# Supplemental Material: IonGate: Circuit Memory and Waveform-Protected Learning for Functional Inference in Cardiac Electrophysiology

Authors: Pengyuan Xu (1), Tianyi Ma (2), Peng Xue (3)*, Jicheng Li (2)*

Affiliations and correspondence: (1) Department of Materials Science and Engineering, Monash University, Clayton, Victoria 3800, Australia.

(2) Department of Neurosurgery, The Second Affiliated Hospital of Xuzhou Medical University, Xuzhou, Jiangsu 221006, China.

(3) Department of Critical Care Medicine, The Second Affiliated Hospital of Xuzhou Medical University, Xuzhou 221006, China.

* Joint corresponding authors: Peng Xue, xuepeng08262024@163.com; Jicheng Li, 3272464246@qq.com

**Author Contributions**

Pengyuan Xu developed the model. Tianyi Ma collected the publicly available data and formatted the manuscript. Peng Xue and Jicheng Li provided overall scientific guidance and supervision.

**Sources of Funding**

None

**Disclosures**

None

### Ethics and Data Reuse

This study was a secondary analysis of publicly available experimental recordings and images. No new participants were recruited, animal experiments performed or biological samples collected. No new ethics approval was required for the present secondary analysis. The original studies describe the approvals governing their data collection.

### Reference Convention

All bibliographic references for this Supplemental Material are included in the main manuscript reference list. Archived historical figure plates retain their original source labels; the reference-number crosswalk below translates original manuscript citation labels to the current main reference list.

Original → current reference numbers: 1 → 1; 2 → 2; 3 → 3; 4 → 4; 5 → 5; 6 → 6; 7 → 7; 8 → 8; 9 → 18; 10 → 17; 11 → 12; 12 → 13; 13 → 14; 14 → 15; 15 → 16; 16 → 19; 17 → 9; 18 → 10; 19 → 11.

## S1 Scope and implementation

This supplement preserves complete development evidence. Version numbers are internal candidate identities, not independent datasets. The main paper evaluates the human-confirmed V30 complete model. Legacy figure citations below map to Supplementary Figures S1–S19 and legacy tables to Supplementary Tables S1–S16. Earlier endpoint conventions are preserved and must not be numerically pooled with the final window-initial functional.

## S2 Exact observer, fitting and historical calibration methods

### Public data and distinct reference roles

We reanalyzed publicly released data and acquired no new animals, participants or biological samples. The primary input was Lei electronic-model-cell file 20230510-Cfast_auto.dat from Figshare 27193878 [5]. Its HEKA export preserved current, command voltage, independently recorded second-amplifier voltage, timestamps and compensation labels across 16 settings. Native sampling was 10 kHz; the inclusive 45–80 ms primary window contained 351 samples per setting. All settings came from one acquisition session. Supplementary Table S1 defines each dataset's role and unit of analysis.

For Figure S2, the scoring reference was measured second-amplifier Vm. For Figs. 3 and 4, it was a specified circuit functional computed from that voltage, nominal elements and a declared initial-state convention. These references address voltage reconstruction and conditional circuit-charge inference, respectively. The second sensor was permitted for training and separate calibration groups, but never as a test-time predictor input.

Additional resources comprised Lei NaV1.5 HEK recordings, Sato Zenodo 8288407, published Ma/Glazer SCN5A summaries, Kemp hERG recordings and replay, He cardiac optical mapping, and the Xie study's public paired neuronal record [1,2,12,13,14]. We retained separate labels for measured voltage, programmed states, source-model voltage, published functional summaries and external image context. Raw inputs, processing manifests and figure-source hashes remain archived; Tables S1–S8 preserve unrounded outputs and provenance.

### First-amplifier inverse observer

IonGate comprises a voltage observer, a state-aware circuit readout and a grouped calibration layer. We evaluated these components through independent-voltage scoring (Figure S2),

fixed-observer ablation (Figure S3) and held-out interval scoring (Figure S4), respectively. Let t index the recorded time grid, Iout denote output current, Vcmd command voltage, and αR and αP the known resistance and supercharging compensation fractions. The observer parameters are θ = (τout, Rs, R̂s, Ĉm, Voffset). Rs represents the series voltage-drop parameter, while R̂s represents the compensation estimate; they are distinct fitted quantities.

The observer inverts the pinned Lei front-end relations [5]. It first corrects the output-filter relation using time constant τout:

$$\hat{I}_{\text{in}} = I_{out} + \tau_{out}\frac{dI_{out}}{dt} \quad \text{(S3)}$$

We computed dIout/dt using a seven-sample, second-order Savitzky–Golay filter, derivative order one, median native sample spacing and interpolation edge handling [15]. Compensation dynamics obey

$$\frac{dV_{est}}{dt} = \frac{V_{cmd} - V_{est}}{(1 - \alpha_P)\hat{C}_m\hat{R}_s}, \quad V_{est}(0) = V_{cmd}(0) \quad \text{(S4)}$$

For each piecewise-constant command interval, Vest,i+1 = Vcmd,i + (Vest,i − Vcmd,i) exp(−Δti/τv), where τv = (1 − αP)ĈmR̂s. This exact interval propagation retains the compensation state instead of replacing it by instantaneous command voltage. The observer then calculates pipette and membrane voltage:

$$V_p = V_{cmd} + \hat{R}_s[\alpha_R I_{out} + \alpha_P\hat{C}_m\frac{dV_{est}}{dt}] \quad \text{(S5)}$$

$$\hat{V}_m = V_p + V_{offset} - R_s[\hat{I}_{\text{in}} + \hat{C}_m\frac{dV_{est}}{dt}] \quad \text{(S6)}$$

Resistances are in GΩ, capacitance in pF, current in pA and time in ms, giving voltage in mV. Inputs are timestamps, Iout, Vcmd and known compensation labels; measured test Vm is excluded. Neighboring samples enter the derivative estimate, making this an offline observer

rather than a demonstrated causal real-time estimator. The compensation state Vest in Eq. (4) describes the amplifier front end, whereas branch current Ik in Eq. (2) describes the model-cell readout. These are separate dynamical variables with different roles. Figure S2 evaluates the observer containing Vest; Figure S3 fixes that observer and changes only the charge contribution of Ik. The ablation therefore does not test removal of all dynamics from IonGate.

**Grouped fitting and prediction freeze**

We partitioned compensation settings rather than individual time samples. Setting identifiers g ran from 0 to 15. In fold f, test settings satisfied g mod 4 = f, calibration settings satisfied g mod 4 = (f + 1) mod 4, and the remaining eight settings trained the observer. Each setting appeared once as a held-out test setting across four folds. Within any fold, its training, calibration and test settings were disjoint. This design tests within-session compensation transfer, not cross-session replication.

The five shared parameters used initial vector θ0 = (0.110 ms, 0.030 GΩ, 0.03004 GΩ, 44 pF, 0 mV). Lower bounds were (0.040, 0.010, 0.010, 10, −20), and upper bounds were (0.300, 0.070, 0.070, 300, 20), in corresponding units. Training voltage residuals were $\hat{V}m - Vm$ at every fourth sample of the inclusive 45–80 ms window, yielding 88 sampled residuals per training setting. We concatenated these with five regularization residuals $2(\theta - \theta 0)/s$, where s = (0.08 ms, 0.015 GΩ, 0.015 GΩ, 45 pF, 8 mV).

SciPy least_squares minimized this concatenated vector using soft_l1 loss, f_scale = 1 and max_nfev = 180. For the concatenated numerical residual vector r, the objective was one half of the sum of $\rho(rj^2)$, with $\rho(z) = 2(\sqrt{(1 + z)} - 1)$. The same robust loss acted on voltage and regularization residuals; the latter were not a separately added quadratic penalty. Tables S8A and S8B record exact frozen parameters and group roles. Training used second-sensor voltage only on training settings; interval calibration used it on separate calibration settings.

The freeze stage wrote test voltage predictions and endpoint estimates before the score stage loaded test Vm. A SHA-256 manifest locked predictions and metadata; the prediction CSV hash is 1719891faea17f29b4f550e0eae529408020c3df934d7b393b8033df3b0ed34a. For each test setting, the frozen observer first generated voltage from permitted first-amplifier inputs. The readout then generated charge, and the stored same-fold calibration radius generated the interval and prediction-width decision. Test Vm entered only subsequent scoring. Subsequent figure generation, robustness analysis and manuscript revision did not refit the observer. Phase and stress analyses were post-freeze descriptive analyses of previously inspected data.

**Comparator implementation and voltage endpoints**

Voltage comparators were measured Vcmd, Vcmd − 0.030Iout, Vcmd − 0.035Iout, and the source-exact published Lei forward model. Static-resistance coefficients are in GΩ. The Lei model received the same recorded command and compensation labels, with its published parameters held fixed. IonGate instead used parameters estimated on disjoint training settings. We therefore compared these specified implementations, not model classes receiving identical optimization. Published-model replay used the source implementation and Myokit infrastructure [16].

We calculated RMSE, MAE, signed bias and the within-setting 95th percentile of absolute voltage residuals. RMSE was the square root of mean squared residuals, MAE the mean absolute residual, and bias the mean signed residual. Across-setting summaries reported medians and P95 of setting-level metrics; P95 setting RMSE was not a pointwise error bound or confidence interval. The primary full window included 45 and 80 ms. Post-freeze explanatory windows were onset [50,52), plateau [52,69) and offset [70,72) ms, with 20, 170 and 20 samples after rounding timestamps to six decimals.

Setting wins compared paired RMSEs. Samples within a waveform and settings within the session were not treated as independent experimental sessions. We reported descriptive differences without deriving cross-session significance. Native close-ups in Figure S2l–n used fixed setting 00 and illustrative display windows; those windows did not alter the scoring definitions.

**State-aware circuit readout and matched ablation**

The circuit readout used nominal gm = 10 pA/mV, gk = 100 pA/mV and Ck = 100 pF, giving τk = 1 ms. For piecewise-linear voltage with slope si = (Vi+1 − Vi)/Δti, we propagated the branch current exactly:

$$I_{k,i+1} = e^{-\Delta t_i/\tau_k} I_{k,i} + C_k s_i \left[1 - e^{-\Delta t_i/\tau_k}\right] \quad \text{(S7)}$$

For Figure S3, Ik started at nominal zero at the beginning of the record and propagated to both scoring endpoints. Each predicted or reference waveform supplied its own mean voltage over 40–49.9 ms as Vhold. We used trapezoidal integration from 45 to 80 ms, including both endpoints, and Eq. (2) with the propagated states. This calculation avoids numerical differentiation of measured Vm for the charge functional. In stored units, pA × ms and pF × mV give fC; plotted pC values divide by 1,000.

We computed L = gm ∫(Vm − Vhold)dt, C = Ck[Vm(tb) − Vm(ta)] and S = −τk[Ik(tb) − Ik(ta)]. Full, no-memory and leak-only readouts were L + C + S, L + C and L, respectively. The ablation retained identical observer predictions, nominal elements, integration samples, hold definition and state trajectory; it changed only the included readout terms. We did not refit the observer, reinitialize state or retune a threshold for an ablation. The reference was always the full functional of measured second-sensor voltage under the same convention. We did not remove its memory term when scoring a no-memory prediction. Thus, all readouts estimated the same target rather than redefining the reference for each ablation.

For signed error analysis, let δL, δC and δS denote predicted minus reference contributions. The full residual was $\delta Q_{full} = \delta L + \delta C + \delta S$, whereas the no-memory residual was $\delta Q_{no\text{-}memory} = \delta L + \delta C - S_{ref}$. This decomposition separates component-estimation errors from omission of reference memory. It explains the readout design; the paired absolute discrepancies establish whether including memory improved a setting. The identity alone does not require every setting to improve.

We normalized discrepancy as $100|\hat{Q} - Q_{ref}|/\max(|Q_{ref}|, 100\ fC)$; the floor prevents unstable division near zero. We retained signed components, all 16 settings and every paired reversal. Relative reductions of median and P95 discrepancies were calculated from the respective summary ratios. Their values are not medians of per-setting percentage reductions. Component closure checked numerical implementation only. The reference depends on topology, nominal elements and the initial-state convention, rather than an independent current-meter integral.

**Conditional intervals and selection policies**

Figure S4 retained the archived center_charge endpoint: each waveform used its own window-initial voltage as Vhold and Ik(45 ms) = 0. The same 45–80 ms samples and nominal elements were used for predictions and references. This window-initial convention differs from Figure S3's hold-based, record-propagated convention; we neither pooled their charge errors nor interpreted one as an improvement over the other.

For each method and fold, we calculated four absolute calibration residuals $|\hat{Q} - Q_{ref}|$ from that method's predictions and the corresponding voltage-derived references. The rank was $k = \text{ceil}[(n_{cal} + 1)(1 - \alpha)]$. When this rank existed, radius r was the kth ordered calibration residual and the test interval was $[\hat{Q} - r, \hat{Q} + r]$. With $n_{cal} = 4$ and $\alpha = 0.2$, r equaled the maximum calibration residual. At $\alpha = 0.05$, rank five exceeded the calibration sample size, so

this rule supplied no finite radius. Every method used its own same-fold calibration residuals, never held-out errors, to set r.

For a nominal 80% interval, the score was 2r + 10 max(|$\hat{Q}$ − Qref| − r,0). We normalized scores and summary full widths by max(|Qref|,100 fC) and multiplied by 100. This score penalizes width plus missed-reference distance. Signed containment margins in Figure S4r used 100(r − |$\hat{Q}$ − Qref|)/max(|Qref|,100 fC); nonnegative values indicated containment.

The operational reporting rule used prediction-normalized width, 100 × 2r/max(|$\hat{Q}$|,100 fC), at budgets 0.25, 0.5, 1, 1.5, 2, 2.5 and 25%. A setting was accepted when its width did not exceed the budget. This denominator differs from the reference-normalized summary width and requires no test reference at selection time. Containment margins and reference-normalized metrics were computed after reveal as evaluation outputs, not inputs to width selection. We retained empty-set statistics as unavailable, reported every miss and compared common accepted subsets. Observed containment describes the fixed compensation groups; split-conformal coverage additionally depends on exchangeability [18]. Width-based abstention alone was not defined as a detector of acquisition mismatch.

**Analytical error bounds are distinct from empirical calibration**

For known circuit parameters and the window-initial convention, the archived implementation writes the piecewise-linear functional as Q = ΣwiVi with Σwi = 0. Constant common voltage offset therefore cancels in this functional. If each voltage sample has an independently justified box error εi, the sharp sample-box radius is Σ|wi|εi: each error can take the sign of its coefficient to attain the extremum. With bounded interpolation error εint, initial-current uncertainty b0 and model discrepancy bm, an additional outer radius is

$$r_{outer} = \varepsilon_{\int}\left[g_m T + C_k\left(1 - e^{-T/\tau}\right)\right] + b_0\tau_k\left(1 - e^{-T/\tau}\right) + b_m \quad \text{(S8)}$$

The implementation also encloses continuous positive parameter boxes using analytical extrema of gk exp(−u gk/Ck), rather than assuming a finite parameter grid captures the extrema. These deterministic bounds are conditional on externally justified voltage, parameter, state and model-error budgets. The current experiment does not supply a traceable budget, so the physical-certification gate remains separate from the empirically calibrated intervals reported in Figure S4. Implementation details and assumptions are retained in the supplementary analysis archive.

**Frozen NaV1.5 application**

The biological application freezes the median τout and Voffset across the four existing model-cell folds, substitutes public per-sweep Rs and Cm metadata for the corresponding resistance/capacitance terms, sets $\hat{R}s = Rs$, and uses $\alpha R = \alpha P$ = compensation/100. This is a declared metadata-adapted transfer, not preservation of every numerical model-cell coefficient. The two cell3/25°C files without per-voltage metadata use the source author's cell-level info.csv fallback. We subtract the first-5-ms current mean, preserve the source 25 kHz time grid, and reconstruct command timing from the published protocol.

The scored comparison uses cell6 at 35°C and the source Gray + VC_MIN fit trained on 0% and 80% compensation. We choose the compatible archived fit by its published training objective, never by the held-out score, and keep it fixed. The 20%, 40% and 60% settings contain 15 voltage sweeps each. Scoring spans [8,14) ms and aggregates sweep RMSE first within condition and then across the three conditions. Unscored application outputs and the 45 scored outputs are separately listed. Fold-parameter sensitivity uses the original four variants without outcome-driven selection.

**Hardware controls, external assays and raw image display**

The Sato experimental-noise subset contains 240 hardware records, each with time, programmed 0/1 state and current [12]. Each record has 500,000 samples at 25 kHz. Within each noise band, low-rate k1 = k2 = 1 records determine closed baseline and open-minus-closed amplitude; a low-rate target excludes itself from its calibration set. We compare baseline-subtracted output area with amplitude × integrated open occupancy, and use inverted labels with the same amplitude and baseline for the negative control. Relative and full-scale denominators are reported separately. This analysis does not involve IonGate predictions.

SCN5A analyses align published site-level feature summaries without training an IonGate variant classifier. hERG analyses preserve source-model protocol holdout/replay and the observed areas of four original genotype/drug traces. Optical displays preserve raw source frames, timestamps, crops, native scale information and common linear contrast. Figure S7 uses the mean of 100 frames for the field and 18 actual nearest-time frames around the first upward 0 mV electrode crossing. It retains all 2,273 optical/electrode pairs and treats repeated events as within-neuron observations, not biological replicates. The primary paper citation is Xie et al.; the public paired dataset lists Rozsa as a data contributor [14]. Tables S7 and S8 record provenance and hashes.

**Post-freeze robustness and reproducibility**

The S18 analysis retains windows [45,80], [49,75], [49,80], [50,75], [50,80] and [45,85] ms; hold windows [40,49.9], [42,49.9] and [45,49.9] ms; trapezoidal/Simpson integration; and propagated, reset or ±100 pA observer start-state challenges. The nominal measured-voltage reference remains unchanged by observer-state perturbation. For perturbation $\delta I_k$ at window start, the induced functional shift is $\tau_k \delta I_k[1 - \exp(-\text{duration}/\tau_k)]$. We saved the child

protocol before computing this grid, while acknowledging that the underlying dataset had already been inspected.

The S19 analysis applies no added filtering or a causal second-order Butterworth low-pass at 2,000, 1,000 or 500 Hz to both first-amplifier inputs, initialized from the first value. Anti-aliased resample_poly with line padding precedes decimation factors 1, 2, 5 and 10. The seven-sample derivative and all model coefficients remain unchanged; its physical support therefore varies with sampling rate. The native second-sensor voltage is only interpolated to score timestamps. Command/static-resistance comparators receive the same challenged inputs, and native Lei predictions are not presented as a matched challenged comparator. Original interval radii remain fixed; reported stress hits have no recalibrated coverage guarantee.

The archive contains exact protocols, frozen predictions, input/output hashes, per-setting/sweep outputs, software versions, figure sources and regeneration instructions. Numeric workbook round-trip checks compare every stored value with its CSV source. Existing recorded data and prediction freezes remain unmodified. This manuscript stage adds writing and document artifacts only; it does not create new performance results.

**Matched-budget optimization and prospective-update safeguards**

The retrospective extension retained the original four outer folds and 45–80 ms endpoint. The five shared free parameters were output-filter time constant, actual and estimated series resistance, estimated membrane capacitance and effective voltage offset. Bounds were [0.040,0.300] ms, [0.010,0.070] GOhm for both resistances, [10,300] pF and [-20,20] mV. Both methods used the original soft-L1 objective, the same prior scale, three deterministic starts and maximum 180 optimizer evaluations per start. Model-cell gm, gk, Ck and actual Cm remained nominal. The initial hold was -40 mV for the forward source model; its

measured command was linearly interpolated using an exact first-order-hold state-space implementation. Numerical equivalence was checked against the original source solver. Equal optimization caps do not imply identical actual run times or identical structural use of measured current: the inverse consumes Iout and Vcmd, whereas the forward model consumes Vcmd.

Ridge corrections used seven fixed features: intercept, current, output-time-constant-scaled current derivative, command displacement, two compensation labels and a current-by-resistance-compensation interaction. Coefficients were fitted on inner training samples only; penalties 1, 10 and 100 were compared on the two inner-validation settings. Corrections were clipped to ±5 mV. After selection, a retained candidate was refitted on the outer training pool and interval radius recalibrated on the disjoint outer calibration pool. No outer test target selected a model. The later multi-endpoint release guard is a future-session policy, not a retroactive explanation of the voltage-only candidate's performance.

The acquisition gate allowed ±1% relative deviation from the native 0.1 ms sampling interval, required unchanged verified additional-filter metadata and rejected a hold-current difference MAD below one quarter of the minimum training value. Native and transformed test signals were evaluated without target access; recorded second-sensor voltage entered only subsequent scoring. A deliberately inaccurate metadata challenge separated metadata-aware rejection from signal-only detection. For downsampling factors four and eight, missing 45 and 80 ms charge-integration endpoints were linearly interpolated, with this clarification stored before mismatch scoring.

For biological activation, all pre-existing four-cell NaIV records were retained. A common observed current minimum between 0.2 and 4 ms after step onset defined each peak. Apparent conductance was -Ipeak/(ENa-Vpeak), using the source 140/10 mM sodium concentrations at the recorded temperature. Positive conductance with at least 10 mV reversal

separation entered a bounded Boltzmann fit over command voltages -60 to +20 mV; -60 to 0 mV was the fixed sensitivity window. At least six points were required. Fits, failures and parameter-bound information were reported for every available condition. A post-score audit separately flagged half-activation not bracketed by two observations on either side, near-boundary solutions, Jacobian rank/conditioning problems or approximate local Vhalf standard error above 10 mV. This audit restricts interpretation and does not replace the original endpoint or create confirmatory subgroup evidence.

The optional update interface requires new independently labelled sessions, disjoint session roles and reference-traceability records. Protected test signal files are not opened. Candidate snapshots are immutable, previously consumed validation batches cannot be reused for repeated tuning, and voltage improvement alone is insufficient for future release. A candidate also needs noninferior median nominal circuit-functional discrepancy and bounded voltage/functional tails. Any released voltage update starts with intervals disabled until fresh separate calibration. No background service or unattended endless retraining was enabled.

**Joint observer, cell-held-out development and four-role growth**

The innovation observer used a source-predicted current residual converted to voltage by series resistance. Fixed coefficients were fitted with training-only feature scales and a group-balanced voltage-plus-functional ridge objective. The memory family added causal exponential states at 0.1, 0.3, 1, 3 and 10 ms, an output-time-constant-scaled derivative and compensation interactions. Matched FIR and affine families received the same wave-hyperparameter budget. A training-calibrated inverse-functional anchor was imposed by the minimum-energy smooth full-record correction for the declared 45–80 ms query. The anchor correction was bounded to 2% of its first-channel estimate and the waveform shift to 5 mV. All three comparator families received the same anchoring opportunity. This is an offline

query-conditioned estimator, not an unconstrained estimate of the complete physical voltage trajectory.

For repeated-inner selection, each of four outer folds contained four internal rotations, with three starts and 180 maximum optimizer evaluations per source/inverse fit on each internal fitting pool. Root fits were reused only when the training indices were identical. Every family received 12 waveform and 36 anchor candidates, or 432 joint combinations. Joint worst normalized median performance and tail guards determined selection; numerical ties were resolved at six decimal places rather than by floating-point residuals around 1e-13. Outer calibration retained the maximum absolute residual over four calibration settings. This rank supports descriptive 80% single-session calibration under exchangeability, not a traceable 95% physical-charge certificate.

Biological fitting retained the previous voltage support and identifiability tests, but each model-defined current proxy supplied its own peak within the common 0.2–4 ms window. IonGate BioRC and the static-dynamic control each tested 54 configurations. Selection minimized equally weighted training-cell losses; temperatures remained nested within cell. A second 27-policy ensemble search was available to both dynamic families. The trust variant used Vhat = Vcmd − gamma Rs[(1−alpha)Iout + tau dIout/dt + (1−alpha)Cm dVest/dt] + Voffset. Gamma is an estimated reliability/shrinkage coefficient, not independently calibrated physical resistance. The apparent ionic-current proxy and the conductance curve remain model-derived.

The new future-session updater requires at least two independent training, four calibration and four validation sessions, plus a protected test role that is never opened. Coefficients and functional calibration learn from training sessions only; interval radii use separate calibration sessions. Validation uses session-level aggregated endpoints and requires at least 1% lower medians for voltage RMSE, nominal functional discrepancy and interval score against every

required strong comparator, without worsening any P95 or reference-hit count. Duplicate content, role overlap and reused validation batches are rejected, and immutable snapshots support rollback. No autonomous background retraining or new physical-session collection occurred in this study.

**Unified RC-State development and release audit**

The passive-state implementation used Cm = 44 pF, gm = 10 pA/mV, gk = 100 pA/mV and Ck = 100 pF as nominal electronic-model-cell parameters. The state equations were Cm dVm/dt = Ie − gm Vm − Ik and dIk/dt = gk dVm/dt − gk Ik/Ck, with Ie = Iout + tau dIout/dt + Ce dVest/dt. Stable first-order-hold state responses were computed around the initial steady input. These nominal quantities define the circuit-functional reference and do not establish component calibration or absolute charge traceability.

For every fold, all extra training traces belonged to the allowed fitting groups only. Fixed stimulus indices were step0, sine0, staircase0, action-potential0 and ramp0/9. Each trace contributed at most 500 uniformly spaced samples after 20 ms, with equal total weight per group and stimulus. The step45–80ms circuit-functional objective remained unchanged. Earlier source/inverse prior fits were reused only for exactly matching training pools. This is group-held-out evaluation on stimulus classes present in training, not a new held-out-protocol replication.

The orthogonal head used a training-only pseudoinverse to fit six static nuisance features and residualize the family-specific features. Residual ridge penalties were 0.01, 0.1, 1, 10, 100 and 1000, crossed with functional weights 0, 0.1, 1 and 10: 24 candidates per family. Four internal4fit/2cal/2validation rotations selected one candidate; eight outer fitting and four calibration settings remained disjoint from the four tests. The radius was the maximum absolute functional residual on the four calibration records. All native and historical strong

comparators were retained. No candidate used held-out sensor labels to fit, but repeated inspection makes this retrospective development rather than independent confirmatory testing.

Biological state derivatives were evaluated with the exact discrete FOH coefficients b0 = 1 − (tau/dt)(1−a), b1 = (tau/dt)(1−a) − a and a = exp(−dt/tau). The matched FIR control truncated the same smoothing response to32taps; compensationVest retained its commanded-step convention. Each family used72equal-grid configurations and the unchanged peak window, activation support and fit-identification rules. Unknown Rseal added zero extra correction for all models. A unit-slope ramp check is a software/mathematical control, not a biological accuracy reference.

**Protected-state development and complete retention audit**

V16 used the existing four outer folds, excluding all stimulus records from four test compensation settings in each fold. Within twelve allowed settings, four base models each fitted nine and calibrated on the remaining three. Five source parameters, their three-restart/180-evaluation caps, six fixed real stimulus records per group, 500 samples per record and group/protocol weights were unchanged. Prior source and affine-backbone fits were reused only from exactly matching roles; no calibration/test label entered their fitting or selection.

The complete affine basis contained six source/compensation nuisance terms and three instantaneous-current innovations. Its fitted coefficients were fixed before augmentation. The family innovation matrix was orthogonalized against all nine affine columns using training-only projection P=A-pseudoinverse E. RC contributed nine filtered KCL-state columns; FIR contributed100lags in addition to the instantaneous backbone; the stronger affine control contributed three compensation-polynomial gain terms, still affine in current conditional on

metadata. Residual ridge values1,10,100,1000,1000000 and an exact-zero option crossed functional weights0,0.1,1,10, giving24candidates each. Innovation voltage was capped at2mV without changing the protected backbone. Inner selection used matched voltage median/P95 ratios and a nominal-functional tail guard.

The readout retained four features [1,alphaR,alphaP,Qwave/25000], ridge1,10,100 and blend0,0.5,1. An optional scale was sqrt(0.1+zHz), with H the inverse training ridge information matrix and z the training-standardized feature vector. Scale choices doubled the readout grid to18per family. Staged coefficient-held-out predictions within nine groups selected the readout using median functional error, functional P95 and an internal leave-one score surrogate. Because backbone/wave hyperparameters were selected within the same nine-group pool, these internal predictions are not fully hyperparameter-held-out estimates. They are development criteria, not unbiased independent validation. The three calibration and four test settings remained excluded from every fitting/selection stage.

The scalar correction was limited to0.5%of the predicted functional and returned to the voltage waveform by a smooth minimum-energy projection of length scale1ms and maximum2mV shift. The final point estimate was the mean of four corrected voltage trajectories, and Q was recomputed from this single waveform. Twelve calibration endpoint pairs used absolute residual divided by the calibration scale and multiplied by the corresponding test scale. Primary alpha0.2 retained the second/eleventh order statistics; alpha0.1 sensitivity used first/twelfth. These are empirically evaluated scale-adapted CV+ intervals [19], not a new theorem or traceable physical-charge calibration. The primary median interval score and all original controls were retained; means remained secondary.

Biological static correction was Vcmd-B*tanh(gamma*Rs*(1-alpha)*Iout/B)+offset. Dynamic corrections included output-filter deconvolution, compensation-capacitor current and membrane-capacitive subtraction with kappa0or1. A common nonnegative signal-only

scalar bounded dynamic voltage by3mV and current increments by eta times observed peak magnitude. Public Rseal metadata supplied optional perturbation-leak subtraction;48records with missing Rseal received no extra correction in every family. Gamma0/0.25/0.5, B10/30mV, tau0/0.03/0.13ms, kappa0/1, eta0/0.2/0.6 and leak0/1 gave216configurations per family.

Whole-cell LOCO maximized the worst training-cell identifiable fraction, then its equal-cell mean, then compared worst/mean normalized penalized losses. Both temperatures of cell3 stayed together. Point eligibility, peak0.2-4ms, command support[-60,+20]mV, minimum six points, Boltzmann bounds, two-point half-activation bracketing, Jacobian condition threshold1e8 and Vhalf SE threshold10mV were unchanged. All450sweeps,24conditionfits and failures were scored. Exactly the same biological algorithm was rerun in V14,V15,V16; no cell-specific best-version selection was allowed.

V14 evaluated a quartic-tail readout with fitted residual scales; V15 evaluated twelve leave-one-group-out models fitted to eleven groups. Both failed complete retention and remain archived with their frozen predictions. V16 underwent same-coefficient RC ablation without refitting. Portable inference accepted only first-channel arrays and compensation metadata; injected second-sensor labels did not change outputs. Source, model and data hashes plus immutable prediction receipts define the candidate. Complete promotion still required all current/historical strong controls and the explicit prior-advantage retention audit. No active model was replaced, and no new physical recording or new-cell confirmation was acquired.

**First-channel tangent adaptation and bounded functional transport**

V19/V20 retained the exact four-outer-fold design and four same-role base models per fold, each using nine fitting, three calibration and four test settings. Parent physical parameters, waveform heads and their selection budgets were inherited from identical legal V13/V16

roles. A five-parameter source-current/voltage Jacobian was computed by central differences, using steps[1e-4,1e-5,1e-5,0.01,0.001] and parameter scales[0.04,0.008,0.008,15,3]. The current fitting window 20-110 ms and noise floor 10 pA were fixed. The augmented design X comprised scaled physical current derivatives, a 100 pAoffset nuisance and optionally a 1% current-gain nuisance. Equation(9) gives the unconstrained local ridge/MAP update; physical coordinates were clipped to source bounds and one prior-scale step, then propagated through the voltage Jacobian with a 5 mVwhole-record cap. These are bounded model corrections, not measurements of the true amplifier parameters.

$$\delta\hat{p} = \left(X^T W X + \Lambda\right)^{-1} X^T W \left(I_{out} - I_{model}\right) \quad \text{(S9)}$$

Each V19 family received 40 ridge/nuisance/gain configurations for each waveform and functional head. The fixed compensation mixer used V16 below 80% alphaP and approached V13 at 95%, without group-ID lookup. V20 inherited its own same-role V19 waveform update and searched 13 functional configurations: zero residual, or ridge 0.1/1/10/100 with 10/25/50 fCcaps. Three features were the query functional of a representative current-tangent correction, multiplied by 1,alphaR-0.5 andalphaP-0.5. Functional coefficients were fitted on six settings and checked on three within the nine-setting pool; selection minimized functional median subject to a training P95 non-inferiority guard. The functional target was projected onto the single waveform by Equation(10), using the existing smooth direction withQRC[psi]=1. Equality to the target holds only when the 2 mVprojection cap does not bind.

$$V_{new} = V_{shape} + \left(Q_{target} - Q_{RC}[V_{shape}]\right)\psi \quad \text{(S10)}$$

V20's scale proxy was the square root of the source-Q Jacobian quadratic form under an inverse current-MAP Hessian, normalized by the training median. Powers 0,0.25,0.5 and 1 were compared on training-only leave-one-out interval scores, with scale clamped to 0.25-4.

This is a nominal sensitivity proxy, not a traceable physical uncertainty budget. Twelve signed relative calibration residuals were transported to each test prediction using this scale, and their minimum/maximum formed a fixed envelope. Alpha 0.2 remained the primary proper interval-score penalty; alpha 0.1 was a secondary penalty on the same bounds, not a separately calibrated 90% interval. The empirical envelope is not asserted to inherit the original CV+ finite-sample guarantee[19]. All three model families received identical new opportunities and information; old stronger controls remained in the audit. Affine and FIR denote the residual backbones: both controls include the shared nonlinear current-driven adapter and bounded functional readout, rather than pure affine/FIR end-to-end operators.

The independent second-voltage sensor supplied legal training/calibration references and post-freeze test scoring, never query-time adaptation. Outer calibration/test labels were excluded from fitting and selection. Parent hyperparameters had already been selected within each nine-setting training pool, so the coefficient-held-out inner features were partial OOF development criteria rather than fully nested independent validation. Repeated inspection of the same hardware session makes all new results retrospective. The biological adapter and its 216 configurations, eligibility rules,450 sweeps and 24 fits were unchanged and actually refitted under whole-cell LOCO; it did not inherit the electronic model-cell's fixed conductance assumption. Two frozen-coefficient interventions separately removed RC heads or current-driven updates, without recalibrating intervals. Portable replay tolerances explicitly propagated 12-significant-digit CSV quantization and finite-difference roundoff, remained below 0.001 fC, and were required not to change coverage classifications; no scientific scores were altered to pass software tests.

**Conditional residual learning and bias-aware interval calibration**

V21 and V22 retained the four outer folds and four same-role base models per fold, with nine fitting, three calibration and four test settings per base model. Every candidate inherited its

own identical-role V20/V19 waveform components; no test group was used to fit a source circuit, residual head, functional learner or selection criterion. V21 searched 15 soft-threshold/gain readouts and 16 linear-bias/profile configurations per family. V22 searched 48 functional configurations: physical, quadratic-compensation or joint features, ridge penalties 0.1/1/10/100 and gains 0/0.25/0.5/1. The physical feature vector was the existing three-component first-channel current-tangent functional. Compensation features were [1, aR-0.5, aP-0.5, (aR-0.5)^2, (aR-0.5)(aP-0.5), (aP-0.5)^2]. Each design column was normalized by its training RMS. The ridge target was (Qref-Qanchor)/100, where Qanchor was the same-family V13 functional prediction, not a test-reference value.

