and related geometries. Examples include the III–V nanowire model of Ganeriwala et al. [6] and recent surface-potential-based compact models with subband quantization [7]. These works motivate the choice of gate capacitance as an output and provided downstream uses for the reliability assessment developed here. They did not address the reliability of reduced numerical models.

Physical and numerical foundations of the semiconductor model: Effective-mass treatment across material interfaces follows the BenDaniel–Duke formulation [16]. Stern [17] established self-consistent quantized inversion-layer electrostatics in silicon, later comprehensively reviewed by Ando, Fowler, and Stern [18]. Trellakis et al. [19] developed robust iterative solution strategies for coupled Schrödinger–Poisson systems. Quantum displacement of inversion charge has long been known to alter intrinsic metal–oxide–semiconductor (MOS) capacitance and hence circuit-level compact-model behavior [20]. These results motivate the fidelity requirements of the full-order model and the need to assess the charge derivative as well as the state.

Goal-oriented reliability for reduced-order models (ROMs): Janon, Nodet, and Prieur [14] developed goal-oriented error bounds for reduced-basis outputs with applications to sensitivity analysis. In that work, "sensitivity" denotes Sobol-index computation from bounds on a linear state functional, a different object from the derivative-error estimation considered here. Blonigan and Parish [11] evaluated dual-weighted-residual (DWR) error estimation for projection-based ROMs of steady partial differential equations (PDEs), including high-consequence decision settings. DWR has also been extended to space–time goal-oriented control of proper orthogonal decomposition (POD) ROMs [12]. In electrical engineering, Rewieński, Lamecki, and Mrozowski [21] demonstrated output-specific a posteriori error estimation for reduced-basis microwave models. Further work developed a posteriori estimators and goal-oriented adaptive sampling for parametric ROMs [22, 23]. Granzow, Seidl, and Bond [13] analyzed linearization errors in discrete goal-oriented estimates of nonlinear functionals, directly relevant to the second-derivative coupling term in our augmented adjoint. These studies provide the general tools. Our online goal is a bias derivative of a self-consistent Schrödinger–Poisson solution, evaluated within the specified compressed shared-operator decision architecture.

Derivative-observable goal functionals: The closest general precedent is by Guo and Zhong [15], who obtained strict upper and lower bounds for quantities of interest associated with response-sensitivity derivative fields in static finite-element analysis. Their bounds concern discretization error in linear static response. We assess reduction error online for a nonlinear self-consistent eigenproblem, using a compressed dual space, shared dual-operator construction across goals, and a decision controller. Derivative-goal error estimation is therefore established in general; our contribution is its realization in

an online ROM reliability architecture for Schrödinger–Poisson electrostatics.

Here, compressed online dual estimation targets the direct bias derivative of a self-consistent Schrödinger–Poisson solution. The charge and derivative goals share operator construction and are assessed by an empirical tolerance controller. The contribution lies in this implementation and the evaluation of its behavior, with the local hypotheses and limitations stated explicitly.

Open boundaries and alternative quantum discretizations: Arnold [24] treated open quantum boundary conditions for injection and extraction. Our coordinate describes transverse equilibrium confinement and does not represent longitudinal transport. Arnold, Ben Abdallah, and Negulescu [25] used Wentzel–Kramers–Brillouin (WKB) filtering to reduce the resolution needed for rapidly oscillatory semiclassical Schrödinger solutions. Wang and Shu [26] combined a WKB local discontinuous Galerkin treatment of the Schrödinger equation with a Poisson discretization for a resonant tunneling diode. These methods resolve spatial oscillations, whereas our ROM reduces the cost of repeated parameter evaluations and estimates output reduction error. The previously mentioned methods could provide an alternative reference discretization, but combining them with the ROM would require a new residual, derivative implementation, bases, and validation.

Aliffi, Nastasi, and Romano [27, 28] studied fourth- and higher-order Schrödinger descriptions of ballistic transport with non-parabolic dispersion and transparent boundary treatments. Our results use a second-order effective-mass Hamiltonian and do not validate those models. For a differentiable higher-order algebraic residual, the augmented state–tangent Jacobian retains its repeated diagonal block. Additional boundary equations, higher-order interface conditions, and any energy- or parameter-dependent boundary operators must enter the state and all Jacobian and curvature derivatives. The current bases, rank, calibration, and cost results cannot be transferred unchanged.

3. Model, estimator, and controller

3.1. Physical model and full-order reference

We consider a symmetric Si(100) double-gate structure. The declared in-domain parameter range is $t_{\mathrm{Si}} \in [3, 8]$ nm, $t_{\mathrm{ox}} \in [0.8, 1.2]$ nm, $V_g \in [0.15, 0.90]$ V, and $E_F \in [-0.30, +0.30]$ eV. Additional states extending to $t_{\mathrm{Si}} = 2.6$–9.0 nm, $V_g = 1.05$ V, and $|E_F| = 0.36$ eV are used only as out-of-domain (OOD) stress probes and are excluded from formal coverage claims.

The electrostatics are governed by the one-dimensional self-consistent Schrödinger–Poisson system. For each conduction-band valley ν, envelope states satisfy the position-dependent effective-mass (BenDaniel–Duke [16]) Schrödinger equation on the extended oxide–silicon–oxide domain,

$$-\frac{\hbar^2}{2}\frac{d}{dx}\left[\frac{1}{m_\nu(x)}\frac{d\psi_{\nu n}}{dx}\right] + [E_c(x) - q\varphi(x)]\psi_{\nu n} = E_{\nu n}\psi_{\nu n}. \tag{3.1}$$

We set $\psi = 0$ at the outer oxide boundaries. The quantum grid contains $N_Q = 121$ uniformly spaced nodes, leaving 119 interior eigenvector unknowns for each valley. Across the silicon/oxide interfaces, the material coefficients change abruptly, and we use the conservative face stencil defined below. The oxide barrier is $\Delta E_c = 3.10$ eV and the oxide effective mass is $0.50\, m_0$. On the chosen energy reference, the unshifted silicon conduction-band edge is $E_{c,\,\mathrm{Si}} = 0.45$ eV. We calculate subband populations with

two-dimensional Fermi–Dirac statistics at 300 K:

$$N_{\nu n} = \frac{g_s g_\nu m_\nu^{\mathrm{dos}} kT}{2\pi\hbar^2} \ln\left[1 + \exp\left(\frac{E_F - E_{\nu n}}{kT}\right)\right]. \tag{3.2}$$

The occupation prefactor in Eq (3.2) separates spin degeneracy $g_s = 2$ from the valley factors $g_{\Delta_2} = 2$ and $g_{\Delta_4} = 4$. All occupations and their derivatives include both multiplicities. The spatial electron density is $n(x) = \sum_{\nu,n} N_{\nu n} |\psi_{\nu n}(x)|^2$. Integrating the density over silicon gives the positive electron sheet-charge magnitude $Q = q \int_{\mathrm{Si}} n(x)\, dx$ in C m^{-2}; the signed electron charge is $-Q$. When the two gates vary together at fixed E_F, the common-gate derivative is $C_{gg} = dQ/dV_g$, in F m^{-2}. This is a derivative of the magnitude of silicon charge, not a separately evaluated terminal-capacitance matrix for each gate. Wavefunctions penetrate the oxide, but oxide probability is excluded from Q. The potential satisfies Poisson's equation on 81 silicon nodes:

$$-\varepsilon_{\mathrm{Si}} \frac{d^2\varphi}{dx^2} = -qn(x). \tag{3.3}$$

At the two silicon interfaces, the boundary condition is $\varepsilon_{\mathrm{Si}}\, \partial\varphi/\partial n = c_{\mathrm{ox}} \left(V_g - \varphi\right)$, where n is outward-facing and $c_{\mathrm{ox}} = \varepsilon_{\mathrm{ox}}/t_{\mathrm{ox}}$. The model describes transverse equilibrium confinement and quasi-static common-gate charge response. The model does not calculate longitudinal injection, extraction, or transport current. E_F is the specified carrier chemical-potential offset; the Hamiltonian potential energy is $E_c - q\varphi$. Physical constants are listed in Table 1.

Table 1. Physical parameters of the full-order reference model.

Quantity	**Value**
Valleys (Si(100), 2-D DOS)	Δ_2: $m_{\mathrm{conf}} = 0.916m_0$, $m_{\mathrm{dos}} = 0.19m_0$, $g = 2$; Δ_4: $m_{\mathrm{conf}} = 0.19m_0$, $m_{\mathrm{dos}} = \sqrt{0.916 \cdot 0.19}m_0$, $g = 4$
Spin degeneracy	$g_s = 2$; $g_\nu = 2, 4$ (valleys only)
Dielectrics	$\varepsilon_{\mathrm{Si}} = 11.7\varepsilon_0$; $\varepsilon_{\mathrm{ox}} = 3.9\varepsilon_0$
Oxide barrier / mass	$\Delta E_c = 3.10$ eV; $m_{\mathrm{ox}} = 0.50m_0$
Temperature / statistics	$T = 300$ K; 2-D Fermi–Dirac per subband, Eq (3.2)
Quantum BCs	$\psi = 0$ at outer oxide boundaries; BenDaniel–Duke interface flux continuity
Poisson BCs	Symmetric Robin gate coupling, $c_{\mathrm{ox}} = \varepsilon_{\mathrm{ox}}/t_{\mathrm{ox}}$
Silicon band-edge reference	$E_{c,\,\mathrm{Si}} = 0.45$ eV before electrostatic shift $-q\varphi$

The full-order reference retains 12 subbands per valley. Starting from the uniform gate potential, the deterministic damped fixed-point solver tries mixing factors 0.10, 0.05 and 0.025 in that order, with up to 4000 iterations per attempt. A solve is retained only if its maximum absolute residual meets the requested tolerance. We record failed attempts and the damping factor chosen. This safeguard changes the iteration, not the equations. One tuple in the corrected 180-tuple training set requires the second solver attempt.

Our reduction-error tolerances refer to the fixed discrete reference ($N_Q = 121$, $N_{\mathrm{Si}} = 81$, $M_{\mathrm{ref}} = 12$). They do not bound spatial discretization or physical-model error. Section 5 examines sensitivity to

spatial and interface sampling through a separate grid study, without providing a continuum-error enclosure. The interval procedure in Section 3.5 instead gives a local statement about selected frozen discrete systems.

All main corrected truth labels and effectivities use Q from the tight reference solve and a Richardson derivative formed from central differences at $h = 10^{-4}$ and 2×10^{-4} V, with residual tolerance 2×10^{-12} V. Section 5 compares this derivative with the production central difference ($h = 2\times10^{-4}$ V, tolerance 2×10^{-10} V) and additional steps. For the diagnostic grid and step sweeps, we started with the smaller mixing factor 0.05 and then tried 0.025 and 0.0125 because the larger mixing factor repeatedly failed. The equations and stopping tolerances remained unchanged. The interval experiments enclosed the exact implicit derivative of their frozen discrete equations.

Algebraic residual and quantum discretization

The algebraic state $u \in \mathbb{R}^{81}$ contains only the silicon nodal potentials in volts. Eigenvalues, eigenvectors, occupations, and density are eliminated by evaluation of the quantum problem; they are not independent components of u. Set $\Delta x = t_{\mathrm{Si}}/80$. The Poisson matrix A_{P} has interior row coefficients $(-\varepsilon_{\mathrm{Si}}/\Delta x^2, 2\varepsilon_{\mathrm{Si}}/\Delta x^2, -\varepsilon_{\mathrm{Si}}/\Delta x^2)$, first-row coefficients $(\varepsilon_{\mathrm{Si}}/\Delta x + c_{\mathrm{ox}}, -\varepsilon_{\mathrm{Si}}/\Delta x)$, and the reflected last row. Let $b_0 = b_{80} = c_{\mathrm{ox}}\ V_g$ and $b_i = -q\ n_i$ for $1 \le i \le 79$. The implemented fixed-point residual is $F_{\mathrm{ref}}(u,\ V_g) = u - A_P^{-1}\ b(n_{\mathrm{ref}}(u,\ V_g),\ V_g)$.

This inverse-Poisson scaling gives every residual component units of volts despite the different physical units of boundary and interior Poisson rows. $J_{\mathrm{ref}} = \partial_u F_{\mathrm{ref}}$ is dimensionless. The stopping test is the maximum absolute fixed-point residual in volts; the damping safeguards are specified above. No additional nondimensionalization is applied inside the solver. Lengths and energies supplied in nm and eV are converted to SI for the operators.

