Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

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

yu an xu
Monash University
tianyi ma
Second Affiliated Hospital of Xuzhou Medical College
Peng Xue
Second Affiliated Hospital of Xuzhou Medical College
Jicheng Li
3272464246@qq.com

Second Affiliated Hospital of Xuzhou Medical College

Research Article

Keywords: cardiac electrophysiology, voltage clamp, circuit memory, functional inference, uncertainty calibration

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11665941/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

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

Authors: Yu An Xu[1], Tianyi Ma[2], Peng Xue[3]*, Jicheng Li[2]*

[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.

Corresponding authors: Peng Xue, xuepeng08262024@163.com; Jicheng Li, 3272464246@qq.com.

Abstract

Background

Functional assessment of SCN5A and hERG channels supports arrhythmia research, but voltage-clamp feedback can distort both recorded waveforms and derived functional quantities. We developed IonGate to coordinate voltage reconstruction and state-aware functional inference in one model.

Results

We retrospectively evaluated a frozen circuit-informed observer, branch-memory readout and waveform-protected learning architecture across 16 compensation settings from one electronic-model-cell session. An independent second-amplifier voltage channel provided the reconstruction reference. Affine and finite-impulse-response controls received matched incremental enhancement budgets. Fixed-waveform ablation isolated the readout contribution; a retained biological application comprised 450 sweeps from four cells. Median/P95 voltage root-mean-square errors were 0.252515/2.565794 mV; nominal resistance–capacitance functional errors were 0.105459%/0.443297%. Median functional error was 49.47% and 52.18% lower than affine and finite-impulse-response controls, respectively; voltage P95 was 28.99% and 28.00% lower. Removing branch memory from identical predicted waveforms increased median functional error to 0.349058%; the full readout improved 14/16 settings. Residual calibration contained 16/16 nominal references, with wider intervals and worse proper scores than matched controls. The unchanged biological adapter retained 20/24 identifiable activation fits, matching enhanced controls.

Conclusions

IonGate improved waveform and nominal circuit-functional point estimates together and separated the contributions of memory, bounded learning and calibration. References used

measured voltage and nominal component values, not independent ionic-charge metrology. Clinical utility and cross-session generalization require further study.

Keywords: cardiac electrophysiology; voltage clamp; circuit memory; functional inference; uncertainty calibration

List of abbreviations

FIR, finite impulse response; HEK, human embryonic kidney; IQR, interquartile range; MAE, mean absolute error; P95, empirical 95th percentile; RC, resistance–capacitance; RMSE, root-mean-square error; SE, standard error; SNR, signal-to-noise ratio; TAC, transverse aortic constriction; WT, wild type.

Background

Ion-channel experiments connect molecular perturbations to electrical function. In cardiac electrophysiology, functional SCN5A assays inform the interpretation of Brugada-syndrome-associated variants, while hERG measurements relate channel gating to repolarization and drug response [1,2]. The useful endpoint is often not the sampled current alone, but a derived quantity such as an activation curve, recovery feature or finite-window current integral. Measurement errors can therefore propagate into quantities used to compare mutations, compounds and recording conditions.

Patch-clamp methods provide high-resolution electrical recordings [3], yet the command voltage is not necessarily the membrane voltage. Access resistance, capacitance, output filtering and compensation form a coupled feedback system. Published artefact models explain why these effects can alter fast-current measurements [4,5], and coupled experimental–mathematical approaches address related problems in sodium-current characterization [6]. The challenge is not simply to remove a fixed filter: changing the

electrical load also changes the effective measurement gain. A useful inference model must retain the relevant circuit states and specify which physical quantity it estimates.

State inference from noisy biophysical recordings and conductance estimation under voltage clamp emphasize this distinction between an observed signal and its underlying state [7,8]. Accurate waveform reconstruction and accurate functional estimation are related but not interchangeable objectives. A small voltage residual can have a substantial endpoint contribution to a circuit functional; conversely, a low integral error can conceal a distorted waveform. Dynamic branch memory adds another dependence on initial and final states. These considerations motivate a model that evaluates waveform and function together rather than optimizing either in isolation.