For each configuration, functional coefficients were fitted on six settings and evaluated on the other three within the nine-setting training pool. Selection minimized the larger of median error normalized by the V13 inner median and P95 error normalized by the V20 inner P95. Both denominators were computed from the legal nine-setting training contexts, never from the outer reported scores. The selected learner was refitted on all nine settings. Its residual was capped at 50 fC before the selected gain was applied (Equation 11), and projected onto the retained current-adapted waveform by Equation 10 with the unchanged 2 mV cap. One final waveform therefore supplied both voltage and nominal functional scores. Parent hyperparameters had already been selected on nine settings: inner coefficient-heldout features were partial OOF development criteria, not fully nested independent validation.

$$Q_{target} = Q_{anchor} + \gamma \; clip\left(100 \; z^{T}\beta, \; -50, \; 50\right) \quad \text{(S11)}$$

V22 calibration searched 18 configurations per family: no, linear or quadratic relative-bias ridge; ridge penalty 1/10; and compensation-sensitivity profile power 0/0.5/1. The relative bias mu(aR,aP) was clipped to +/-0.005. The profile was s=max(1-aP,0.05)^(-power). Candidate selection used six/three fitting/validation of the bias model inside the training pool, followed by leave-one-setting envelopes scored by mean proper interval score. This training

selection objective emphasized costly misses; the primary reported endpoint remained the median proper interval score at alpha=0.2. The selected bias model was refitted on nine training settings, then each of the three disjoint calibration settings supplied ecal=(Qref-Qhat)/max(|Qhat|,100). No outer calibration outcome was used to select the bias/profile configuration.

For each test setting and each of the four base models, the interval center was c=Qhat+mu*D and candidate radius was r=D*s*|ecal-mucal|/scal, with D=max(|Qhat|,100). The three calibration residuals per model generated 12 candidate centers/radii. Equation 12 selected the second lower and eleventh upper order statistics for alpha=0.2; alpha=0.1 used the first and twelfth. These empirical bias/profile rank intervals replaced V20's fixed signed envelope; the scoring formula, test windows and reference definition did not change. Model fitting, adaptively learned profiles, repeated-session development and nonexchangeable compensation conditions preclude asserting a nominal physical coverage guarantee. Neither observed 16/16 hits nor software replay establishes a traceable ionic-charge certificate.

$$L_\alpha = ord(k_L;\ c - r), \quad U_\alpha = ord(k_U;\ c + r) \quad \text{(S12)}$$

All families received identical new configuration counts, role boundaries and information, while historical stronger controls remained in the release gate. Affine/FIR refer to residual backbones with the same nonlinear current adapter and readout, not pure linear end-to-end systems. The unchanged 216-configuration biological adapter was actually refitted with whole-cell LOCO, retaining all 450 sweeps, 24 conditions, boundary failures and missing-metadata flags. It is a separate biological application, not independent validation of the electronic-model-cell current-tangent learner. V21/V22 were two bounded retrospective development trials on the repeatedly inspected session. No candidate was automatically activated, and no cross-version metric stitching was permitted. Portable interval replay propagated the actual 12-significant-digit CSV charge quantization through candidate

bounds/order statistics, enforced a 0.001 fC hard budget and required unchanged hit classifications.

**Signed residual updates and consensus-centered calibration**

V23/V24 used the same four outer folds, four nine-training/three-calibration/four-test base fits and identical-role V22 parents. V23 compared gains 0/0.1/0.25/0.5/1 on the inherited signed relative-bias estimate. Within the nine-setting training pool, six/three coefficient-heldout predictions supplied the reference median/P95, and a minimax normalized criterion selected the gain. V24 compared the same gains with uncertainty thresholds kappa=0/0.5/1 (15 configurations); training selection minimized P95 among candidates noninferior to the V22 inner median and P95. The uncertainty proxy used the ridge-bias sandwich covariance with residual variance divided by max(n-effective_df,1). This is a model-based shrinkage proxy, not a traceable error budget. Equation 13 defines the selected soft update; V23 corresponds to kappa=0. Each update was projected onto the one output waveform with the inherited 2 mV cap. All positive gains selected kappa=0 in V24, and its IonGate point waveforms exactly matched V23.

$$Q_{new} = Q_{22} + \gamma \; sign(\mu) \; max(|\mu| - \kappa \; s_{\mu}, \; 0) \; D \quad \text{(S13)}$$

V23 interval learning used conditional residual quantiles at 0.1 and 0.9, fitted by linear programming with pinball loss and an L1 slope penalty; the intercept was unpenalized. Feature families were linear [1,x,y], quadratic [1,x,y,x^2,xy,y^2] and hinge [1,x,y,max(aR-0.8,0),max(aP-0.8,0),|aR-aP|], where x=aR-0.5 and y=aP-0.5. Four penalties 0.01/0.1/1/10 gave 12 equal-budget configurations per family. Features were training-RMS normalized; residual predictions were clipped to +/-2 percentage points. Inner six/three fits and leave-one-setting calibration envelopes selected the minimum mean proper score. Outer calibration nonconformity was max(lower-e,e-upper,0), with e=100(Qref-Qhat)/max(|

Qhat|,100). Equation 14 generated 12 adjusted candidate bounds, ranked second/eleventh at alpha=0.2 or first/twelfth at alpha=0.1.

$$L_k = Q_k + D_k(l_k - a_{cal})/100, \quad U_k = Q_k + D_k(u_k + a_{cal})/100 \quad \text{(S14)}$$

V24 used conditional median pinball fitting with the same feature/penalty choices plus a zero-trend option (13 configurations). It selected minimum inner median proper score subject to mean score no greater than 1.05 times the zero-trend baseline. After freezing the point/trend learner, each disjoint outer calibration record supplied its signed residual minus its own predicted trend. Adding the query trend transported these 12 relative residuals around the mean point prediction (Equation 15). There was no inverse compensation-margin scale and no base-prediction spread term in the center. This consensus-centered empirical interval is not the original CV+ construction and is not asserted to have a finite-sample or physical coverage guarantee. The primary score formula and alpha=0.2, test windows and nominal functional reference remained unchanged; mean/P95 scores were secondary endpoints.

$$L = \bar{Q} + \bar{D}\, r_{(kL)}/100, \quad U = \bar{Q} + \bar{D}\, r_{(kU)}/100 \quad \text{(S15)}$$

The new configuration counts, role exclusions, first-channel inputs and calibration opportunities were identical for IonGate, affine and FIR backbones; all controls retained the shared nonlinear adaptation/readout and historical stronger comparators remained available. Parent hyperparameters were already selected within nine training settings, so coefficient-heldout inner criteria remain partial OOF, not fully nested independent validation. Outer calibration/test labels were excluded from fitting and selection, but the repeatedly inspected single hardware session makes both trials retrospective development. The unchanged biological adapter was actually refitted with the same 216 choices, whole-cell LOCO and all failures retained. Software replay propagated actual CSV quantization through candidate bounds/order statistics with a 0.001 fC hard limit and unchanged coverage labels. None of

these checks establishes independent ionic-charge accuracy or perpetual autonomous improvement.

**Derivative-free dynamic current balance and endpoint controls**

V27 identified Equation 16 on baseline-centered training signals using trapezoidal integrals, with Vk a low-pass filtered membrane voltage and Vest the command estimator state. The five coefficients [Cm,gm,Ck,tauout,Ce] had prior [44,10,100,0.1,44], scales [30,3,100,0.05,30] and bounds [5,1,10,0,0] to [300,50,500,1,300]. Time constants tauk=0.3/1/3/10 ms and tv0=0.5/1.32/3 ms, with ridge penalties 0.01/0.1/1, gave 36 physical candidates; four blending gains gave 144 configurations per family. Fitting sampled 25-110 ms every fourth native sample and weighted each training setting equally by dividing its residual block by 100 times the square root of its sample count. Bounded linear least squares used only the legal six-setting inner or nine-setting final training role. Query-time passive-state inference used only current/command inputs; nominal scoring coefficients were not redefined by the learned physical coefficients. Source files were hash-frozen before fitting.

$$A_I = C_m \Delta V_m + g_m A_V + C_k \Delta V_k - \tau_{out} \Delta I_{out} - C_e \Delta V_{est} \quad \text{(S16)}$$

V25/V26 inherited identical-role V23 models, retaining four outer folds, four base fits per fold and nine fitting/three calibration/four test settings per base fit. V25 used the nominal passive load Cm=44 pF, gm=10 pA/mV, Ck=100 pF and gk=100 pA/mV, together with its own training-derived output/compensation parameters. Writing the passive two-state system as dx/dt=A x+B[u+tauout du/dt+Ce dVest/dt], the state transformation in Equation 17 eliminated numerical current differentiation. Signals were baseline-centered using 20-40 ms and initialized at 20 ms; the same nominal RC functional used for scoring was applied to the reconstructed perturbation. Sixty ridge/basis/gain candidates learned a capped innovation from this current-balance functional, with at most 1% functional correction and the

unchanged 2 mV waveform projection cap. All three families received the same new opportunities.

$$z = x - \tau_{out} Bu; \quad dz/dt = Az + (B + \tau_{out} AB)u + BC_e dVest/dt \quad \text{(S17)}$$

V26 tested the steady-endpoint proxy in Equation 18, using two baseline windows (20-40 or 40-45 ms), endpoint averages of 1 or 5 ms, two calibrated/hybrid feature bases, four ridge penalties and gains 0/0.5/1 (96 configurations per family). It approximated endpoint voltage changes by current changes divided by nominal gm and replaced Vest endpoint changes by command changes. This approximation was not assumed to hold after the independent endpoint check: the scoring window 45-80 ms spans a pulse and an incompletely settled return transient. V25/V26 remained unchanged. One additional V27 trial was explicitly recorded as a post-initial-plan amendment, not preregistration or fresh blind validation.

$$Q_{mass} := A_I + (\tau_{out} - \frac{C_m}{g_m})\Delta I_{out} + C_e \Delta V_{cmd} \quad \text{(S18)}$$

Functional selection minimized the worse of training-inner median discrepancy relative to V22 and P95 discrepancy relative to V23. The same 12 conditional-quantile interval options and primary alpha=0.2 score were retained. Parent hyperparameters were already selected within nine training settings, so the inherited inner criteria remain partial OOF development criteria, not fully nested independent validation. Outer calibration/test labels were excluded from fitting and selection. The endpoint diagnostic used independent voltage only after prediction freezing and did not revise V25/V26 outcomes. Every current control received the new physics and budget, including the stronger FIR25 tail. All trials reused the inspected one-session dataset; no nominal physical coverage or independent ionic-charge certification is claimed. The biological adapter and 216-choice whole-cell LOCO analysis were unchanged and actually rerun with all 450 sweeps and failed fits retained.

**Bounded residual self-adaptation and conditional signed calibration**

V28 inherited each same-role V27 backbone. A radial-basis kernel on the two compensation coefficients used length scales 0.25, 0.6 or 1.5, ridge penalties 0.1, 1 or 10, and gains 0.25, 0.5 or 1, plus a zero-gain option (28 candidates). The target was signed relative nominal-functional discrepancy; output was capped at 0.5 percentage points before the gain and projected onto the inherited waveform with the unchanged 2 mV cap. Median/P95 minimax selection used coefficient-held-out six-setting/three-setting inner predictions relative to the own V27 inner reference. Selected coefficients were refitted on nine training settings. These remain partial OOF criteria because inherited hyperparameters were selected within that nine-setting pool. No outer calibration/test labels entered fitting or selection. One waveform supplied both voltage and functional endpoints.

The interval layer used a conditional median trend with zero/linear/quadratic/hinge bases, slope L1 penalties 0.1/1/10 and compensation scale max(1-alphaP,0.05)^(-p), p=0/0.5/1 (36 candidates, including duplicate zero-trend penalties). Training-inner leave-one-setting signed-rank intervals were selected by mean proper interval score. Final bounds transported each legal calibration signed residual after subtracting its fitted trend and dividing by its scale, then added the query trend and scale around the mean of four whole base predictions. Twelve candidates used ranks 2/11 at alpha=0.2 and 1/12 at alpha=0.1. The primary median score, scoring window, nominal RC parameters and query-only inputs were unchanged. These are empirical signed-rank intervals, not a CV+ or physical metrology coverage guarantee. The unchanged biological implementation was reused with byte-identical V27 result and held-cell receipts, rather than described as new fitting. All earlier strong baselines, including the FIR25 functional tail, remain in the full benchmark. Repeated use of this inspected hardware session makes V28 retrospective development, not fresh external validation.

**Synchronized active-coefficient refitting and signed-envelope calibration**

For each outer fold, twelve non-test settings formed the fitting/calibration pool. For each setting g, the algorithm selected the unique earlier V27 crossfit whose three-setting calibration set contained g; that model's nine-setting hyperparameter source excluded g and all four outer tests. All active coefficients were then refitted on the other eleven settings: the family-specific V13 waveform/readout, protected V16 waveform/readout, V22 functional and bias layers, and bounded V27 dynamic-current-balance parameters. The own-family selected architecture, gains, ridge penalties and physical time constants were inherited by the same deterministic rule, without a new outer-test hyperparameter search. Inner three-way coefficient fits used seven/eight fitting and four/three validation settings; inherited hyperparameters mean these inner development residuals are not fully nested validation. Current-only source priors remained the legal earlier priors. No calibration-setting label entered its own active or hyperparameter fits.

Each eleven-setting model supplied both a first-channel-only calibration prediction at g and the four outer-query predictions. Its signed relative nominal-functional residual was transported to its own query center, yielding twelve candidates. The primary bounds were their full minimum/maximum, fixed in advance at both alpha=0.2 and alpha=0.1; alpha changed only the proper-score penalty. The mean of twelve complete waveforms supplied both voltage and functional point endpoints. These synchronized signed envelopes are empirical intervals, not a claimed CV+ theorem or guaranteed physical-charge coverage. Retrospective reuse of the inspected hardware session, one-session sampling, nominal RC scoring parameters, full failed-setting retention and no automatic activation remain unchanged. An initial technical attempt stopped before any prediction freeze because an inherited point wrapper unnecessarily evaluated a removed interval head; the recorded

preflight fix used the mathematically identical point-only entry, preserved the failed source/protocol and did not change roles, targets or model mathematics.

**Protected convex readout and residual-kernel calibration**

The own-family, same-role V22, V25 and V27 functionals were actual input-derived predictions, not historical performance summaries. A convex readout learned non-negative expert weights summing to one, an optional offset bounded by 25 fC, and a shrinkage gain. Linear programming minimized a weighted mean/max absolute training discrepancy with a penalty on departing from V27. Tail weights 0.25/0.75, trust penalties 0.01/0.1/1, offset off/on and gains 0.25/0.5/1, plus a zero-gain option, gave 37 configurations per family. Coefficient-held-out six-setting/three-setting inner predictions selected a median/P95-noninferior minimax option relative to the own V27 inner reference, followed by refitting on nine settings. Parent hyperparameters had already been selected in that nine-setting pool, so this remains partial OOF development. Outer calibration/test labels were excluded from selection. Equation 19 defines the functional update, further projected onto the frozen V27 waveform with both 50 fC and 0.15 mV limits; the same output waveform supplied voltage and functional endpoints. No backbone or physical parameter was refitted.

$$Q_{30} = Q_{27} + \gamma\ clip\left(w_{22}Q_{22} + w_{25}Q_{25} + w_{27}Q_{27} + b - Q_{27},\ -50,\ 50\right) \quad \text{(S19)}$$

V30 fitted a regularized residual-kernel mean and geometry scale using only the nine training settings and two known compensation coefficients. Length scales 0.25/0.6/1.5, ridge penalties 0.1/1/10 and geometry powers 0/1 gave 18 options per family. The mean was clipped at one percentage point and the posterior geometry variance floored at 0.05; neither is an independently calibrated physical uncertainty. Training-only inner/leave-one-setting proper-score selection was followed by fitting on nine settings. Each base model's three held-out calibration errors were centered and scaled by its own residual model, then transferred to that

same base query prediction. Twelve adjusted candidates used ranks 2/11 for alpha=0.2 and 1/12 for alpha=0.1. V31 instead retained its own frozen V27 quantile head and recomputed nonconformity from actual new readout predictions. Both alternatives were recorded before their outer scoring. V32 added a separately documented, later role-disjoint calibration-adaptation trial with twelve gain/kernel settings selected only on the nine fitting settings. All three families received identical new opportunities. Complete historical controls and all failed candidates were retained, and the biological module stayed byte-identical to V27 (all 450 sweeps, 24 conditions, four failures). These are retrospective empirical nominal-RC-functional intervals, not guarantees for independently metered ionic charge.

## S3 Biological identifiability criteria

Command support[-60,+20]mV;peak window0.2–4ms;at least6points;half-activation bracketed by2points on each side;condition number<=1e8;SE(Vhalf)<=10mV;original boundary rules;current ratio0.2–3. All failures retained. Exact source is soft_biology_v16.py, inherited unchanged by V27 and V30. These are model-fit acceptance criteria, not measured physiological truth.

## S4 Complete historical results and rejected candidates

### Feedback changes effective gain even when the instrument is unchanged

A simple access-resistance circuit already breaks a load-independent gain assumption. Let ionic current satisfy Iion = gVm and membrane voltage satisfy Vm = Vcmd − RsIion. At steady state, without additional output filtering, substitution gives

$$I_{out} = \frac{g V_{cmd}}{1 + R_s g} \quad \text{(S1)}$$

The effective gain depends on g as well as the instrument parameter Rs. With Rs = 30 MΩ and conductances of 10 and 100 pA/mV, an ideal current ratio of 10 becomes an observed

ratio of 3.25, a relative deviation of −67.5% (Figure S1; S2 Fig). This is an analytical counterexample, not a new measurement. It identifies the prerequisite that a common-filter cancellation argument would need to establish: equivalence of the effective transfer relation after accounting for the load and feedback.

The implication is constructive. IonGate uses circuit state and compensation metadata explicitly instead of assuming that identical hardware alone makes two recordings comparable. Supplementary Table S1 separates the independent voltage experiment from the known-drive, model-reference and external assay analyses. Those reference roles determine the claims evaluated below.

**First-amplifier inference reconstructs independently recorded membrane voltage**

IonGate reconstructed membrane voltage in held-out compensation settings using only first-amplifier current, command voltage, timestamps and compensation metadata. We evaluated 16 settings from one electronic-model-cell session in four grouped partitions, each with eight training, four calibration and four test settings. Predictions were frozen before the corresponding test-group second-amplifier voltage was read for scoring. Figure S2 pairs all 16 predicted waveforms with this independent sensor; S3, S7 and S14 Figs provide complementary diagnostics.

Across the complete 45–80 ms window, IonGate's median setting-level root-mean-square error (RMSE) was 1.802 mV and its across-setting 95th percentile (P95) was 5.396 mV. Command voltage and fixed 30 MΩ correction gave median RMSEs of 32.725 and 47.004 mV, respectively. IonGate improved all 16 settings against both proxies and against the 35 MΩ correction (Supplementary Table S2). These comparisons established a direct reconstruction benefit over treating the stimulus or a static voltage drop as membrane voltage.

The comparison with the fixed published Lei forward model revealed a temporal distinction. Plateau median RMSE decreased from 1.315 mV for Lei to 0.659 mV for IonGate, a 49.9% reduction calculated from the two medians. IonGate improved all 16 plateau settings. The full-window median remained lower for Lei, at 1.617 mV, although IonGate won nine settings and had lower median MAE, median absolute bias and P95 setting RMSE. Onset and offset RMSE favored IonGate in six and one settings, respectively. Thus, the strongest independently measured voltage advantage was sustained plateau reconstruction, not uniformly improved transition tracking.

The phase summaries, residual maps and native close-ups connect this result to the recorded signals. The onset, plateau and offset close-ups retain every displayed sample from the fixed original setting 00, selected by identifier rather than prediction error (Figure S2l–n). They illustrate temporal error structure without replacing the all-setting comparison. Phase scoring was a post-freeze explanatory analysis; its windows were not prospectively selected primary endpoints. These results established first-amplifier reconstruction and located its observed advantage within the voltage trajectory. We next held those voltage predictions fixed to isolate the downstream contribution of branch memory.

**Dynamic branch memory improves the charge readout on identical voltage predictions**

Retaining dynamic memory improved IonGate's finite-window charge readout without changing its voltage predictions. We fixed the observer and compared three readouts on identical frozen waveforms. Integrating the model-cell branch equation dIk/dt = gk dVm/dt − (gk/Ck)Ik gives [5]

$$Q_{RC} = g_m \int_{t_a}^{t_b} (V_m - V_{hold})\,dt + C_k \Delta V_m - \tau_k \Delta I_k, \quad \tau_k = \frac{C_k}{g_k} \quad \text{(S2)}$$

We denote the integrated leak-like, endpoint-capacitor and state-memory contributions by L, C and S, respectively, so that QRC = L + C + S. Here, $S = -\tau k[Ik(tb) - Ik(ta)]$ is a charge contribution, not the branch-current state itself. The full readout retained all three terms; no-memory retained L + C, and leak-only retained L. All used the same 40–49.9 ms hold mean, record-propagated state and 45–80 ms integration window. The reference applied the full nominal circuit functional to independently recorded second-amplifier voltage (Figure S3; S4 and S16 Figs).

Full RC reduced median normalized discrepancy from 0.456% for no-memory to 0.213%, and P95 from 1.423% to 0.522% (Supplementary Table S3). These were reductions of 53.3% and 63.3%, calculated from the corresponding summary statistics rather than from medians of paired percentage reductions. Leak-only gave a median of 0.551% and P95 of 4.303%. Full RC improved 13/16 settings against each ablation, while retaining the three reversals in each comparison. This was an advantage of IonGate's readout mechanism on fixed predictions, rather than an advantage obtained by replacing the observer. All three readouts were scored against the same full circuit functional, so the comparison tested approximations to one declared endpoint.

The signed decomposition and cumulative component curves connected this improvement to retained circuit state. Finite-window charge depended on the voltage trajectory and the branch states at both endpoints. Capacitor and memory contributions could oppose one another, so dropping S altered the magnitude and sign of the net endpoint correction. Memory therefore contributed a physically specified correction, not an additional fit to the test reference. The native L, C and S trajectories for fixed setting 00 expose these contributions over the complete window (Figure S3k–m), alongside all-setting errors and component comparisons. Numerical decomposition closed within $1.6 \times 10{-}14$ pC, confirming the implemented identity rather than a hardware measurement precision. These paired results

isolate the benefit of retaining branch memory for a nominal circuit-charge functional supported by independently recorded voltage.

**Conditional intervals expose the accuracy–width trade-off**

IonGate's conditional intervals improved the descriptive balance between reference containment and interval width. Under the original window-initial endpoint convention, IonGate contained 14/16 second-sensor-derived circuit references, compared with 10/16 for the fixed Lei comparator (Figure S4; S5 and S17 Figs). The median 80% interval score was 1.131% for IonGate and 1.316% for Lei, with lower values indicating a smaller width-and-miss penalty. The corresponding median full widths were 1.084% and 0.911% (Supplementary Table S3). The score therefore made the cost of IonGate's additional containment explicit instead of ranking intervals by hit count alone.

Point estimates gave a different comparison on this endpoint. Median normalized discrepancy was 0.260% for IonGate and 0.162% for Lei; IonGate won 6/16 settings. Baseline-subtracted direct Iout integration gave 1.602%, while command-voltage, 30 MΩ and 35 MΩ circuit proxies gave 5.716%, 22.611% and 27.332%, respectively. IonGate improved the aggregate point discrepancy over these simpler proxies, whereas Lei retained the better aggregate point estimate. Supplementary Table S3 separates this comparator stratum from the matched readout ablation. In particular, the 0.260% original-convention discrepancy and the 0.213% matched-ablation discrepancy are different endpoints, not competing estimates of one error.

Both methods calibrated interval radii using their own four calibration settings in the same fold assignments. Figure S4p–r displays all calibration and held-out residuals, unchanged fold radii and all 16 signed containment margins. These panels linked each hit and miss to its calibration radius and held-out residual. They distinguished larger intervals from smaller residuals as explanations for containment, rather than attributing the count alone to improved

point prediction. The original endpoint used window-initial voltage and zero branch state, whereas Figure S3 used a preceding hold and record-propagated state. The reference and prediction followed the same convention within each analysis.

With four calibration settings, the order-statistic rule supplied a finite radius at nominal 80% but not 95%. We reported observed containment of the circuit reference, with the calibration assumptions specified in Methods [18].

The width policy further exposed the reporting trade-off. At the original 25% budget neither method abstained. At a 1% budget, IonGate accepted seven settings and contained five references; Lei accepted 11 and contained six. Figure S4 retained complete width, miss and common-accepted-subset comparisons rather than discarding difficult settings. The result supported conditional interval utility beyond point accuracy, with containment gains, width cost and accepted population reported together. Figs. 2–4 thus evaluated voltage reconstruction, state-dependent functional estimation and the final reporting decision as separate outputs of IonGate.

**Historical equal-budget comparison separates waveform and functional utility**

The original fixed-parameter comparison was supplemented with a retrospective matched-budget analysis, rather than silently replacing its comparator. For each of four outer folds, both methods used eight training, four calibration and four test compensation settings. Both fitted the same five amplifier parameters with identical bounds, prior, robust training-voltage objective, three starts and a limit of 180 optimizer evaluations per start. Actual evaluation counts and run times were retained. The source-equivalent forward implementation reproduced all 16 frozen Myokit traces within 0.002 mV. Test predictions were saved and hashed before loading their corresponding second-sensor voltage for scoring. The historical

recordings had already been inspected, so this extension is grouped retrospective evaluation, not a new prospective blind experiment (Figure S8; Supplementary Table S5).

The tuned Lei model achieved a median full-window voltage RMSE of 0.615 mV, compared with 1.802 mV for IonGate, and was better in all 16 settings. In contrast, IonGate's median nominal circuit-functional discrepancy was 0.260% versus 0.479% for tuned Lei, with P95 values of 0.723% and 1.176%; IonGate had lower functional discrepancy in 9/16 settings. The median paired difference was 0.032 percentage points in IonGate's favor. Thus, the lower ratio of summary medians does not imply uniform or statistically established superiority across independent experiments.

Separate four-group calibration yielded 14/16 reference hits for IonGate and 12/16 for tuned Lei. Median full widths were 1.084% and 1.703%, and median interval scores were 1.131% and 2.013%, respectively. These results locate an observed IonGate advantage in finite-window circuit-function estimation and descriptive interval utility, not in universally superior voltage reconstruction. The comparator is a tuned five-parameter source model, not an exhaustive search over all Lei-model parameters.

**Window sensitivity preserves the median benefit of retaining circuit state**

The readout advantage persisted across six declared windows without selecting a window for the best result. Holding the original baseline, propagated state and trapezoidal integration fixed, full-RC median discrepancy ranged from 0.140% to 0.476% and stayed below both ablations in all six windows (S18 Fig). The full-versus-no-memory setting win counts ranged from 11 to 16, demonstrating that a favorable median need not imply improvement in every condition.

The complete sensitivity grid contains 6 windows × 3 hold choices × 2 quadrature rules × 4 state cases × 16 settings × 3 readouts, or 6,912 rows. Reset and ±100 pA state perturbations

test an assumed initial-state challenge; they are not measured metrological uncertainty. We retained every case and reproduced the original endpoint before expanding the analysis. This post-freeze analysis supports the readout mechanism across window choices while making the dependence on baseline and initial state visible.

**Acquisition mismatch can defeat unchanged interval calibration**

Sampling and input bandwidth produced distinct responses in voltage reconstruction and the integrated circuit functional. We processed only the first-amplifier inputs with declared low-pass and anti-aliased sampling challenges; the scoring voltage remained the native independently recorded signal. No observer parameter, interval radius or alignment was retuned. The 4 sampling rates × 4 filter conditions × 16 settings × 4 methods produce 1,024 score rows (S19 Fig).

Without an additional low-pass filter, median voltage RMSE rose from 1.802 mV at 10 kHz to 2.737, 5.782 and 10.508 mV at 5, 2 and 1 kHz. Containment under the unchanged original intervals was 14/16, 14/16, 13/16 and 2/16, respectively. The width policy still abstained in zero settings. This failure shows that an apparently narrow interval can survive a change that invalidates its calibration. Added filtering produced different voltage and charge trends, reinforcing the need to score both endpoints. Acquisition compatibility must therefore be checked independently of interval width; the present analysis identifies this requirement rather than claiming an already validated automatic detector.

**Acquisition compatibility is a separate decision from interval width**

A target-free compatibility guard compared test sampling with the training acquisition contract, required verified filtering provenance and checked a training-only current-noise floor. For the retained default IonGate, it accepted 16/16 native settings, preserving 14/16 reference hits, and rejected all 560 non-native declared configurations in a 6-by-6

sampling/filtering grid. These configurations were software transformations of real hardware recordings, not 560 independent acquisitions. A separate deliberately inaccurate-metadata challenge rejected only 9/80 native-rate filtered cases, leaving 11 accepted interval misses. The guard therefore closes the previously identified failure for declared acquisition changes but is not a universal blind filter detector (Figure S9; Supplementary Table S6).

**Controlled learning is reversible and endpoint-specific**

IonGate-Grow added bounded ridge residual correction [17] without altering the circuit equations. Each outer training pool was split into six fitting and two inner-validation settings. Three fixed penalties were compared without reading outer calibration or test labels; a candidate required at least 2% lower inner-validation median voltage RMSE and no more than 5% P95 deterioration. One of four candidates passed; three retained the unmodified observer. Model snapshots, validation decisions and frozen predictions were archived immutably.

On outer test settings, the selected voltage-targeted candidate lowered median voltage RMSE from 1.802 to 1.726 mV. Its median circuit-functional discrepancy increased from 0.260% to 0.274%, reference hits decreased from 14/16 to 12/16 and median interval score increased from 1.131% to 1.345%. Accordingly, the original IonGate remains the default; the update is a research candidate, not an overall performance upgrade. An implemented future-session update interface additionally requires noninferior circuit-functional validation, protects test-session identities and disables interval release until separate recalibration. This interface has software tests but no longitudinal new-session performance evidence.

**Frozen biological application supports transfer as a model-reference comparison**

IonGate applied frozen model-cell τout and Voffset to 450 NaV1.5 voltage sweeps spanning four cells, two temperatures and 24 groups. It adapted resistance and capacitance using public

acquisition metadata without refitting to biological outcomes. A compatible frozen published Gray + VC_MIN fit was available for one cell at 35°C. That author model used the 0% and 80% compensation conditions for its original fitting, leaving 20%, 40% and 60% for the present 45-sweep comparison (Figure S6; S9 and S13 Figs).

The median of the three condition-level median voltage RMSEs was 20.793 mV for IonGate, 22.337 mV for command voltage and 23.343 mV for the static residual-resistance proxy (Supplementary Table S4). IonGate improved the condition-group summary by 6.9% and 10.9%, respectively, winning two of three conditions against command voltage and all three against static resistance. Four original frozen-fold parameter variants produced aggregate errors of 20.757–20.822 mV. The 60% comparison with command voltage and all sweep-level reversals remain in the supplied tables.

The result demonstrates metadata-adapted biological application without refitting to biological outcomes. IonGate generated outputs for 450 sweeps, with 45 sweeps from one cell scored against a compatible frozen author-model voltage reference. This extends the workflow beyond the electronic model cell while retaining the reference type and analysis unit specified in Supplementary Table S1.

**A biological functional endpoint reveals both compensation sensitivity and identification limits**

The frozen observer was applied to all 450 acquired NaV sweeps from four cells, representing five cell-temperature cohorts and 24 compensation records. Apparent activation curves used the same observed peak current for every method; command voltage, static-resistance voltage and IonGate voltage supplied alternative driving potentials. The nominal sodium reversal potential came from the source-model concentrations. This endpoint

measures compensation sensitivity of an output-current-derived apparent conductance, not independently measured channel activation.

Across four methods and two fixed voltage-support windows, 192 curve fits were attempted and 18 lacked sufficient eligible points. Additional identification auditing marked only 103/192 fits usable; optimization boundaries, absent half-activation bracketing and weak local identification restricted interpretation. For example, cell3 at 35 °C had an apparent Vhalf range of 12.976 mV with IonGate versus 21.869 mV with static correction. It was the only completely identified paired cohort for that comparison in the audit; it is one cell-condition example, not evidence of a general cross-cell advantage. All missing, boundary and unfavorable fits remain visible (Figure S10; Supplementary Table S7).

**External hardware and image data clarify what each reference contributes**

Sato's 240 known-drive hardware records across 48 kinetic–noise strata test area transfer and label specificity. We compared recorded current area with an occupancy-based reference calibrated from other low-rate records and evaluated correct versus inverted labels under identical calibration (Figure S5; S8 Fig). These programmed-state controls characterize the hardware reference design separately from the IonGate prediction benchmark.

The biological context panels retain 49 paired SCN5A peak-current-density summaries, 30 paired recovery summaries and 48 hERG author-model protocol comparisons. Tissue images and a paired neuronal optical–electrical record show complementary measurement scales (Figs. 6 and 7; S6 and S10–S12 Figs). Figure S7 displays actual frames synchronized with an independently recorded electrode voltage, which anchors the distinction between seeing a signal and measuring its electrical counterpart. These analyses document reference types and assay variability rather than attributing source-study biological effects to IonGate. Their role

is to guide validation design for applications where recorded currents support functional interpretation.

**Joint innovations improve waveform reconstruction while preserving a circuit-functional anchor**

The joint observer adds measured-current innovations, stable exponential memory and a smooth functional constraint to the published forward circuit prior. Training used only first-channel inputs and the permitted training references; the held-out second voltage channel entered scoring after prediction freezing. For the original 45–80 ms window-initial nominal functional, the retained anchored candidate had median voltage RMSE 0.464 mV, functional discrepancy 0.256% and interval score 1.131%, compared with 0.615 mV, 0.479% and 2.013% for tuned Lei (Supplementary Table S8; Figure S11a–c). The point discrepancy is a functional of the independently recorded voltage under nominal circuit parameters, not independent ionic-charge metrology.

The voltage improvement was not specific to the memory feature family. Matched anchored FIR and affine controls achieved median RMSE 0.371 and 0.304 mV, respectively; their functional and interval medians were numerically tied to the common anchor. The IonGate memory observer instead had the lower voltage P95: 2.968 mV versus 3.759 and 3.566 mV. This comparison separates the observed median gain of calibrated residual models from the tail behavior of memory-based innovations. The earlier matched readout ablation remains the direct evidence for the benefit of retaining branch-state memory on identical predictions.