On $x_j = -t_{\mathrm{ox}} + j\ \Delta x_q$, $\Delta x_q = (t_{\mathrm{Si}} + 2\ t_{\mathrm{ox}})/120$, define $\alpha_{j+1/2} = \left(1/m_j + 1/m_{j+1}\right)/2$ and $\tau_{j+1/2} = \hbar^2\alpha_{j+1/2}/\left(2\Delta x_q^2\right)$. The interior Hamiltonian has the diagonal $\tau_{j-1/2} + \tau_{j+1/2} + E_{c,j} - q\varphi_{q,j}$ and off-diagonal $-\tau_{j+1/2}$. Nodes in the closed silicon interval use the silicon coefficients; others use oxide coefficients. This conservative flux stencil uses an arithmetic average of inverse mass and does not fit every interface to a grid face. Interface sampling contributes to the non-monotone grid sensitivity in Section 5.

We linearly interpolate the silicon potential to the quantum nodes and extend it linearly through each oxide to the gate value. To normalize over the extended domain, we divide the Euclidean-normalized interior eigenvectors by $\sqrt{\Delta x_q}$ and add zero endpoints, giving $\Delta x_q \sum_j \left|\psi_j\right|^2 = 1$. We then linearly interpolate the wave amplitudes back to silicon nodes before squaring them. Silicon integrals use trapezoidal quadrature without renormalizing the probability inside silicon. The reference diagonalizes the 119×119 tridiagonal matrix and retains its lowest 12 eigenstates per valley.

The reliability residual F_{rel} corresponds to a distinct electrostatic model that retains all silicon potential nodes. Its quantum Hamiltonian is projected onto the lowest 12 zero-field finite-barrier eigenvectors W_{12} per valley, and all 12 eigenstates of this projected matrix are used. Thus $H_{12} = \Delta x_q\ W_{12}^{\mathrm{T}}\ H\ W_{12}$ and $F_{\mathrm{rel}}(u,\ V_g) = u - A_P^{-1} b(n_{12}(u,\ V_g),\ V_g)$, with the Jacobian $J_{\mathrm{rel}} = \partial_u F_{\mathrm{rel}}$. The primal model analogously uses W_{10}. Retaining 12 states of the full-grid Hamiltonian differs from using a fixed 12-vector quantum space. Section 3.3.1 explicitly accounts for this model discrepancy.

We compute residual and charge derivatives by centered differences, re-diagonalizing at each perturbation. This differentiates the occupied density matrix rather than individual eigenvector phases

or eigenvalue labels. Within the fixed reduced space, $\Gamma = C\ \mathrm{diag}(f(\lambda))C^{\mathrm{T}}$ is invariant to rotations of an exactly degenerate eigenspace because f is the same smooth Fermi occupation there. The implementation therefore requires no division by the gap between near-coincident eigenvalues. Smoothness of the truncated full-grid reference additionally requires an isolated retained spectral cluster at the cutoff. We do not claim a uniform spectral-gap or Jacobian-invertibility result over the parameter domain.

3.2. Reduced-order model and derivative output

The primal potential is $u = \varphi_0 1 + B_p a$, with the $81 \times K_p$ matrix B_p containing Euclidean-orthonormal POD vectors, $K_p \in \{1, 2, 3, 4\}$. The reduced coefficient vector a is measured in volts. The reduced residual is $B_p^{\mathrm{T}} F_{10}$. One or two chord updates use the Jacobian at the initialized coefficients. The tangent instead uses the refreshed reduced Jacobian at the final coefficients: $J_p a_V = -r_V$. It defines the implicit reduced tangent, not the derivative of the finite-count chord algorithm. The output is $C_{\mathrm{gg}} = Q_V + Q_a a_V$ for the $M = 10$ charge. The bias dependence of φ_0 is included in r_{V} and Q_{V}. All derivatives are centered differences: $\mathrm{FD_V} = 2{\times}10^{-5}$ V for bias derivatives in the tangent; $\mathrm{FD_{small}} = 2{\times}10^{-6}$ V for coefficient, dual-coordinate, and bias-direction perturbations in output/reliability evaluation; $\mathrm{FD}_\varphi = 2{\times}10^{-6}$ V for nodal perturbations in offline full-node adjoints. The primal Jacobian uses a separate coefficient step of 10^{-6} V.

Here φ_0 is not an independent model parameter. For each $\left(t_{\mathrm{Si}},\ V_g,\ t_{\mathrm{ox}},\ E_F\right)$, we obtain it from the uniform-silicon charge balance $2c_{\mathrm{ox}}\left(V_g - \varphi_0\right) = q\sum_{v,\,n} N_{vn}(\varphi_0)$, evaluated with the finite-barrier zero-field subband energies and 12 states per valley, so the gate bias enters the feature vector through φ_0. We use normalized coordinates $z = ((t_{\mathrm{Si}} - 5.5)/2.5,\ \ (t_{\mathrm{ox}} - 1.0)/0.2,\ \ (\varphi_0 - 0.50)/0.40,\ \ E_F/0.30)$ and the full quadratic expansion $\{1,\ z_i,\ z_i z_j\}$, which gives 15 features. The coefficients were trained on the 180-state tensor design $t_{\mathrm{Si}} \in \{3,\ 5.5,\ 8\}$ nm, $V_g \in \{0.15,\ 0.40,\ 0.65,\ 0.90\}$ V, $t_{\mathrm{ox}} \in \{0.8,\ 1.0,\ 1.2\}$ nm, $E_F \in \{-0.30,\ -0.15,\ 0,\ 0.15,\ 0.30\}$ eV.

The feature order is $1, z_1, \ldots, z_4$, followed by $z_i\, z_j$ in lexicographic order with $i \le j$. The 81×4 primal basis is obtained by an unweighted Euclidean singular value decomposition (SVD) of the corrected reference potentials minus their uniform baselines. Ordinary least squares maps the 15 features to the four projected coefficients; its 15×4 matrix initializes a. The full training procedure and resulting numerical arrays are supplied. No evaluation or validation output enters this construction. The uniform baseline root is solved on $[-1.5, V_g]$ V with the absolute root tolerance 10^{-13} V. Its charge balance uses zero-field subband occupations, including both spin and valley factors, while the final Q includes silicon probability weights.

3.3. Compressed reliability model and augmented derivative-goal estimation

Charge-adjoint snapshots $J_{\mathrm{rel}}^{-\mathrm{T}}(\partial_u Q_{12})^{\mathrm{T}}$ are evaluated at the 180 corrected $K_p = 4$, two-stage training states. Their standard, unweighted Euclidean SVD gives an $81 \times r$ matrix Z with $Z^{\mathrm{T}}Z = I$; $r = 9$ is the retained choice. Full adjoints $z_1, z_2 \in \mathbb{R}^{81}$ and compressed coefficients $\widehat{z_1}, \widehat{z_2} \in \mathbb{R}^r$ are distinguished throughout. Reconstructed compressed adjoints are $Z\widehat{z_j}$; these Galerkin solutions need not equal orthogonal projections of the full adjoints. The charge equation is:

$$\eta_Q = -\widehat{z_Q}^{\mathrm{T}} r_Z, \qquad A^{\mathrm{T}}\widehat{z_Q} = g_Q \tag{3.4}$$

For the reliability model only, we write $F = F_{\mathrm{rel}}$, $Q = Q_{12}$, $J = F_u$, and s for the reconstructed reduced tangent. The augmented goal is $C(u, s; V) = Q_u s + Q_V$, with its u derivative taken at fixed s. The full adjoints satisfy (3.5) and (3.6):

$$J^{\mathrm{T}} z_2 = Q_u^{\mathrm{T}}, \tag{3.5}$$

$$J^{\mathrm{T}} z_1 = (Q_{uu}[s] + Q_{uV})^{\mathrm{T}} - (F_{uu}[s] + F_{uV})^{\mathrm{T}} z_2. \tag{3.6}$$

Define $K = F_{uu}[\cdot, s] + F_{Vu}$, meaning $K\delta u = F_{uu}[\delta u, s] + F_{Vu}\delta u$. Then $K_2 = Z^{\mathrm{T}} K Z$, $g_Q = Z^{\mathrm{T}} Q_u^{\mathrm{T}}$, and $g_C = Z^{\mathrm{T}} C_u^{\mathrm{T}} = Z^{\mathrm{T}}(Q_{uu}[\cdot, s] + Q_{Vu})^{\mathrm{T}}$. Also $r_Z = Z^{\mathrm{T}} F$ and $t_Z = Z^{\mathrm{T}}(Js + F_V)$. With $A = Z^{\mathrm{T}} J Z$, the compressed equations reuse the charge coefficients $\widehat{z_2} = \widehat{z_Q}$:

$$A^{\mathrm{T}} \widehat{z_1} = g_C - K_2^{\mathrm{T}} \widehat{z_2} \tag{3.7}$$

and the derivative-goal error estimate is:

$$\eta_C = -\left(\widehat{z_1}^{\mathrm{T}} r_Z + \widehat{z_2}^{\mathrm{T}} t_Z\right) \tag{3.8}$$

The complete oracle-free online workflow, including the shared dual solve and its separation from post hoc full-order evaluation, is summarized in Figure 1.

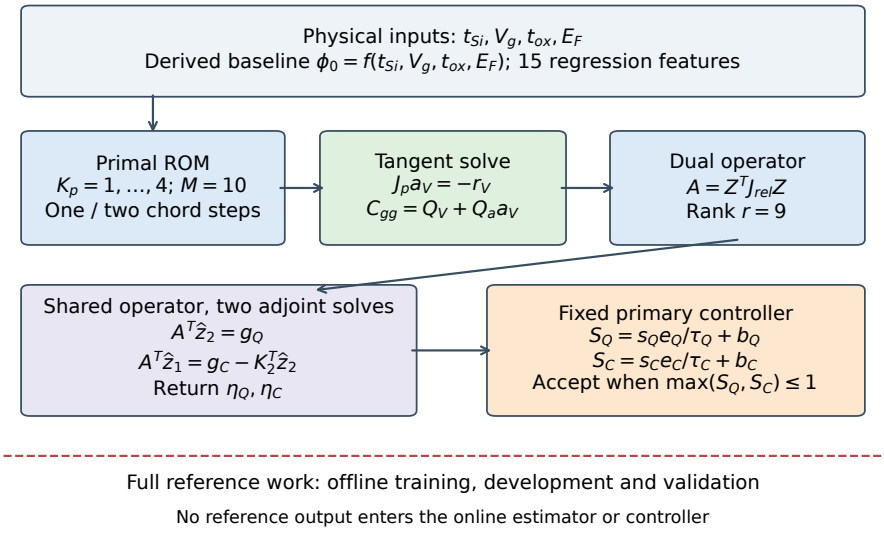


Figure 1. Online workflow. The uniform baseline φ_0 is derived from the four physical inputs. The charge and derivative goals share one compressed operator; the timed implementation of Section 6.7 also reuses its lower–upper (LU) factorization. Reference outputs are excluded from online evaluation; they are used in offline construction, controller development, and post hoc validation. Primary controller parameters are fixed before both validation campaigns in Section 6.1.

The algebra permits one LU factorization of A^{T} and two solves, with additional derivative-specific curvature and goal-gradient assembly. The released campaign evaluator shares the assembled dual operator but makes separate general linear-solver calls for the two adjoints. The optimized

implementation timed in Section 6.7 factors A^{T} once and reuses its LU factors. Both implementations solve the same adjoint equations. We supply the numerical solver and verification source code, rebuilt offline arrays, input designs, and corrected row-level results. The timing comparison in Section 6.7 includes curvature assembly and the full online path. Historical results from the model with the omitted spin factor are excluded from the corrected numerical evidence.