Here we introduce IonGate, which separates voltage reconstruction from state-aware functional readout and reconnects them through waveform-protected learning. A circuit-informed observer reconstructs membrane voltage from first-amplifier signals; a dynamic RC readout retains the branch-state contribution to finite-window charge. A bounded learned correction then improves the functional estimate while leaving the inherited observer coefficients fixed and limiting changes to the returned waveform. The contribution is this integrated inference architecture, not a new RC law or a claim that existing voltage-clamp models lack dynamic states [4,5,6].

We evaluate the architecture at three distinct levels. Independently measured second-amplifier voltage tests reconstruction; fixed-waveform ablation tests the contribution of branch-memory readout; and identical-point comparisons isolate interval calibration. Affine and finite-impulse-response (FIR) controls receive the same incremental readout and calibration opportunities. The retrospective primary evaluation includes 16 compensation settings from one electronic-model-cell session, alongside a retained biological application with 450 sweeps from four cells. This design tests whether one complete model can improve

both waveform and nominal RC-functional estimates, rather than assembling the best outcomes from different versions.

Methods

Data and software

The primary model-cell and NaV1.5 recordings are publicly archived in Figshare (10.6084/m9.figshare.27193878) and the source model in Zenodo (10.5281/zenodo.15200038). Sato known-drive recordings are archived in Zenodo (10.5281/zenodo.8288407). Cardiac and paired neuronal imaging datasets are available in the cited public repositories [9,10,11]. Processed figure data, complete numerical tables and attribution records are included in Data S2. Data S3 contains frozen inference sources, serialized models, calibration records, processed model-cell inputs, expected predictions and integrity manifests. It supports primary frozen-model inference and scoring without the original IBEX project, not historical training or biological refitting. Raw DAT/TIFF files remain in their public source repositories. Data S1 supplies full supplementary methods, tables and original figure plates.

Software project: IonGate V30 frozen replay. Tested operating system: Linux; language: Python 3.9.18; requirements: NumPy 2.0.2, pandas 2.2.3 and SciPy 1.13.1. The frozen inference and scoring archive is supplied as Data S3. A permanent public identifier for the original software has not been assigned. Third-party licenses and notices remain in the archive.

Public recordings and reference hierarchy

The primary recording was 20230510-Cfast_auto.dat from the Lei electronic-model-cell dataset, Figshare DOI 10.6084/m9.figshare.27193878 [5]. The HEKA conversion preserves output current, command voltage, independent second-amplifier membrane voltage,

timestamps and compensation labels. Native sampling is 10 kHz. The primary scoring window is 45–80 ms inclusive, containing 351 samples per setting. The 16 compensation settings belong to one acquisition session, not 16 independent experiments. The biological application comprises 450 NaV1.5 sweeps from four cells and 24 cell–temperature–compensation conditions. Table 1 distinguishes the references and units of analysis.

Only timestamps, first-amplifier output current, command voltage and known compensation metadata enter a held-out prediction. Measured voltage is permitted in designated fitting and calibration roles and is used to score test predictions after their freeze. The primary functional reference is calculated from measured voltage and nominal circuit elements. Programmed-state hardware, published SCN5A summaries, hERG replay and paired optical/electrical recordings provide additional, separately labeled contexts [1,2,12,13,14]; they are shown in distinct modules in Figures 5–7 and preserved in the Supplementary Material, rather than pooled into the primary performance estimates.

Dynamic observer and the protected backbone

The inverse observer follows the front-end circuit relations underlying the published voltage-clamp model [5]. The core parameters describe output-filter time constant, access resistance, compensation resistance estimate, estimated capacitance and voltage offset. Output-filter inversion uses Iin = Iout + tauout dIout/dt; the derivative is estimated offline by a seven-sample second-order Savitzky–Golay filter with interpolation edge handling [15]. An exponentially propagated compensation-voltage state accounts for the known resistance and supercharging compensation fractions. Exact front-end equations and frozen coefficients are supplied in Supplementary Methods S1 and the model manifests.