**Repeated internal validation exposes a non-dominating update**

Four rotations of the eight outer-training settings replaced the single internal split. Each rotation used four fitting, two calibration and two validation settings; each outer-training setting entered internal validation once. The numerically stabilized candidate achieved

median voltage RMSE 0.598 mV and P95 2.716 mV, but functional discrepancy rose to 0.381% and interval score to 1.341% (Supplementary Table S8; Figure S11d,e,h). Coverage counts remained 14/16. A separate normalized-residual interval experiment also worsened the retained candidate's score to 1.442% while maintaining 14/16 hits. These updates were retained as failed joint upgrades rather than relabelled improved models.

**Whole-cell transfer improves identifiability but not every functional comparison**

All 450 biological sweeps entered leave-one-entire-cell-out development; both temperatures of cell3 stayed in the same fold. Each dynamic family received 54 parameter candidates. The BioRC candidate yielded lower penalized compensation-consistency loss than fixed static resistance in all four cells and identified 17/24 activation conditions versus 14/24. Its loss was lower than leave-cell-out tuned static resistance in only two cells, and that comparator identified 19/24 conditions (Supplementary Table S9; Figure S11f,g).

A trust-regularized variant scaled the complete voltage correction rather than only the negative electrode-current contribution. This prevented the effective leak correction from changing sign merely because the trust factor was smaller than the compensation fraction. The new variant identified 19/24 conditions, including 4/5 rather than 2/5 in cell6, but surpassed tuned static resistance on penalized loss in only one of four cells. Robust top-configuration ensembles did not resolve the trade-off. All failed and boundary fits remain in the audit. These are functional-consistency and identifiability endpoints, not independently referenced biological activation accuracy.

**Earlier V10 development: Additional stimuli improve point estimation without establishing joint dominance**

The RC-State extension replaced empirical memory features with passive-circuit conservation. The electronic-model-cell adapter used nominal capacitance and branch

parameters to filter recorded current into voltage-state contributions; the output-filter derivative was evaluated as the derivative of the filtered state rather than by differentiating noisy current samples. All families retained the same published forward prior, first-channel inputs and compensation metadata. The strengthened affine comparator included instantaneous-current innovations and compensation interactions, and the FIR comparator included 101 native lags spanning 10 ms.

Training then added six predetermined real stimulus records per training compensation group: one step, one sine wave, one staircase, one action-potential train and two ramp sweeps. Each family received the same 500 post-20 ms samples per record, group/protocol weighting and internal candidate budget. All stimulus records from each outer-test compensation group remained excluded. The multistimulus RC-State candidate reached median voltage RMSE 0.340 mV, nominal functional discrepancy 0.159% and interval score 1.620%, with 15/16 reference hits. These are scores of one candidate; its lower functional discrepancy was not substituted into later versions (Supplementary Table S10).

**Orthogonalized physical residuals approach affine point accuracy but retain an interval trade-off**

A subsequent candidate first learned six common source/compensation nuisance features, then orthogonalized each family's residual features against that training-only nuisance space. The RC-State family included passive-state and instantaneous-innovation terms; the affine and FIR families received the same 24-candidate residual search and the same real training records. No inverse-functional anchor was imposed. This frozen orthogonal candidate achieved median voltage RMSE 0.338807 mV, functional discrepancy 0.178483% and interval score 1.154764%. The matched affine scores were 0.337156 mV, 0.179084% and 0.926462%; the matched FIR scores were 0.530359 mV, 0.236658% and 1.019535%. The voltage and functional medians were therefore effectively tied with affine at the predefined

1% practical-improvement threshold, rather than evidence of joint superiority (Figure S12b–e; Supplementary Table S10).

Voltage P95 was lower for RC-State: 2.297 mV versus 3.533 mV for affine and 4.115 mV for FIR. However, separate calibration produced only 11/16 reference hits for RC-State, compared with 13/16 and 15/16. A same-coefficient ablation, removing only the nine orthogonalized RC-state residual terms, increased voltage P95 from 2.297 to 3.665 mV and median functional discrepancy from 0.178% to 0.184%; median voltage RMSE slightly improved from 0.339 to 0.335 mV. Thus the state features changed tail behavior and functional estimation, but their contribution was not uniformly beneficial across metrics (Figure S12f).

**The biological adapter does not close the unified release criterion**

The biological adapter used electrode feedback and capacitive-current conservation without treating electronic-model-cell passive conductances as known biological parameters. Each dynamic family received the same 72 settings of output-filter time constant, resistance-correction trust, capacitance subtraction, derivative smoothing and optional metadata-based leak correction. Every cell was held out in its entirety, with both temperatures of cell3 kept together. Missing Rseal information was never imputed: additional leak correction was set to zero for those records in every method, and the missingness flag remained in the frozen predictions.

The initial causal derivative adapter identified 16/24 activation conditions. A matched first-order-hold discretization correction gave a unit-slope response to a linear-ramp software check, but the resulting biological adapter identified only 14/24 conditions. The earlier tuned-static and Trust-RC comparators each identified 19/24. The corrected RC-State adapter did not have lower penalized consistency loss than tuned static resistance in any of the four cells,

and its FIR counterpart was nearly numerically equivalent. Analytical discretization consistency therefore did not translate into better biological functional fits (Figure S12g,h; Supplementary Table S11).

The final release record combines the orthogonal multistimulus electronic head and the FOH biological adapter under one version, with domain-specific physical equations and frozen role assignments. It evaluates voltage, circuit function, interval score, reference hits and whole-cell consistency together. Both domain gates failed, so the candidate was not promoted and no active observer was replaced (Figure S12i).

**Protected circuit-state innovations improve the median waveform error**

The earlier V16 candidate fixed the complete nine-feature affine backbone and fitted RC innovations only after training-only orthogonalization against that backbone. Matched FIR received 100 lagged innovations, and the strengthened affine control received three compensation-dependent current-gain terms that remained affine in measured current. All families used the same 24 innovation and 18 readout candidates, extra-stimulus records, weights, caps and outer splits. Four nine-fit/three-calibration base models generated each outer-fold point estimate, with all four test settings excluded. V14 conditional-scale and V15 leave-one candidates were retained as unsuccessful development paths (Supplementary Table S12; Supplementary Table S12).

**Current matched comparisons do not supersede the stronger historical charge-tail benchmark**

V16 voltage RMSE median/P95 was 0.264573/3.454898 mV, compared with 0.292931/3.610832 mV for affine and 0.291699/3.610766 mV for FIR. Nominal functional discrepancy median/P95 was 0.179171%/0.465936%, versus 0.207329%/0.470942% and 0.201527%/0.462014%. Thus its functional tail was slightly below the current affine control,

but remained above the historical V13 affine value 0.447167%; the FIR functional P95 was also lower. Historical controls were retained, so expanding the comparator family could not by itself establish superiority (Supplementary Table S12; Supplementary Table S12).

**Lower median interval scores entail a retained coverage trade-off**

At unchanged primary alpha=0.2, median interval scores were 0.792837%, 0.821173% and 0.828084%. This improved on both current controls and the historical affine median 0.850080%. Mean scores, kept as secondary summaries, were 1.021804%, 0.987521% and 1.015334%. Reference hits were 15/16, 14/16 and 14/16. The candidate median width was 0.781583%. Its 15/16 hits did not retain V13's 16/16, and median functional discrepancy also worsened from V13's 0.131053%. These full-version trade-offs failed the complete retention gate; no active model was replaced (Supplementary Table S12).

**Same-coefficient intervention identifies a model-specific waveform contribution**

Removing only the nine RC innovation coefficients, while retaining the affine backbone and readout, changed voltage median/P95 from 0.264573/3.454898 mV to 0.289996/3.610595 mV and nominal functional median/P95 from 0.179171%/0.465936% to 0.196396%/0.445588%. Full-model voltage error was lower in 14/16 settings and functional error in 9/16. Four of sixteen base fits selected exactly zero innovations; the adaptive model could therefore retain the protected baseline when its training criterion did not favor extra states. Functional-tail behavior remained adverse. No recalibrated interval claim is made for the ablation (Supplementary Table S12).

**Whole-cell identifiability exceeds nineteen conditions without changing criteria**

The capacitive/leak-aware adapter was actually refitted in V14, V15 and V16, rather than assembling cell-wise best results. The same 216 configurations were available to the state, matched-static and FIR families. Whole-cell LOCO selection first protected training-cell

identification fractions, then compared normalized consistency losses. No activation-fit oracle was used at inference. The 450 original sweeps, 24 fits, boundary checks and requirement for two points on either side of half activation were unchanged. Twenty of twenty-four fits were identifiable in each current family, compared with seventeen for V13 and nineteen for the earlier tuned-static control (Supplementary Table S12; Figure S14; Supplementary Table S13).

**Biological recovery is strongest in cell6, but is not a unique RC-memory effect**

For cells3,6,9,18, V16 identified 9/10,5/5,4/5,2/4 conditions. Penalized losses were 19.992,3.671,43.263,56.513 mV, versus V13's 17.603,78.816,46.826,56.751 mV. Thus cell6 recovered all five compensation conditions and identification did not regress in any cell, but cell3's loss worsened. Enhanced static/FIR controls also reached20/24, so the biological recovery supports the shared acquisition-aware capacitance/leak correction rather than a unique RC-state advantage. Four failed/boundary conditions remain visible in the full24-panel figure. This endpoint measures apparent functional consistency and identifiability, not independently referenced biological voltage accuracy.

**Conditional residual learning improves functional median without losing matched point advantages**

The complete V22 model retained the V20 current-adapted waveform shape and learned a bounded signed functional correction from legal training data. Voltage RMSE median/P95 was 0.252199/2.565983 mV; nominal circuit-functional median/P95 discrepancy was 0.117432%/0.438437%. The functional median improved from V20 0.142470% and beyond the V13 target 0.131053%, while voltage median/P95 also remained numerically below V20. The functional P95 remained below the old strong affine 0.447167%, but was slightly above V20 0.437989%. Both voltage and functional metrics were computed from the same single

predicted waveform, without substituting historical per-endpoint scores (Figure S13a-c,m,n; Supplementary Table S12).

**Equal enhancement of affine and FIR controls isolates a current model advantage**

Matched affine voltage median/P95 was 0.285646/3.613352 mV and functional median/P95 0.144994%/0.449336%; matched FIR values were 0.285855/3.563499 mV and 0.195087%/0.455742%. Thus all four point summaries favored IonGate over both current controls, including a stronger newly enhanced affine functional median. Each family inherited its own same-role parent models and received 48 new functional and 18 bias/profile configurations. Affine and FIR name residual backbones rather than purely linear end-to-end systems: both controls include the shared nonlinear current adapter and bounded functional readout. These are retrospective summary advantages, not universal per-setting or independent-session superiority (Figure S13o; Supplementary Table S12).

**Circuit memory and current-driven readout make separable, nonuniform contributions**

With all other frozen coefficients retained, removing RC terms increased voltage median/P95 from 0.252199/2.565983 to 0.274075/3.739402 mV and functional median from 0.117432% to 0.135684%. Full-model voltage error was lower in 14/16 settings. However, RC removal improved functional P95 to 0.350477%. Separately removing current-driven waveform and conditional functional updates increased functional median/P95 to 0.131053%/0.460049%; the full model reduced functional error in 12/16 settings. This update still worsened voltage P95 relative to the disabled variant, 2.527397 mV. The interventions identify endpoint-specific contributions, not universally beneficial components; neither intervention recalibrated intervals (Figure S13g,h,k,l; Supplementary Table S12.3).

**Separate complete candidates reveal the remaining accuracy-coverage trade-off**

At the unchanged primary interval-score penalty alpha=0.2, V22 achieved 15/16 reference hits, median score 1.204046%, mean score 1.290131% and median width 1.198954%. V21 had restored 16/16 hits with median score 1.090428% and functional median 0.140360%; its coverage cannot be combined with V22 point scores. Neither candidate retained V18 median score 0.732700%. V22 passed five of eight frozen checks, failing exact V20 functional-tail retention, interval-score retention and 16/16 coverage. Actual whole-cell reruns of the unchanged biological adapter retained all 450 sweeps, 24 conditions and four failed fits, with 20/24 identifiable conditions in both candidates and enhanced controls. Neither candidate passed the joint release gate or replaced an active model (Figure S13d-f,i,j,p; Figure S14; Tables 12-13).

**A complete update combines tail recovery, full empirical coverage and matched point advantages**

V23 added a bounded training-selected correction from the V22 signed functional-bias learner, while inheriting its own same-role parent waveforms. It achieved voltage RMSE median/P95 0.252106/2.565980 mV and functional median/P95 0.127688%/0.428057%, with 16/16 reference hits. Functional P95 was below V20 0.437989%, and all four point summaries remained lower than current equally enhanced affine and FIR controls. Thus tail recovery and full empirical reference coverage occurred in the same complete candidate, not by substituting another version's bounds. This gain did not preserve V22's 0.117432% functional median; V23 median interval score was 1.398749%, not the V18 0.732700% target. Six of eight frozen checks passed (Figure S15a-c,e-g,l; Supplementary Table S12).

**Consensus-centered calibration improves mean and tail interval utility, but not the primary joint target**

V24 used the same frozen IonGate point waveforms as V23: all 16 serialized waveforms were identical, so changes in interval metrics cannot be attributed to a new point predictor. Signed calibration residuals were transported by a training-only conditional median trend and ranked around the ensemble point, rather than dispersed base-model centers. Median/mean/P95 interval scores were 1.023436%/1.010881%/1.445136%, versus affine 0.977817%/1.258009%/3.161869% and FIR 0.835032%/1.060520%/2.068190%. V24 improved both secondary score summaries over the two current controls, but not the primary median. Reference hits were 14/16, versus affine 11/16 and FIR 12/16; settings 01 and 05 missed. No positive uncertainty-threshold coefficient was selected in any of the 16 IonGate base models, so the new soft-threshold option has no demonstrated contribution to the point gains. Five of eight checks passed (Figure S15d,h,i,k,l; Supplementary Table S12).

**Whole-cell reruns retain all fits and the joint gate prevents cross-version score assembly**

The unchanged 216-configuration acquisition-aware adapter was actually refitted in both trials, preserving all 450 sweeps, 24 conditions and four failed activation fits. IonGate and current enhanced static/FIR controls again reached 20/24 identifiable conditions. The biological result is therefore retained identifiability, not evidence that the new electronic-model-cell interval learner improves biological voltage or ionic-charge accuracy. V22 retains the best functional median among these three candidates; V23 combines the stronger functional tail and 16/16 hits; V24 has better mean/tail interval scores. None simultaneously retains V22 functional median, V18 primary score and all other checks. No active or retained model was replaced (Figure S15j,l; Figure S14; Tables 12-13).

**Current-balance challenges separate tail improvement from joint superiority**

V25 reconstructed a passive-load voltage perturbation from first-channel current and command compensation without numerically differentiating current, then learned a bounded functional innovation. It retained 16/16 reference hits and reduced functional P95 from V23 0.428057% to 0.391394%. However, functional median increased to 0.145089% and interval-score median to 1.508475%; the equally enhanced FIR control had the stronger functional P95 0.353553%. The full four-summary advantage was therefore not retained. V26 tested an integral-charge shortcut with a steady-endpoint approximation; its functional median/P95 0.236164%/0.543826% and 15/16 hits failed key retentions. These outcomes remain in the benchmark, rather than being removed as inconvenient controls (Figure S16c-f; Supplementary Table S12).

**Return-transient diagnosis motivates a bounded dynamic-identification amendment**

A post-freeze diagnostic found that the command pulse returned near 70 ms while low-compensation membrane voltage was still relaxing at the 80 ms scoring endpoint. The absolute mismatch between delta Iout/nominal gm and independently recorded delta Vm had median/P95 7.751/41.040 mV across the 16 settings. This is a finite-transient endpoint mismatch, not evidence of an incorrectly calibrated DC current gain. V27 was one explicitly recorded amendment to the initial two-trial budget: five bounded load/instrument coefficients were identified from integrated current balance on legal training records, including capacitor memory and return transients. V27 retained 16/16 hits and all four point-summary advantages over its current matched controls; functional median was 0.119793%, closer to V22 0.117432%, but functional P95 0.442318% did not retain V23 and interval-score median 1.461672% did not retain V18. No candidate passed all eight checks. Each biological adapter was actually rerun and retained 20/24 fits from all 450 sweeps, as did enhanced controls.

Earlier V22/V23 results and all negative updates were preserved (Figure S16a,b,g,h; Tables 12-13).

**Bounded residual learning preserves point advantages but exposes a calibration trade-off**

V28 added a regularized compensation-conditioned residual learner to the whole V27 predictor. All 16 settings were retained, and each matched backbone received the same 28 point-head and 36 interval configurations. Voltage median/P95 were 0.252437/2.565788 mV. The four point summaries remained below current enhanced affine and FIR controls, but nominal functional median/P95 were 0.123099%/0.443292%, missing the V22/V25 targets of 0.117432%/0.391394%. Recentring signed calibration ranks on the final ensemble point reduced proper-score median by 25.98%, from 1.461672% to 1.081955%; reference hits fell from 16/16 to 13/16, so this reduction was not a joint calibration improvement. All three misses are displayed in Figure S17g. The enhanced FIR had a lower score median 0.791466% but only 12/16 hits (Supplementary Table S14).

Frozen-parameter ablation separated the two modifications. Disabling the residual head left 13/16 hits and score median 1.097313%, close to full V28; the modest learning head therefore did not explain the three coverage losses. Replacing the interval layer with the fixed V27 envelope, while recomputing calibration on the legal three settings per base fit, recovered 16/16 hits but score median 1.456769%, still far above V18 0.732700%. This post-freeze diagnostic was not selected as a new model or merged with the best historical point scores. The result identifies uncertainty transport, rather than a missing point-prediction correction alone, as the immediate bottleneck. V28 was not activated. The unchanged biological adapter retained the exact V27 450-sweep, 20/24-fit records, including all four failures; no additional biological improvement is claimed (Figure S17i-l).

**Whole-model leave-setting-out refitting changes the calibration trade-off**

V29 replaced the four nine-fit/three-calibration base fits with twelve eleven-fit/one-calibration active-coefficient refits per outer fold. Each calibration setting was excluded from both its active coefficient fits and its nine-setting hyperparameter prior; all four outer test settings remained excluded throughout. The same structure and number of fits were given to IonGate, enhanced affine and enhanced FIR. IonGate voltage P95 improved from V27 2.565818 to 2.467720 mV, and interval-score median/mean/P95 changed from 1.461672/1.538137/2.309190% to 1.037998/1.075982/1.638706%. However, empirical reference hits were 13/16, functional median/P95 were 0.146481%/0.466615%, and voltage median rose to 0.289890 mV, above both current matched controls. V29 therefore did not retain four-summary dominance or meet the joint V22/V25/V18 targets. The matched affine itself achieved 16/16 hits and score median 1.031743%; this stronger calibration result is retained rather than excluded (Figure S18; Supplementary Table S15).

Two fixed diagnostics separated active refitting from interval tightening. Keeping the old nine-setting coefficients and applying the signed full envelope gave 15/16 hits with score median 1.092070%. Keeping V29 point predictions but recentering ranks 2/11 on their ensemble mean reduced score median to 0.871646% while hits fell to 10/16; it was not selected as a new release. An explicit post-score negative control used test references only to calculate the minimum common absolute padding needed for 16/16 hits within the V29 interval shape. That oracle padding was 31.254014 fC and would yield score median 1.232855%, still above V18. It was never used for inference or model selection. This narrow result rules out uniform widening as a solution for this frozen interval family, not the feasibility of better conditional calibration generally. The unchanged biological adapter retains the byte-identical V27 450-sweep, 20/24-fit evidence and four failures; no new biological fitting or independent ionic-charge measurement is claimed.

**Waveform-protected learning improves median functional accuracy without losing matched point advantages**

Locking the V27 waveform and physical coefficients, V30 changed only a bounded functional readout and the interval layer. The functional median discrepancy fell from 0.119793% to 0.105459% (11.97% relative reduction), below the V22 target 0.117432%. Voltage median/P95 were 0.252515/2.565794 mV; all four point summaries remained below both equally enhanced matched controls. Empirical reference hits were 16/16, with unchanged 20/24 biological fits. Functional P95 was 0.443297%, however, exceeding both V25 0.391394% and V27 by 0.000979 percentage points. Primary interval-score median was 1.362887%, improved relative to V27 but above V18 0.732700%. Thus the same complete candidate combines the median target, full empirical hits and matched point advantages, not all requested endpoints (Figure S19; Supplementary Table S16).

V31 used exactly the same frozen point waveforms as V30, but recalibrated the inherited V27 quantile head on the new point predictions. Both achieved 16/16 hits; V31 score median was 1.483720%, compared with V30 1.362887%. The 8.14% reduction is attributable to the changed calibration layer under identical point predictions. Separately, the V30-vs-V27 functional-median improvement is attributable to the new readout because the backbone coefficients were untouched and query waveforms changed only through the bounded functional projection. Neither comparison establishes a general cross-session advantage or independent ionic-charge accuracy.

One disclosed post-initial-plan amendment tested V32: one of each base fit's three originally held-out settings adapted the readout, and the other two remained exclusively for interval calibration. This produced twelve members and twenty-four correlated interval candidates per fold, not twenty-four independent settings. V32 functional median/P95 were 0.114725%/0.448022%, score median 1.210591% and 15/16 hits. The original strict rule,

requiring no functional-median or P95 regression relative to V27 in addition to core protections, recommends V27. A separately labelled engineering-tolerance sensitivity recommends V30; its one-percent tail allowance is not a metrological uncertainty bound. The recorded final preference is balanced, retaining the entire V30 version. No outcome is assembled by mixing metric winners, and the complete joint goal remains false.

## S5 Selection and endpoint accounting

The strict zero-functional-regression rule retainsV27. The human-confirmed balanced research recommendation selectsV30, accepting the disclosed0.000979percentage-point Q-P95 increase. V30 has16/16hits;itsQ-P95 target0.391394% and interval-median target0.732700% remain unmet. V31 has identical point predictions but worse interval scores. The laterV32 trial has15/16hits and is not selected. No activation or autonomous learning is reported. FINAL_RECOMMENDATION.json and FINAL_SUMMARY_TABLE.csv preserve all complete-version comparisons.

## S6 Reproduction and supplementary inventory

Run reproduce/reproduce_ibex.py --verify-only to validate immutable parents;run without that flag to recalculate primary tables, figures and descriptive ablations without fitting. This is an IBEX dependency-aware analysis release, not a standalone raw-data redistribution. New relative-charge ablation is post-freeze and requires no optimized parameters. All16historical table families below and20supplementary figures are retained. The existing227test result is archived, not represented as newly executed by manuscript assembly.

**Supplementary Table S1. Datasets, reference types and analysis units**

Table1 Datasets

| Dataset | Extent / analysis unit | Reference | Role |
|---|---|---|---|
| Lei electronic model cell | 16 compensation settings; 1 acquisition session | Independent second-amplifier measured Vm | IonGate independent-Vm reconstruction and circuit-functional evaluation |
| Lei NaV1.5 HEK | 450 sweeps;4 cells;2 temperatures;24 groups | 45 scored sweeps /3CP groups of cell6: fixed author-model Vm | Frozen IonGate application and held-out model-reference comparison |
| Sato hardware | 240 records;48 rate/noise strata | Programmed0/1 state;cross-record output calibration | Physical measurement-chain area transfer and state specificity |
| SCN5A assays | 49 pCD pairs;30 recovery-feature pairs | Published paired summary features | Variant functional-assay reproducibility and platform variation |
| hERG WT/R56Q | 12 cell-condition records;4 protocols;48 comparisons | Observed currents;published source-model replay | Cross-channel electrophysiology and protocol-specific replay |
| Cardiac imaging | One displayed Sham slice and one TAC slice | Original fluorescence frames/ROI flux | Tissue-scale fluorescence dynamics and measurement context |
| Voltron imaging | One neuron/session;2273 frames | Simultaneous electrode Vm and optical ROI | Same-cell optical/electrical pairing and temporal response |

Each row has its own reference and analysis unit. Applied sweeps, scored sweeps, compensation settings, images and independent sessions are not interchangeable sample counts.

## Supplementary Table S2. Independent membrane-voltage benchmark

Table2 Voltage

| Window | Predictor | Median RMSE (mV) | P95 setting RMSE (mV) | Median MAE (mV) | Median abs bias (mV) | IonGate wins /16 |
|---|---|---|---|---|---|---|
| full | IonGate | 1.802 | 5.396 | 0.797 | 0.448 | — |
| full | Lei | 1.617 | 5.705 | 1.149 | 0.561 | 9 |
| full | Vcmd | 32.725 | 44.983 | 16.842 | 0.350 | 16 |
| full | Rs30 | 47.004 | 85.234 | 34.288 | 3.730 | 16 |
| full | Rs35 | 55.189 | 99.574 | 40.828 | 4.401 | 16 |
| onset | IonGate | 4.254 | 15.413 | 2.419 | 1.646 | — |
| onset | Lei | 3.357 | 19.456 | 3.115 | 3.115 | 6 |
| onset | Vcmd | 86.536 | 98.474 | 84.409 | 84.409 | 16 |
| onset | Rs30 | 93.741 | 222.248 | 84.757 | 65.771 | 16 |
| onset | Rs35 | 98.373 | 259.971 | 88.472 | 70.268 | 16 |
| plateau | IonGate | 0.659 | 0.939 | 0.610 | 0.536 | — |
| plateau | Lei | 1.315 | 2.194 | 0.803 | 0.595 | 16 |
| plateau | Vcmd | 10.638 | 29.026 | 6.640 | 6.640 | 16 |
| plateau | Rs30 | 29.554 | 46.738 | 21.635 | 21.635 | 16 |
| plateau | Rs35 | 36.013 | 54.434 | 25.706 | 25.706 | 16 |
| offset | IonGate | 4.795 | 14.036 | 3.053 | 1.683 | — |
| offset | Lei | 2.263 | 11.915 | 2.022 | 1.914 | 1 |
| offset | Vcmd | 90.739 | 101.352 | 87.027 | 87.027 | 16 |
| offset | Rs30 | 90.562 | 214.796 | 82.182 | 66.632 | 16 |
| offset | Rs35 | 94.957 | 252.446 | 85.630 | 61.670 | 16 |

Median/P95 aggregate setting-level RMSEs. Onset [50,52), plateau [52,69), offset [70,72) ms; full 45–80 ms inclusive. All 16 settings originate in one session. A blank win entry denotes the IonGate row itself.

## Supplementary Table S3. Circuit readout, conditional intervals and selection

Table3A RC

| Readout | Median error (%) | P95 error (%) | Full RC wins /16 |
|---|---|---|---|
| Full RC | 0.213 | 0.522 | — |
| No memory | 0.456 | 1.423 | 13 |
| Leak only | 0.551 | 4.303 | 13 |

Original window-initial point comparison (separate convention)

| Predictor | Median discrepancy (%) | P95 discrepancy (%) | IonGate wins |
|---|---|---|---|
| IonGate original full RC | 0.260 | 0.723 | — |
| Lei original full RC | 0.162 | 0.587 | 6/16 |
| Direct Iout integral | 1.602 | 2.850 | 14/16 |
| Vcmd original RC | 5.716 | 21.874 | 16/16 |
| Rs30 original RC | 22.611 | 28.708 | 16/16 |
| Rs35 original RC | 27.332 | 33.720 | 16/16 |

Table3B Intervals

| Predictor | Hits /16 | Median full width (%) | Median interval score (%) | Abstentions /16 |
|---|---|---|---|---|
| IonGate | 14 | 1.084 | 1.131 | 0 |
| Lei_frozen | 10 | 0.911 | 1.316 | 0 |

Table3C Acceptance

| method | budget_pct | n_accepted | n_conditional_hits | empirical_covered_among_accepted | median_error_pct |
|---|---|---|---|---|---|
| IonGate | 0.250 | 0 | 0 | — | — |

| method | budget_pct | n_accepted | n_conditional_hits | empirical_covered_among_accepted | median_error_pct |
|---|---|---|---|---|---|
| Lei_frozen | 0.250 | 4 | 0 | 0 | 0.382 |
| IonGate | 0.500 | 0 | 0 | — | — |
| Lei_frozen | 0.500 | 4 | 0 | 0 | 0.382 |
| IonGate | 1 | 7 | 5 | 0.714 | 0.246 |
| Lei_frozen | 1 | 11 | 6 | 0.545 | 0.165 |
| IonGate | 1.500 | 12 | 10 | 0.833 | 0.260 |
| Lei_frozen | 1.500 | 15 | 9 | 0.600 | 0.165 |
| IonGate | 2 | 16 | 14 | 0.875 | 0.260 |
| Lei_frozen | 2 | 16 | 10 | 0.625 | 0.162 |
| IonGate | 2.500 | 16 | 14 | 0.875 | 0.260 |
| Lei_frozen | 2.500 | 16 | 10 | 0.625 | 0.162 |
| IonGate | 25 | 16 | 14 | 0.875 | 0.260 |
| Lei_frozen | 25 | 16 | 10 | 0.625 | 0.162 |

Matched ablation uses the 40–49.9 ms hold mean and record-propagated state. Original point/interval results use Vm(45 ms) and Ik(45 ms)=0. Do not pool their errors. The displayed width summary is reference-normalized; budget acceptance uses prediction-normalized widths. Conditional reference containment is not independent physical-charge coverage.

## Supplementary Table S4. Frozen biological application and separate external endpoints

Table4A Frozen NaV

| CP (%) | Predictor | Sweeps | Median RMSE (mV) | P95 sweep RMSE (mV) |
|---|---|---|---|---|
| 20 | IonGate_frozen | 15 | 19.999 | 20.793 |
| 20 | Vcmd | 15 | 22.996 | 25.791 |

| CP (%) | Predictor | Sweeps | Median RMSE (mV) | P95 sweep RMSE (mV) |
|---|---|---|---|---|
| 20 | static_Rs | 15 | 23.290 | 27.048 |
| 40 | IonGate_frozen | 15 | 20.793 | 21.322 |
| 40 | Vcmd | 15 | 22.337 | 25.190 |
| 40 | static_Rs | 15 | 23.343 | 26.586 |
| 60 | IonGate_frozen | 15 | 21.609 | 26.385 |
| 60 | Vcmd | 15 | 21.242 | 24.634 |
| 60 | static_Rs | 15 | 23.821 | 26.719 |

Table4B Hardware

| Noise band | Nominal SNR | Records | Median fullscale error (%) | P95 fullscale error (%) | P95 programmed-relative error (%) |
|---|---|---|---|---|---|
| high | 10.200 | 80 | 3.280 | 4.052 | 675.989 |
| low | 1.820 | 80 | 1.292 | 3.633 | 880.745 |
| med | 5.140 | 80 | 2.901 | 4.002 | 459.885 |

Table4C SCN5A

| Published feature | Paired variants | Median VUMC-VCCRI | Median absolute site difference |
|---|---|---|---|
| pCD (% WT) | 49 | 0.950 | 9.310 |
| RFI (source scale) | 30 | 1.056 | 2.197 |

Table4D hERG AP

| genotype | condition | native_samples | recorded_current_area_nA_s |
|---|---|---|---|
| R56Q | Control | 88240 | 4.949 |

| genotype | condition | native_samples | recorded_current_area_nA_s |
|---|---|---|---|
| R56Q | RPR260243 | 88240 | 4.817 |
| WT | Control | 88240 | 5.785 |
| WT | RPR260243 | 88240 | 3.219 |

The NaV voltage reference is the frozen author model in one cell, not independently measured biological Vm. Hardware state, SCN5A summaries and hERG output areas are distinct endpoints, not pooled IonGate performance.

## Supplementary Table S5. Matched-budget retrospective benchmark

Table5.1

| Method | Vm RMSE median (mV) | Q discrepancy median (%) | Reference hits | Interval score median (%) |
|---|---|---|---|---|
| IonGate | 1.802 | 0.260 | 14/16 | 1.131 |
| Lei_forward_tuned | 0.615 | 0.479 | 12/16 | 2.013 |
| IonGate_Grow | 1.726 | 0.274 | 12/16 | 1.345 |

One electronic-model-cell session; 16 nested settings. Same five free parameters and optimizer caps; different structural input use. Q is a nominal circuit reference, not independent absolute charge.

## Supplementary Table S6. Compatibility and controlled-update accounting

Table6.1

| Evaluation | Accepted / detected | Total | Interpretation |
|---|---|---|---|
| Native settings | 16 accepted | 16 | No native false rejection in this test |
| Declared transformed configurations | 560 rejected | 560 | Metadata-aware; software transformations |
| Native-rate undeclared filter | 9 rejected | 80 | 11 accepted interval misses remain |

| Evaluation | Accepted / detected | Total | Interpretation |
|---|---|---|---|
| Inner-validation candidates | 1 retained | 4 | Three candidates rolled back |
| Default production model | Original IonGate | -- | Grow is research-only; interval score did not improve |

Counts of transformed conditions are not independent hardware sessions. Future multi-endpoint release is implemented and software-tested, not longitudinally validated.

## Supplementary Table S7. Biological functional fit completeness

Table7.1

| Cohort | Method | Fits / attempted | Active-bound fits | Vhalf range (mV) |
|---|---|---|---|---|
| cell18/35C | IonGate_frozen | 4/4 | 0 | 119.998 |
| cell18/35C | Vcmd | 4/4 | 1 | 64.636 |
| cell18/35C | static_Rs | 4/4 | 0 | 78.505 |
| cell3/25C | IonGate_frozen | 5/5 | 0 | 8.880 |
| cell3/25C | Vcmd | 5/5 | 0 | 33.580 |
| cell3/25C | static_Rs | 5/5 | 0 | 4.871 |
| cell3/35C | IonGate_frozen | 5/5 | 0 | 12.976 |
| cell3/35C | Vcmd | 5/5 | 1 | 18.954 |
| cell3/35C | static_Rs | 5/5 | 0 | 21.869 |
| cell6/35C | IonGate_frozen | 1/5 | 1 | NA |
| cell6/35C | Vcmd | 5/5 | 0 | 18.990 |
| cell6/35C | static_Rs | 5/5 | 4 | 31.931 |
| cell9/35C | IonGate_frozen | 5/5 | 4 | 2.703 |

| Cohort | Method | Fits / attempted | Active-bound fits | Vhalf range (mV) |
|---|---|---|---|---|
| cell9/35C | Vcmd | 5/5 | 1 | 40.937 |
| cell9/35C | static_Rs | 5/5 | 0 | 66.966 |

All predefined cohorts retained. Active-bound counts alone do not establish identifiability. Full audit:103/192 usable fits across four methods and two support windows;18 insufficient-point fits. Only one complete identifiable IonGate/static paired cohort. Ranges are not accuracy errors.