3.3.1. Conditional first-order representation and compression

Proposition 1. *At fixed bias, let $F = F_{\mathrm{rel}}$ and $Q = Q_{12}$ be C^3 near the segment joining $U_a = (u_a, s_a)$ and an exact augmented reliability solution $U^* = (u^*, s^*)$. Define $R(U) = (F(u, V), F_u(u, V)s + F_V(u, V))$, and $C(U) = Q_u(u, V)s + Q_V(u, V)$. Assume $J = F_u(u_a, V)$ is nonsingular. All vectors use Euclidean inner products in fixed numerical units: potentials in volts, tangents in V/V, and C in* $\mathrm{F\,m^{-2}}$. *Equivalently, we use $u/(1\,\mathrm{V})$, $s/(1\,\mathrm{V/V})$, and $C/(1\,\mathrm{F\,m^{-2}})$ to define dimensionless product norms. The operator and multilinear norms below are induced by these fixed scales.*

$$R_U(U_a) = \begin{pmatrix} J & 0 \\ K & J \end{pmatrix}, \qquad R_U(U_a)^T z = C_U(U_a)^T$$

Let $e = U^* - U_a$, $f = F(u_a, V)$, and $t = Js_a + F_V(u_a, V)$, and let M_C and M_R bound the induced norms of C'' and R'' on that segment. Then,

$$C(U^*) - C(U_a) = -z_1^T f - z_2^T t + R_{nl}, \quad |R_{nl}| \le \frac{1}{2}(M_C + \| z \| M_R) \| e \|^2 .$$

Proof. Taylor expansion gives $0 = R(U_a) + R_U(U_a)e + \rho_R$ and $C(U^*) - C(U_a) = C_U(U_a)e + \rho_C$, with $\|\rho_R\| \le M_R\|e\|^2/2$ and $|\rho_C| \le M_C\|e\|^2/2$. Multiplication by the adjoint yields $R_{\mathrm{nl}} = \rho_C - z^{\mathrm{T}}\rho_R$ and the stated inequality. This is an error identity with a conditional remainder estimate, not a computable enclosure of the target full-order-model (FOM) error. □

Assume also that $A = Z^{\mathrm{T}}JZ$ is nonsingular. Let $z_Z = (\widehat{Zz_1}, \widehat{Zz_2})$, where the coefficients solve (3.4) and (3.7), and define $\eta_Z = -\widehat{z_1}^{\mathrm{T}} r_Z - \widehat{z_2}^{\mathrm{T}} t_Z$. The compression term is $R_Z = (z_Z - z)^{\mathrm{T}} R(U_a)$. If η_{FD} is the implemented finite-difference estimate, set $R_{\mathrm{FD}} = \eta_Z - \eta_{\mathrm{FD}}$. These definitions also fix the signs of the terms in the decomposition.

The output actually scored in the experiments is $C^*_{\mathrm{ref}} - C_{10}(U_a)$, whereas the proposition concerns $C^*_{\mathrm{rel}} - C_{\mathrm{rel}}(U_a)$. This requires an explicit model-bridge term:

$$R_{\mathrm{bridge}} = (C^*_{\mathrm{ref}} - C^*_{\mathrm{rel}}) + (C_{\mathrm{rel}}(U_a) - C_{10}(U_a)),$$

$$C^*_{\mathrm{ref}} - C_{10}(U_a) = \eta_{\mathrm{FD}} + R_{\mathrm{nl}} + R_Z + R_{\mathrm{FD}} + R_{\mathrm{bridge}}.$$

In this notation, ${C_{\mathrm{ref}}}^*$ denotes the exact derivative for the full-grid discrete reference, ${C_{\mathrm{rel}}}^*$ the derivative at the reliability-model solution, and $C_{10}(U_a)$ the $M = 10$ derivative output at the approximate state–tangent pair. The reported derivatives are numerical rather than exact, so finite differences in C_{10} and the converged reference may introduce a further output-evaluation term. We examine that contribution in the reference-step study. A corresponding bridge exists for the charge goal because its reliability model also differs from the full-grid reference. The controller floors were chosen empirically and cannot be interpreted as proven bounds on either bridge.

For nonzero signed target error $\Delta C = C^*_{\mathrm{ref}} - C_{10}(U_a)$, signed effectivity satisfies

$$\eta_{FD}/\Delta C = 1 - \left(R_{nl} + R_Z + R_{FD} + R_{bridge}\right)/\Delta C,$$

before accounting for numerical evaluation error in the output and reference. Effectivity can approach one only when the combined remainder is small compared with $|\Delta C|$. If the local remainder constants and adjoint norms stay bounded as $e \to 0$, an additional condition $|C_{\mathrm{rel}}{}^* - C_{\mathrm{rel}}(U_a)| \geq \gamma\|e\|$ with $\gamma > 0$ yields an $O(\|e\|)$ defect in ideal Taylor effectivity. This nondegeneracy of the scalar output is a separate assumption; accurate states alone do not imply it. We have not established uniform constants or bounds on the model bridges.

The charge-adjoint field $z_Q(u,\ V) = F_u(u,\ V)^{-\mathrm{T}} Q_u(u,\ V)^{\mathrm{T}}$ explains why a shared space can be useful. Differentiating its defining equation in direction $(s, 1)$ gives $z_2 = z_Q$ and $z_1 = D_{(s,\ 1)} z_Q$. The identity is evaluated at the same state and tangent. It helps explain why the charge-trained space may be useful, but it does not guarantee accurate projection of every derivative adjoint.

More precisely, in a smooth constant-stage region, let $q(V) = z_Q(u_{\mathrm{path}}(V), V)$, $P_Z = ZZ^{\mathrm{T}}$, and $\varepsilon_\rho = \sup_{|\xi - V| \leq \rho} \|(I - P_Z)q(\xi)\|$, and assume additionally that q is C^3 with $\|q'''\| \leq C_3$ on that interval. A centered-secant argument gives $\mathrm{dist}(q'(V), \mathrm{range}\, Z) \leq \varepsilon_\rho/h + C_3 h^2/6$ for $0 < h \leq \rho$. For $\varepsilon_\rho > 0$ and $C_3 > 0$, if $h^* = (3\varepsilon_\rho/C_3)^{1/3} \leq \rho$, the minimum is $(3/2)(C_3/3)^{1/3}\varepsilon_\rho^{2/3}$; otherwise, use $h = \rho$. If $\varepsilon_\rho = 0$, the bound tends to zero as $h \downarrow 0$; if $C_3 = 0$ and $\varepsilon_\rho > 0$, choose $h = \rho$. The implemented tangent adds at most $\|D_u z_Q\|\|s_{\tan} - s_{\mathrm{path}}\|$. These are local, conditional estimates involving an out-of-sample envelope, not training residuals. They do not apply across a stage-count switch. Direct projection measurements in Section 6.5 test the actual fields.

3.3.2. Reproducible computational procedure

Algorithm 1. *Offline construction, online estimation, and post hoc evaluation*

Inputs: the training design, physical tuple $\left(t_{Si}, V_g, t_{ox}, E_F\right)$*, primal rank* K_p*, selected chord-stage count, and frozen controller parameters.*

1) Offline reference solves. *For each of the 180 training tuples, solve the corrected full reference and subtract the derived uniform baseline from the reference potential. Store the resulting snapshots.*

2) Offline reduced spaces. *Compute the unweighted Euclidean POD, retain four primal modes, and fit the 15-feature initializer by ordinary least squares. At the two-stage,* $K_p = 4$ *training states, form the full-node reliability charge adjoints and compute their SVD to obtain Z. The supplied builder reproduces every offline array.*

3) Online initialization. *For each requested physical tuple, construct the geometry-dependent quantum and projected Poisson tensors if the geometry is new. Derive* φ_0*, evaluate the quadratic feature row, and truncate the initialized coefficients to* K_p*.*

4) Online primal updates. *Form* J_0 *at the initialized coefficients. Apply the selected one or two chord updates using this same Jacobian.*

5) Online tangent and output. *Refresh* J_p *at the final coefficients, compute* r_V *by centered bias differences, and solve* $J_p a_V = -r_V$*. Reconstruct* $s = \varphi_{0,V} 1 + B_p a_V$ *and evaluate the* $M = 10$ *derivative output.*

6) Online reliability. *From the* $M = 12$ *projected quantum model, evaluate the reliability charge and assemble* r_Z, A, g_Q*. Compute* t_Z, K_2, g_C *using the specified centered perturbations. Using the same*

assembled operator A^T, solve for $\widehat{z_2}$ and then $\widehat{z_1}$, and evaluate η_Q and η_C using Eqs (3.4) and (3.8). The timing implementation in Section 6.7 additionally factors A^T once and reuses its LU factors for both solves.

7) Online decision. *Form the branch scores with the frozen controller. If* $\max(S_Q, S_C) \leq 1$*, accept; otherwise, reject. Return the outputs, estimates, and decision. No reference output or truth label enters steps 3–7.*

8) Post hoc reference evaluation. *Solve the full-grid reference at the requested and neighboring biases to obtain the converged derivative. Compute signed errors, effectivity, and truth labels. Keep repeated algorithmic states distinct from independent physical tuples, and record row-level inputs, outputs, and environment information.*

3.4. Guarded tolerance controller

The primary controller uses nonnegative normalized offsets $b_Q,\ b_C$ and multipliers $s_Q,\ s_C \geq 1$. Its parameters are selected by a declared finite-grid rule using development data, and are then held fixed for the independent validation. The historical rule is retained as a comparator in Section 6.8.

The online relative indicators $e_Q = |\eta_Q|/|Q_{M12}|$ and $e_C = |\eta_C|/|C_{\mathrm{gg},\,M10}|$ use the returned reliability charge and primal tangent capacitance, respectively. They define the branch scores:

$$S_Q = s_Q \frac{e_Q}{\tau_Q} + b_Q, \qquad S_C = s_C \frac{e_C}{\tau_C} + b_C. \tag{3.9}$$

The prescribed reduction-error budgets are $\tau_Q = 0.005$ and $\tau_C = 0.01$, relative to the fixed discrete reference. A truth label is based on the corresponding reference-normalized errors, whereas the online indicators use online output magnitudes. This denominator distinction is retained in calibration and validation. Neither budget is a continuum-accuracy or physical-device guarantee.

Section 5 treats subband truncation and spatial/interface resolution as separate questions. The reference and reliability models both retain 12 subbands, but the reliability Hamiltonian uses a fixed projected quantum space. Empirical offsets account for discrepancies observed in the tests; they have not been proved to bound the model bridge in Proposition 1.

We round summary percentages and effectivities for readability, while retaining additional digits for small-error diagnostics and near-threshold examples. The underlying values are supplied in the supplementary CSV files. Prescribed tolerances, controller constants, and population limits retain their declared precision. All accept/reject decisions use the full stored numerical values, never the rounded displayed values.

The grid contains ten values for each multiplier and eleven for each normalized offset, yielding 12,100 candidate rules. Selection first requires zero development false acceptances and no unresolved references; among feasible candidates, it first minimizes false rejections, then minimizes the total tightening $(s_Q - 1) + (s_C - 1) + b_Q + b_C$ among candidates tied on false rejections, and finally applies lexicographic ordering to the parameter tuple to break any remaining ties. The resulting guarded joint score is:

$$G = \max(S_Q,\ S_C). \tag{3.10}$$

The state is accepted if and only if $G \leq 1$. The selected constants are $s_{\mathrm{Q}} = 1$, $s_{\mathrm{C}} = 1$, $b_{\mathrm{Q}} = 0.0025$, and $b_{\mathrm{C}} = 0$. The equivalent online indicator cutoffs are 0.49875% for charge and 1.00000% for capacitance. These parameters are fixed before validation. Section 6.1 reports their selection and observed errors.

An offset can be written as $b_Q = \delta_Q/\tau_Q$ or $b_C = \delta_C/\tau_C + m$. The decision depends only on the effective offsets; separating an offset into a floor and a margin does not create a mathematical error bound. Table 2 summarizes the frozen model and primary-controller specification.

Table 2. Model and primary controller specification. Historical constants are studied separately in Section 6.8.