The protected model also retains learned current-innovation and dynamical-state components developed on the fitting pools. Its inherited memory basis includes exponential states at 0.1,

0.3, 1, 3 and 10 ms, derivative information and compensation interactions. Source-model residuals and regularized state features contribute to the voltage estimate; therefore IonGate is a circuit-informed learned observer, not a model-free inversion. The final update leaves every inherited backbone and physical coefficient fixed. Neither source-model replay nor held-out measured voltage is substituted as the model output.

The simplified steady-state relation illustrates why this structure is necessary. For Iion = gVm and Vm = Vcmd − RsIion, with no additional filtering,

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

Thus the effective gain depends on the load g as well as the instrument resistance. This analytical relation motivates retaining feedback physics; it is not an empirical performance result (Figure 1).

State-aware circuit functional

The nominal model-cell readout uses gm = 10 pA/mV, gk = 100 pA/mV and Ck = 100 pF, giving tauk = Ck/gk = 1 ms. Its branch state obeys tauk dIk/dt + Ik = Ck dVm/dt. Integrating this relation gives

$$Q_{RC} = g_m \int_{t_a}^{t_b} (V_m - V_{hold})dt + C_k \Delta V_m - \tau_k \Delta I_k, \ \tau_k = \frac{C_k}{g_k} \quad (2)$$

For the primary final-model endpoint, each waveform uses Vhold = Vm(45 ms) and Ik(45 ms) = 0, with both endpoints included. Piecewise-linear voltage supplies an analytic convolution-weight calculation, so the charge functional does not require numerical differentiation of the measured reference. The same convention is used for predictions and references. Older hold-mean, record-propagated-state analyses remain separate in the supplement; their numbers are not combined with the final endpoint.

The term nominal RC charge functional denotes Equation (2) under these component and state conventions throughout this paper. Setting the initial branch current to zero defines the scored functional; it is not an independent observation that the physical branch had exactly that state. Voltage accuracy and this circuit-derived functional accuracy are consequently reported as separate endpoints.

The fixed-waveform ablation uses the final IonGate voltage without refitting. Full RC, no-memory and leak-only predictions are respectively L + C + S, L + C and L, where L is the conductance integral, C the endpoint capacitance term and S the branch-state term in Equation (2). Every ablation is scored against the same full reference computed from measured voltage. This changes the readout, not its target, and tests branch memory rather than removal of every dynamical component from the observer.

Waveform-protected functional learning

The final readout combines actual, same-family, same-role functional predictions from three frozen internal experts. Expert labels and their implementation history are provided in Supplementary Methods S2; they denote complete fitted models rather than selected historical scalar scores. Nonnegative weights sum to one, an optional offset is bounded by 25 fC, and a shrinkage gain limits departure from the protected anchor:

$$Q_{target} = Q_{anchor} + \gamma\, clip(\Sigma_j w_j Q_j + b - Q_{anchor}, -50, 50) \quad (3)$$

A smooth minimum-energy direction psi transports this correction to the waveform:

$$V_{new} = V_{anchor} + s(Q_{target} - Q_{anchor})\psi,\ Q_{RC}[\psi] = 1 \quad (4)$$

The scale s is reduced when necessary to enforce both a 50 fC functional correction limit and a 0.15 mV maximum pointwise waveform change. The actual returned waveform, after this

scaling, supplies both voltage and functional scores. Functional results are not assembled from a separate unreported predictor.

Linear programming balances mean and maximum absolute fitting residuals with a penalty on departure from the anchor. 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, give 37 configurations per family. Six-setting/three-setting coefficient-held-out inner predictions select a median/P95-noninferior minimax option before refitting on nine settings. Inherited hyperparameters were already selected within those nine settings; this is partial out-of-fold development, not fully nested validation.