## Supplementary Table S8. Joint observer and repeated-inner development audit

Table8.1

| Method | Vm median mV | Vm P95 mV | Q median % | Score median % | Hits |
|---|---|---|---|---|---|
| IonGate_v7 | 1.801602 | 5.395687 | 0.259825 | 1.130579 | 14/16 |
| Lei_tuned | 0.614709 | 3.567413 | 0.479473 | 2.012682 | 12/16 |
| IonGate_MemoryInnovation_anchored | 0.463789 | 2.968216 | 0.256178 | 1.130579 | 14/16 |
| Generic_FIR_innovation_anchored | 0.370988 | 3.758790 | 0.256178 | 1.130579 | 14/16 |
| Lei_affine_anchored | 0.304415 | 3.566187 | 0.256178 | 1.130579 | 14/16 |
| IonGate_MemoryInnovation_stable | 0.597633 | 2.716449 | 0.380671 | 1.340981 | 14/16 |
| Generic_FIR_innovation_stable | 0.392254 | 3.758393 | 0.380671 | 1.340981 | 14/16 |
| Lei_affine_stable | 0.280249 | 3.566024 | 0.380671 | 1.340981 | 14/16 |

All scores use the same16compensation settings from one session and45–80ms window-initial nominal functional. Query-conditioned anchored estimates are offline. More flexible controls receive the same head/anchor search budget; they are not the five-parameter Lei model. No overall dominance or95percent physical-charge coverage is established.

## Supplementary Table S9. Whole-cell-held-out functional consistency and identification

Table9.1

| Heldout cell | Method | Penalized loss mV | Identifiable / attempted |
|---|---|---|---|
| cell3 | IonGate_v7 | 19.79564 | 9/10 |
| cell3 | static_Rs | 23.67817 | 9/10 |
| cell3 | LOCO_static_Rs | 17.52586 | 9/10 |
| cell3 | Matched_static_dynamic_control | 17.57946 | 9/10 |
| cell3 | IonGate_BioRC | 20.09700 | 9/10 |
| cell3 | IonGate_TrustRC | 17.57549 | 9/10 |
| cell6 | IonGate_v7 | 220.00000 | 0/5 |
| cell6 | static_Rs | 181.00089 | 1/5 |
| cell6 | LOCO_static_Rs | 48.35631 | 4/5 |
| cell6 | Matched_static_dynamic_control | 62.54666 | 3/5 |
| cell6 | IonGate_BioRC | 87.66511 | 2/5 |
| cell6 | IonGate_TrustRC | 45.22972 | 4/5 |
| cell9 | IonGate_v7 | 201.66123 | 0/5 |
| cell9 | static_Rs | 30.61448 | 4/5 |
| cell9 | LOCO_static_Rs | 40.32837 | 4/5 |
| cell9 | Matched_static_dynamic_control | 39.25316 | 4/5 |
| cell9 | IonGate_BioRC | 28.06303 | 4/5 |
| cell9 | IonGate_TrustRC | 42.24381 | 4/5 |

| Heldout cell | Method | Penalized loss mV | Identifiable / attempted |
|---|---|---|---|
| cell18 | IonGate_v7 | 176.90019 | 1/4 |
| cell18 | static_Rs | 200.98297 | 0/4 |
| cell18 | LOCO_static_Rs | 58.81528 | 2/4 |
| cell18 | Matched_static_dynamic_control | 59.71649 | 2/4 |
| cell18 | IonGate_BioRC | 55.23448 | 2/4 |
| cell18 | IonGate_TrustRC | 61.65430 | 2/4 |

One heldout cell per fold; both temperatures ofcell3remain in the samefold. Loss combines identifiable Vhalf range, failurefraction andfitresidual; lower is better, not a reference-accuracy error. Each dynamic family has54base candidates. All450sweeps and24conditionfits are retained per selected method, including failures.

## Supplementary Table S10. Unified-version development, including unsuccessful candidates

Table10.1

| Candidate / control | Vm median mV | Vm P95 mV | Q median % | Score median % | Hits |
|---|---|---|---|---|---|
| IonGate_RCState | 0.442912 | 2.565909 | 0.537172 | 3.134596 | 13/16 |
| Matched_affine | 0.294217 | 3.567189 | 0.268175 | 1.097151 | 14/16 |
| Matched_FIR | 0.399494 | 4.011766 | 0.328944 | 1.836958 | 14/16 |
| IonGate_RCState_multi | 0.340088 | 2.662927 | 0.158951 | 1.619652 | 15/16 |
| Matched_affine_multi | 0.303641 | 3.533578 | 0.222058 | 1.354927 | 14/16 |
| Matched_FIR_multi | 0.524920 | 4.887157 | 0.206882 | 1.060885 | 15/16 |
| IonGate_RCState_orth | 0.338807 | 2.297167 | 0.178483 | 1.154764 | 11/16 |
| Matched_affine_orth | 0.337156 | 3.533258 | 0.179084 | 0.926462 | 13/16 |

| Candidate / control | Vm median mV | Vm P95 mV | Q median % | Score median % | Hits |
|---|---|---|---|---|---|
| Matched_FIR_orth | 0.530359 | 4.115095 | 0.236658 | 1.019535 | 15/16 |

Single-step, multistimulus and orthogonal are distinct complete candidate records. The final unified candidate uses the orthogonal head; its metrics are not replaced by better entries from other rows. Functional references are nominal circuit quantities from independentVm. One acquisition session; no confirmatory significance claim.

## Supplementary Table S11. Whole-cell audit of the unified FOH adapter

Table11.1

| Heldout cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell3 | IonGate_RCState | 20.04102 | 9/10 |
| cell3 | Matched_static_dynamic | 19.99111 | 9/10 |
| cell3 | Matched_FIR | 20.04102 | 9/10 |
| cell3 | LOCO_static_Rs | 17.52586 | 9/10 |
| cell3 | IonGate_TrustRC | 17.57549 | 9/10 |
| cell6 | IonGate_RCState | 181.23010 | 1/5 |
| cell6 | Matched_static_dynamic | 181.23936 | 1/5 |
| cell6 | Matched_FIR | 181.23010 | 1/5 |
| cell6 | LOCO_static_Rs | 48.35631 | 4/5 |
| cell6 | IonGate_TrustRC | 45.22972 | 4/5 |
| cell9 | IonGate_RCState | 71.07610 | 3/5 |
| cell9 | Matched_static_dynamic | 51.39338 | 4/5 |
| cell9 | Matched_FIR | 71.07610 | 3/5 |
| cell9 | LOCO_static_Rs | 40.32837 | 4/5 |
| cell9 | IonGate_TrustRC | 42.24381 | 4/5 |

| Heldout cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell18 | IonGate_RCState | 175.95768 | 1/4 |
| cell18 | Matched_static_dynamic | 60.39035 | 2/4 |
| cell18 | Matched_FIR | 175.95767 | 1/4 |
| cell18 | LOCO_static_Rs | 58.81528 | 2/4 |
| cell18 | IonGate_TrustRC | 61.65430 | 2/4 |

All24conditions per method, including failures; cell3contains both temperatures. Loss is a penalized compensation-consistency endpoint, not independently referenced accuracy. MissingRseal causes no extra leak correction for any family.

## Supplementary Table S12. Complete development through V27, including physical-readout failures

Table12.1 Point estimates

| Version / method | Vm median mV | Vm P95 mV | Q median % | Q P95 % |
|---|---|---|---|---|
| IonGate v13 | 0.286491 | 2.527397 | 0.131053 | 0.460049 |
| FIR v13 | 0.422341 | 3.553636 | 0.202974 | 0.498797 |
| Affine v13 | 0.289945 | 3.610683 | 0.198315 | 0.447167 |
| IonGate v14 | 0.286322 | 2.527158 | 0.236435 | 0.512486 |
| FIR v14 | 0.422160 | 3.553643 | 0.221489 | 0.483579 |
| Affine v14 | 0.289946 | 3.610705 | 0.245599 | 0.477264 |
| IonGate v15 | 0.327083 | 2.430736 | 0.161065 | 0.503880 |
| FIR v15 | 0.439779 | 3.672841 | 0.206172 | 0.515130 |
| Affine v15 | 0.290738 | 3.612266 | 0.202887 | 0.464330 |
| IonGate v16 | 0.264573 | 3.454898 | 0.179171 | 0.465936 |

| Version / method | Vm median mV | Vm P95 mV | Q median % | Q P95 % |
|---|---|---|---|---|
| FIR v16 | 0.291699 | 3.610766 | 0.201527 | 0.462014 |
| Affine v16 | 0.292931 | 3.610832 | 0.207329 | 0.470942 |
| IonGate v17 | 0.284986 | 3.355038 | 0.198286 | 0.480813 |
| FIR v17 | 0.291846 | 3.610698 | 0.205669 | 0.446388 |
| Affine v17 | 0.292998 | 3.610789 | 0.215792 | 0.497494 |
| IonGate v18 | 0.291058 | 2.829484 | 0.185305 | 0.451386 |
| FIR v18 | 0.291500 | 3.610790 | 0.174549 | 0.464065 |
| Affine v18 | 0.292763 | 3.610815 | 0.176087 | 0.442399 |
| IonGate v19 | 0.252096 | 2.565919 | 0.179384 | 0.429866 |
| FIR v19 | 0.285565 | 3.563382 | 0.226084 | 0.454935 |
| Affine v19 | 0.285518 | 3.613301 | 0.207699 | 0.485676 |
| IonGate v20 | 0.252394 | 2.566111 | 0.142470 | 0.437989 |
| FIR v20 | 0.285893 | 3.563583 | 0.190046 | 0.484810 |
| Affine v20 | 0.285536 | 3.613390 | 0.188883 | 0.447414 |
| IonGate v21 | 0.252393 | 2.566122 | 0.140360 | 0.439048 |
| FIR v21 | 0.285893 | 3.563583 | 0.190046 | 0.484810 |
| Affine v21 | 0.285536 | 3.613390 | 0.191288 | 0.447345 |
| IonGate v22 | 0.252199 | 2.565983 | 0.117432 | 0.438437 |
| FIR v22 | 0.285855 | 3.563499 | 0.195087 | 0.455742 |
| Affine v22 | 0.285646 | 3.613352 | 0.144994 | 0.449336 |

| Version / method | Vm median mV | Vm P95 mV | Q median % | Q P95 % |
|---|---|---|---|---|
| IonGate v23 | 0.252106 | 2.565980 | 0.127688 | 0.428057 |
| FIR v23 | 0.285739 | 3.563486 | 0.204331 | 0.450802 |
| Affine v23 | 0.285646 | 3.613352 | 0.145504 | 0.449336 |
| IonGate v24 | 0.252106 | 2.565980 | 0.127688 | 0.428057 |
| FIR v24 | 0.285728 | 3.563484 | 0.204331 | 0.449662 |
| Affine v24 | 0.285646 | 3.613350 | 0.145504 | 0.449336 |
| IonGate v25 | 0.251650 | 2.565630 | 0.145089 | 0.391394 |
| FIR v25 | 0.285171 | 3.563279 | 0.205451 | 0.353553 |
| Affine v25 | 0.285810 | 3.613308 | 0.198259 | 0.539603 |
| IonGate v26 | 0.253013 | 2.565805 | 0.236164 | 0.543826 |
| FIR v26 | 0.286440 | 3.563469 | 0.274705 | 0.563086 |
| Affine v26 | 0.286398 | 3.613335 | 0.224816 | 0.633659 |
| IonGate v27 | 0.252581 | 2.565818 | 0.119793 | 0.442318 |
| FIR v27 | 0.285681 | 3.563486 | 0.206635 | 0.480635 |
| Affine v27 | 0.285734 | 3.613352 | 0.193687 | 0.521251 |

Table12.2 Intervals

| Version / method | Score median % | Score mean %* | Score P95 % | Hits | Width median % |
|---|---|---|---|---|---|
| IonGate v13 | 0.918820 | 0.915260 | 1.106286 | 16/16 | 0.919731 |
| FIR v13 | 0.991162 | 1.104718 | 1.853556 | 14/16 | 0.948981 |
| Affine v13 | 0.850080 | 1.013413 | 1.849032 | 13/16 | 0.741332 |
| IonGate v14 | 0.905330 | 1.055203 | 1.820638 | 14/16 | 0.900644 |

| Version / method | Score median % | Score mean %* | Score P95 % | Hits | Width median % |
|---|---|---|---|---|---|
| FIR v14 | 1.050596 | 1.132148 | 1.897317 | 14/16 | 0.954907 |
| Affine v14 | 0.886640 | 1.035187 | 1.613956 | 14/16 | 0.850641 |
| IonGate v15 | 0.877120 | 1.081122 | 1.997049 | 14/16 | 0.840945 |
| FIR v15 | 1.008431 | 1.123881 | 1.980229 | 14/16 | 0.988899 |
| Affine v15 | 0.817181 | 1.005569 | 1.791391 | 14/16 | 0.797693 |
| IonGate v16 | 0.792837 | 1.021804 | 1.802858 | 15/16 | 0.781583 |
| FIR v16 | 0.828084 | 1.015334 | 1.691032 | 14/16 | 0.787269 |
| Affine v16 | 0.821173 | 0.987521 | 1.641072 | 14/16 | 0.777991 |
| IonGate v17 | 0.991023 | 1.247275 | 2.903515 | 15/16 | 0.951877 |
| FIR v17 | 0.930025 | 1.169735 | 2.057993 | 14/16 | 0.897233 |
| Affine v17 | 0.927627 | 1.076115 | 1.982680 | 14/16 | 0.847906 |
| IonGate v18 | 0.732700 | 1.050184 | 2.095874 | 14/16 | 0.688609 |
| FIR v18 | 0.915368 | 1.089455 | 2.076510 | 15/16 | 0.903279 |
| Affine v18 | 0.850483 | 1.044356 | 2.015639 | 14/16 | 0.848728 |
| IonGate v19 | 1.003316 | 1.040734 | 1.348052 | 15/16 | 0.992029 |
| FIR v19 | 0.953942 | 0.964247 | 1.156940 | 16/16 | 0.953552 |
| Affine v19 | 0.884087 | 0.973474 | 1.519215 | 14/16 | 0.859213 |
| IonGate v20 | 1.057109 | 1.060761 | 1.313072 | 14/16 | 1.043293 |
| FIR v20 | 1.067459 | 1.182106 | 1.854000 | 15/16 | 1.048086 |
| Affine v20 | 0.983430 | 1.103353 | 1.892719 | 14/16 | 0.952314 |
| IonGate v21 | 1.090428 | 1.227414 | 2.017813 | 16/16 | 1.092269 |
| FIR v21 | 0.989486 | 1.083225 | 1.467850 | 14/16 | 0.974229 |

| Version / method | Score median % | Score mean %* | Score P95 % | Hits | Width median % |
|---|---|---|---|---|---|
| Affine v21 | 0.983618 | 1.072963 | 1.565676 | 15/16 | 0.974797 |
| IonGate v22 | 1.204046 | 1.290131 | 2.153996 | 15/16 | 1.198954 |
| FIR v22 | 1.003707 | 1.052569 | 1.500478 | 15/16 | 0.922476 |
| Affine v22 | 1.002241 | 1.120695 | 1.815450 | 14/16 | 0.891691 |
| IonGate v23 | 1.398749 | 1.480329 | 2.350902 | 16/16 | 1.397034 |
| FIR v23 | 1.201393 | 1.312484 | 1.826442 | 15/16 | 1.190483 |
| Affine v23 | 1.016785 | 1.275001 | 2.219439 | 15/16 | 1.002328 |
| IonGate v24 | 1.023436 | 1.010881 | 1.445136 | 14/16 | 0.957380 |
| FIR v24 | 0.835032 | 1.060520 | 2.068190 | 12/16 | 0.720063 |
| Affine v24 | 0.977817 | 1.258009 | 3.161869 | 11/16 | 0.751697 |
| IonGate v25 | 1.508475 | 1.557093 | 2.103442 | 16/16 | 1.509064 |
| FIR v25 | 1.168093 | 1.230224 | 1.843147 | 16/16 | 1.165239 |
| Affine v25 | 1.128654 | 1.160753 | 1.645607 | 15/16 | 1.058970 |
| IonGate v26 | 1.227672 | 1.605108 | 3.169915 | 15/16 | 1.228209 |
| FIR v26 | 1.421136 | 1.425874 | 2.112535 | 15/16 | 1.324803 |
| Affine v26 | 1.133005 | 1.285350 | 2.383743 | 15/16 | 1.090599 |
| IonGate v27 | 1.461672 | 1.538137 | 2.309190 | 16/16 | 1.462931 |
| FIR v27 | 1.240918 | 1.390712 | 2.093412 | 15/16 | 1.190460 |
| Affine v27 | 1.250991 | 1.320318 | 1.918796 | 15/16 | 1.166779 |

Table12.3 Fixed-coefficient memory ablation

| Counterfactual | Vm median mV | Vm P95 mV | Q median % | Q P95 % |
|---|---|---|---|---|
| V16full | 0.264573 | 3.454898 | 0.179171 | 0.465936 |

| Counterfactual | Vm median mV | Vm P95 mV | Q median % | Q P95 % |
|---|---|---|---|---|
| V16noRCfixedcoeff | 0.289996 | 3.610595 | 0.196396 | 0.445588 |
| V18full | 0.291058 | 2.829484 | 0.185305 | 0.451386 |
| V18noRCfixedcoeff | 0.289528 | 3.652745 | 0.190838 | 0.464667 |
| V20 full | 0.252394 | 2.566111 | 0.142470 | 0.437989 |
| V20 RC disabled | 0.274231 | 3.739478 | 0.149091 | 0.377137 |
| V20 adaptation disabled | 0.265174 | 2.527397 | 0.131053 | 0.460049 |
| V22 full | 0.252199 | 2.565983 | 0.117432 | 0.438437 |
| V22 RC disabled | 0.274075 | 3.739402 | 0.135684 | 0.350477 |
| V22 current/readout disabled | 0.265174 | 2.527397 | 0.131053 | 0.460049 |
| V23 full | 0.252106 | 2.565980 | 0.127688 | 0.428057 |
| V23 RC disabled | 0.274022 | 3.739401 | 0.199046 | 0.378981 |
| V23 current/readout disabled | 0.265174 | 2.527397 | 0.131053 | 0.460049 |
| V24 full | 0.252106 | 2.565980 | 0.127688 | 0.428057 |
| V24 RC disabled | 0.274022 | 3.739401 | 0.199046 | 0.378981 |
| V24 current/readout disabled | 0.265174 | 2.527397 | 0.131053 | 0.460049 |
| V25 full | 0.251650 | 2.565630 | 0.145089 | 0.391394 |
| V25 inherited RC heads disabled | 0.273671 | 3.739201 | 0.193411 | 0.385431 |

| Counterfactual | Vm median mV | Vm P95 mV | Q median % | Q P95 % |
|---|---|---|---|---|
| V25 balance update disabled | 0.252106 | 2.565980 | 0.127688 | 0.428057 |
| V26 full | 0.253013 | 2.565805 | 0.236164 | 0.543826 |
| V26 inherited RC heads disabled | 0.275184 | 3.739306 | 0.258550 | 0.534859 |
| V26 balance update disabled | 0.252106 | 2.565980 | 0.127688 | 0.428057 |
| V27 full | 0.252581 | 2.565818 | 0.119793 | 0.442318 |
| V27 inherited RC heads disabled | 0.274410 | 3.739303 | 0.172065 | 0.427950 |
| V27 balance update disabled | 0.252106 | 2.565980 | 0.127688 | 0.428057 |

All16settings from one hardware session. Primaryalpha=.2. Mean score is secondary and added after inspection; median remains the promotion endpoint. Ablation has no calibrated interval claims. Functional references are nominal circuit quantities, not ionic-charge metrology. Nohistoricalbest-value stitching. V17/V18 failed the exact new retention gate and were not activated; V18 lower median score came with 14/16 hits. V19/V20 fixed signed residual envelopes use alpha only as the score penalty, without a nominal coverage guarantee. Point improvements did not pass the full new retention target. V21/V22 use bias/profile-adjusted rank intervals with the same primary alpha=0.2 score. V21 restores 16/16, while V22 improves functional median; these are separate complete candidates. No nominal physical coverage guarantee is asserted. V23 conditional-quantile and V24 consensus-centered intervals are distinct empirical constructions, not interchangeable or guaranteed CV+ certificates. V24 secondary mean/P95 score gains do not replace the primary median requirement. V25/V26/V27 have 60/96/144 new point configurations per family and 12 interval choices. V27 was one disclosed amendment after transient-endpoint diagnosis; no new candidate passed joint qualification.

## Supplementary Table S13. All whole-cell outcomes, including failed development candidates

Table13.1

| Held-out cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell3 | Retained V13 | 17.603387 | 9/10 |
| cell3 | IonGate V16 | 19.992710 | 9/10 |
| cell3 | Matched static V16 | 19.991513 | 9/10 |
| cell3 | FIR V16 | 19.992710 | 9/10 |
| cell3 | Old tuned Rs | 17.525864 | 9/10 |
| cell6 | Retained V13 | 78.816366 | 2/5 |
| cell6 | IonGate V16 | 3.671017 | 5/5 |
| cell6 | Matched static V16 | 4.301502 | 5/5 |
| cell6 | FIR V16 | 3.671066 | 5/5 |
| cell6 | Old tuned Rs | 48.356313 | 4/5 |
| cell9 | Retained V13 | 46.825813 | 4/5 |
| cell9 | IonGate V16 | 43.262557 | 4/5 |
| cell9 | Matched static V16 | 43.292177 | 4/5 |
| cell9 | FIR V16 | 43.262557 | 4/5 |
| cell9 | Old tuned Rs | 40.328373 | 4/5 |
| cell18 | Retained V13 | 56.751200 | 2/4 |
| cell18 | IonGate V16 | 56.513307 | 2/4 |
| cell18 | Matched static V16 | 56.501295 | 2/4 |
| cell18 | FIR V16 | 56.513307 | 2/4 |
| cell18 | Old tuned Rs | 58.815283 | 2/4 |

| Held-out cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell3 | IonGate V17 | 19.992710 | 9/10 |
| cell6 | IonGate V17 | 3.671017 | 5/5 |
| cell9 | IonGate V17 | 43.262557 | 4/5 |
| cell18 | IonGate V17 | 56.513307 | 2/4 |
| cell3 | IonGate V18 | 19.992710 | 9/10 |
| cell6 | IonGate V18 | 3.671017 | 5/5 |
| cell9 | IonGate V18 | 43.262557 | 4/5 |
| cell18 | IonGate V18 | 56.513307 | 2/4 |
| cell3 | IonGate V19 | 19.992710 | 9/10 |
| cell6 | IonGate V19 | 3.671018 | 5/5 |
| cell9 | IonGate V19 | 43.262557 | 4/5 |
| cell18 | IonGate V19 | 56.513307 | 2/4 |
| cell3 | IonGate V20 | 19.992710 | 9/10 |
| cell6 | IonGate V20 | 3.671018 | 5/5 |
| cell9 | IonGate V20 | 43.262557 | 4/5 |
| cell18 | IonGate V20 | 56.513307 | 2/4 |
| cell3 | IonGate SoftState v21 | 19.992710 | 9/10 |
| cell6 | IonGate SoftState v21 | 3.671017 | 5/5 |
| cell9 | IonGate SoftState v21 | 43.262557 | 4/5 |
| cell18 | IonGate SoftState v21 | 56.513307 | 2/4 |

| Held-out cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell3 | Matched static bio v21 | 19.991513 | 9/10 |
| cell6 | Matched static bio v21 | 4.301502 | 5/5 |
| cell9 | Matched static bio v21 | 43.292177 | 4/5 |
| cell18 | Matched static bio v21 | 56.501295 | 2/4 |
| cell3 | Matched FIR bio v21 | 19.992710 | 9/10 |
| cell6 | Matched FIR bio v21 | 3.671066 | 5/5 |
| cell9 | Matched FIR bio v21 | 43.262557 | 4/5 |
| cell18 | Matched FIR bio v21 | 56.513307 | 2/4 |
| cell3 | IonGate SoftState v22 | 19.992710 | 9/10 |
| cell6 | IonGate SoftState v22 | 3.671017 | 5/5 |
| cell9 | IonGate SoftState v22 | 43.262557 | 4/5 |
| cell18 | IonGate SoftState v22 | 56.513307 | 2/4 |
| cell3 | Matched static bio v22 | 19.991513 | 9/10 |
| cell6 | Matched static bio v22 | 4.301502 | 5/5 |
| cell9 | Matched static bio v22 | 43.292177 | 4/5 |
| cell18 | Matched static bio v22 | 56.501295 | 2/4 |
| cell3 | Matched FIR bio v22 | 19.992710 | 9/10 |
| cell6 | Matched FIR bio v22 | 3.671066 | 5/5 |
| cell9 | Matched FIR bio v22 | 43.262557 | 4/5 |
| cell18 | Matched FIR bio v22 | 56.513307 | 2/4 |

| Held-out cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell3 | IonGate SoftState v23 | 19.992710 | 9/10 |
| cell6 | IonGate SoftState v23 | 3.671017 | 5/5 |
| cell9 | IonGate SoftState v23 | 43.262557 | 4/5 |
| cell18 | IonGate SoftState v23 | 56.513307 | 2/4 |
| cell3 | Matched static bio v23 | 19.991513 | 9/10 |
| cell6 | Matched static bio v23 | 4.301502 | 5/5 |
| cell9 | Matched static bio v23 | 43.292177 | 4/5 |
| cell18 | Matched static bio v23 | 56.501295 | 2/4 |
| cell3 | Matched FIR bio v23 | 19.992710 | 9/10 |
| cell6 | Matched FIR bio v23 | 3.671066 | 5/5 |
| cell9 | Matched FIR bio v23 | 43.262557 | 4/5 |
| cell18 | Matched FIR bio v23 | 56.513307 | 2/4 |
| cell3 | IonGate SoftState v24 | 19.992710 | 9/10 |
| cell6 | IonGate SoftState v24 | 3.671017 | 5/5 |
| cell9 | IonGate SoftState v24 | 43.262557 | 4/5 |
| cell18 | IonGate SoftState v24 | 56.513307 | 2/4 |
| cell3 | Matched static bio v24 | 19.991513 | 9/10 |
| cell6 | Matched static bio v24 | 4.301502 | 5/5 |
| cell9 | Matched static bio v24 | 43.292177 | 4/5 |
| cell18 | Matched static bio v24 | 56.501295 | 2/4 |

| Held-out cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell3 | Matched FIR bio v24 | 19.992710 | 9/10 |
| cell6 | Matched FIR bio v24 | 3.671066 | 5/5 |
| cell9 | Matched FIR bio v24 | 43.262557 | 4/5 |
| cell18 | Matched FIR bio v24 | 56.513307 | 2/4 |
| cell3 | IonGate SoftState v25 | 19.992710 | 9/10 |
| cell6 | IonGate SoftState v25 | 3.671017 | 5/5 |
| cell9 | IonGate SoftState v25 | 43.262557 | 4/5 |
| cell18 | IonGate SoftState v25 | 56.513307 | 2/4 |
| cell3 | Matched static bio v25 | 19.991513 | 9/10 |
| cell6 | Matched static bio v25 | 4.301502 | 5/5 |
| cell9 | Matched static bio v25 | 43.292177 | 4/5 |
| cell18 | Matched static bio v25 | 56.501295 | 2/4 |
| cell3 | Matched FIR bio v25 | 19.992710 | 9/10 |
| cell6 | Matched FIR bio v25 | 3.671066 | 5/5 |
| cell9 | Matched FIR bio v25 | 43.262557 | 4/5 |
| cell18 | Matched FIR bio v25 | 56.513307 | 2/4 |
| cell3 | IonGate SoftState v26 | 19.992710 | 9/10 |
| cell6 | IonGate SoftState v26 | 3.671017 | 5/5 |
| cell9 | IonGate SoftState v26 | 43.262557 | 4/5 |
| cell18 | IonGate SoftState v26 | 56.513307 | 2/4 |

| Held-out cell | Method | Penalized loss mV | Identified / attempted |
|---|---|---|---|
| cell3 | Matched static bio v26 | 19.991513 | 9/10 |
| cell6 | Matched static bio v26 | 4.301502 | 5/5 |
| cell9 | Matched static bio v26 | 43.292177 | 4/5 |
| cell18 | Matched static bio v26 | 56.501295 | 2/4 |
| cell3 | Matched FIR bio v26 | 19.992710 | 9/10 |
| cell6 | Matched FIR bio v26 | 3.671066 | 5/5 |
| cell9 | Matched FIR bio v26 | 43.262557 | 4/5 |
| cell18 | Matched FIR bio v26 | 56.513307 | 2/4 |
| cell3 | IonGate SoftState v27 | 19.992710 | 9/10 |
| cell6 | IonGate SoftState v27 | 3.671017 | 5/5 |
| cell9 | IonGate SoftState v27 | 43.262557 | 4/5 |
| cell18 | IonGate SoftState v27 | 56.513307 | 2/4 |
| cell3 | Matched static bio v27 | 19.991513 | 9/10 |
| cell6 | Matched static bio v27 | 4.301502 | 5/5 |
| cell9 | Matched static bio v27 | 43.292177 | 4/5 |
| cell18 | Matched static bio v27 | 56.501295 | 2/4 |
| cell3 | Matched FIR bio v27 | 19.992710 | 9/10 |
| cell6 | Matched FIR bio v27 | 3.671066 | 5/5 |
| cell9 | Matched FIR bio v27 | 43.262557 | 4/5 |
| cell18 | Matched FIR bio v27 | 56.513307 | 2/4 |

450sweeps,4cells,5cell/temperaturecohorts,24conditions per method. Failed and boundary fits remain in scores. Apparentactivation consistency/identification, not independentbiologicalVm orionicchargeaccuracy. SameV13model-cell andbiological adapters areboundinone release record.

Same216configurationbiologicalalgorithmwasactuallyrefitV14,V15,V16; point-levelpredictionsandallfailedfitsarearchived. V17/V18 used the unchanged algorithm and were refitted, again reaching 20/24; these are not new independent biological replications. The V19/V20 biological adapter was unchanged and refitted; these reruns do not independently validate current-driven adaptation in living cells. V21/V22 include all current enhanced biological controls, again 20/24. The unchanged adapter is not a newly validated current-tangent learner in living cells. V23/V24 rerun the unchanged biological adapter; all 450 point records and four failures are retained.

## Supplementary Table S14. Bounded residual learning and frozen mechanism ablations

Table14.1 Point endpoints: complete frozen candidates

| **Model** | **Vm median** | **Vm P95** | **Q median %** | **Q P95 %** |
|---|---|---|---|---|
| IonGate v27 | 0.252581 | 2.565818 | 0.119793 | 0.442318 |
| FIR v27 | 0.285681 | 3.563486 | 0.206635 | 0.480635 |
| Affine v27 | 0.285734 | 3.613352 | 0.193687 | 0.521251 |
| IonGate v28 | 0.252437 | 2.565788 | 0.123099 | 0.443292 |
| FIR v28 | 0.285665 | 3.563475 | 0.205503 | 0.479949 |
| Affine v28 | 0.285671 | 3.613353 | 0.191958 | 0.523703 |
| IonGate head_disabled | 0.252581 | 2.565818 | 0.119793 | 0.442318 |
| IonGate legacy_envelope | 0.252437 | 2.565788 | 0.123099 | 0.443292 |
| FIR head_disabled | 0.285681 | 3.563486 | 0.206635 | 0.480635 |
| FIR legacy_envelope | 0.285665 | 3.563475 | 0.205503 | 0.479949 |
| Affine head_disabled | 0.285734 | 3.613352 | 0.193687 | 0.521251 |

| Model | Vm median | Vm P95 | Q median % | Q P95 % |
|---|---|---|---|---|
| Affine legacy_envelope | 0.285671 | 3.613353 | 0.191958 | 0.523703 |

Table14.2 Primary interval endpoints

| Model | IS median % | IS mean % | IS P95 % | Hits | Width % |
|---|---|---|---|---|---|
| IonGate v27 | 1.461672 | 1.538137 | 2.309190 | 16/16 | 1.462931 |
| FIR v27 | 1.240918 | 1.390712 | 2.093412 | 15/16 | 1.190460 |
| Affine v27 | 1.250991 | 1.320318 | 1.918796 | 15/16 | 1.166779 |
| IonGate v28 | 1.081955 | 1.300414 | 2.616709 | 13/16 | 0.944073 |
| FIR v28 | 0.791466 | 1.031393 | 1.842072 | 12/16 | 0.742326 |
| Affine v28 | 1.043329 | 1.493089 | 4.328708 | 11/16 | 0.749796 |
| IonGate head_disabled | 1.097313 | 1.324881 | 2.606319 | 13/16 | 0.925665 |
| IonGate legacy_envelope | 1.456769 | 1.560380 | 2.330661 | 16/16 | 1.461926 |
| FIR head_disabled | 0.791466 | 1.025541 | 1.874881 | 12/16 | 0.767878 |
| FIR legacy_envelope | 1.195176 | 1.383146 | 2.093412 | 15/16 | 1.183654 |
| Affine head_disabled | 1.049635 | 1.645206 | 4.352778 | 10/16 | 0.753969 |
| Affine legacy_envelope | 1.221726 | 1.330625 | 1.918796 | 15/16 | 1.144334 |

All 16 settings from one hardware session. Lower errors/scores are better. Primary alpha=0.2. Ablations are post-freeze diagnostics, not selected new releases. Biology is byte-identical V27 evidence (20/24 fits, 450 sweeps). Historical stronger controls remain in Table12 and full machine-readable benchmark.