Component	Setting
Full reference	N_Q= 121; N_{Si} = 81; 12 subbands per valley; spin 2; valley counts 2 and 4
Comparison derivative	Richardson combination of $h = 10^{-4}$ and 2×10^{-4} V; residual tolerance 2×10^{-12} V
Primal ROM	Four trained POD modes; K_p = 1–4 online; M = 10; one or two chord updates
Internal differences	$FD_V = 2\times10^{-5}$ V; $FD_{small} = FD_\varphi = 2\times10^{-6}$ V; primal coefficient Jacobian 10^{-6} V
Dual space	r = 9; Euclidean charge-adjoint SVD on 180 corrected training states
Reduction budgets	$\tau_Q = 0.005$; $\tau_C = 0.01$
Primary controller	$s_Q = 1$; $s_C = 1$; $b_Q = 0.0025$; $b_C = 0$

3.5. Optional full-model-assisted point certification

The empirical error estimate leaves terms that Proposition 1 does not bound. For selected outputs, we therefore carry out a separate interval calculation using established methods [29, 30]. It works with full-model information, so it is a post hoc check rather than part of the oracle-free online controller. The quantity being verified is an exact real solution of the nonlinear system assembled from binary64 coefficients at fixed geometry and Fermi energy. We freeze those coefficients before beginning interval arithmetic. The resulting enclosure consequently does not account for their assembly error relative to ideal real constants.

For verification only, expand the unknown to $x = (u, \{w_{vj}, E_{vj}\})$, where u has 81 potentials and each of the 24 eigenpairs has 119 interior eigenvector components and one energy in eV. Thus x has 2961 components. The Hamiltonian is the full 119×119 tridiagonal matrix, not the projected reliability Hamiltonian. The equations $G(x, V) = 0$ comprise the Poisson equations divided by their diagonal coefficients, the eigenvalue equations $(H_v(u, V) - E_{vj}I)w_{vj} = 0$, and $(w_{vj}^{\mathrm{T}} w_{vj} - 1)/2 = 0$. This enlarged formulation and row scaling are used only by the verifier; the eliminated 81-variable residual in Section 3.1.1 remains unchanged. Interpolation, occupations, and charge quadrature follow Section 3.1.1. The accompanying Supplementary Material describes this verification target and provides the coefficient assembly and Jacobian blocks.

Proposition 2. *Let G be continuously differentiable on a neighborhood of a closed, bounded convex box X centered at x_0. Let $J(X)$ be a componentwise interval enclosure of all Jacobians on X, and suppose every matrix in $J(X)$ is nonsingular. Let Δ enclose $A^{-1}G(x_0, V)$ for every $A \in J(X)$. If $x_0 - \Delta$ lies strictly inside X, then $G(\cdot, V)$ has exactly one zero x^* in X. This is a local interval-verification criterion, not a new general interval Newton theorem.*

Proof. For $x \in X$, set $B(x) = \int_0^1 G_x(x_0 + t(x - x_0), V)\, dt$. Componentwise convexity implies $B(x) \in J(X)$, hence $B(x)$ is nonsingular. The continuous map $T(x) = x_0 - B(x)^{-1}G(x_0, V)$ maps X into itself. Brouwer's fixed-point theorem supplies $x^* = T(x^*)$; the identity $G(x, V) = G(x_0, V) + B(x)(x - x_0)$ then gives

$G(x^*, V) = 0$. For any two zeros in X, the segment-averaged Jacobian belongs to $J(X)$ and annihilates their difference. Its nonsingularity proves uniqueness in X. □

The Jacobian has an arrow structure: its top row contains the scaled Poisson block M and 24 coupling blocks B_k; each lower row contains C_k and a bordered eigenpair block

$$D_k = \begin{pmatrix} H_v - EI & -w \\ w^{\mathrm{T}} & 0 \end{pmatrix}.$$

Verified inversion of all D_k and $S = M - \sum_k B_k D_k^{-1} C_k$ proves regularity of the full interval family and encloses the Newton correction. For every component, the absolute upper bound of the entire Newton-correction interval must be strictly smaller than the corresponding nominal box radius. We also use interval tridiagonal Sturm counts at 12 separators per valley and separated energy boxes to identify the retained eigenpairs as the first 12 eigenvalues. Any pivot interval containing zero, a failed inversion, or a failed inclusion leaves the state without a certificate.

At the verified root, the implicit function theorem gives $G_x t = -G_V$ for $t = dx^*/dV$, holding geometry and Fermi energy fixed. The same verified block solve encloses t. Evaluating Q and its analytic directional derivative $C = Q_x t + Q_V$ over the state and tangent boxes yields I_{Q} and I_{C}. Both output intervals must exclude zero. For the actual returned binary64 ROM outputs, compute the interval errors $R_{\mathrm{Q}} = |Q_{\mathrm{ROM}} - I_{\mathrm{Q}}|/|I_{\mathrm{Q}}|$ and $R_{\mathrm{C}} = |C_{\mathrm{ROM}} - I_{\mathrm{C}}|/|I_{\mathrm{C}}|$. Acceptance is certified only if the whole R_{Q} is at most 0.005 and the whole R_{C} is at most 0.01. A violation is certified only if the whole of either interval exceeds its threshold; otherwise, the result is inconclusive. Interval inclusion proves these statements about the true reference-relative errors.

This direct comparison bounds the net discrepancy between each supplied ROM output and the verified target without separately bounding the Taylor, projection, finite-difference, and model-bridge terms in Proposition 1. The verifier requires explicit spin-degeneracy metadata and rejects earlier frozen assemblies that lack it. All reported assemblies were rebuilt for Eq (3.2). The certificate relies on validated arithmetic and the supplied implementation; it is not a proof-assistant formalization. Uniqueness is established only within X. The certificate does not cover continuum error, omitted higher subbands, physical uncertainty, and untested parameter values.

4. Computational protocol and evaluation designs

We regenerated every result after repairing the occupation prefactor in Eq (3.2). Only the 180-tuple tensor training design contributes to the primal POD, regression initializer, and charge-adjoint basis. The numerical dimensions remain $K_p = 1$–4, with $M = 10$ in the primal model, $M_{\mathrm{rel}} = 12$ for the reliability calculation, and nine modes in the compressed dual space. All 180 reference solves meet the tolerance, with one training tuple requiring the smaller-damping fallback. The feature matrix has a rank of 15. Source code, rebuilt arrays, and the training records specify the implementation used in the results below.

A physical tuple consists of $(t_{\mathrm{Si}}, V_g, t_{\mathrm{ox}}, E_{\mathrm{F}})$, whereas an algorithmic state also specifies the primal rank and chord-stage count. Four ranks evaluated at two stage counts give eight correlated states for one tuple, as indicated in Table 3. We reused the coordinates of the earlier 240-state structured and 190-state targeted designs, but recalculated their outputs under the corrected model. The structured set contains

16 identified training-support corner states; another 16 states are outside the declared domain and are used only as stress probes. Neither group is an independent test of in-domain generalization.

Table 3. Roles of the physical designs in the corrected study. Reused designs do not become new independent tests when their numerical outputs are recomputed.

Design	Physical tuples	States	Role
Tensor training	180	180	POD, regression, and dual snapshots
Reused structured design	96	240	Diagnostics; 16 training corners and 16 OOD states
Reused targeted design	190	190	Diagnostics; earlier adaptive coordinates
Calibration	256	2048	Margin selection; primary development
Earlier test design	512	4096	Margin validation; primary development
Online prescreen	10,000	10,000	Online indicators only; no reference labels
Selected confirmation	128	128	Online-selected; primary development
Reused validation design	1024	8192	Corrected-model validation on examined coordinates
New independent validation	1024	8192	Fresh uniform physical tuples; no selection use

Calibration, the earlier test, and the prescreen use independent uniform coordinates over the declared four-dimensional box, with seeds 3133201, 3133202, and 3133203. For the primary rule, the 512-tuple earlier test is explicitly reclassified as development data. The primary development set contains 2048 + 4096 + 128 = 6272 states. The reused 1024-tuple validation has seed 3133301. A separate 1024-tuple holdout with seed 3133302 is fixed before corrected outcomes and excludes all supplied training, historical, calibration, earlier-test, prescreen, and reused-validation coordinates. Both validation campaigns start after the primary parameters are fixed. No outcome from either campaign is used to retune them.

The 128 confirmation states are selected from the corrected 10,000-state prescreen using online quantities only. The first 64 are closest to the unit joint score under the historical multipliers and floors without its added margin; the next 64 remaining C-dominant online states are closest to the unit C score. State identifiers break ties, with a declared joint-score fallback if needed. Full references are computed only after this selection. All prescreen rows, choices, unresolved states, and nonlinear solve attempts are retained.

Truth acceptance means that both reference-normalized errors, $|Q_{\mathrm{M12}}-Q_{\mathrm{ref}}|/|Q_{\mathrm{ref}}|$ and $|C_{\mathrm{gg,\,M10}}-C_{\mathrm{ref}}|/|C_{\mathrm{ref}}|$, satisfy 0.005 and 0.01. C dominance means that its tolerance-normalized true error exceeds the corresponding charge value. The joint near-threshold band is $0.8 \le \max(e_{\mathrm{Q,\,true}}/\tau_{\mathrm{Q}}, e_{\mathrm{C,\,true}}/\tau_{\mathrm{C}}) \le 1.2$. A strict C-dominated near failure additionally requires charge within tolerance and capacitance error in (1%, 1.2%]. Nonfinite outputs or an unmet reference tolerance remain unresolved and are not silently counted as acceptable.

Effectivity is $|\eta_{\mathrm{C}}|/|C_{\mathrm{ref}}-C_{\mathrm{gg,\,M10}}|$; signed agreement compares η_{C} with $C_{\mathrm{ref}}-C_{\mathrm{gg,\,M10}}$. The predefined capacitance-error populations are the 0.5–2% decision band, the 0.8–1.2% tolerance band, and C-dominated true failures. The tables give both state counts and physical-tuple counts. The absolute-error floor $|C_{\mathrm{ref}}-C_{\mathrm{gg,\,M10}}| \ge 10^{-6}$ F m^{-2} is used only as a reporting convention for ratio diagnostics. It is neither a controller input nor a physical regime boundary. Results without this floor, together with absolute

metrics for every row, are also supplied.

5. Reference convergence and numerical sensitivity

We evaluated derivative-step size and nonlinear-solver tolerance on 96 physical tuples. The central-difference steps were $h = 0.5$, 1, 2, 4, and 8×10^{-4} V, with the tighter Richardson estimate serving as the comparison reference. In a separate study of 24 tuples, we changed the number of full-grid subbands from 12 to 16 and 20, and refined the quantum/silicon meshes from (121, 81) to (181, 121) and (241, 161). Each setting is solved self-consistently with the corrected occupations. The supplied damping comparisons test agreement between the two documented iteration policies on their overlapping cases. Table 4 collects the resulting reference sensitivities.

Table 4. Corrected reference sensitivity. Relative shifts are observations on the stated designs, not continuum-error bounds.

Perturbation	Tuples	Largest observed relative change	Interpretation
Production central difference vs. tight Richardson	96	0.000999%	0/240 truth-label changes
Residual tolerance $2\times10^{-10} \rightarrow 2\times10^{-12}$ V at $h = 2\times10^{-4}$ V	96	0.000003%	Nonlinear stopping sensitivity
Full-grid subbands 12 → 20	24	1.27×10^{-7}%	Capacitance truncation sensitivity
Production → either refined grid	24	3.584%	Spatial/interface sensitivity
Production → finest grid	24	2.054%	Spatial/interface sensitivity

The central-difference step-halving ratio has a median of 0.250048 and an interquartile range of 0.249202–0.252385, consistent with second-order behavior on these tuples. None of the three supplied derivative comparisons changes a truth label on the 240 corresponding states. The effectivity ratio is more sensitive when the true error is small: replacing the tight Richardson reference with the production central difference changes effectivity by at most 0.018964 on the 171 states above the absolute-error reporting floor. All headline controller and effectivity results use the tight Richardson reference.

The largest production-to-refined capacitance change is 3.584%, and the largest production-to-finest change is 2.054%. Refinement is non-monotone on 19 of 24 tuples. Material interfaces are sampled on a uniform grid and do not remain fitted to the same faces as the mesh changes. A formal second-order Richardson construction reaches 4.365%, although its assumed asymptotic regime has not been established. We retain this construction as a diagnostic, not as a continuum-error estimate, bound, or guarantee.

Figure 2 reports the separate internal-difference and dual-rank checks on the 240-state design. FD_V, FD_{small}, and the offline nodal step FD_φ are varied by factors 1/4, 1/2, 2, and 4. Each change to FD_φ requires rebuilding the charge-adjoint basis from the training states. Ranks 7–11 use the same corrected SVD. An online tangent-step change can affect both the returned ROM output and its truth label, so the records distinguish changes in decisions from changes in truth labels. These diagnostics assess the sensitivity of the stated implementation; they are not used to optimize its settings on validation outcomes.