## Supplementary Table S15. Whole-model synchronized recalibration and frozen diagnostics

Table15.1 Complete point endpoints

| Model | Vm median | Vm P95 | Q median % | Q P95 % |
|---|---|---|---|---|
| IonGate v29 | 0.289890 | 2.467720 | 0.146481 | 0.466615 |
| FIR v29 | 0.287680 | 3.668293 | 0.213289 | 0.469307 |
| Affine v29 | 0.283808 | 3.613016 | 0.196448 | 0.544171 |
| IonGate frozen9_signed_envelope | 0.252581 | 2.565818 | 0.119793 | 0.442318 |
| IonGate recentered_rank2_11 | 0.289890 | 2.467720 | 0.146481 | 0.466615 |
| FIR frozen9_signed_envelope | 0.285681 | 3.563486 | 0.206635 | 0.480635 |
| FIR recentered_rank2_11 | 0.287680 | 3.668293 | 0.213289 | 0.469307 |
| Affine frozen9_signed_envelope | 0.285734 | 3.613352 | 0.193687 | 0.521251 |
| Affine recentered_rank2_11 | 0.283808 | 3.613016 | 0.196448 | 0.544171 |

Table15.2 Primary proper interval endpoints

| Model | IS median % | IS mean % | IS P95 % | Hits | Width % |
|---|---|---|---|---|---|
| IonGate v29 | 1.037998 | 1.075982 | 1.638706 | 13/16 | 1.018497 |
| FIR v29 | 1.215037 | 1.228555 | 1.737177 | 15/16 | 1.120607 |
| Affine v29 | 1.031743 | 1.018382 | 1.265792 | 16/16 | 1.027804 |

| Model | IS median % | IS mean % | IS P95 % | Hits | Width % |
|---|---|---|---|---|---|
| IonGate frozen9_signed_envelope | 1.092070 | 1.080048 | 1.332597 | 15/16 | 1.029739 |
| IonGate recentered_rank2_11 | 0.871646 | 1.031792 | 1.817929 | 10/16 | 0.719653 |
| FIR frozen9_signed_envelope | 1.180516 | 1.225988 | 1.772040 | 16/16 | 1.179287 |
| FIR recentered_rank2_11 | 0.719874 | 0.931073 | 1.534852 | 12/16 | 0.715901 |
| Affine frozen9_signed_envelope | 1.010499 | 1.183664 | 2.295414 | 14/16 | 0.969830 |
| Affine recentered_rank2_11 | 0.787582 | 1.065723 | 2.443887 | 13/16 | 0.777982 |

All16settings/onehardware_session;nominalRCfunctional. Equal12refits/fold/family. Frozen diagnostics are not selected releases. The affine control attains16/16 and is retained. Biological results remain byte-identical V27 evidence. Oracle padding is diagnostic only, never used at inference.

## Supplementary Table S16. Waveform-protected readout and complete calibration alternatives

Table16.1 Complete point endpoints

| Model | Vm median | Vm P95 | Q median % | Q P95 % |
|---|---|---|---|---|
| IonGate v27 | 0.252581 | 2.565818 | 0.119793 | 0.442318 |
| FIR v27 | 0.285681 | 3.563486 | 0.206635 | 0.480635 |
| Affine v27 | 0.285734 | 3.613352 | 0.193687 | 0.521251 |
| IonGate v30 | 0.252515 | 2.565794 | 0.105459 | 0.443297 |

| Model | Vm median | Vm P95 | Q median % | Q P95 % |
|---|---|---|---|---|
| FIR v30 | 0.285681 | 3.563487 | 0.220536 | 0.450802 |
| Affine v30 | 0.285783 | 3.613359 | 0.208699 | 0.547341 |
| IonGate v31 | 0.252515 | 2.565794 | 0.105459 | 0.443297 |
| FIR v31 | 0.285681 | 3.563487 | 0.220536 | 0.450802 |
| Affine v31 | 0.285783 | 3.613359 | 0.208699 | 0.547341 |
| IonGate v32 | 0.252482 | 2.565840 | 0.114725 | 0.448022 |
| FIR v32 | 0.285681 | 3.563486 | 0.220536 | 0.450802 |
| Affine v32 | 0.285781 | 3.613359 | 0.206923 | 0.546054 |

Table16.2 Proper interval endpoints

| Model | IS median % | IS mean % | IS P95 % | Hits | Width % |
|---|---|---|---|---|---|
| IonGate v27 | 1.461672 | 1.538137 | 2.309190 | 16/16 | 1.462931 |
| FIR v27 | 1.240918 | 1.390712 | 2.093412 | 15/16 | 1.190460 |
| Affine v27 | 1.250991 | 1.320318 | 1.918796 | 15/16 | 1.166779 |
| IonGate v30 | 1.362887 | 1.377353 | 2.091487 | 16/16 | 1.364174 |
| FIR v30 | 1.069416 | 1.194838 | 1.790461 | 15/16 | 1.042015 |
| Affine v30 | 1.159797 | 1.266840 | 1.840088 | 15/16 | 1.117981 |
| IonGate v31 | 1.483720 | 1.560441 | 2.339500 | 16/16 | 1.485024 |
| FIR v31 | 1.194897 | 1.379835 | 2.093412 | 15/16 | 1.177501 |
| Affine v31 | 1.128703 | 1.308477 | 1.920036 | 15/16 | 1.121917 |
| IonGate v32 | 1.210591 | 1.127197 | 1.541311 | 15/16 | 1.194509 |
| FIR v32 | 0.924269 | 1.043019 | 1.667880 | 14/16 | 0.889162 |

| Model | IS median % | IS mean % | IS P95 % | Hits | Width % |
|---|---|---|---|---|---|
| Affine v32 | 0.950162 | 1.114967 | 1.862160 | 15/16 | 0.933157 |

All16hardware_settings/one_session. V30/V31pointCSVfilesarebyte-identical; no metric stitching. V32 uses24correlated interval candidates, not24independent settings. Biology unchanged V27 all450sweeps20/24fits. Strict and engineering-tolerance recommendations are reported separately; neither implies all requested targets passed.

## S7 Claim–evidence and attribution index

Voltage reconstruction: main Figure 2 and Table 2; independent second-amplifier reference, with group roles and frozen provenance in source/INPUT_LINEAGE.json and evidence/FINAL_MODEL_V30.json. Primary values are recomputed by reproduce/reproduce_ibex.py --verify-only.

Branch-memory readout: main Figure 3, figure_data/V30_SAME_WAVE_MEMORY_ABLATION.csv and Equation (2). The returned voltage, window, initialization and full reference are fixed; only the terms used for prediction change.

Bounded functional learning: main Figure 3 and the protected-anchor comparison in Results 3.3. The 50 fC functional and 0.15 mV pointwise limits constrain the update, not experimental uncertainty. Both metrics are calculated from the returned waveform.

Calibration: main Figure 4, the alpha = 0.2 setting-level results and the identical-point alternative. Full reference containment, interval width and proper score are separate endpoints; the 8.14% internal score improvement is not strong-baseline superiority.

Biological applicability: main Figure 6 and Table 3; all 450 sweeps, four cells, 24 conditions and four unidentifiable cases remain included. Enhanced controls share the 20/24 count. Hardware and optical contexts in Figures 5–7 are not pooled V30 accuracy trials.

Implementation assurance: source lineage, model manifests, retained software tests and reproduction scripts establish traceability and computational consistency. They are not additional biological or independent-session replications. This writing revision changes no prediction, calibration endpoint or figure pixels.

## S8 Review access, reuse licenses and frozen replay

DataS3.zip is the primary numerical review route. Install its requirements.txt and execute python replay.py. The program reconstructs predictions and alpha = 0.2 intervals for 16 settings and all three matched families, then checks voltage RMSE, nominal RC-functional error, reference hits, prediction-normalized relative width and reference-normalized proper score. A read guard rejects undeclared access to the original project, and fitting routines are disabled. This reproduces frozen primary inference and scoring, not historical fitting or biological retraining.

DataS2.zip contains complete table sources, main-figure numerical data, public dataset identifiers, metadata snapshots and image-processing provenance. The main image panels were rendered from dataset pixels: Sham data [9] and neuronal data [11] are CC BY 4.0, whereas TAC data [10] are CC0. These dataset licenses are distinct from the license of the accompanying journal article. No publisher article-image pixels were reused. Cropping, false-colour scaling and selected-frame display are documented. BSD-3-Clause notices for source circuit-model repositories are retained.

The stored numerical results are unchanged. A submission audit corrected the wording of the width denominator: relative interval width is normalized by the predicted functional magnitude; proper score is normalized by the reference magnitude. This matches the frozen scoring code and all existing table values.

## Supplementary figure captions

**Figure S1. IonGate separates measurement physics, functional readout and calibration.**

The preserved framework contrasts requested voltage, amplifier output and the independently observed membrane voltage. Analytical panels show load-dependent gain and the 10-to-3.25 current-ratio counterexample, distinct from recorded measurements. Added panels i–k show native command voltage, current and independent Vm with frozen IonGate Vm for fixed setting 00 over 45–80 ms. The waveform example is chosen by identifier rather than performance. Dynamic state and reference role define the downstream readout and interval. Complete sources and conditions accompany the figure assets.

**Figure S2. First-amplifier predictions are scored against an independent voltage sensor.**

All 16 native compensation-setting waveform panels are retained in two rows of eight, including 351 samples in the scoring window. Purple denotes second-sensor Vm, teal IonGate and dashed grey the fixed Lei comparator. Full-window and phase metrics retain all wins and reversals. Phase windows are onset [50,52), plateau [52,69) and offset [70,72) ms; illustrative zoom windows in l–n are distinct from these scoring windows. Distribution bands denote descriptive setting summaries, not confidence intervals. Signed residual matrices preserve transition error. All settings belong to one physical session.

**Figure S3. Retaining dynamic RC memory improves the functional readout.**

The 16 original native residual curves, paired full/no-memory/leak-only endpoints, cumulative distributions and component maps use identical frozen IonGate voltage. L, C and S denote leak, capacitor and state-memory charge contributions. All full-versus-ablation reversals remain visible. Panels k–m give cumulative contribution trajectories for fixed setting 00, compared with the nominal readout from measured Vm. The hold is the 40–49.9 ms mean, state is propagated from record start, and the integration window is 45–80 ms. The

reference is an independently voltage-supported nominal circuit functional, not independently metered charge.

**Figure S4. Conditional interval utility depends on calibration support and width policy.**

All 16 frozen nominal-80% intervals retain 14 IonGate hits, 10 Lei hits and every miss. Widths, interval scores, fold radii, level sensitivity, assumed parameter-budget cases and budget-dependent decisions accompany the interval grid. Panels p–r expose calibration residuals, test residuals and signed containment margins. Width budgets never use held-out errors; common accepted subsets are reported separately. Four calibration groups cannot supply the required fifth order statistic at a nominal 95% level. This figure uses the original window-initial functional convention, not the matched hold-based Figure S3 convention. Assumed bounds are sensitivity cases, not calibration certificates.

**Figure S5. Programmed-state hardware recordings test area transfer and specificity.**

The Sato experimental-noise hardware subset contains 240 records, with 80 per nominal SNR band of 10.2, 5.14 and 1.82. Recorded area is compared with a cross-record amplitude/baseline-calibrated occupancy reference. Panels j–l retain empirical full-scale error distributions, all area–occupancy pairs and correct/inverted-label controls using identical calibration. Programmed states define the drive sequence; they are not independently metered current. The figure characterizes the real measurement chain and is not a scored IonGate advantage experiment.

**Figure S6. Frozen NaV1.5 application is linked to distinct functional and imaging contexts.**

The preserved 24-image grid uses actual source frames in three rows of eight and fixed source-specific contrast/scale information. Quantitative panels distinguish 450 applied NaV sweeps from 45 scored sweeps in one cell against an unchanged author-model Vm reference. Additional panels m–o retain all 45 paired comparisons with voltage proxies, all 49 paired

SCN5A peak-density summaries, and signed areas of the four original hERG genotype/treatment traces. SCN5A and hERG panels are source-assay/replay context, not additional IonGate classifiers or charge certificates. Cardiac optical images derive from He et al. [13]; paired neuronal material derives from the Xie study dataset [14]. Exact frame IDs, crops, scales and pixel hashes appear in Table S7.

**Figure S7. Real fluorescence is paired with independently recorded electrode voltage.**

The field is the mean of 100 raw frames; the cell-detail crop preserves native context. The 18-frame grid spans fixed target offsets −25 to +60 ms in 5 ms steps, in three rows of six, using actual nearest frames around the first upward 0 mV electrode crossing. Labels give rounded observed times. Optical/electrode traces, all 2,273 paired observations, the recorded command, event-centered trajectories and frame-time fluorescence values retain their native observations. Common crops, linear intensity limits and scale bars preserve comparability. One neuron/session and its dependent repeated events provide external measurement context, not an IonGate optical benchmark. Source: Xie et al. [14]; paired public dataset doi:10.25378/janelia.14143118, collection doi:10.25378/janelia.c.5325254.

**Figure S8. Matched-budget comparison separates waveform and functional endpoints.**

(a,b) All 16 setting-level voltage and nominal circuit-functional errors. (c,d) Interval scores and full widths after separate same-fold calibration. (e) Native onset from fixed setting 00, chosen by identifier, against the independent voltage sensor. (f) Paired circuit-functional discrepancies; the identity line does not indicate a null hypothesis test. (g) Actual residual calls under identical three-start, 180-evaluation caps. (h) Changes from the voltage-targeted Grow candidate; positive coordinates indicate improvement. All settings belong to one physical session.

**Figure S9. Acquisition contracts and bounded learning expose distinct failure modes.**

(a,b) Accepted settings and ungated reference hits in software sampling/filtering transformations of real recordings. (c) Rejections when added filtering is deliberately omitted from metadata; missed cases remain visible. (d) Inner-validation voltage gain, with accepted candidate in purple. (e) Reference hits and width after separate calibration; original IonGate remains default. (f) Implemented future-session release workflow, not evidence of new-session learning performance.

**Figure S10. Apparent biological activation is reported with identification failures.**

(a–e) All available compensation fits in five cell-temperature cohorts. Crosses mark fits flagged by the post-score identification audit; NA denotes insufficient points. (f) Unfiltered cohort-level Vhalf ranges, including unfavorable and incomplete cases; these are not activation-accuracy errors. (g) Identifiable fits relative to all 24 attempted conditions per method in the primary support window. (h) Analysis units and interpretation. All methods share the observed current peak; assumed ENa and model-reconstructed voltage do not supply independent activation truth.

**Figure S11. Joint optimization and whole-cell transfer expose endpoint-specific gains.**

a–c, All 16 outer-held-out setting scores for voltage RMSE, nominal circuit-functional discrepancy and interval score. d, Median versus P95 voltage error. e, Setting-paired voltage improvement over tuned Lei; positive values favor IonGate. f, Penalized compensation-consistency loss in every held-out cell. g, Identifiable/attempted activation conditions, retaining all failures. h, Voltage–functional median trade-off. i, At least 1% median-improvement checks for the V9 candidate against each comparator; the complete release gate additionally checks tails and reference hits. Joint denotes the retained anchored V8 candidate; CV denotes the non-promoted V9 repeated-inner candidate. Sixteen settings came from one

hardware session, not 16 independent acquisitions. Biological data comprise 450 sweeps from four cells, with cell3 temperatures kept together. No p-values or independent physical-charge coverage are implied.

**Figure S12. One-version RC-State evaluation against enhanced matched controls.**

a, Real first-amplifier reconstruction and independent membrane-voltage reference, with representative setting selected nearest candidate median RMSE, not largest gain. b,c, All16heldout setting voltage/function errors. d, Interval scores; open rings mark reference misses. e, Median andP95voltage errors. f, Fixed-coefficient RC-state feature ablation; positive values mean improvement with states, and the negative median-voltage result is retained. g,h, Whole-cell-heldout biological identification and penalized consistency, with all450sweeps and failed fits retained. i, Unified release checks. The same orthogonal electronic head and FOH biological adapter generate the complete version record; no best-value stitching is allowed. Model-cell observations derive from one session, and biology from four cells. No independently measured ionic-charge coverage is implied.

**Figure S13. Conditional functional learning, circuit-state attribution and complete-version trade-offs.**

a, Real waveform nearest the V22 median voltage RMSE; independent second sensor withheld from inference. b-d, Every hardware setting: voltage, nominal functional and interval-score distributions. e, All reference intervals and point estimates. f, Empirical coverage versus median width; retained V18 included. g,h, Frozen-coefficient RC and current/readout interventions; negative contributions remain visible. i,j, Entire held-out biological cells; enhanced controls also reach 20/24. k, Current-driven waveform-update RMS and four-base-model min/max, not a confidence interval. l, Training-selected functional basis (0 physical, 1 quadratic, 2 joint), gain and profile power. m, Fixed high-compensation

setting 05 with retained V16/V18 and V22. n, Complete-version median frontier including V21 and V22. o, All four current matched point-summary ratios. p, Eight unchanged joint checks; no complete column passes. One hardware session, four biological cells; nominal RC functional, not independently metered ionic charge. Frozen interventions were not recalibrated. Affine/FIR controls include the shared nonlinear adapter/readout.

**Figure S14. All24activationconditions,with20identifiedfitsandfourretainedfailures.**

Points are apparent conductance and estimated peak voltage derived from actual first-amplifier currents. Lines are per-condition Boltzmannfits. Normalization is to the maximum eligible conductance in that condition, not a shared physiological truth. Brown panels retain failed/boundary cases. The predictor has no independently measured biologicalVmreference; enhanced matchedcontrolsalsoreach20/24. The original all-condition image is preserved: V20 reruns the same unchanged biological adapter and point traces, not a newly acquired biological experiment; current rerun scores are in Supplementary Table S13. V21/V22 reruns use the same unchanged biological adapter; all 450 point inputs were checked against the preserved plate.

**Figure S15. Complete-version tail recovery, empirical coverage and interval utility.**

a, Representative real voltage nearest the V23 median RMSE; V24 IonGate point waveforms are identical. b, All 16 paired functional discrepancies, including adverse changes. c, Functional error distributions and the V20 tail threshold. d, Full primary interval-score distributions; current controls share new opportunities. e, Empirical coverage-width trade-off of complete versions. f, Every V23/V24 reference-centered interval. g, Actual functional updates from V22. h, Selected gains/thresholds; no positive threshold selected. i, Every coverage outcome. j, Whole-cell fit retention, all current enhanced families 20/24. k, V24 ratios against current matched controls for four point summaries and two secondary score

summaries; the primary median remains worse. l, Eight unchanged current targets; none passes all. One hardware session and four biological cells, not 16 independent experiments. No nominal physical coverage guarantee or independently metered charge claim.

**Figure S16. Dynamic current-balance challenges and finite-endpoint validation.**

a, Fixed zero-compensation setting: return transient, command and current-equivalent voltage. b, Finite endpoint delta Iout/nominal gm versus independent delta Vm; not a DC gain calibration. c, Whole-version functional median/P95 frontier. d, All-setting discrepancies, including the stronger FIR25 tail. e, Primary interval-score distributions and retained V18 threshold. f, Complete-version empirical coverage/width. g, Four V27 point ratios against its current equal-budget affine/FIR. h, Eight common current targets; none passes all. All 16 hardware settings from one session, all biological reruns retaining 20/24 fits, and all negative candidates preserved. Endpoint diagnostics were post-freeze; no independent ionic-charge or guaranteed coverage claim.

**Figure S17. Bounded residual learning and the coverage-width trade-off.**

a,b, Fixed zero- and high-compensation real waveforms, with independent Vm. c,d, All-setting functional errors and paired changes. e,f, Whole-version functional frontier and proper interval scores. g, All 16 intervals and all three misses. h, Coverage-width frontier. i, Frozen head and interval ablations; legal calibration recomputed. j, Four matched point-summary ratios. k, Unchanged biological identifiable fits by cell (all 450 sweeps). l, All eight common current targets. One hardware session; nominal RC functional, not independently metered ionic charge; retrospective development. No activation; all older figures retained.

**Figure S18. Synchronized model recalibration and complete qualification audit.**

a,b, Fixed real waveforms at settings 00/15. c,d, All sixteen functional discrepancies and paired V29-V27 differences. e,f, Whole-version functional and proper-score comparisons.

g,h, All intervals/misses and the coverage-width frontier. i, Old nine-setting coefficient and interval-rank diagnostics, not selected new releases. j, Matched point ratios, including adverse voltage-median comparison. k, Unchanged biological fit evidence from all 450 sweeps. l, All eight common current targets. Twelve panels; one real hardware session, no independent ionic-charge metrology claim. All earlier figures and adverse matched controls are preserved.

**Figure S19. Protected functional learning, full coverage and module attribution.**

a,b, Fixed V30 real waveforms at settings00/15. c,d, All-setting functional discrepancies and V30-V27 paired changes. e,f, Whole-version functional and proper-score comparisons, including V32. g,h, All sixteen V30 intervals and complete-version coverage-width trade-offs. i, V30/V27 functional-median and identical-point V30/V31 interval-score ratios attribute the readout and calibration contributions separately. j, Four current V30 matched point ratios. k, Unchanged biological fits from all450sweeps; matched controls also retain20/24. l, Eight common current targets for complete versions. One hardware session, nominal RC functional, no independently metered ionic-charge claim; all original18figures retained.

**Figure S20 All16 frozen final-model waveforms**

Independent measured Vm, final IonGate and equally enhanced FIR are shown for every setting. Labels identify known compensation fractions. One session; no setting removed.

# Fig. 1 | Physics-constrained functional inference

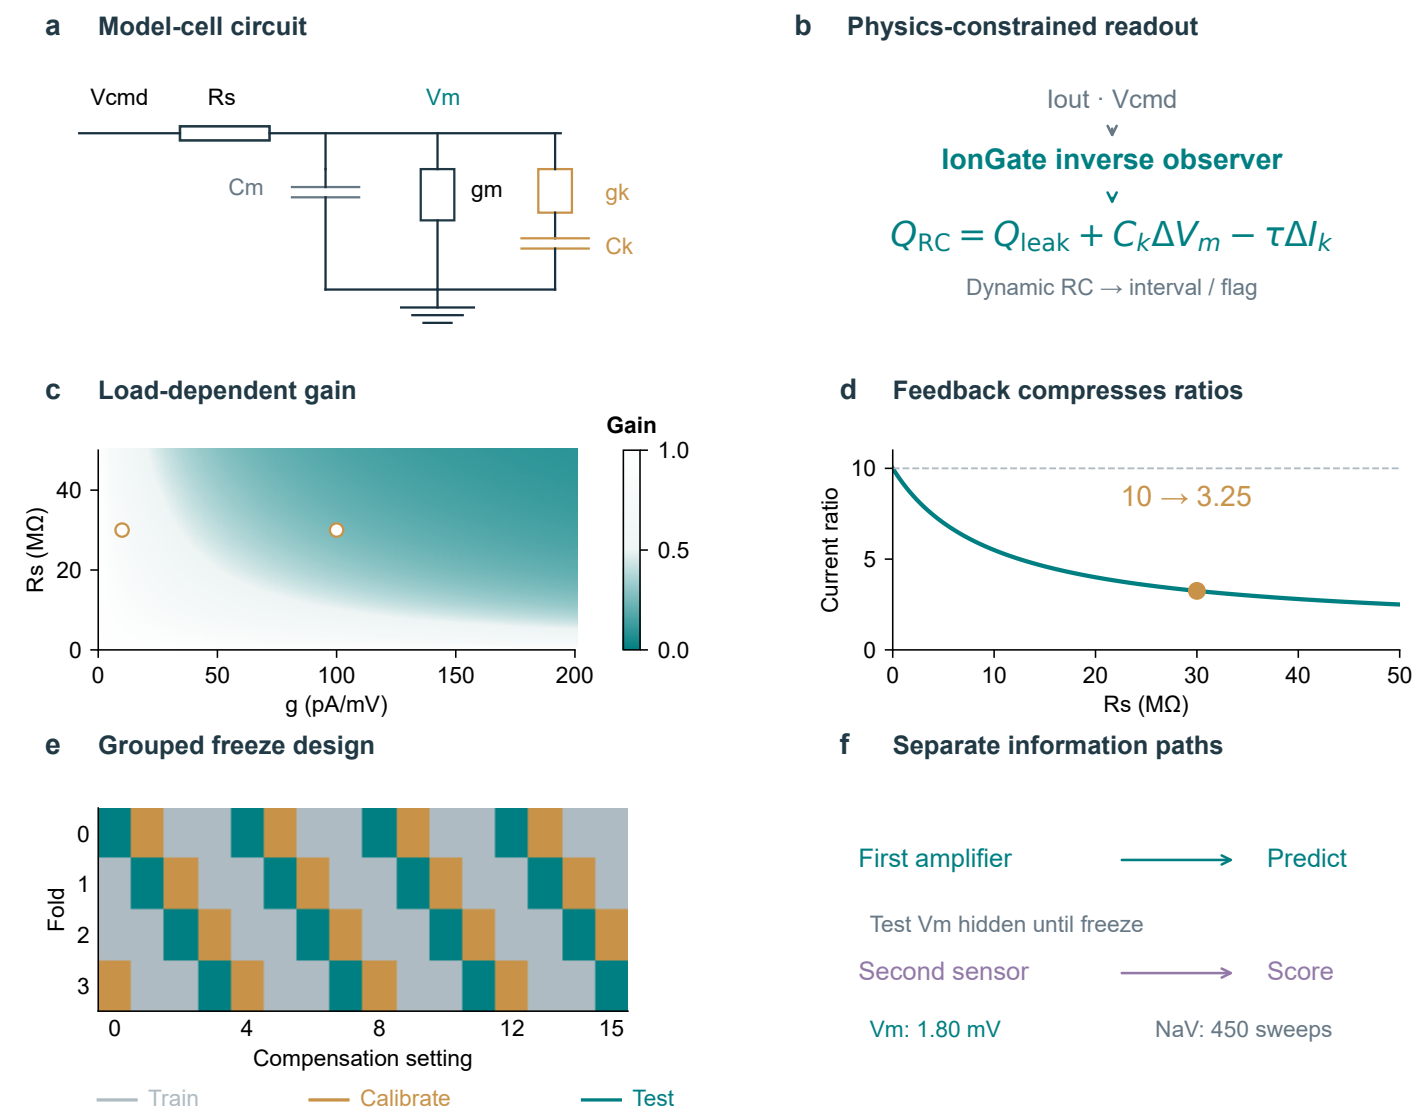

# Fig. 2 | Independent Vm recovery and phase-specific advantage

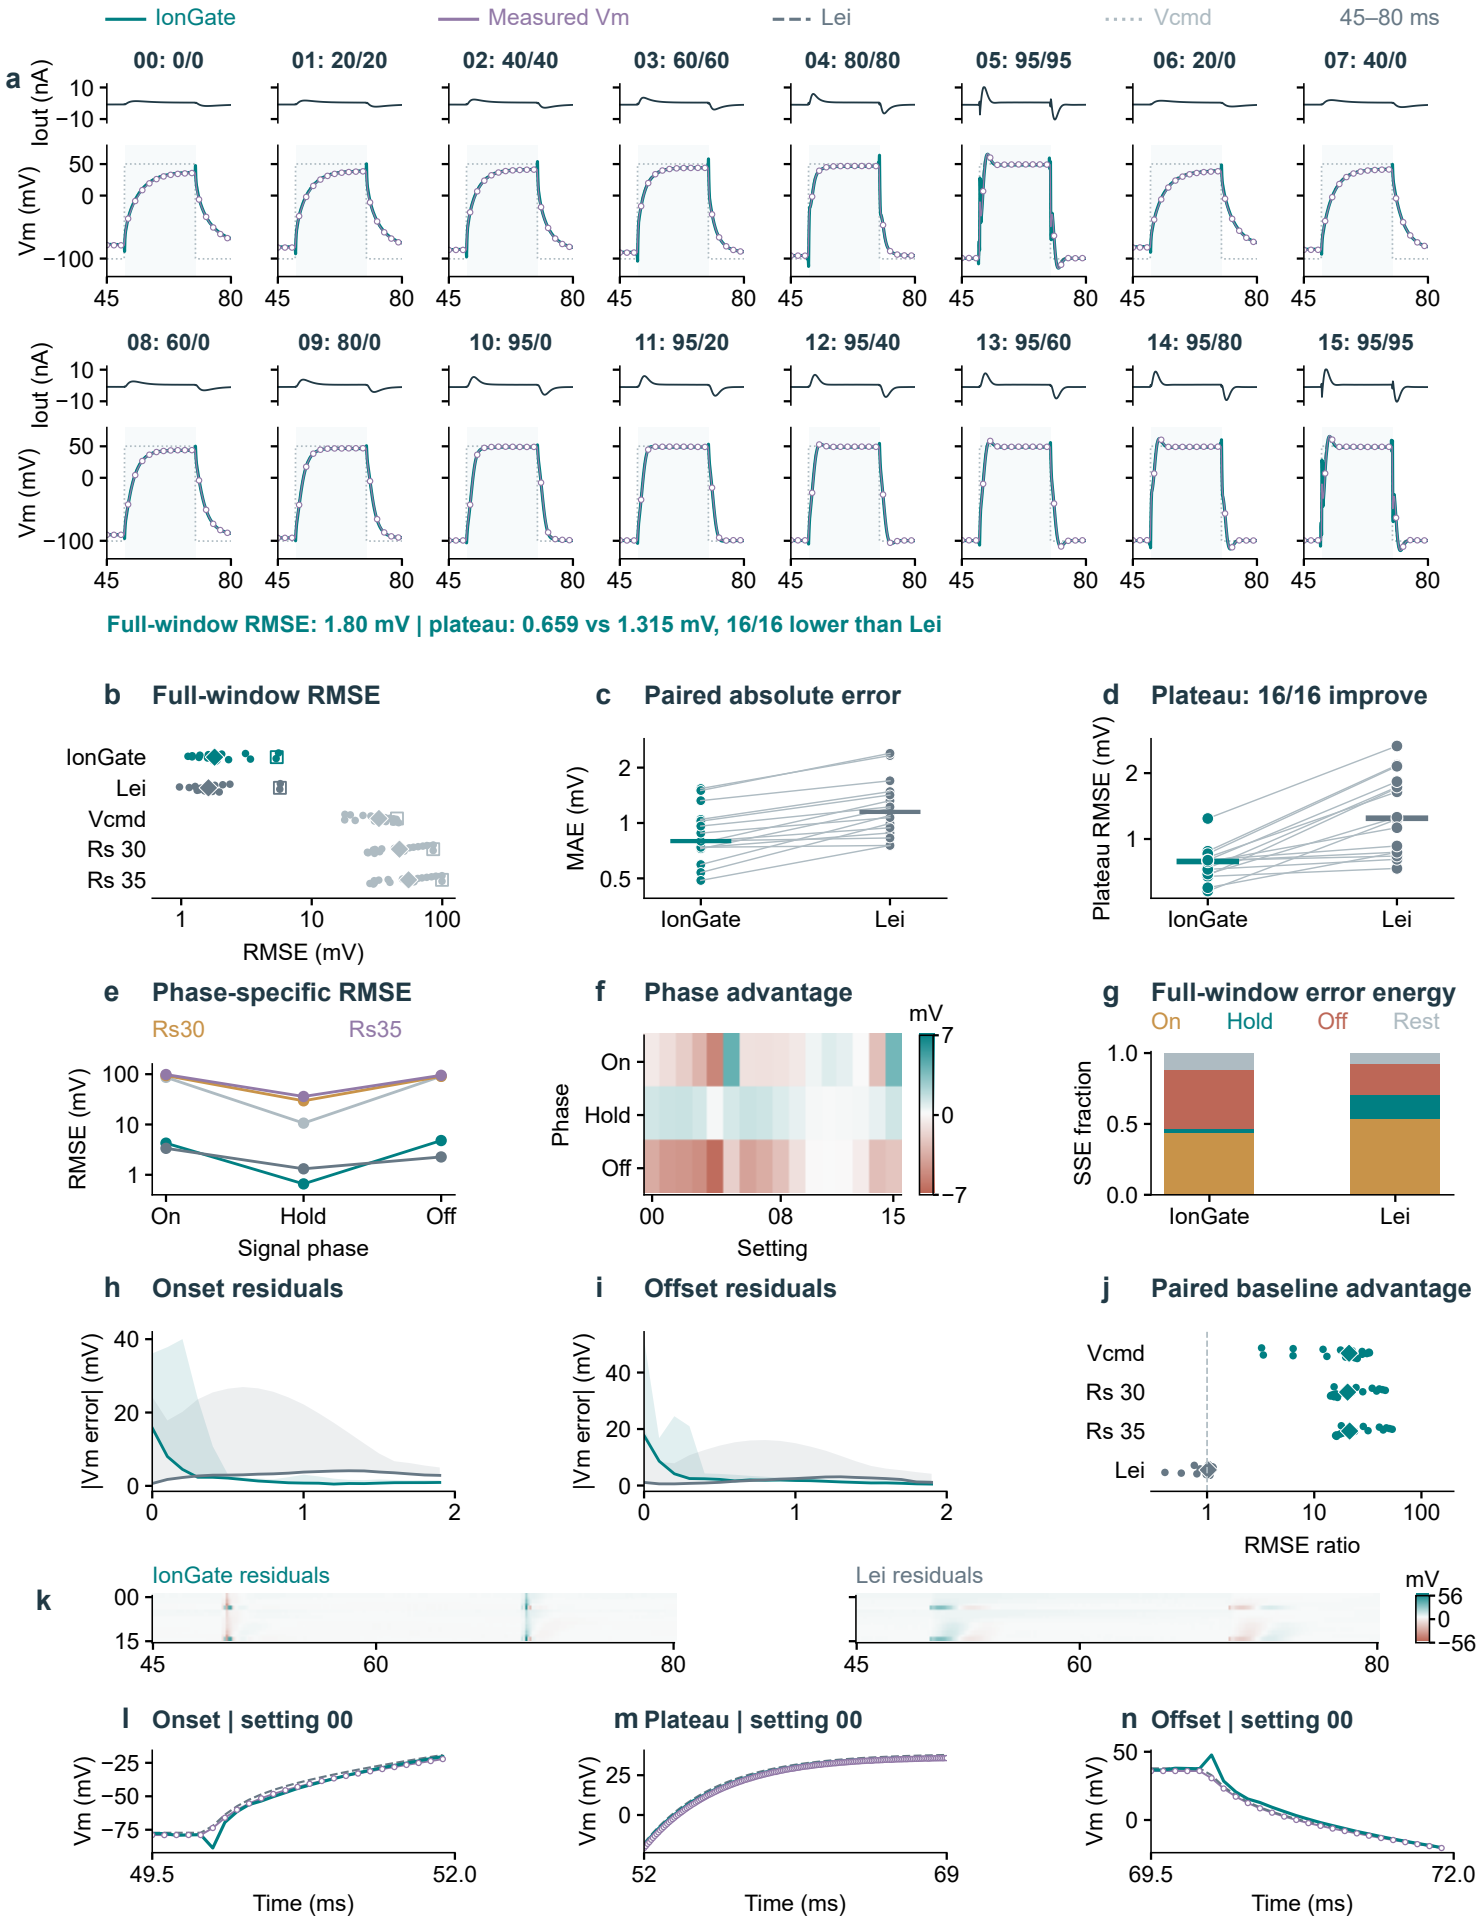

# Fig. 3 | Dynamic RC and state memory explain the IonGate gain

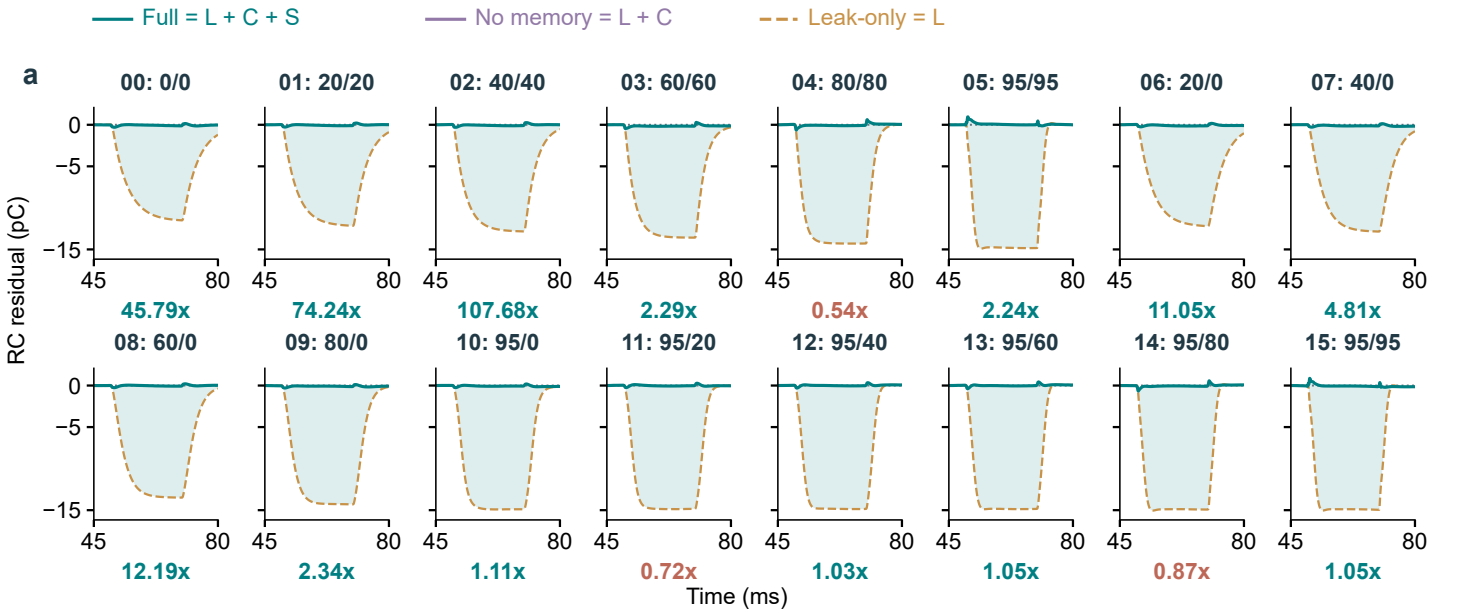


**Matched median errors: full 0.213%, no-memory 0.456%, leak-only 0.551% | 16 settings**

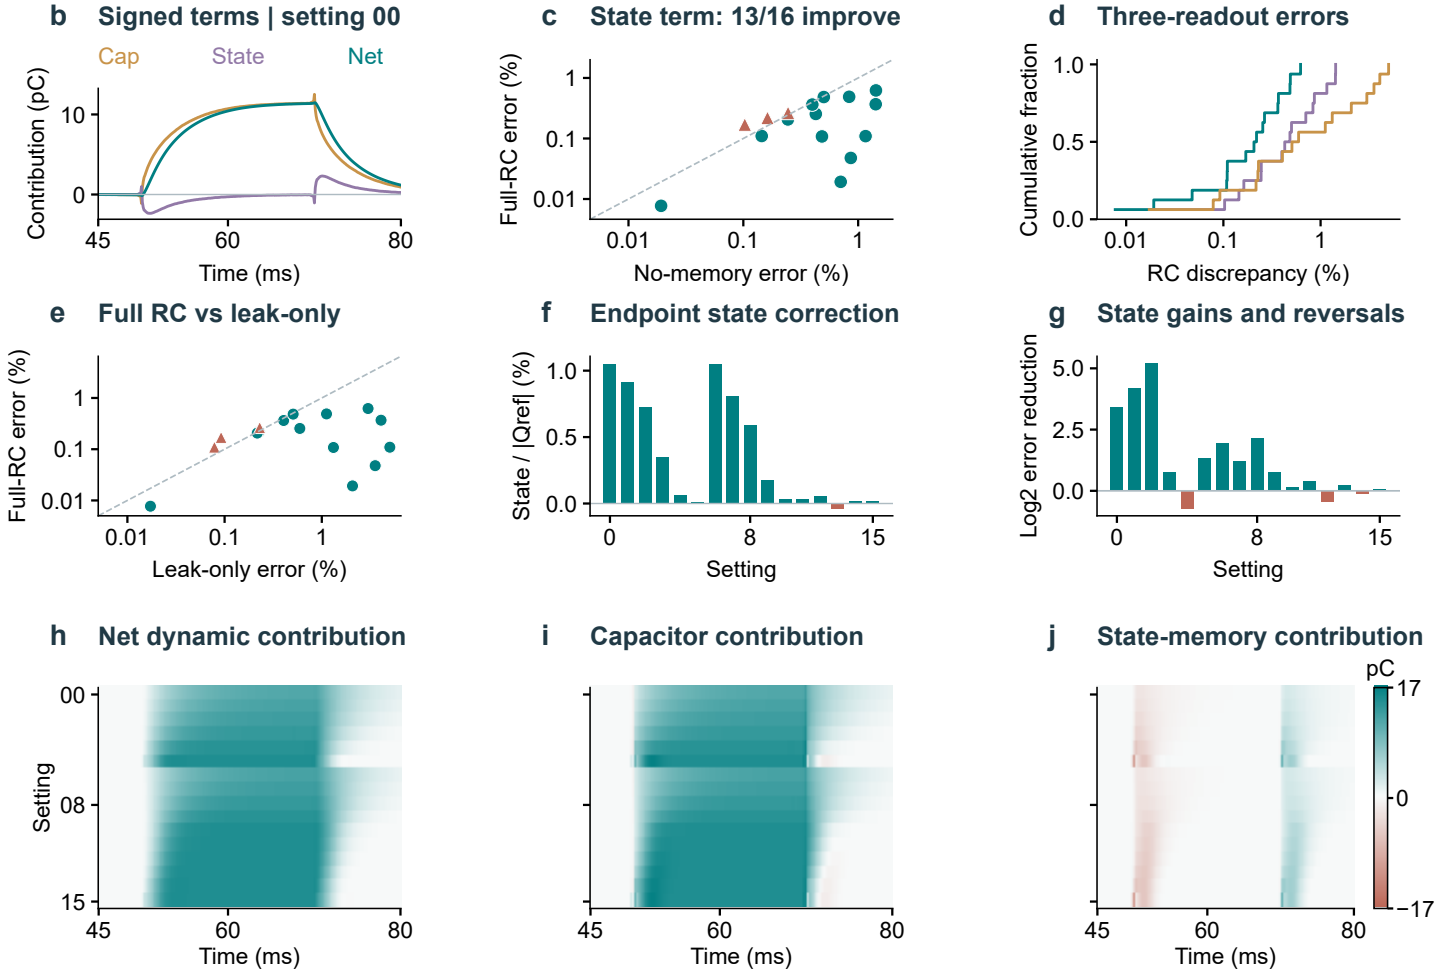


Same frozen observer and state | all reversals retained | conditional RC reference, one session

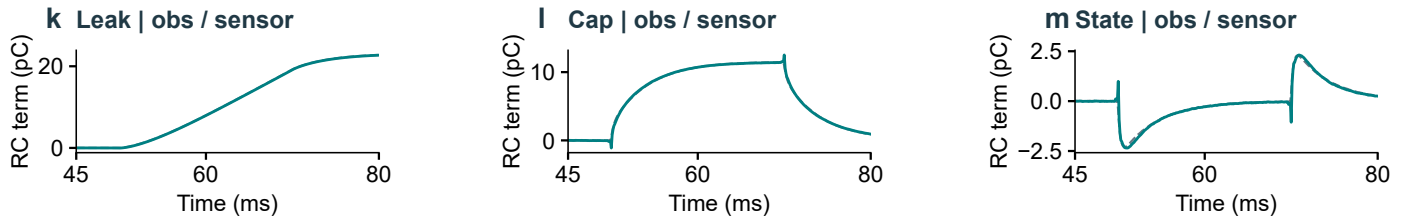

# Fig. 4 | Frozen intervals: reliability, utility and acceptance

Same-fold 80% intervals | 16 settings, one session | interval score: 1.131% vs 1.316%