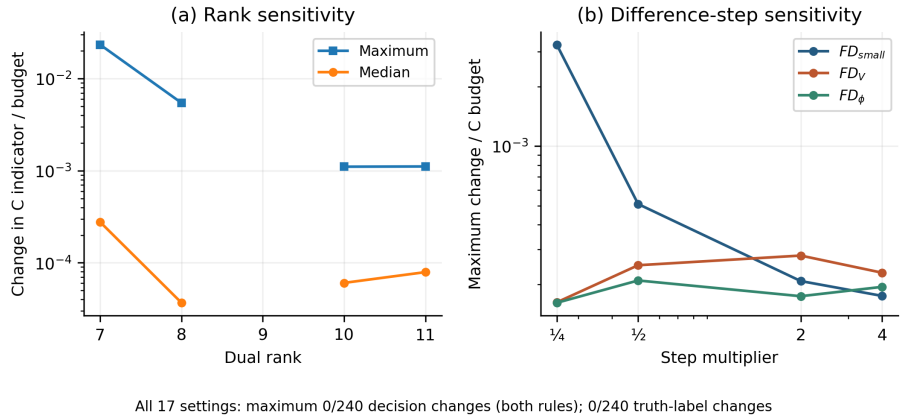


Figure 2. Internal finite-difference and dual-rank sensitivity after rebuilding the spin-degenerate model. Counts compare each perturbation with the nominal implementation on the same 240 states. Zero indicator changes are omitted on the logarithmic axes. Full row-level outputs and both historical and primary decisions are supplied.

6. Numerical results

6.1. Primary controller selection and independent validation

For controller selection, we evaluated ten multiplier values, {1, 1.005, 1.01, 1.015, 1.02, 1.03, 1.05, 1.075, 1.1, 1.2}, against eleven normalized offsets, {0, 0.0025, 0.005, 0.01, 0.02, 0.04, 0.06, 0.08, 0.1, 0.15, 0.2}. The full Cartesian grid contains 12,100 candidates. We fixed the objective and tie-breaking order in Section 3.4 before selection. The released candidate table records every rule, including those eliminated by the constraints.

On the development data, the selected parameters $(s_Q, s_C, b_Q, b_C) = (1, 1, 0.0025, 0)$ produce no false acceptances and five false rejections. In terms of the online relative indicators, the rule uses cutoffs of 0.49875% for charge and 1.00000% for capacitance. These cutoffs apply to the estimators, so they must not be read as bounds on errors normalized by the full-order reference. Table 5 lists the classifications of both rules on the development and validation populations.

Table 5. Controller classifications for the corrected model. FA denotes a false acceptance; FR denotes a false rejection. All rows use the same truth budgets.

Population	Rule	States	True failures	FA	FR
Development	Primary	6272	2558	0	5
Development	Historical	6272	2558	0	102
Reused validation	Primary	8192	3225	2	8
Reused validation	Historical	8192	3225	0	41
New independent validation	Primary	8192	3207	1	6
New independent validation	Historical	8192	3207	0	58

We kept the selected rule fixed when evaluating the new independent design of 1024 physical tuples. It made one false acceptance and six false rejections, compared with zero and 58 under the

historical comparator. On the reused 1024-tuple design, the primary rule gives 2 false acceptances and 8 false rejections. These finite-design results concern the corrected implementation. Since the reused coordinates were examined previously, their evidentiary role differs from that of the new holdout. The supplement includes individual failing and rejected rows, full precision scores, and parameter-perturbation tables. Figure 3 compares the primary online score with the true joint error divided by its budget on the new independent holdout, including a closer view around the decision thresholds.

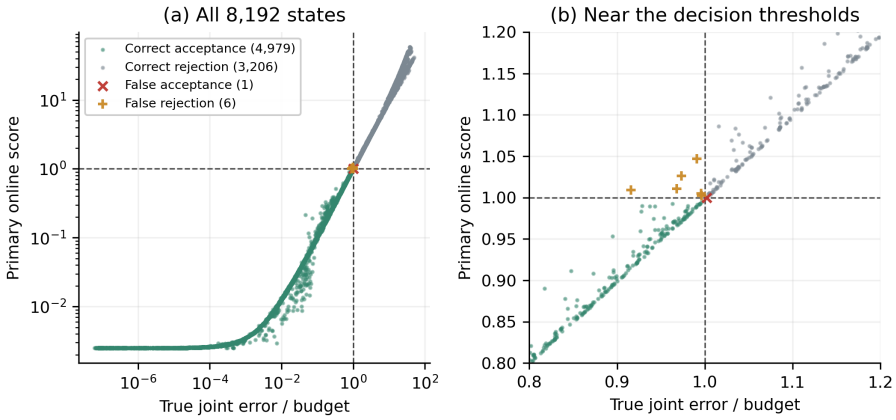


Figure 3. Frozen primary controller on the independent validation. A false acceptance lies to the right of the true-error threshold while remaining below the online-score threshold. The second panel enlarges the joint decision region; the complete first panel retains all states.

All three false acceptances observed across the two validation designs are charge-only failures; none violates the capacitance budget or both budgets. The counts serve as diagnostics and are not a pooled independent risk estimate. For each such state, the records retain the actual unrounded errors, online indicators, and score. These failures show that zero false acceptances during development does not transfer uniformly to unseen states.

In all three charge-only false acceptances, the ROM overpredicts the positive reference charge; the true charge errors range from 0.500211% to 0.500949%. All three have $|\eta_Q|$ at least as large as the observed absolute charge error, so these cases are not explained by an underestimated absolute correction. If $Q_{\mathrm{ROM}} = Q_{\mathrm{ref}}(1+\kappa)$ with $\kappa > 0$, exact error divided by the online charge is $\kappa/(1+\kappa)$, whereas the truth-normalized error is κ. The $\kappa = 0.005$ boundary therefore corresponds to an online value $0.005/1.005 = 0.004975124378$, below the frozen primary cutoff 0.0049875. Finite calibration has not fully compensated for this normalization difference. This post hoc explanation does not change the controller or turn its estimator into a rigorous bound.

The new holdout contains 382 states in the joint near-threshold band; the primary rule has 1 false acceptance and 6 false rejections there. No validation-derived setting replaces the frozen primary parameters. One-parameter sensitivity tables are descriptive post hoc analyses, not a further selection step.

At least one false acceptance occurs at 1 of the 1024 physical tuples when their eight states are evaluated. Under the stated independent uniform tuple-draw model, the exact one-sided 95% Clopper–Pearson upper confidence limit for that tuple-level event probability is 0.462%. The limit inverts the binomial tail [31]: for $0 < k < N$, it is the 0.95 quantile of $\mathrm{Beta}(k+1, N-k)$, with limits $1 - 0.05^{1/N}$

for $k = 0$ and 1 for $k = N$. The one-sided reporting rule is fixed before validation outcomes. This statement concerns the declared artificial sampling distribution and the fixed implementation. It is not a deployment-risk guarantee. The eight states at one tuple are not treated as eight independent Bernoulli observations. No corresponding binomial interpretation is assigned to the structured 224 in-domain states or to adaptively selected confirmation points.

6.2. *Effectivity and the size of the evaluation populations*

The corrected structured campaign contains only 1 state in the narrow 0.8–1.2% band, so its effectivity cannot establish broad near-threshold performance. Across all 190 reused targeted states, effectivity ranges from 0.953 to 2.056, with a median of 1.011. Effectivity varies more widely here than in the selected decision band. Table 6 therefore gives the population sizes and reports the broader targeted population alongside the independent validation. A favorable subset does not establish a uniform bound.

When the true error is very small, the effectivity ratio can vary substantially even if the absolute estimator discrepancy is negligible. Figure 4 separates the full structured range from the independently sampled decision band. The accompanying files retain every state, and Section 6.6 reports metrics that remain finite as the true error approaches zero. Correct sign and near-unit ratios observed in a sample do not control the unbounded remainder terms or model bridge in Proposition 1 at other states.

Table 6. Capacitance effectivity with absolute true error at least 10^{-6} F m^{-2}. The Count column gives states/physical tuples; Sign is agreement with the signed true error.

Population	Count	Min.	Median	Max.	Sign
Structured; 0.5–2%	17/9	0.986	1.009	1.065	100%
Structured; 0.8–1.2%	1/1	0.990	0.990	0.990	100%
Targeted; all states	190/190	0.953	1.011	2.056	100%
Targeted; 0.8–1.2%	6/6	0.953	1.005	1.150	100%
Independent; 0.5–2%	737/326	0.533	1.004	1.363	100%
Independent; 0.8–1.2%	207/121	0.710	1.003	1.216	100%
Independent; C-dominated failures	624/312	0.978	0.997	1.015	100%

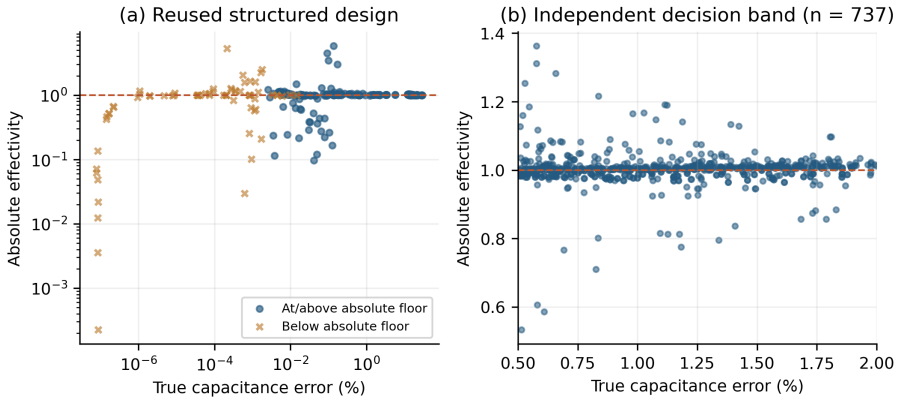


Figure 4. Corrected capacitance effectivity. Left: All 240 structured states, including small-error cases; marker style identifies the absolute-error reporting floor. Right: The independent 0.5–2% capacitance-error population. The horizontal line denotes unit effectivity.

6.3. *Threshold coverage and targeted confirmation*

Table 7 summarizes the coverage of capacitance-dominated failures in each design. Under the corrected model, the reused 190-state targeted design has 141 true failures. This design was selected using an earlier model; it is not a new adaptive search for the corrected model. The separately regenerated prescreen and online-selected confirmation provide boundary evidence for the current model. The supplied records contain all 10,000 prescreen rows and all 128 selected references, including selection reasons and ties. Figure 5 relates the online selections to the confirmed reference errors.

Table 7. Coverage of capacitance-dominated failures. The strict near window requires charge within 0.5% and capacitance error above 1% and at most 1.2%.

Population	States	True failures	C-dominated failures	Strict near failures
Reused structured	240	72	18	0
Reused targeted	190	141	1	0
Corrected confirmation	128	24	0	0
New independent validation	8192	3207	624	0

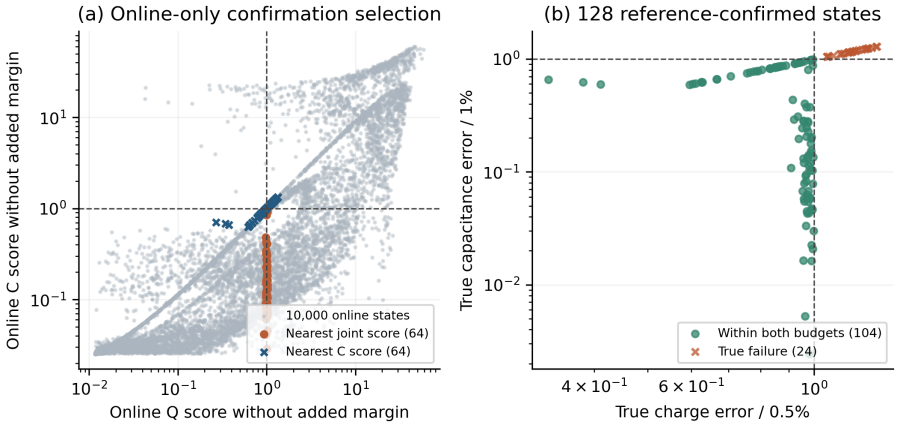


Figure 5. Corrected online prescreen and reference confirmation. Left: All prescreen states and the 128 selected states, using the historical multipliers and floors without the added margin for selection. Right: Confirmed charge and capacitance errors normalized by their respective budgets.