IonGate  Lei  Miss

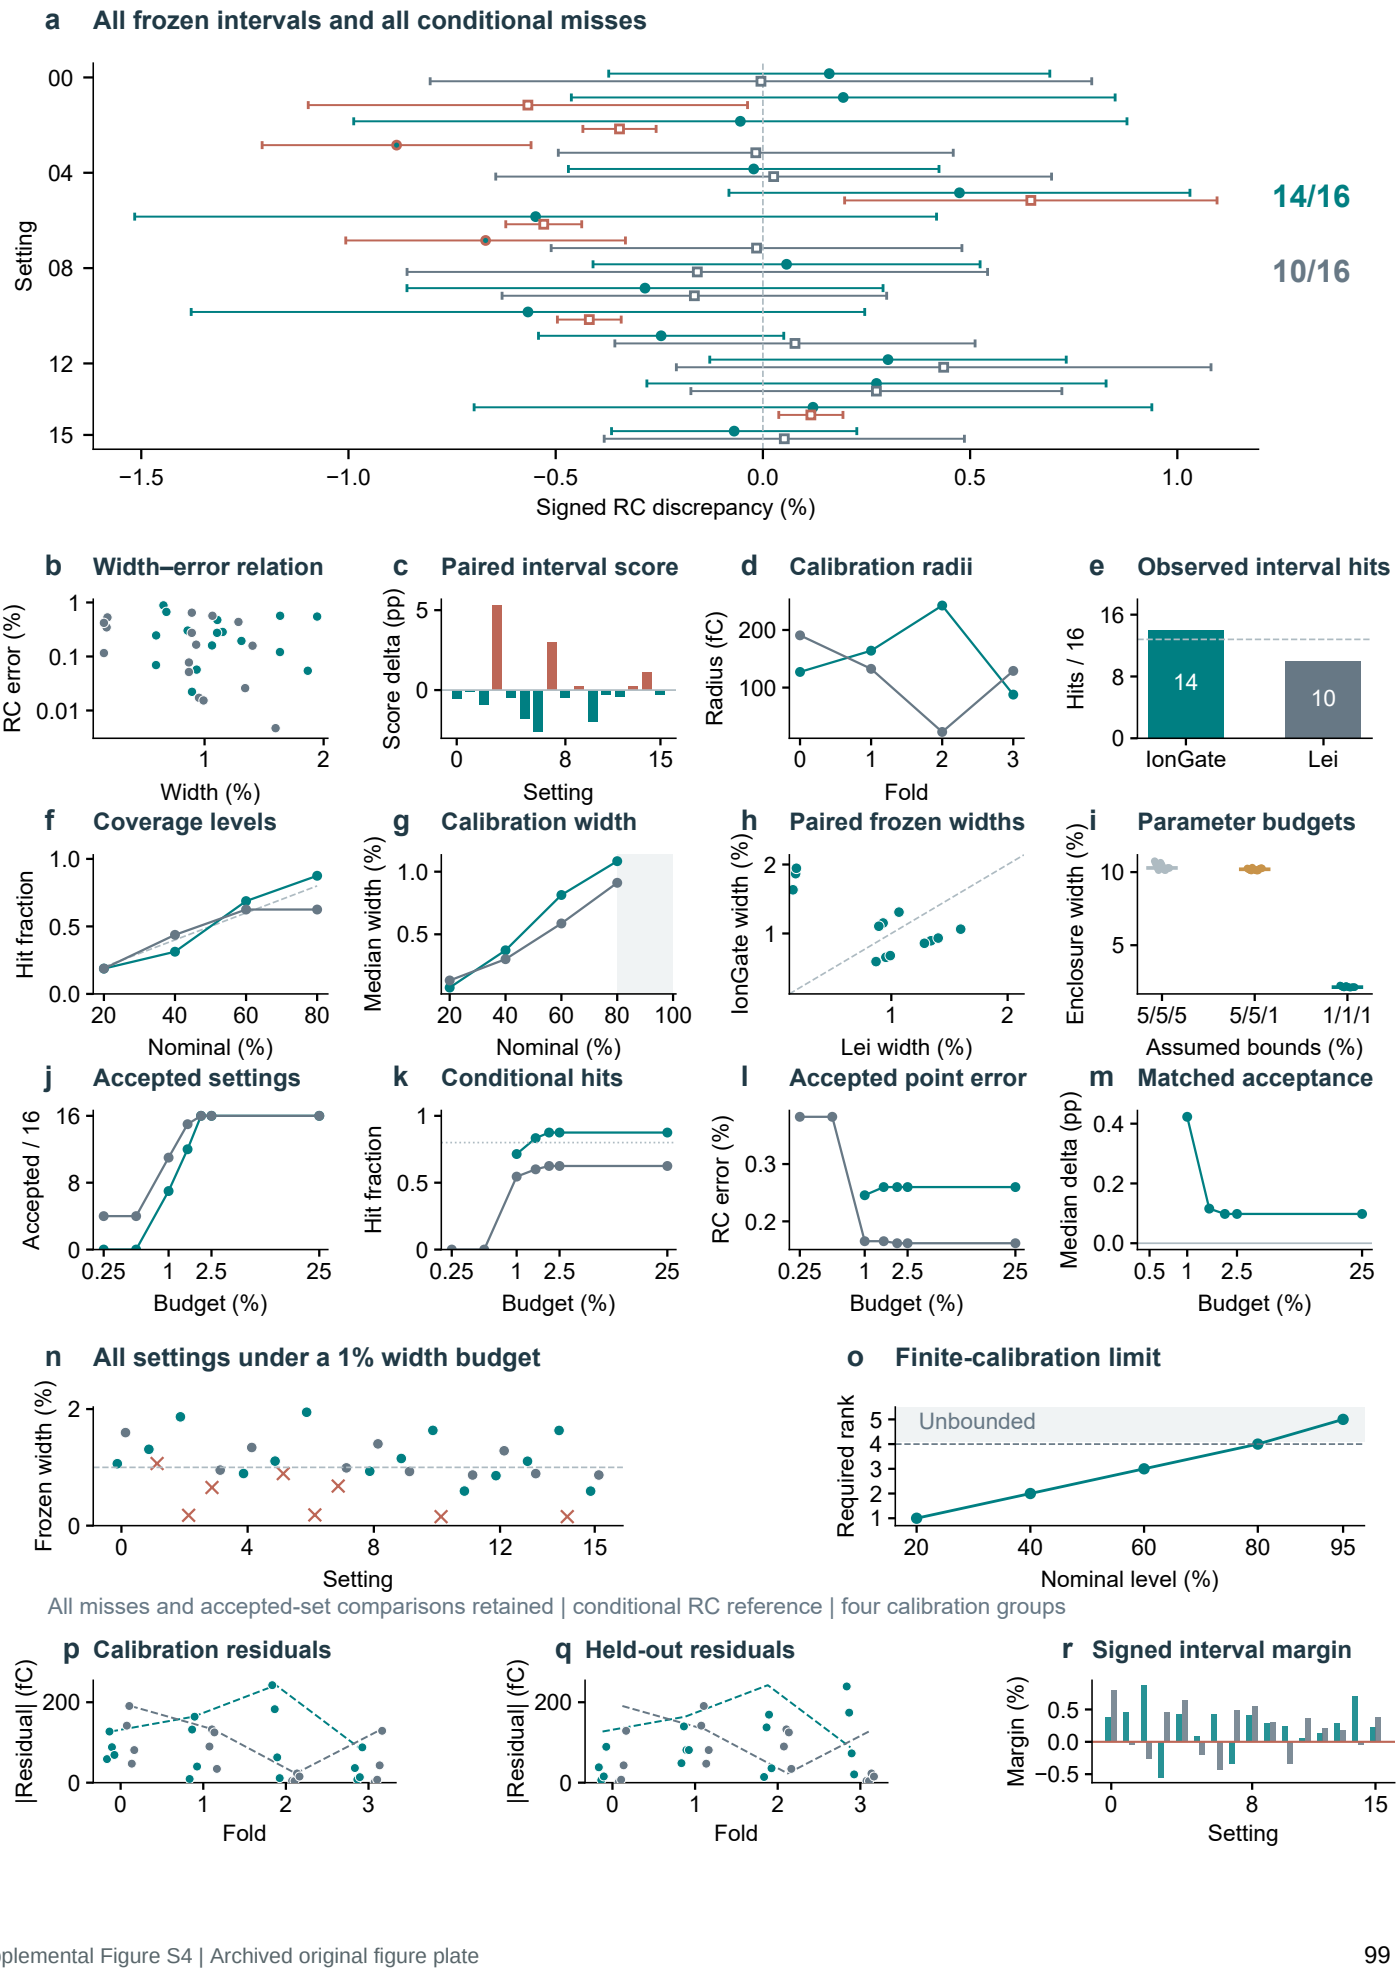

# Fig. 5 | Native hardware noise and switching challenge

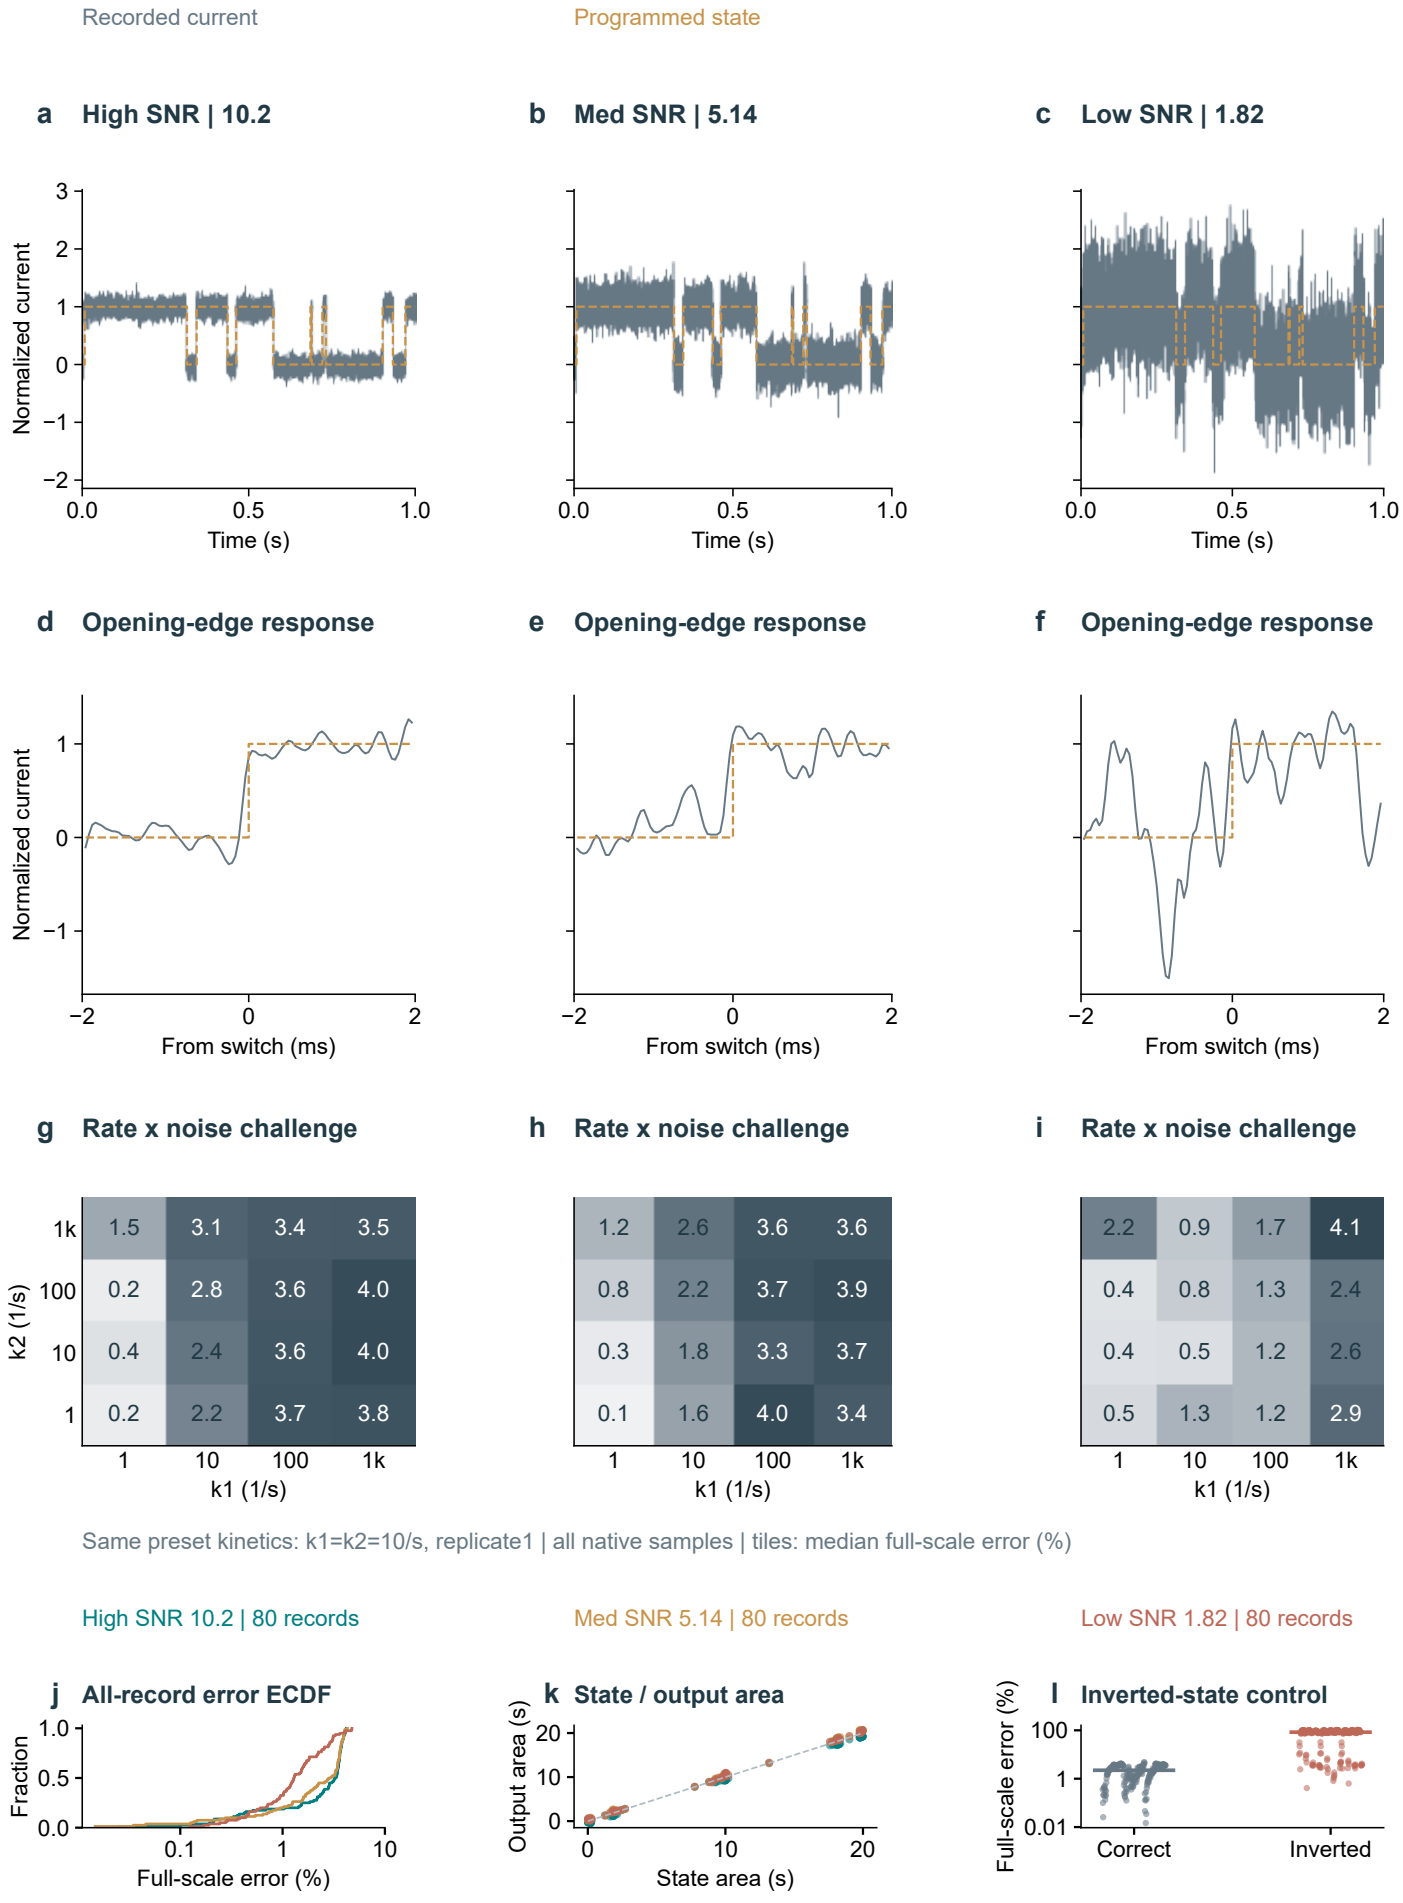

# Fig. 6 | Real imaging and frozen IonGate transfer

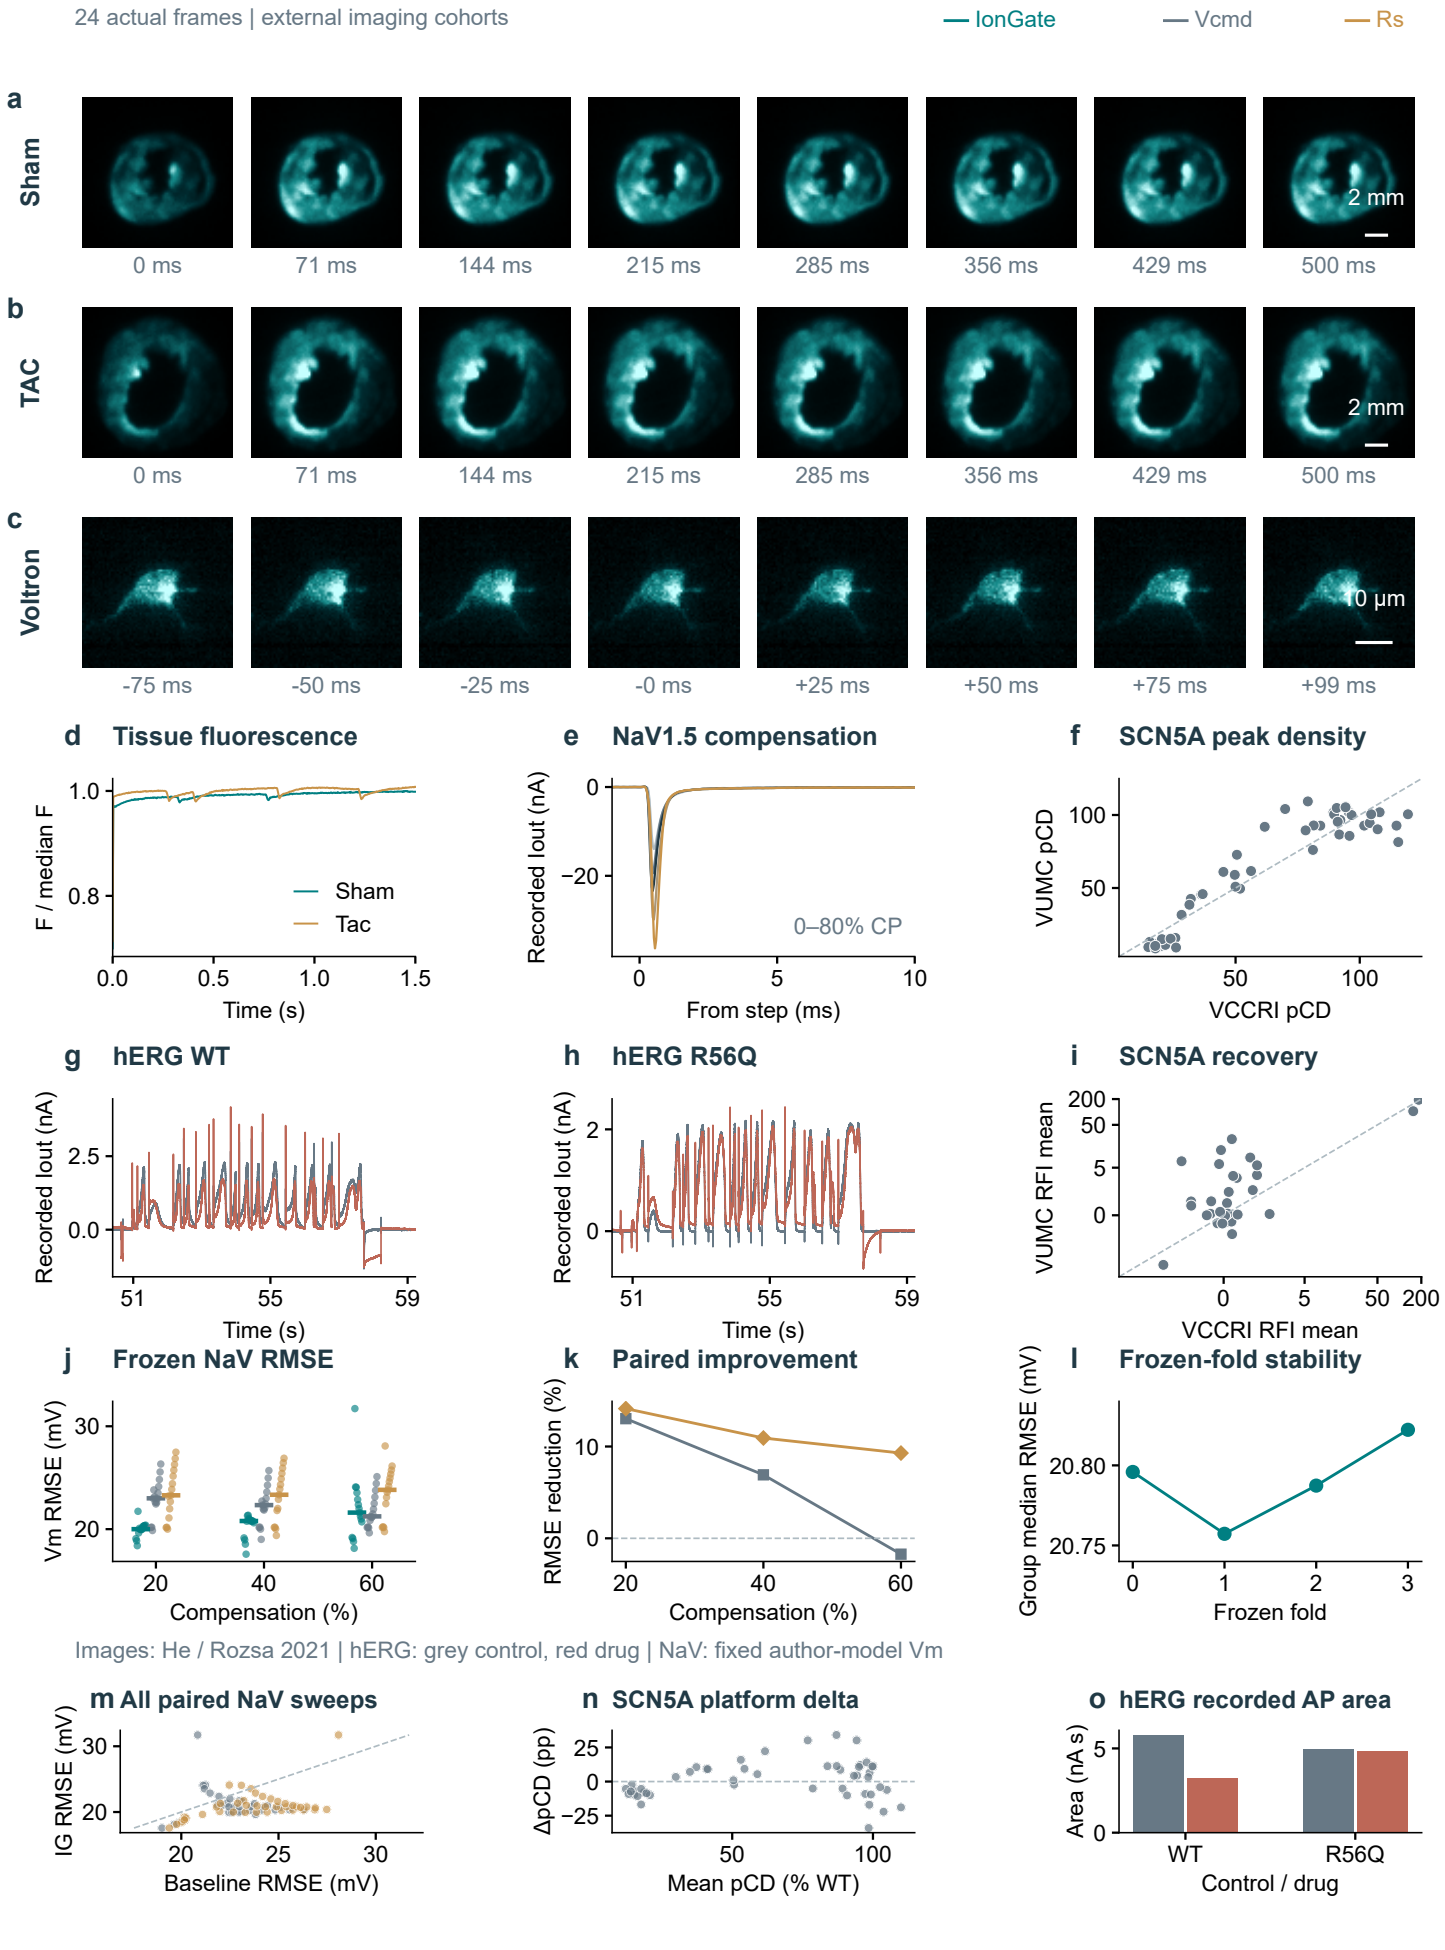

# Fig. 7 | Seeing voltage and recording it independently

Real Voltron images • simultaneous whole-cell electrode • one in-vivo neuron

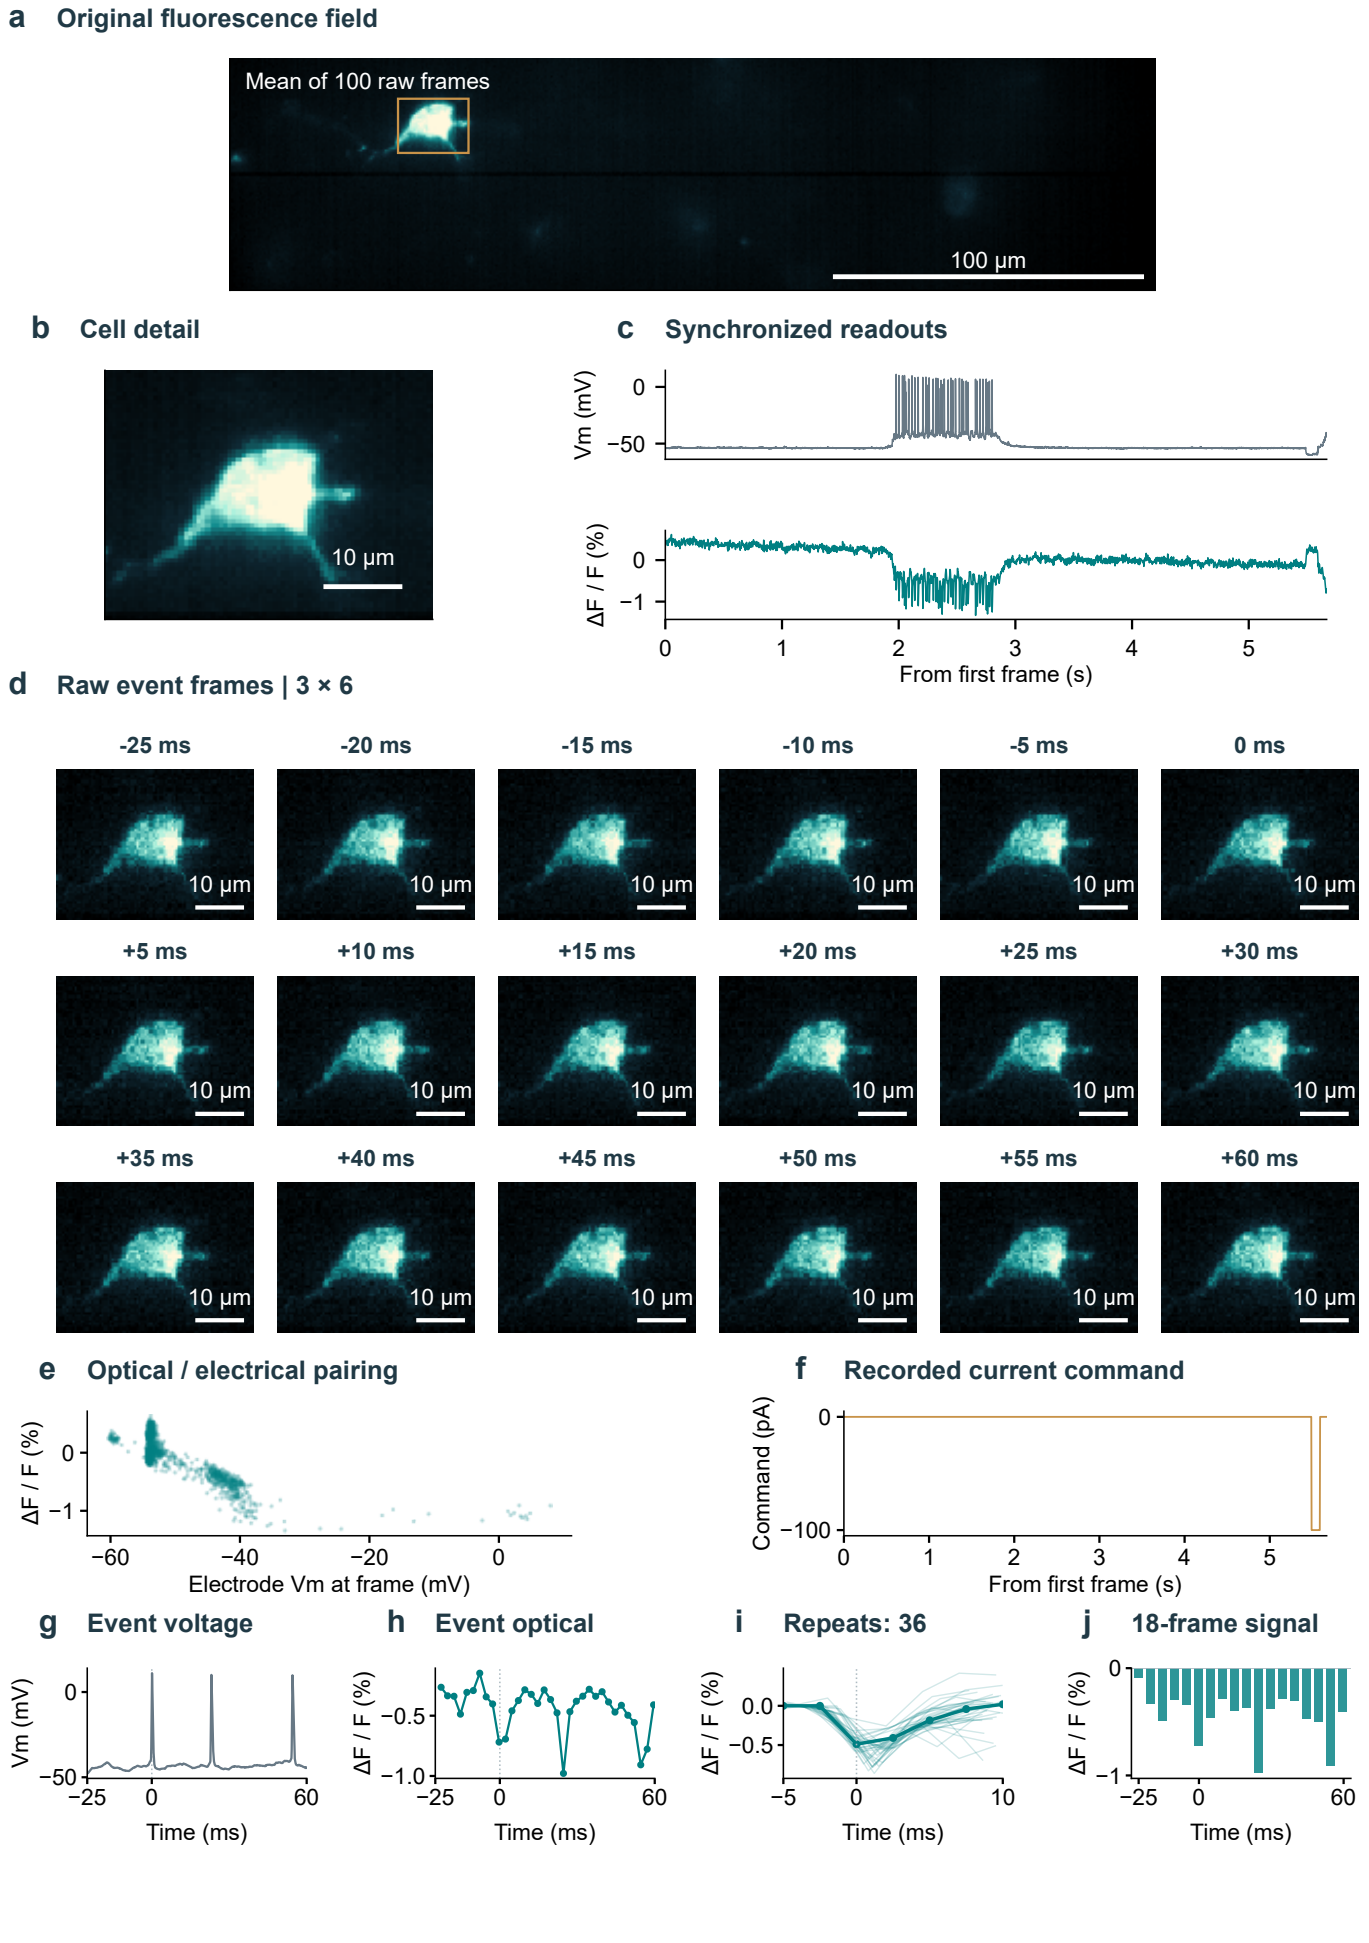

# Fig. 8 | Fair baselines separate voltage and function

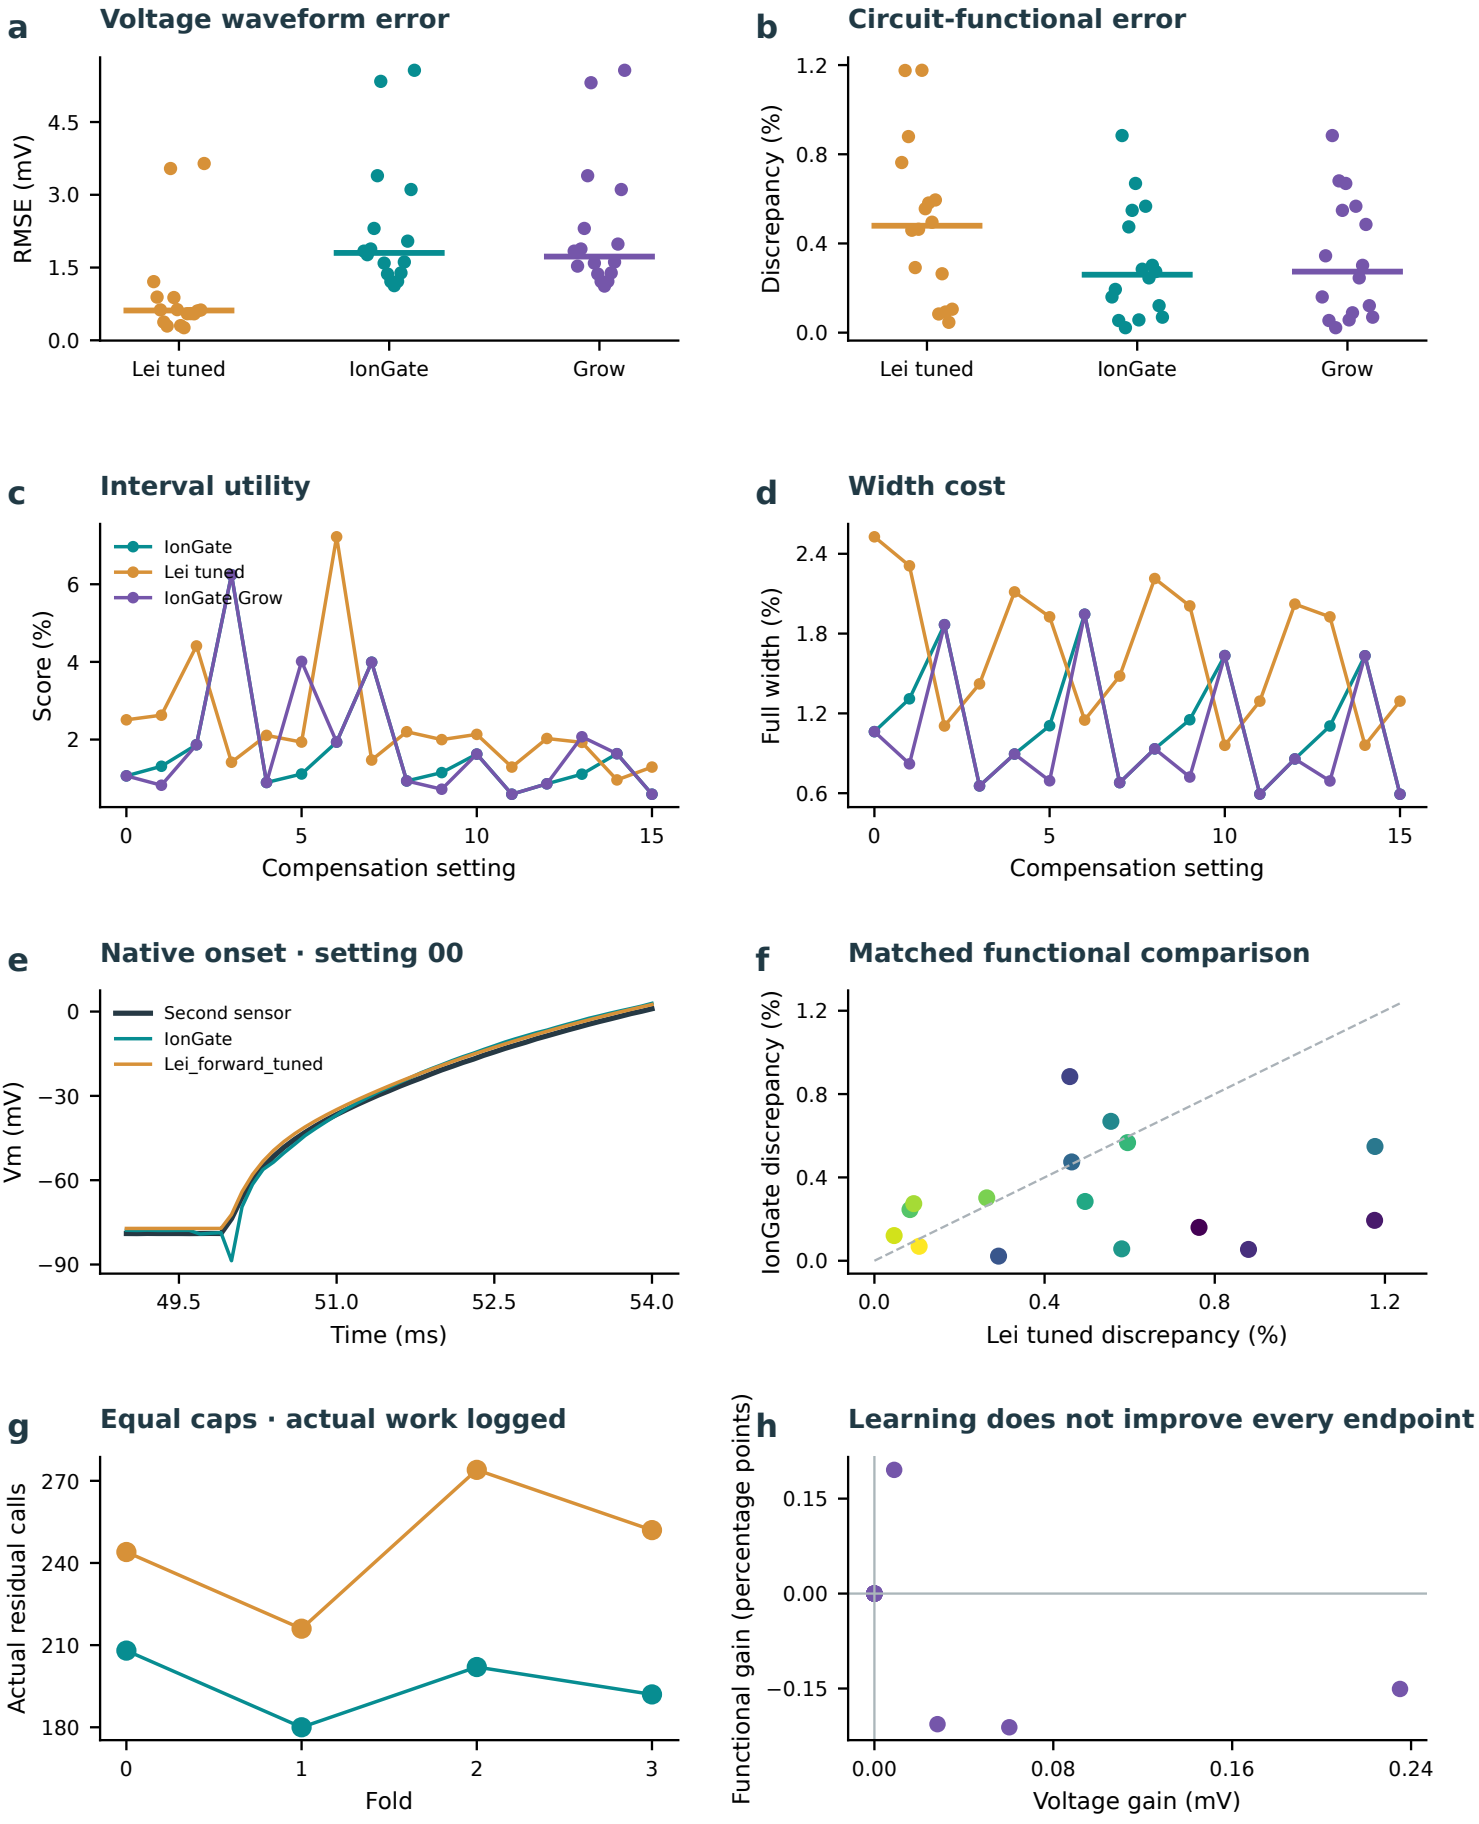

# Fig. 9 | Compatibility checks and reversible growth

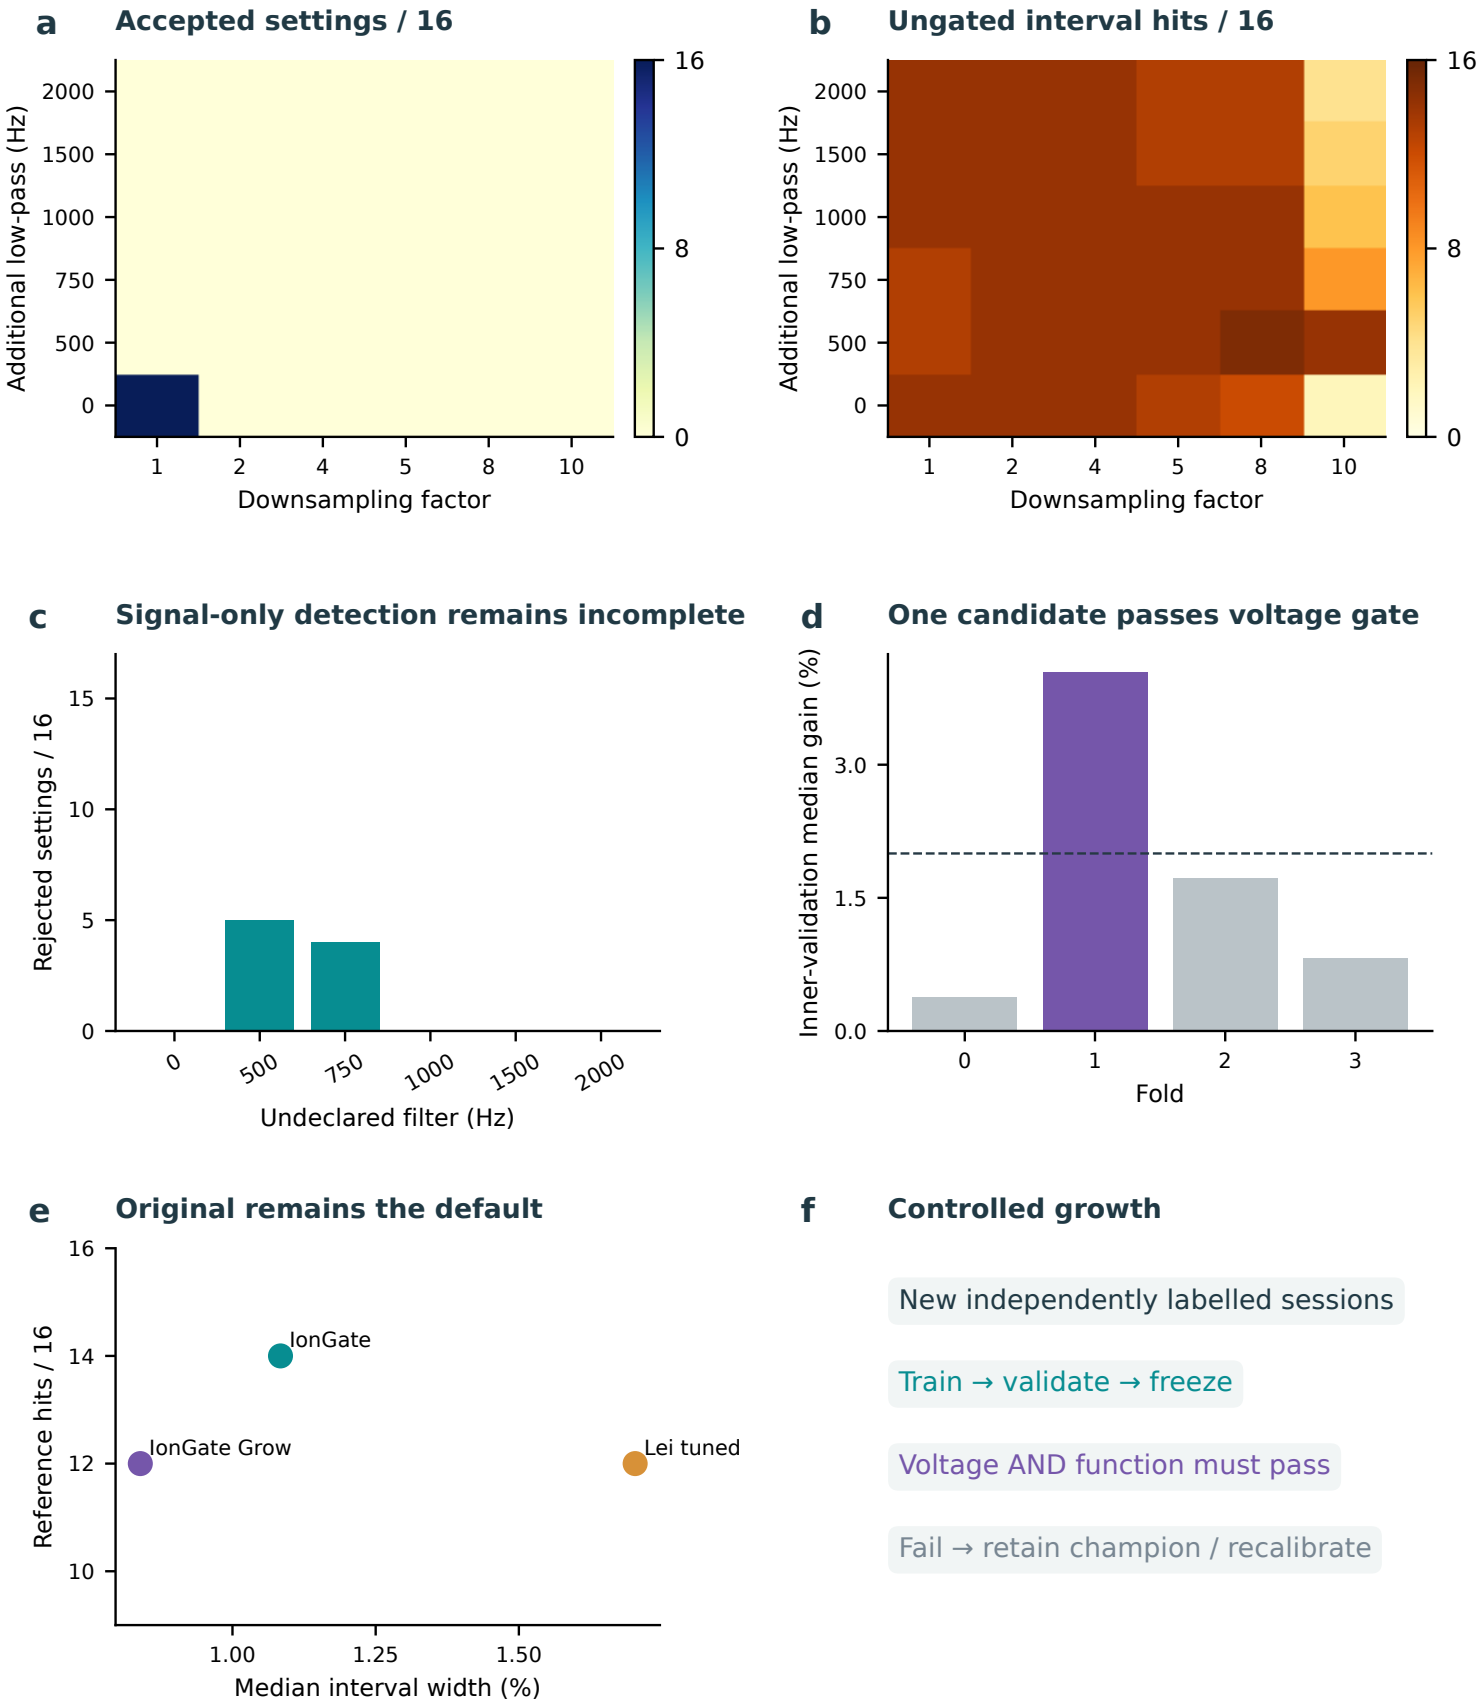

# Fig. 10 | Biological function with fit-failure accounting

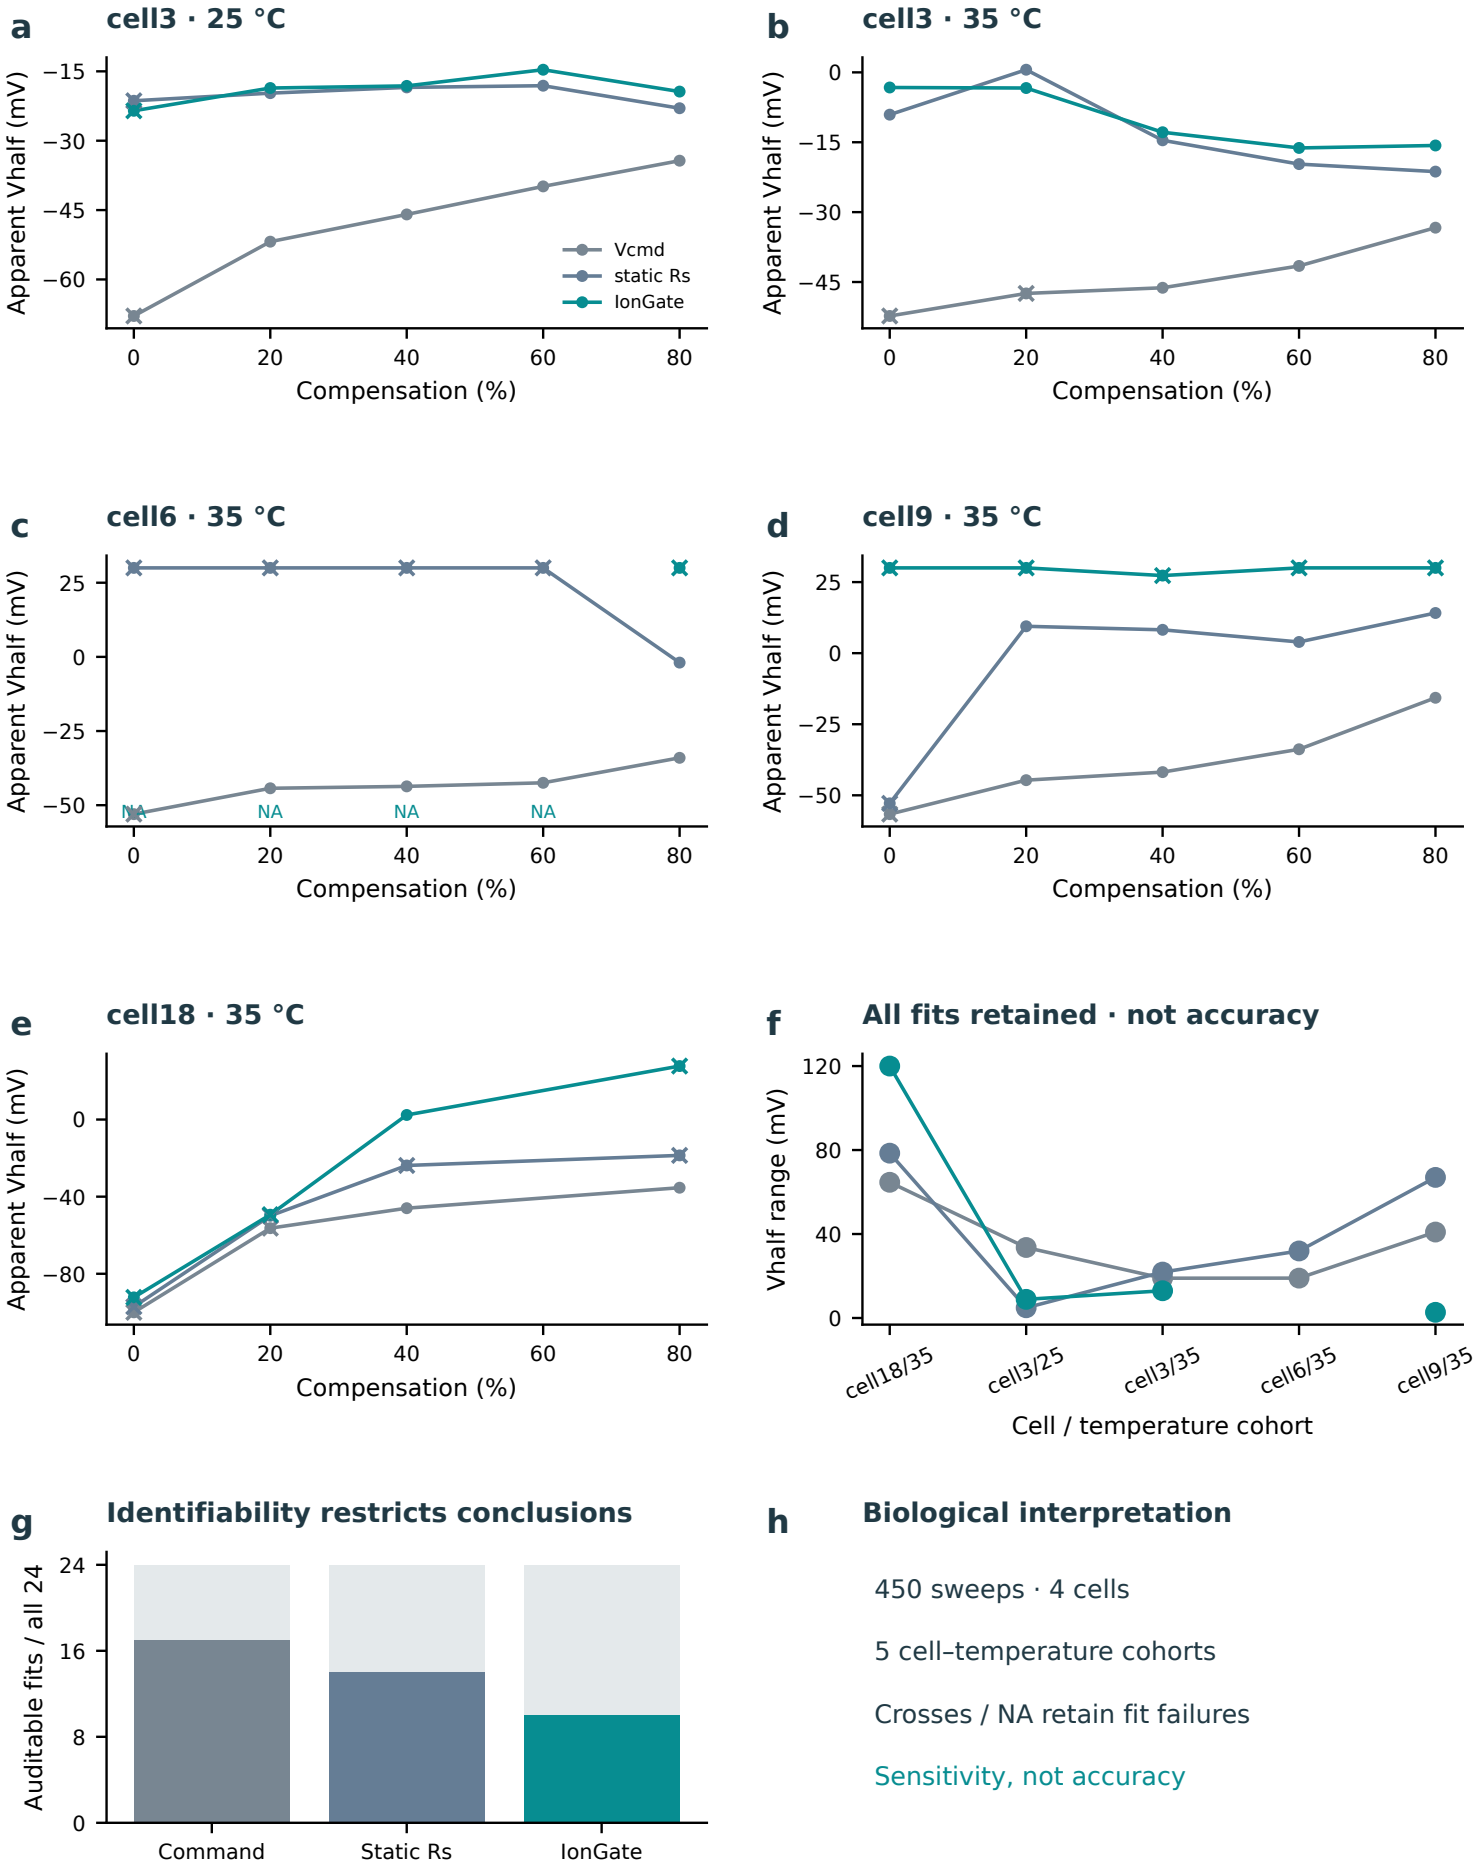

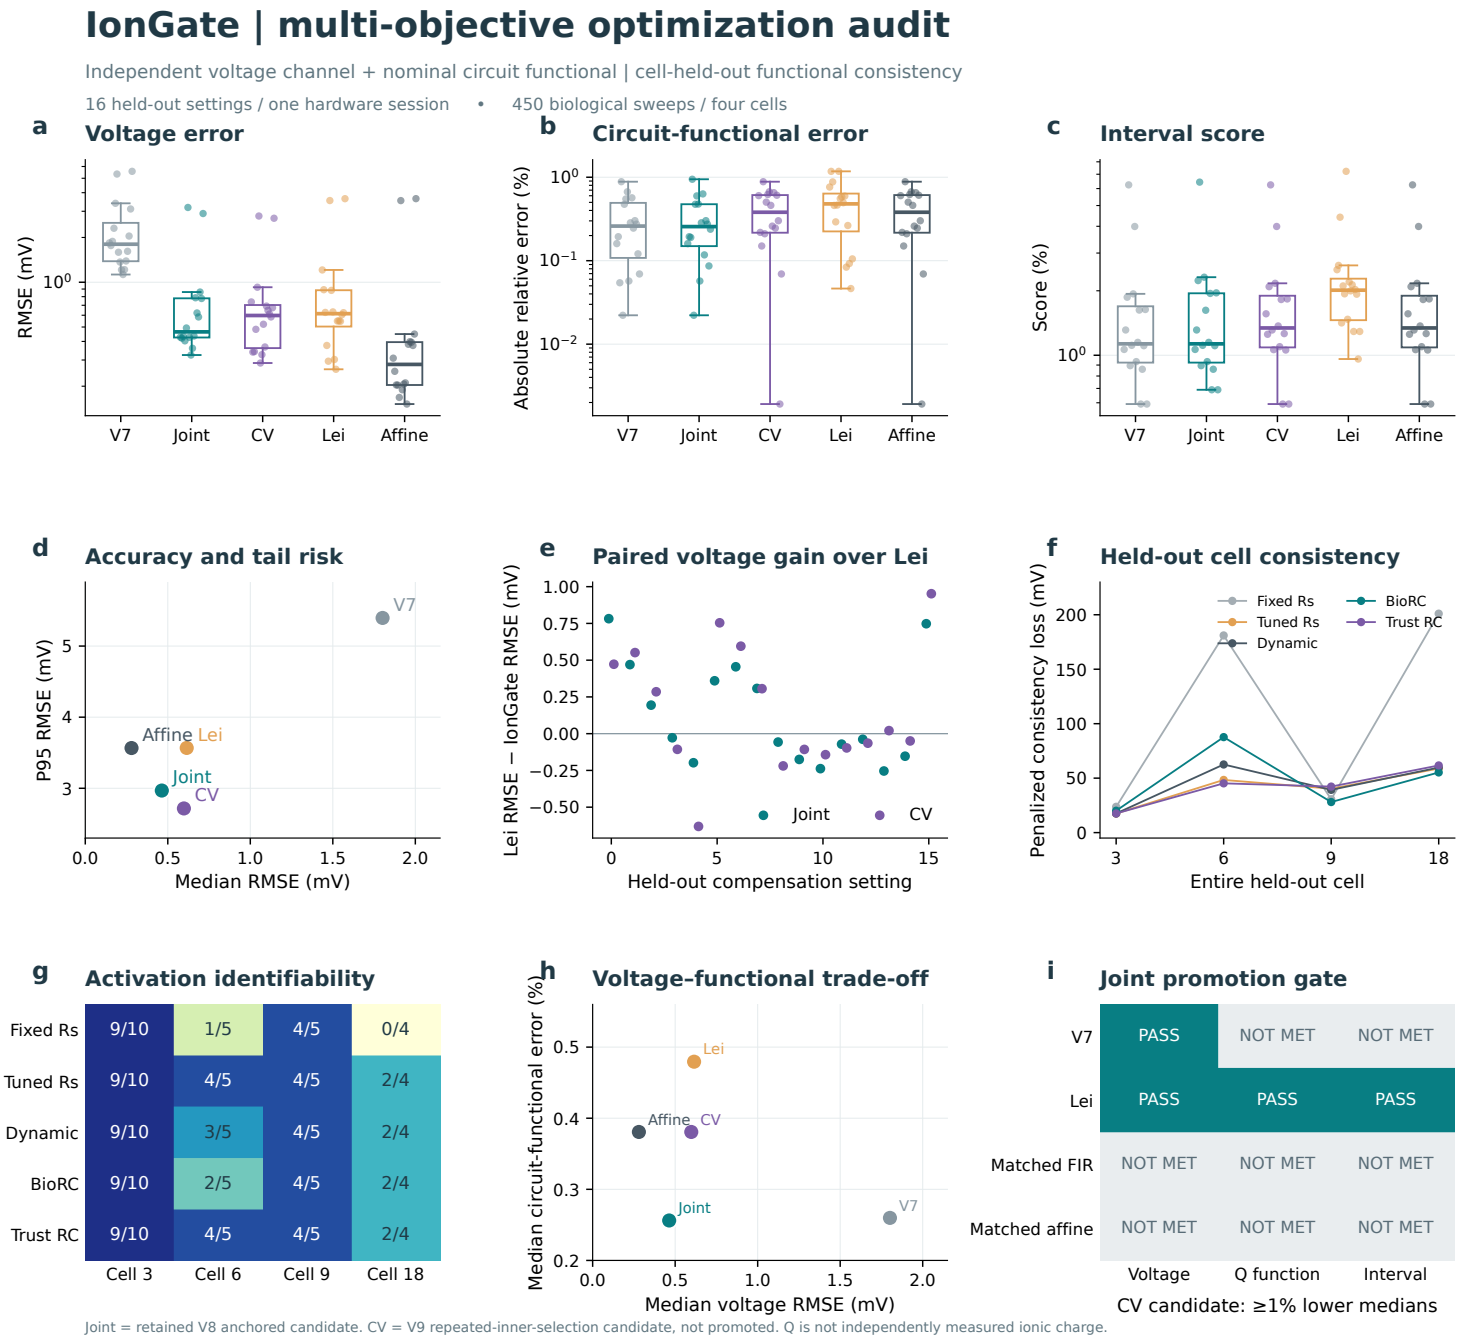
IonGate | multi-objective optimization audit
Independent voltage channel + nominal circuit functional | cell-held-out functional consistency
16 held-out settings / one hardware session • 450 biological sweeps / four cells
a
Voltage error
RMSE (mV)
V7
Joint
CV
Lei
Affine
b
Circuit-functional error
Absolute relative error (%)
c
Interval score
Score (%)
d
Accuracy and tail risk
P95 RMSE (mV)
Median RMSE (mV)
e
Paired voltage gain over Lei
Lei RMSE − IonGate RMSE (mV)
Held-out compensation setting
f
Held-out cell consistency
Penalized consistency loss (mV)
Fixed Rs
Tuned Rs
Dynamic
BioRC
Trust RC
Entire held-out cell
g
Activation identifiability
Fixed Rs 9/10 1/5 4/5 0/4
Tuned Rs 9/10 4/5 4/5 2/4
Dynamic 9/10 3/5 4/5 2/4
BioRC 9/10 2/5 4/5 2/4
Trust RC 9/10 4/5 4/5 2/4
Cell 3
Cell 6
Cell 9
Cell 18
h
Voltage–functional trade-off
Median circuit-functional error (%)
Median voltage RMSE (mV)
i
Joint promotion gate
V7 PASS NOT MET NOT MET
Lei PASS PASS PASS
Matched FIR NOT MET NOT MET NOT MET
Matched affine NOT MET NOT MET NOT MET
Voltage
Q function
Interval
CV candidate: ≥1% lower medians
Joint = retained V8 anchored candidate. CV = V9 repeated-inner-selection candidate, not promoted. Q is not independently measured ionic charge.

# IonGate RC-State | one-version evaluation

Conservation-driven memory • multi-stimulus training • matched affine and 101-tap FIR controls

16 held-out settings / one hardware session | 450 biological sweeps / four held-out cells

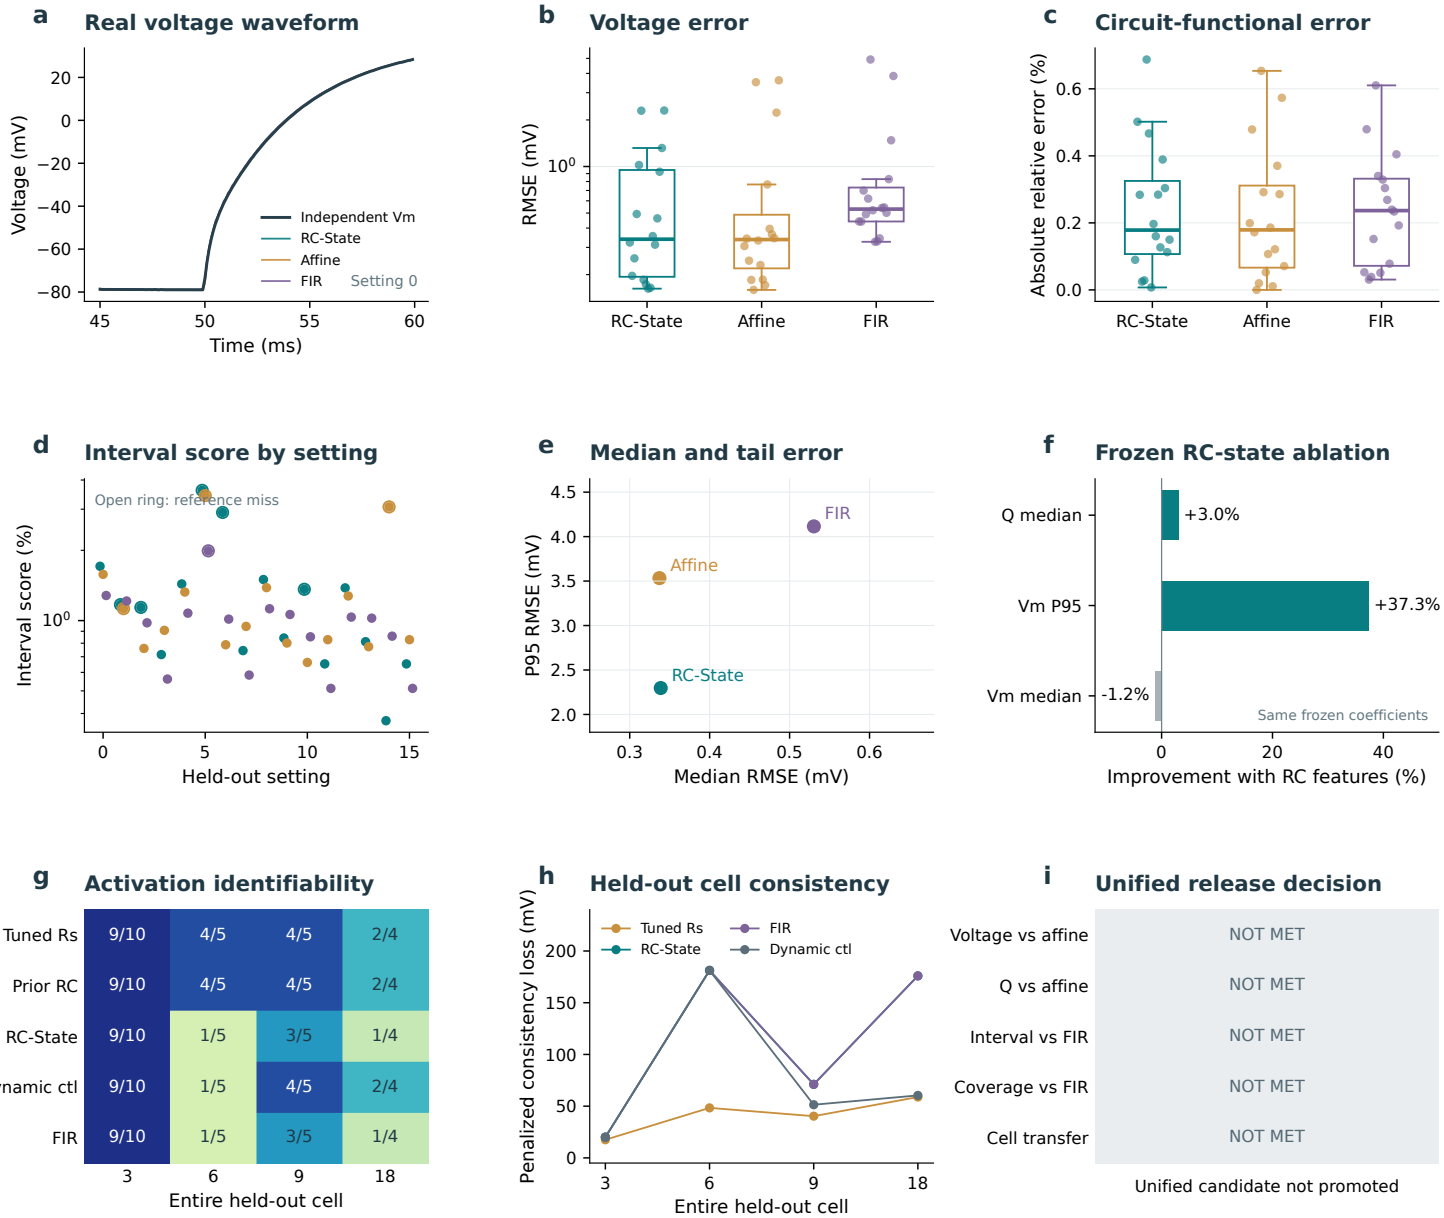


RC-State = orthogonal multi-stimulus head + FOH biological adapter. All endpoints refer to this one version; historical best values are not substituted.
Q is a nominal circuit functional. Biological panels assess consistency/identifiability, not independently measured biological voltage or ionic charge.

# IonGate | conditional residuals and circuit memory

One output waveform / first-channel-only inference / equal-budget affine and FIR controls

16 hardware settings from one session • 450 NaV1.5 sweeps from four cells • complete V22 results, including failures

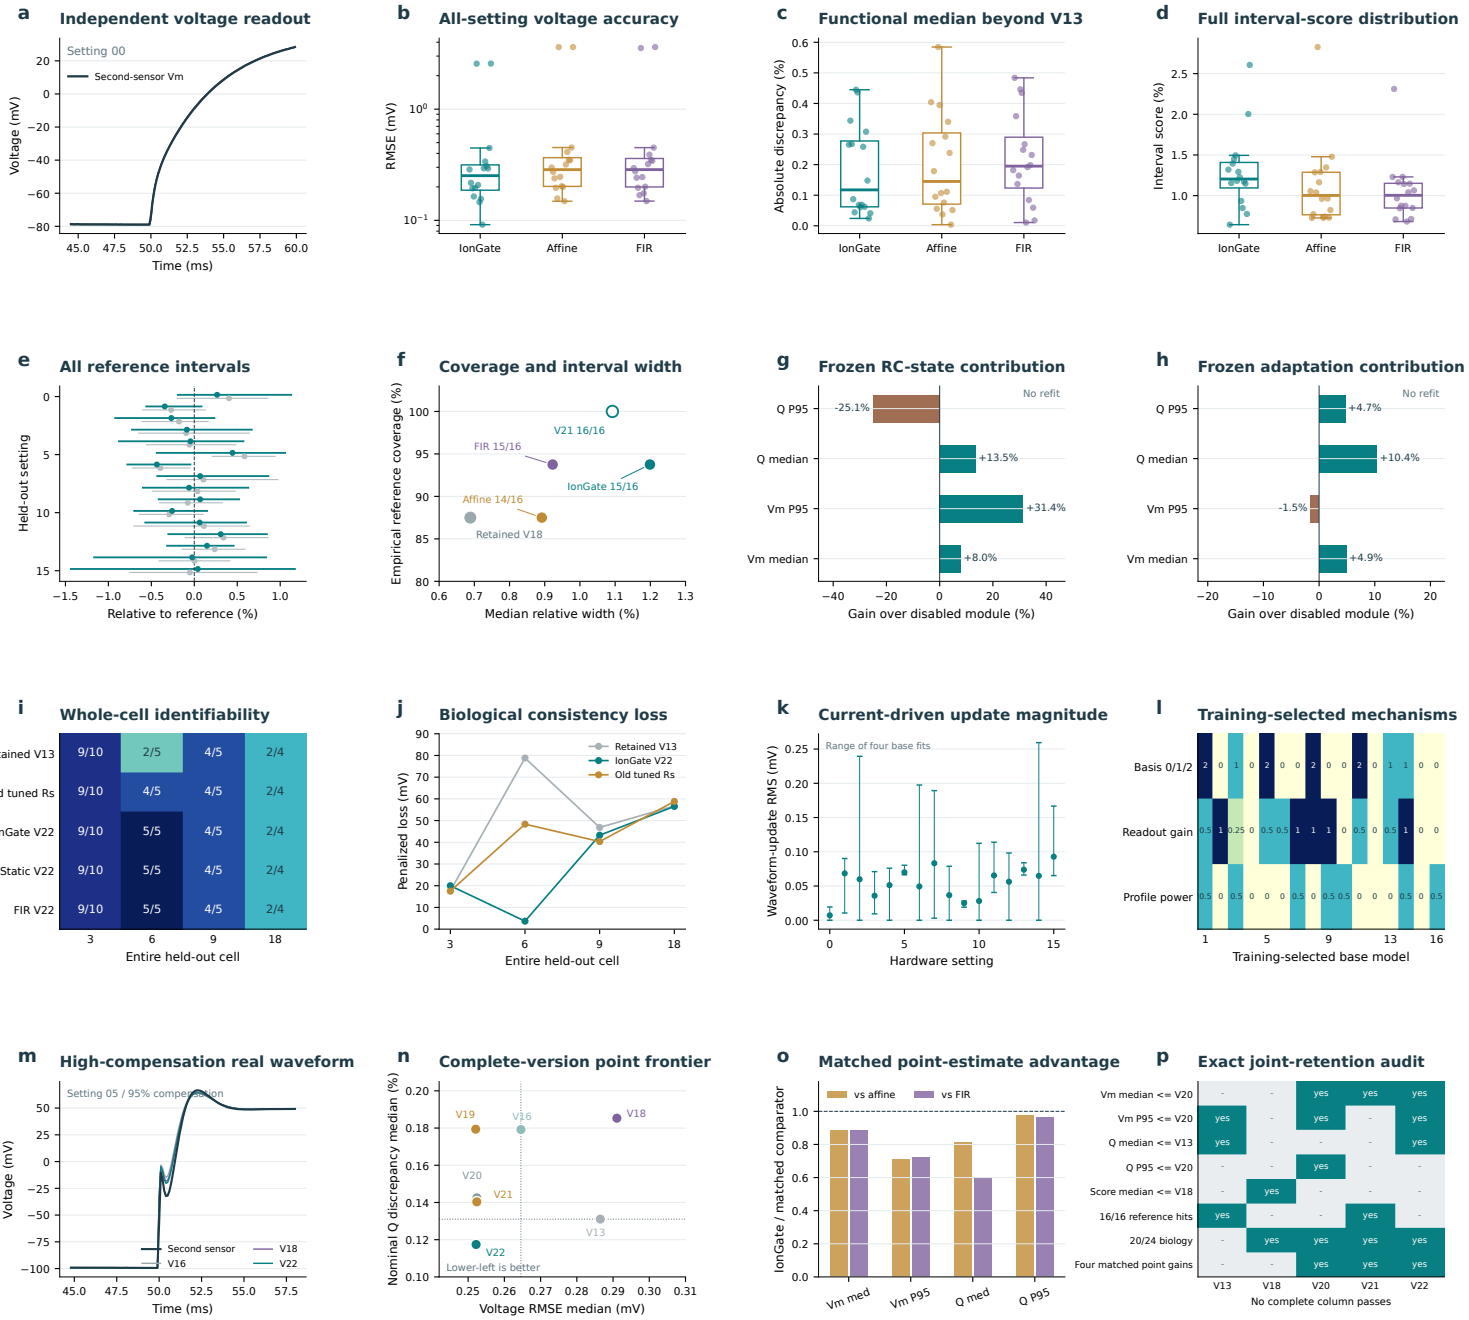


Complete V22: functional median improves beyond V13; four point summaries beat current controls. V21 restores 16/16, but neither closes the interval-score target.

One-session retrospective development. Q is a nominal RC functional, not metered ionic charge. Four-fit ranges are not independent experimental confidence intervals.

## All 24 activation conditions | 20 identifiable, four retained failures

450 real NaV1.5 sweeps • whole-cell LOCO • same peak window, voltage support and identification rules

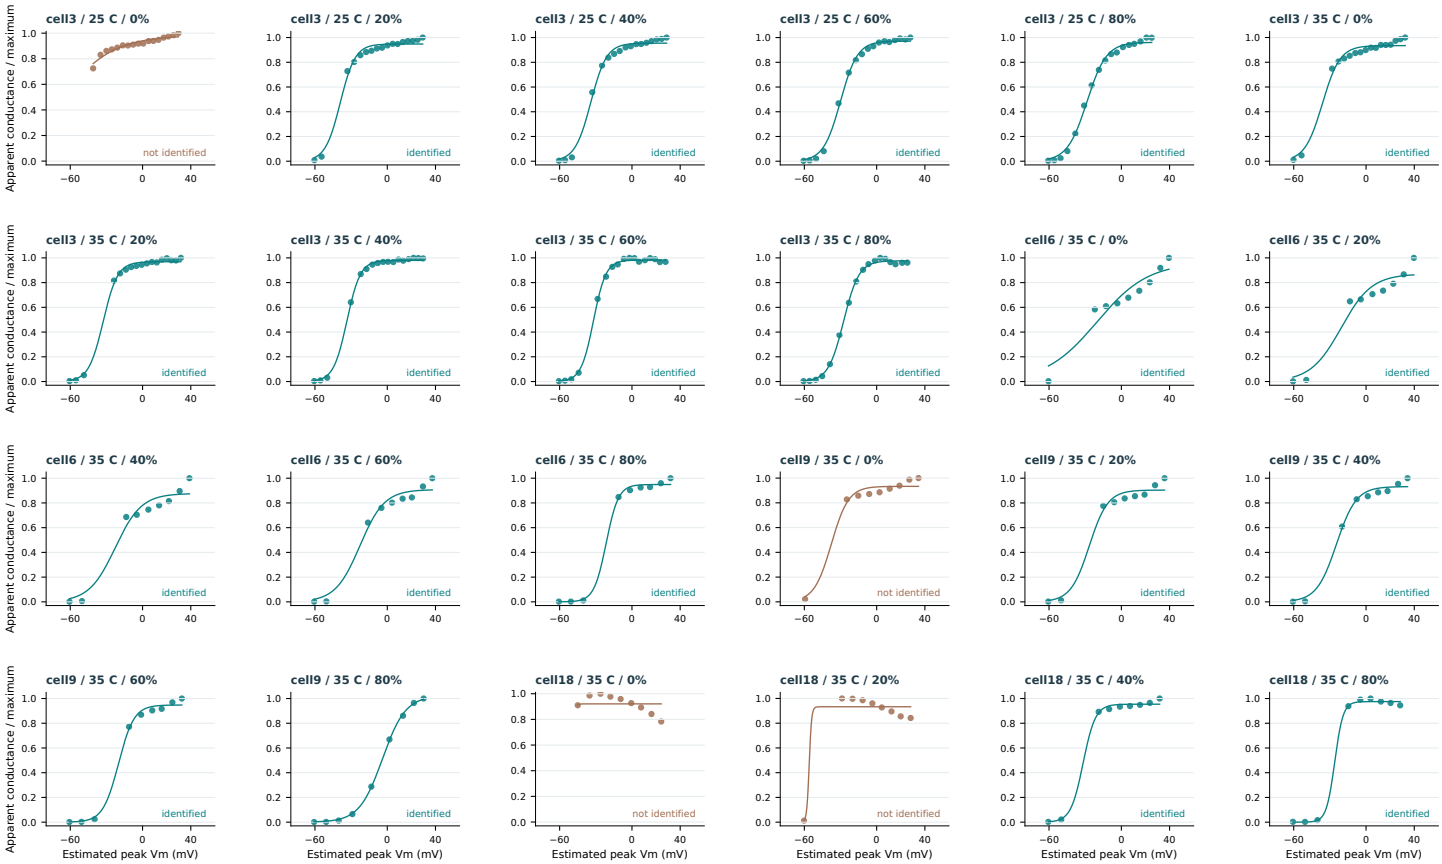