Finding no C-dominated near failures in a particular finite design would not establish that such failures are structurally impossible. The budgets constrain different outputs, the errors are correlated, and a uniform or score-targeted sample can miss a narrow set. Section 6.10 gives interval-verified examples found in a separate search. That adaptive search is excluded from controller selection and cannot measure how prevalent these failures are.

6.4. *Physical profiles and boundary sensitivity*

Figure 6 shows how the reduced approximation is built for the same physical device. The reduction changes the coordinates used for the solution; it leaves the layer thicknesses, interfaces, and boundary

conditions intact. The coordinate represents transverse equilibrium confinement. Setting the wavefunction to zero at the outer boundary truncates the oxide tails with a hard wall; it does not provide an open boundary for longitudinal transport. The quasi-static calculation therefore makes no statement about injection, extraction, or current. Table 8 compares the FOM and ROM profiles for the fixed cases.

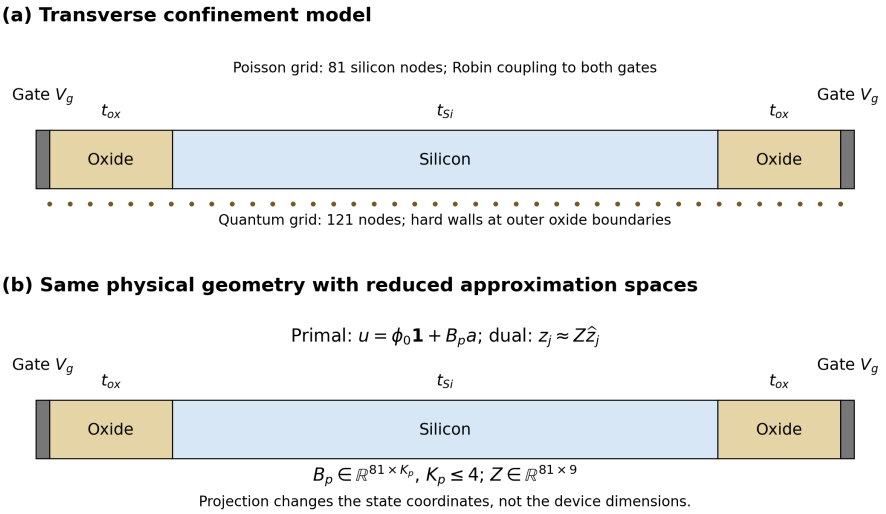


Figure 6. The same double-gate device in the full and reduced descriptions. The quantum grid extends across oxide–silicon–oxide; the Poisson unknowns occupy silicon with Robin gate coupling. POD and dual projection change the approximation spaces while preserving the geometry.

Table 8. FOM–ROM profile comparisons. Density error is the relative spatial L^2 norm and potential error is the maximum absolute difference. Q here is the $M = 10$ primal charge used in the profile calculation; controller truth uses the returned M = 12 reliability charge.

Case	V_g V	Potential V	Density	Q_{M10}	C_{gg}
Low bias	0.15000	1.11×10^{-5}	0.010%	0.009%	0.018%
Intermediate bias	0.35000	3.72×10^{-3}	1.915%	1.352%	1.227%
High bias	0.80000	3.64×10^{-3}	0.981%	0.056%	0.162%
Reused targeted tuple	0.74860	1.43×10^{-3}	2.340%	2.066%	3.218%

The three fixed profiles use (t_{Si}, t_{ox}, E_{F}) = (4 nm, 0.9 nm, 0.11 eV), $K_p = 2$, and two chord stages, at $V_g = 0.15$, 0.35, and 0.80 V. The intermediate case now has 1.2270% capacitance error and exceeds the 1% budget. It is not relabeled as an acceptable near-threshold example. The fourth tuple uses (7.580246014520526 nm, 0.7486026995815337 V, 1.0051022045314313 nm, 0.096555833145976 eV), $K_p = 1$, and one stage; its corrected capacitance error is 3.2183%. Its earlier numerical characterization is superseded. Figure 7 shows the corrected potential and density profiles at the three gate biases, together with the conduction-band edges and the first Δ_2 subband.

At each of the three fixed biases, boundary-placement tests append 0, 4, 8, or 16 oxide grid nodes at each end while keeping the original mesh spacing and interface sampling, gate coupling, and physical device unchanged; only the quantum hard wall moves. The largest observed Q or C shift is 0.000917%.

Other physical sweeps vary oxide thickness over 0.8, 0.9, and 1.2 nm and barrier height over 2.6, 3.1, and 3.6 eV. Relative to the 0.9 nm, 3.1 eV baseline, the oxide-thickness changes shift capacitance by −20.43% to 8.74%, and the barrier sweep shifts it by −2.52% to 3.34%. These vary physical inputs and are not hard-wall truncation errors. The complete 30-case records specify each perturbation. Open transport or a different Hamiltonian requires new residual and derivative implementations, as discussed in Section 2.

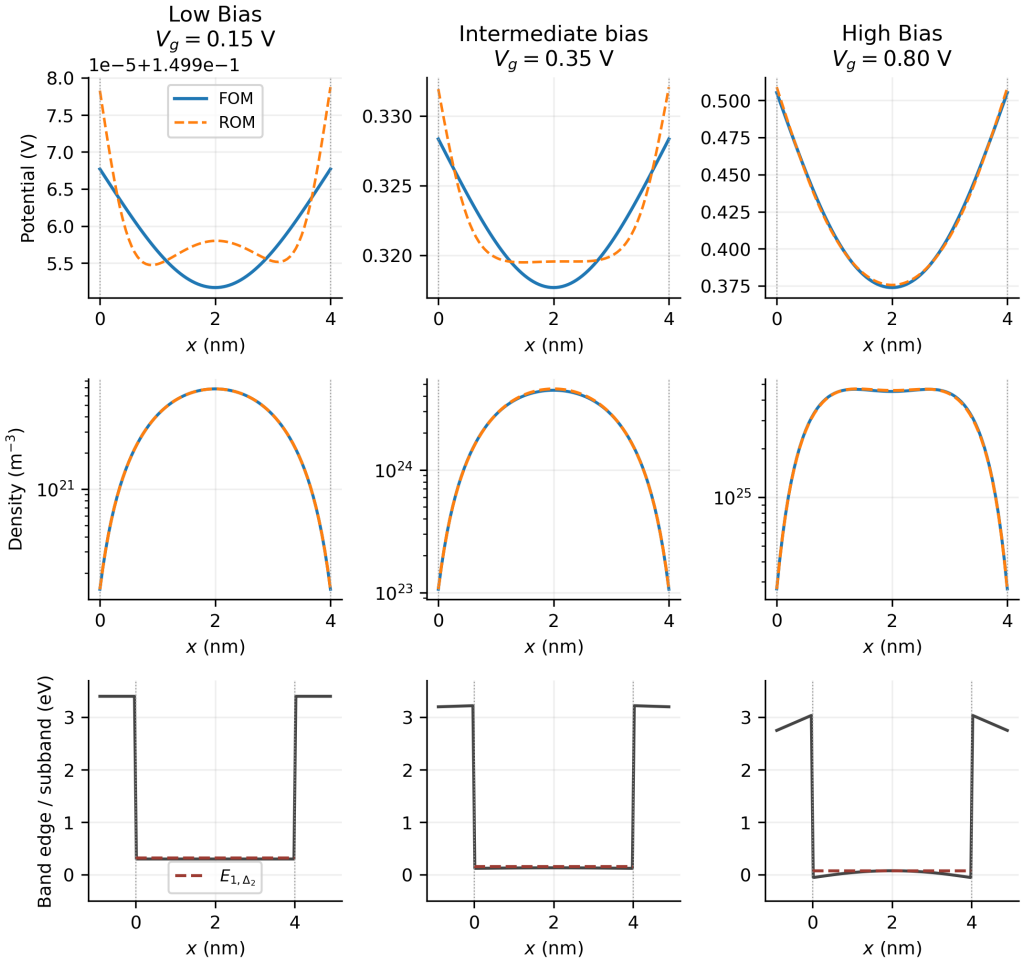


Figure 7. Corrected potential and logarithmic density profiles at low, intermediate, and high gate bias, followed by conduction-band edges and the first Δ_2 subband. Vertical dotted lines mark the silicon interfaces. The fixed intermediate case exceeds the capacitance tolerance after the physical correction.

6.5. Direct adjoint projections and internal-step checks

Full adjoints at 25 diagnostic states are compared with their Euclidean orthogonal projections onto the charge-trained spaces. These diagnostic states do not enrich the basis. The relative field error is $\|(\mathrm{I}-\mathrm{ZZ}^{\mathrm{T}})z_j\|_2/\|z_j\|_2$. It measures representation of the full field; it is distinct from the compressed Galerkin solve and from the residual-weighted output correction. Table 9 reports the resulting projection errors for ranks 7–11.

Table 9. Relative full-adjoint projection errors on 25 cases. z_1 is the derivative adjoint; z_2 is the charge adjoint.

Rank	z_1 median	z_1 maximum	z_2 median	z_2 maximum
7	5.207%	8.476%	2.748%	4.590%
8	5.036%	8.139%	2.367%	4.482%
9	4.360%	7.555%	2.231%	4.425%
10	4.325%	7.405%	2.109%	4.359%
11	4.218%	7.386%	2.093%	4.298%

At rank nine, derivative-adjoint projection error has a median of 4.360% and a maximum of 7.555%; charge-adjoint error has a median of 2.231% and a maximum of 4.425%. Accurate output corrections can therefore coexist with appreciable adjoint-field projection error. These measurements assess the retained rank only on the tested cases. Proposition 1 gives a possible local explanation for compression, but establishes neither a domain-wide envelope nor uniform sufficiency of the retained rank. Figure 8 illustrates how the full-adjoint projection errors vary with the retained rank for the derivative and charge goals.

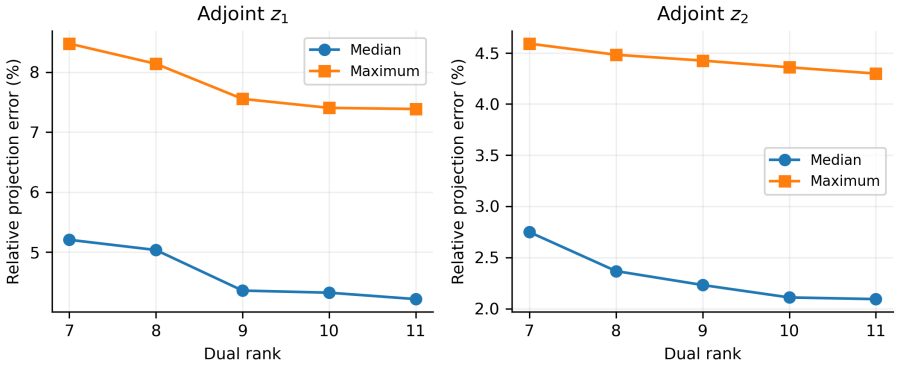


Figure 8. Direct Euclidean projection errors for the full derivative and charge adjoints on 25 diagnostic states. Increasing the retained rank reduces field error modestly; no uniform sufficiency threshold is inferred.

Across the corrected 240-state sensitivity runs, changing the dual rank from 7 to 11 or varying any tested factor of the tangent bias step, online small step, or offline nodal step changes neither the primary controller's decision nor its truth label. Full values, both controllers, and every perturbed output are retained. These checks use the originally specified step factors, so their stability cannot be extrapolated to arbitrary step sizes or other devices.

6.6. Stable absolute and budget-normalized metrics

Define the signed true capacitance error $\Delta C = C_{\mathrm{ref}} - C_{\mathrm{gg,\,M10}}$ and the absolute estimator discrepancy $D_{\mathrm{abs}} = |\eta_C - \Delta C|$. This discrepancy remains meaningful as ΔC tends to zero. For comparison with the engineering budget, define $D_b = D_{\mathrm{abs}}/(\tau_C|C_{\mathrm{ref}}|)$. This normalization avoids division by the true reduction error, although it can still be poorly scaled when the reference capacitance itself approaches zero. Absolute values should therefore be reported alongside budget-normalized values in depletion

regimes. Table 10 summarizes these discrepancies for the structured and targeted designs.

Table 10. Estimator discrepancies over all states, without an effectivity-floor exclusion. D_{abs} is in F m^{-2}; D_b is dimensionless.

Design	D_{abs} median	D_{abs} 95th	D_{abs} max	D_b median	D_b max
Structured	3.27×10^{-7}	6.25×10^{-5}	5.64×10^{-4}	0.00094	0.9173
Targeted	2.74×10^{-6}	4.01×10^{-5}	8.79×10^{-5}	0.01580	0.1679

The fixed 10^{-6} F m^{-2} effectivity floor has no universal physical interpretation. Its ratio to the allowed capacitance error, $10^{-6}/(\tau_{\mathrm{C}}|C_{\mathrm{ref}}|)$, is supplied for every state and varies strongly across the domain. In the structured design, the floor can exceed the tolerance budget by many orders of magnitude when the reference capacitance is extremely small. We therefore use it only to delimit the interpretation of effectivity ratios; no row is removed from controller evaluation or from the absolute-error records. Figure 9 plots the absolute estimator discrepancy against the absolute true capacitance error for both reused designs and retains every state, regardless of the effectivity reporting floor.

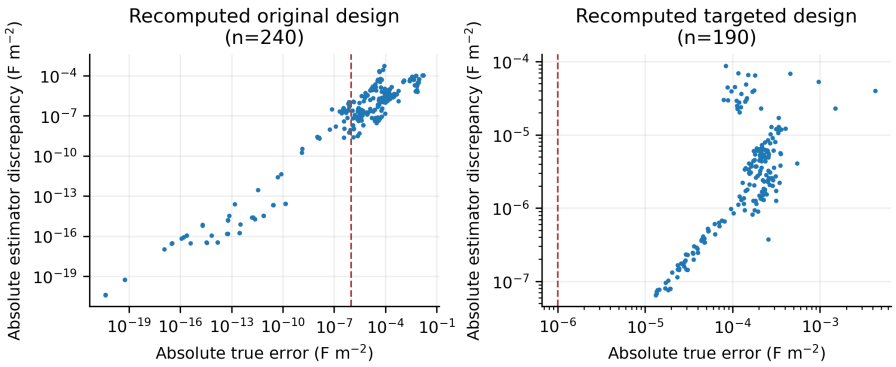


Figure 9. Absolute estimator discrepancy against absolute true capacitance error for both reused designs. The vertical line is the effectivity reporting floor. All points, including those below the floor, enter Table 10 and the controller classifications.

6.7. Measured online cost of operator reuse

The timing comparison uses the same corrected outputs and fixed primary decision rule in both paths. The shared path constructs the compressed dual operator once and reuses a single LU factorization for both adjoint solves; the separate comparator repeats that work for the second goal. Both include the primal solve, tangent, residual and goal assembly, derivative-specific curvature, and final decision. Each geometry is set up afresh. We report operator-assembly, curvature, factorization, and linear-solve components separately. The timing study uses ten paired repetitions per case, 30 pairs overall, in a balanced randomized execution order with seed 3133303 after warm-up. Full-order validation and offline training are excluded from both online timing paths. Table 11 presents the measured timings.

The measured median paired reduction spans 5.84%–13.80% across these cases. Reusing operator construction is more consequential than saving the factorization of a 9×9 matrix. The shared primal/tangent work and derivative-specific curvature limit the savings. These timings depend on the implementation and execution environment; they establish no uniform speedup. We report end-to-end

latency instead of substituting a restricted scalar-operation count.

Table 11. Online timings in milliseconds on three fixed cases. Ranges are observed minima–maxima. Reduction is the median of paired relative reductions, not the ratio of the two marginal medians.

V_g V	Shared median [range]	Separate median [range]	Paired reduction	Curvature median
0.15	100.03 [84.45–113.37]	108.99 [99.07–131.23]	13.80%	18.59
0.35	93.69 [86.74–114.36]	103.14 [96.18–116.22]	9.80%	17.97
0.80	97.06 [91.59–124.66]	106.81 [98.67–125.90]	5.84%	17.16

6.8. Historical constants and the separate margin study

The historical comparator keeps $s_Q = 1.015$, $s_C = 1.020$, $\delta_Q = 5.867181\times10^{-5}$ and $\delta_C = 2.54\times10^{-4}$, with an added C-branch margin $m = 0.2$. The resulting normalized offsets are $b_Q = \delta_Q/0.005$ and $b_C = \delta_C/0.01 + m$. These constants come from the earlier design. Since its original reason for choosing the extra margin was not recorded, we do not reconstruct a rationale from the new results. References [32,33] discuss other ways to treat estimator uncertainty, but neither method is used here or provides a derivation of these constants.

A separate corrected study varies only m over $\{0, 0.01, \ldots, 0.50\}$; the remaining historical constants are held fixed. The selection rule requires no false acceptances on the 256-tuple calibration design, and then chooses the smallest number of false rejections and, in a tie, the smallest margin. It selects $m = 0$. We evaluated that margin on the earlier 512-tuple test, which was independent for this margin-only comparison, although those tuples were later included in primary-controller development. Table 12 keeps these roles distinct.

Table 12. Margin comparison with historical multipliers and floors held fixed. The primary rule in Table 5 is selected separately over all four effective parameters.

Population	Margin	States	FA	FR
Calibration	$m = 0$	2048	0	7
Calibration	$m = 0.2$	2048	0	15
Earlier test	$m = 0$	4096	0	14
Earlier test	$m = 0.2$	4096	0	23
Selected confirmation	$m = 0$	128	0	39
Selected confirmation	$m = 0.2$	128	0	64
New independent validation	$m = 0$	8192	0	35
New independent validation	$m = 0.2$	8192	0	58

After correction, the historical rule produces 0 false acceptances and 2 false rejections on the 240-state design and 0 false acceptances and 0 false rejections on the reused 190-state design. Within the 224 in-domain structured states, there are 0 false acceptances and 2 false rejections. For comparison, the primary rule gives 0 false acceptances and 0 false rejections on both the 224 in-domain structured states and the reused 190-state design. These replace the earlier model's zero-error and 21-false-rejection statements. With historical multipliers and floors fixed, the C indicator cutoff is 0.95549% at $m = 0$

and 0.75941% at $m = 0.2$. The stricter cutoff can reject additional acceptable states, but its value alone supplies no guarantee of eliminating unseen false acceptances.

6.9. Local interval certificates at four physical points

At the four selected physical points, Table 13 records 21 output certificates within both budgets and 11 certificates of a budget violation. No output classification was inconclusive. We regenerated the root, tangent, and output enclosures at both 128 and 192 bits. Each full-grid binary64 assembly was fixed before the interval calculation. The assembly records spin and valley multiplicities; before evaluating it, the verifier checks the spin-degeneracy field explicitly. The supplementary validation file also contains independent comparisons of the kernel, analytic derivatives, and directional Jacobian.

Table 13. Direct interval comparison for eight ROM states at each selected physical point. The same classifications are required at 128 and 192 bits.

Point	States	Within both budgets	Proved violation	Inconclusive
Low bias	8	8	0	0
Intermediate bias	8	2	6	0
High bias	8	6	2	0
Reused targeted tuple	8	5	3	0

These certificates compare the actual binary64 ROM outputs with exact real solutions of the frozen discrete equations. At a specified point, they bound the net reduction discrepancy, including the model bridge and derivative evaluation. The interval procedure requires full-model work and is not included in online latency. It proves neither global uniqueness nor a controller guarantee at untested tuples, and excludes assembly error relative to ideal real constants and continuum or physical-model uncertainty.

Additional interval checks address the three validation states accepted incorrectly by the primary controller. For all three, both 128- and 192-bit proofs certify a violation of at least one budget. The archive gives the exact ROM setting, assembly, and interval bounds for these three cases separately from the four pilot points. These checks establish why the observed decisions were wrong. They do not permit us to adjust the frozen controller using validation outcomes or to make claims about states that were never certified.

6.10. Search for capacitance-dominated near failures

We first searched around the historical target using a finite coordinate scan. All 260 single-state evaluations resolved, but none met the strict near-capacitance-failure conditions. A separately frozen bounded V_g/t_{ox} least-squares attempt retained 97 evaluations (97 resolved, 0 admissible); it was stopped after nonzero-residual stagnation near a sampled interface change, and optimizer termination was not interpreted as success. The subsequent separately frozen smooth-bias search used two fixed geometries, 40 scan tuples with eight rank/stage states each, and 432 retained bisection-state evaluations. That stage retained 752 states (752 resolved, 0 qualifying). All search stages were exploratory and excluded from primary-controller fitting and both validation designs. Since that stage selected no admissible examples, one final conditionally frozen search used the first three distinct development-only seed tuples, two rank/stage states per bias, a fixed global/local bias grid, and at most four 18-step bisection brackets per

capacitance target. It retained 302 states (302 resolved, 2 qualifying) but yielded no target-refined case. The finite search ended at its declared seed set and stopping budget. After examining the completed fixed scan, we used a separate explicit post hoc rule to select all qualifying states for interval illustration: two rank/stage states at one physical tuple. The full 86-state scan snapshot and the post hoc selection record are retained. This selection introduced no new searched physical input and left the planned target-refinement success count unchanged.

Each exploratory stage has a declared finite design or stopping budget, fixed before its evaluations. Later stages respond to earlier search outcomes, so their results provide adaptive evidence. Every evaluated point, failed attempt, and solver trace is retained. No search result enters controller selection or either validation campaign. The earlier model's near-threshold examples and certificates are superseded; a corrected existence claim requires a new successful interval proof.

Table 14. Two interval-verified, C-dominated near-failure configurations at one physical tuple. Interval errors are fractions, so the thresholds are 0.005 for Q and 0.01 for C. Displayed coordinates are rounded; full binary64 coordinates and higher-precision enclosures are supplied.

Case and state	t_{Si} nm	V_g V	t_{ox} nm	E_{F} eV	Q error interval	C error interval
1; $K_p = 1$, one	7.427242	0.806250	0.951610	0.065746	$[0.0015247364 \pm 9.02 \times 10^{-11}]$	$[0.01023 \pm 9.37 \times 10^{-6}]$
2; $K_p = 1$, two	7.427242	0.806250	0.951610	0.065746	$[0.0014758405 \pm 6.94 \times 10^{-11}]$	$[0.0119 \pm 3.11 \times 10^{-5}]$

Both selected configurations satisfy the strict near-failure definition under interval arithmetic at 128 and 192 bits. In each case, the entire charge-error enclosure is below 0.005, the capacitance-error enclosure is above 0.01 but no greater than 0.012, and the dominance of capacitance error is certified. The frozen primary controller rejected both states. These are therefore verified budget violations, not false acceptances. Their existence rules out a structural claim that this type of failure cannot occur in the corrected discrete model and ROM. Two examples do not tell us how often it occurs, cover the entire tolerance boundary, or guarantee either empirical controller. All other ROM states at the certified tuples, including candidates that failed one or more conditions, remain in the supplement.

7. Discussion

7.1. Reduction accuracy and reference-model uncertainty

The largest capacitance change across the tested grids was 3.584%. It is a reference-grid sensitivity, distinct from the 1% budget assigned to ROM reduction error. The triangle inequality gives $|C_{\mathrm{ROM}} - C_{\mathrm{continuum}}| \le |C_{\mathrm{ROM}} - C_{\mathrm{FOM}}| + |C_{\mathrm{FOM}} - C_{\mathrm{continuum}}|$. To turn that inequality into a justified total relative-error allowance, both absolute terms need bounds and the sum must be normalized consistently. The measured inter-grid differences are not bounds on the reference-to-continuum term, so adding 3.584% to the 1% budget would not create a certificate. Physical-model discrepancy, including the effective-mass approximation and uncertain material inputs, adds a further source of error.

Correcting the spin prefactor alters the self-consistent potential, occupations, and differential response. Doubling a previously calculated charge cannot implement the correction; the entire reduced-model and

verification pipeline must be rebuilt. The resulting changes in classifications, effectivity populations, profiles, and certificates show how a shared error in the reference and reduced models can survive internal agreement checks. The explicit occupation convention and executable training pipeline now make this physical assumption reviewable.