Points derive from real first-amplifier currents. Lines are per-condition Boltzmann fits. Brown panels retain all failed/boundary outcomes; no success-based exclusions.

# IonGate | tail recovery and interval calibration

Frozen whole versions, not a collage of endpoint winners

16 real hardware settings / one session • 450 biological sweeps / four cells • V22 retained; V23/V24 not promoted

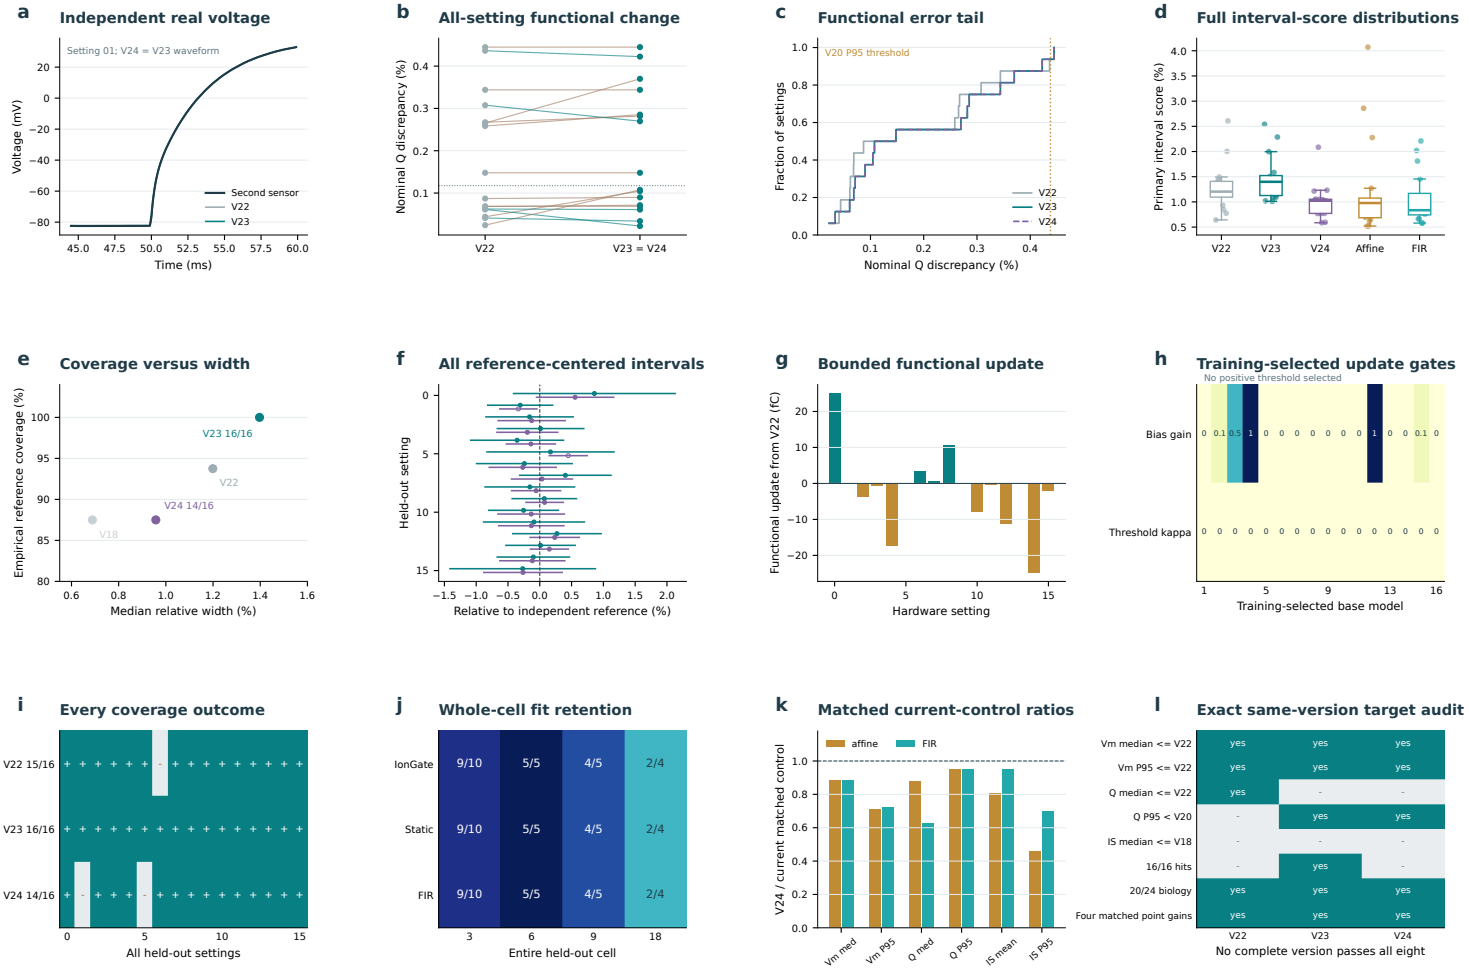


V23: functional tail + 16/16 hits. V24: better mean/tail interval scores, but 14/16 hits. Neither retains the V22 functional median and V18 primary score together.

Nominal RC functional of independent voltage, not metered ionic charge. Retrospective one-session development; controls share adaptation and calibration budgets.

# IonGate | dynamic current-balance challenge

Real return transients explain a failed steady-endpoint shortcut

One hardware session / 16 settings • equal new opportunities for all backbones • whole frozen versions; no score stitching

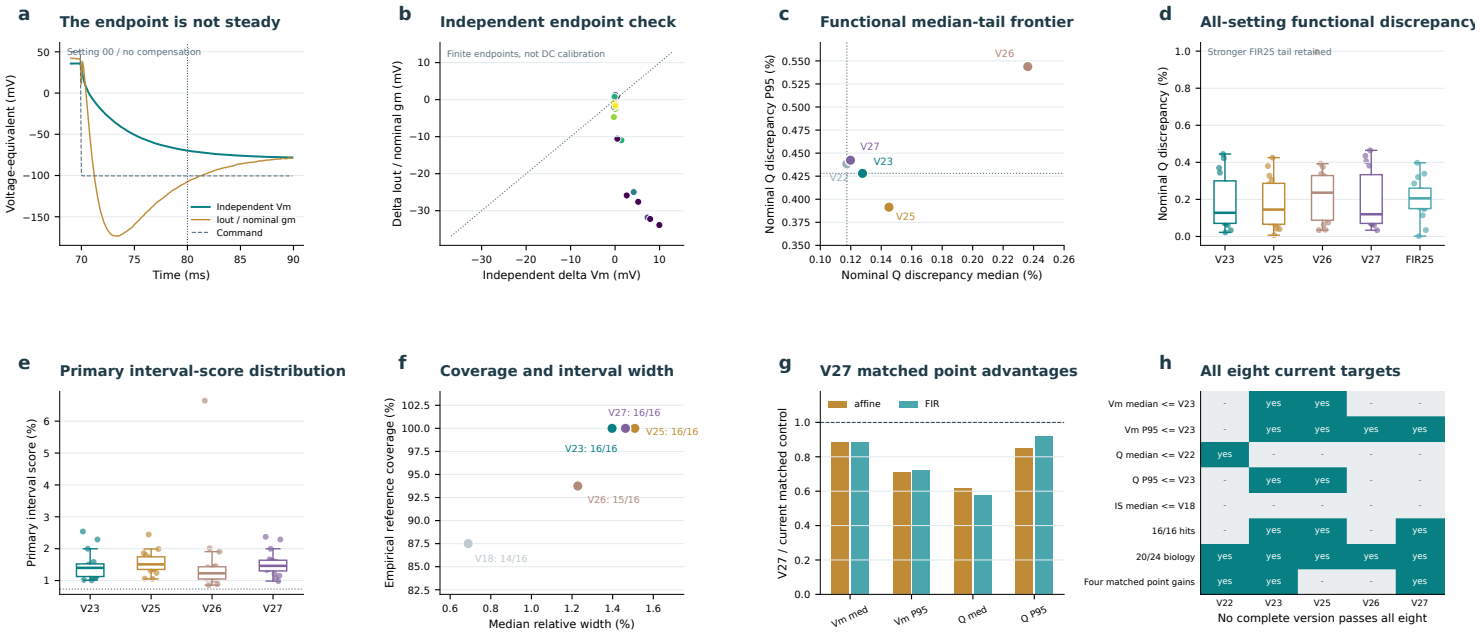


V25 improves Q tail with 16/16 hits, but loses the matched FIR tail comparison. V27 approaches the V22 median, but does not retain the V23 tail or V18 score.

Post-freeze endpoint diagnostic does not estimate DC current gain. Nominal RC functional, not metered ionic charge; retrospective development, no activation.

# IonGate | bounded residual learning

Matched point advantages retained; interval compression loses reference coverage

Frozen whole versions / 16 hardware settings, one session / equal 28-point + 36-interval options per backbone




Nominal RC functional of independent recorded Vm, not independently metered ionic charge. Retrospective development; no candidate activation.

Ablations retain fitted coefficients and recompute calibration within legal roles. Biological data/results are byte-identical V27 evidence, not new fitting.

# IonGate | synchronized model recalibration

Whole-model refitting improves voltage tails, but does not close joint qualification

Frozen whole versions / 16 hardware settings, one session / equal 12 leave-setting-out refits per backbone




Nominal RC functional of independent recorded Vm, not independently metered ionic charge. Retrospective development; no candidate activation.

Ablations compare fixed nine-setting parents and eleven-setting active refits; no test-based model selection. Biological data/results are byte-identical V27 evidence, not new fitting.

# IonGate | waveform-protected functional learning

Lower functional median, full reference coverage, and matched point advantages in one frozen candidate

Frozen whole versions / 16 hardware settings, one session / same readout / calibration opportunities for every backbone




Nominal RC functional of independent recorded Vm, not independently metered ionic charge. Retrospective development; no candidate activation.

V30/V31 share identical point waveforms; V32 adds a separately recorded role-disjoint adaptation trial. Biological data/results are byte-identical V27 evidence, not new fitting.

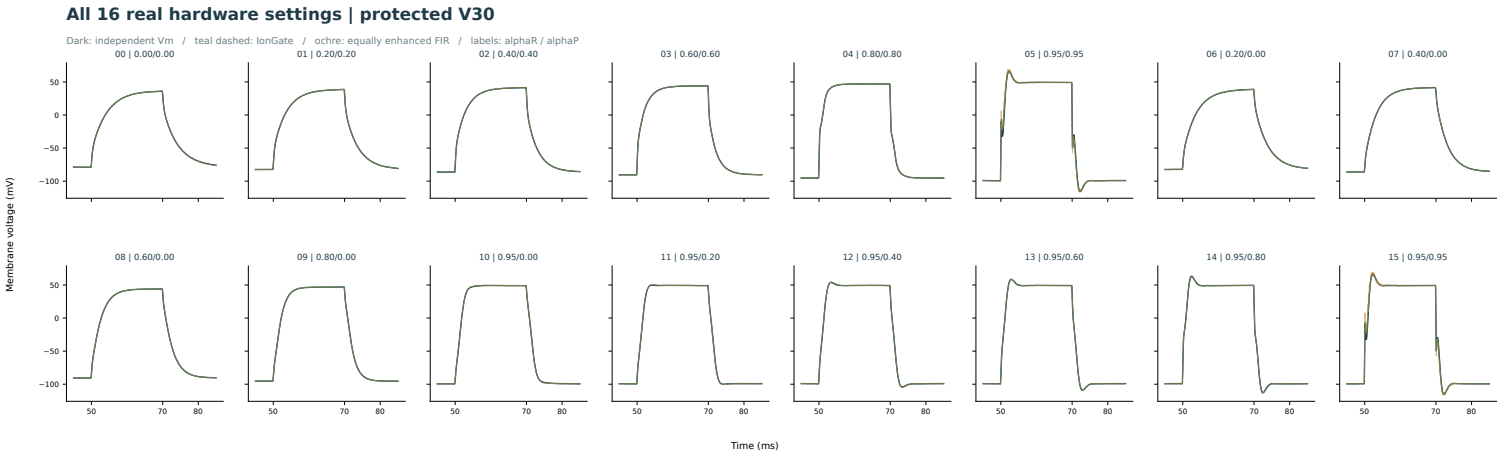
All 16 real hardware settings | protected V30
Dark: independent Vm / teal dashed: lonGate / ochre: equally enhanced FIR / labels: alphaR / alphaP
00 | 0.00/0.00
01 | 0.20/0.20
02 | 0.40/0.40
03 | 0.60/0.60
04 | 0.80/0.80
05 | 0.95/0.95
06 | 0.20/0.00
07 | 0.40/0.00
08 | 0.60/0.00
09 | 0.80/0.00
10 | 0.95/0.00
11 | 0.95/0.20
12 | 0.95/0.40
13 | 0.95/0.60
14 | 0.95/0.80
15 | 0.95/0.95
Membrane voltage (mV)
Time (ms)