7.2. Mathematical conditions and empirical evidence

The repeated diagonal Jacobian in the augmented state–tangent system is standard. The methodological contribution is its application to a direct derivative-goal reliability architecture for a nonlinear self-consistent eigenproblem, with a charge-trained compressed space, shared operator construction, and explicit decision evaluation. Proposition 1 requires local smoothness, invertible linearizations, bounded remainders, and, for ideal effectivity, a scalar-output nondegeneracy condition. Favorable sample statistics cannot establish any of these assumptions.

The model bridge matters because the projected quantum reliability residual and the full-grid reference retain the same number of subbands but use different spaces. Compression and numerical differentiation add further terms. The observed relationship between corrections and true errors supports the method's practical usefulness within the tested setting; the effectivity tails and controller failures show its limits. The independent decision-band values include appreciable underestimation and overestimation, so a central concentration near unity does not imply uniform accuracy even within that band. Interval verification provides a local enclosure when its independent checks succeed. Its target remains narrower than the continuum or device accuracy.

7.3. Controller interpretation and generalization

The primary controller is selected by a reproducible procedure and evaluated on the new holdout, which yields 1 false acceptance and 6 false rejections. Zero development false acceptances are a calibration constraint, not a guarantee about future tuples. The historical comparator yields 0 false acceptances and 58 false rejections on the same independent states. The tradeoff depends on the intended costs of false acceptances and false rejections; this study does not infer an application-specific deployment loss or operating distribution.

The structured design, repeated algorithmic states, targeted confirmation, and independent uniform tuples serve different purposes. Repeated states reveal sensitivity to algorithmic choices, although states at one physical tuple are statistically dependent. A verified example from an adaptive search establishes existence, whereas failure to find an example in a finite search cannot prove nonexistence. Neither outcome measures prevalence. The tuple-level uncertainty calculation applies only to the new independent sampling design. A device workflow with a different input distribution requires a separate assessment.

The augmented tangent–adjoint construction and repeated diagonal block also apply to other differentiable nonlinear parametric residuals. Their implementation depends on the geometry, quantum discretization, material models, regression features, snapshot spaces, numerical differences, controller constants, and cost balance. Higher dimensions, open boundaries, larger spaces, additional parameter derivatives, or higher-order Hamiltonians require new derivative calculations, bases, calibration, and independent validation. The measured rank and timing results cannot be transferred unchanged.

8. Conclusions

The corrected study implements online derivative-goal error estimation for reduced transverse Schrödinger–Poisson electrostatics with a charge-trained augmented adjoint and one shared compressed dual operator. After correcting the omitted spin factor while retaining the valley multiplicities, we regenerated the offline assets and numerical results. We selected the controller on development data and evaluated it on 1024 new physical tuples: the primary rule made one false acceptance and six false rejections, compared with zero and 58 for the historical rule. The projection studies, finite-difference checks, measured timings, and reference-grid tests show both what the approach achieves and where its accuracy is limited. Interval-verified, C-dominated near failures establish that such configurations exist just above the capacitance tolerance. Selected-point certificates prove local discrete output statements, while the controller remains empirical. The results do not supply a uniform reliability bound, a continuum-accuracy bound, or a prediction of deployment risk.

Author contributions

Ahmed Alqurashi: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Visualization, Writing and revision. The author is responsible for the scientific content and final approval of the submitted manuscript.

Use of Generative-AI tools declaration

Generative AI tools, including ChatGPT, were used to assist with manuscript drafting, language editing, and proofreading. The author carefully reviewed, revised, and verified the AI-assisted text and takes full responsibility for the accuracy, integrity, and final content of the manuscript.

Conflict of interest

The author declares no conflict of interest.

References

1. P. Benner, S. Gugercin, K. Willcox, A survey of projection-based model reduction methods for parametric dynamical systems, *SIAM Rev.*, **57** (2015), 483–531. https://doi.org/10.1137/130932715
2. J. S. Hesthaven, G. Rozza, B. Stamm, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, Cham: Springer, 2016. https://doi.org/10.1007/978-3-319-22470-1
3. C. Prud'homme, D. V. Rovas, K. Veroy, L. Machiels, Y. Maday, A. T. Patera, et al., Reliable real-time solution of parametrized partial differential equations: Reduced-basis output bound methods, *J. Fluids Eng.*, **124** (2002), 70–80. https://doi.org/10.1115/1.1448332
4. G. S. H. Pau, Reduced basis method for simulation of nanodevices, *Phys. Rev. B*, **78** (2008), 155425. https://doi.org/10.1103/PhysRevB.78.155425

5. H. Li, G. Li, Component mode synthesis approaches for quantum mechanical electrostatic analysis of nanoscale devices, *J. Comput. Electron.*, **10** (2011), 300–313. https://doi.org/10.1007/s10825-011-0366-7
6. M. D. Ganeriwala, C. Yadav, F. G. Ruiz, E. G. Marin, Y. S. Chauhan, N. R. Mohapatra, Modeling of quantum confinement and capacitance in III–V gate-all-around 1-D transistors, *IEEE Trans. Electron Devices*, **64** (2017), 4889–4896. https://doi.org/10.1109/TED.2017.2766693
7. B. Peng, Y. Jiao, H. Zhong, Z. Rong, Z. Wang, Y. Xiao, et al., Compact modeling of quantum confinements in nanoscale gate-all-around MOSFETs, *Fundam. Res.*, **4** (2024), 1306–1313. https://doi.org/10.1016/j.fmre.2022.09.035
8. R. Becker, R. Rannacher, An optimal control approach to a posteriori error estimation in finite element methods, *Acta Numer.*, **10** (2001), 1–102. https://doi.org/10.1017/S0962492901000010
9. M. Yano, Goal-oriented model reduction of parametrized nonlinear partial differential equations: Application to aerodynamics, *Int. J. Numer. Methods Eng.*, **121** (2020), 5200–5226. https://doi.org/10.1002/nme.6395
10. M. K. Sleeman, M. Yano, Goal-oriented model reduction for parametrized time-dependent nonlinear partial differential equations, *Comput. Methods Appl. Mech. Eng.*, **388** (2022), 114206. https://doi.org/10.1016/j.cma.2021.114206
11. P. J. Blonigan, E. J. Parish, Evaluation of dual-weighted residual and machine learning error estimation for projection-based reduced-order models of steady partial differential equations, *Comput. Methods Appl. Mech. Eng.*, **409** (2023), 115988. https://doi.org/10.1016/j.cma.2023.115988
12. H. Fischer, J. Roth, T. Wick, L. Chamoin, A. Fau, MORe DWR: Space–time goal-oriented error control for incremental POD-based ROM for time-averaged goal functionals, *J. Comput. Phys.*, **504** (2024), 112863. https://doi.org/10.1016/j.jcp.2024.112863
13. B. N. Granzow, D. T. Seidl, S. D. Bond, Linearization errors in discrete goal-oriented error estimation, *Comput. Methods Appl. Mech. Eng.*, **416** (2023), 116364. https://doi.org/10.1016/j.cma.2023.116364
14. A. Janon, M. Nodet, C. Prieur, Goal-oriented error estimation for the reduced basis method, with application to sensitivity analysis, *J. Sci. Comput.*, **68** (2016), 21–41. https://doi.org/10.1007/s10915-015-0127-z
15. M. Guo, H. Zhong, Strict upper and lower bounds for quantities of interest in static response sensitivity analysis, *Appl. Math. Model.*, **49** (2017), 17–34. https://doi.org/10.1016/j.apm.2017.04.029
16. D. J. BenDaniel, C. B. Duke, Space-charge effects on electron tunneling, *Phys. Rev.*, **152** (1966), 683–692. https://doi.org/10.1103/PhysRev.152.683
17. F. Stern, Self-consistent results for n-type Si inversion layers, *Phys. Rev. B*, **5** (1972), 4891–4899. https://doi.org/10.1103/PhysRevB.5.4891
18. T. Ando, A. B. Fowler, F. Stern, Electronic properties of two-dimensional systems, *Rev. Mod. Phys.*, **54** (1982), 437–672. https://doi.org/10.1103/RevModPhys.54.437
19. A. Trellakis, A. T. Galick, A. Pacelli, U. Ravaioli, Iteration scheme for the solution of the two-dimensional Schrödinger–Poisson equations in quantum structures, *J. Appl. Phys.*, **81** (1997), 7880–7884. https://doi.org/10.1063/1.365396

20. W. Liu, X. Jin, Y. King, C. Hu, An efficient and accurate compact model for thin-oxide-MOSFET intrinsic capacitance considering the finite charge layer thickness, *IEEE Trans. Electron Devices*, **46** (1999), 1070–1072. https://doi.org/10.1109/16.760418

21. M. Rewieński, A. Lamecki, M. Mrozowski, A goal-oriented error estimator for reduced basis method modeling of microwave devices, *IEEE Microw. Wirel. Compon. Lett.*, **25** (2015), 208–210. https://doi.org/10.1109/LMWC.2015.2400937

22. L. Feng, S. Chellappa, P. Benner, A posteriori error estimation for model order reduction of parametric systems, *Adv. Model. Simul. Eng. Sci.*, **11** (2024), 5. https://doi.org/10.1186/s40323-024-00260-8

23. D. Blais, S. Nadarajah, C. Biondic, Goal-oriented adaptive sampling for projection-based reduced-order models, *Comput. Fluids*, **290** (2025), 106568. https://doi.org/10.1016/j.compfluid.2025.106568

24. A. Arnold, Mathematical concepts of open quantum boundary conditions, *Transp. Theory Stat. Phys.*, **30** (2001), 561–584. https://doi.org/10.1081/TT-100105939

25. A. Arnold, N. Ben Abdallah, C. Negulescu, WKB-based schemes for the oscillatory 1D Schrödinger equation in the semi-classical limit, *SIAM J. Numer. Anal.*, **49** (2011), 1436–1460. https://doi.org/10.1137/100800373

26. W. Wang, C. W. Shu, The WKB local discontinuous Galerkin method for the simulation of Schrödinger equation in a resonant tunneling diode, *J. Sci. Comput.*, **40** (2009), 360–374. https://doi.org/10.1007/s10915-008-9237-1

27. G. E. Aliffi, G. Nastasi, V. Romano, Ballistic electron transport described by a fourth-order Schrödinger equation, *Z. Angew. Math. Phys.*, **76** (2025), 155. https://doi.org/10.1007/s00033-025-02535-5

28. G. E. Aliffi, G. Nastasi, V. Romano, Ballistic electron transport described by a generalized Schrödinger equation, *J. Phys. A: Math. Theor.*, **59** (2026), 135201. https://doi.org/10.1088/1751-8121/ae5516

29. S. M. Rump, Verification methods: rigorous results using floating-point arithmetic, *Acta Numer.*, **19** (2010), 287–449. https://doi.org/10.1017/S096249291000005X

30. Python-flint developers, *Python-flint 0.9.0 documentation: arb and arb_mat*, 2026. `https://python-flint.readthedocs.io/en/latest/arb_mat.html` (accessed 29 October 2026).

31. C. J. Clopper, E. S. Pearson, The use of confidence or fiducial limits illustrated in the case of the binomial, *Biometrika*, **26** (1934), 404–413. https://doi.org/10.1093/biomet/26.4.404

32. M. Drohmann, K. Carlberg, The ROMES method for statistical modeling of reduced-order-model error, *SIAM/ASA J. Uncertain. Quantif.*, **3** (2015), 116–145. https://doi.org/10.1137/140969841

33. K. Smetana, O. Zahm, A. T. Patera, Randomized residual-based error estimators for parametrized equations, *SIAM J. Sci. Comput.*, **41** (2019), A900–A926. https://doi.org/10.1137/18M120364X

Supplementary

The supplementary package (math-11-11-1502-s001) contains the corrected physics and reliability implementation, executable 180-tuple training procedure, rebuilt offline arrays, frozen input designs,

complete controller candidate sweep, row-level calibration and validation results, diagnostic sweeps, all nine figures, and selected-point interval proofs. The current explicit implementation and rebuilt assets define the corrected reproducible target; byte identity with the unavailable original v1.x source is not claimed. Earlier numerical outcomes are superseded. A small set of earlier CSV files is retained solely to reproduce input coordinates and design provenance; their old output columns are not corrected evidence. The README file identifies canonical corrected results and provides commands, dependency versions, and checksums.



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0)