Waveform protection bounds the update to the returned prediction; it does not guarantee that every error metric decreases. The 50 fC functional-correction limit and 0.15 mV pointwise voltage-change limit define an update trust region, not metrological uncertainty bounds. All reported voltage and functional metrics use the same constrained waveform after any shrinkage, rather than scoring the unachieved target in Equation (3).

Grouped evaluation and matched comparators

Four outer folds hold out four compensation settings each. Within each outer fitting pool, four base-model rotations use nine fitting and three calibration settings; test settings remain outside both roles. The four base predictions are averaged to form the outer-fold waveform. All 16 settings receive one outer-fold test prediction. Frozen model manifests identify every group role and checksum. Repeated model development has inspected this hardware session, so these are retrospective grouped results rather than a new untouched external test.

Affine and FIR comparisons use their own frozen backbones and actual own-family expert predictions. Each receives the same 37 new readout and 18 residual-calibration options and the same group roles. This matches the incremental enhancement budget; it does not imply

identical total historical model-development cost or identical architecture complexity. Source-exact Lei replay, source-model tuning, static resistance, earlier stronger tail controls and all prior candidates are retained as historical comparisons in Supplementary Tables, with endpoint conventions identified. Published-model replay uses the archived implementation and Myokit infrastructure [16].

Voltage metrics are calculated per setting before median, mean and P95 aggregation. Functional discrepancy is

$$e_Q = 100 \frac{|Q - Q_{ref}|}{max(|Q_{ref}|, 100\ fC)} \quad (5)$$

P95 is the empirical 95th percentile of 16 setting-level metrics, not a pointwise error guarantee. Reduction percentages compare summary values, not averages of per-setting reductions. We report descriptive comparisons without treating waveform samples or within-session settings as independent experimental replicates and without a cross-session significance claim.

Empirical intervals and module attribution

A regularized residual kernel uses the two known compensation coefficients, with ridge regularization [17]. Length scales 0.25/0.6/1.5, ridge penalties 0.1/1/10 and geometry powers 0/1 give 18 fitting-only options per family. The residual mean is clipped at one percentage point; geometry variance is floored at 0.05. Neither quantity is independently calibrated instrument uncertainty. Inner/leave-one-setting proper-score selection precedes fitting on nine settings.

Each base model supplies three separate calibration residuals; their own fitted residual mean and scale transfer those residuals to the query. Twelve adjusted endpoint candidates yield lower/upper ranks 2/11 at alpha = 0.2. The 12 candidates are correlated cross-fit residuals, not

12 independent sessions. We report containment, normalized full width and the proper interval score

$$IS_\alpha = U - L + \frac{2}{\alpha}(L - Q_{ref})1(Q_{ref} < L) \quad (6)$$

$$+ \frac{2}{\alpha}(Q_{ref} - U)1(Q_{ref} > U)$$

The interval score is normalized by 100/max(abs(Qref),100 fC), whereas relative width uses 100/max(abs(Qhat),100 fC). Predictive-inference literature motivates separating calibration and evaluation [18,19], but the present empirical intervals are not claimed to inherit an independent-session or ionic-charge coverage guarantee. A frozen conservative calibration alternative has byte-identical point predictions; comparing it with the final residual-kernel layer isolates calibration effects from point-estimation effects.

Biological application and identification accounting

The final package retains the previously frozen biological adapter unchanged. It combines public recording metadata with a model-derived apparent-current proxy and activation fitting. Cell-held-out development and its admissible candidate policies, peak window, voltage support, fit bounds and identifiability criteria are recorded in Supplementary Methods S3 and the archived scripts. Temperature and compensation conditions remain nested within cell. No final hardware-readout update is presented as a new biological refit.

All 450 applied sweeps and all 24 condition-level fits are retained, including boundary and unidentifiable cases. Apparent half-activation voltage and fit uncertainty characterize the model-derived activation curves, not independently measured cellular membrane voltage or clinical variant pathogenicity. Enhanced static and FIR biological controls are reported alongside IonGate, including equal identification counts. The four cells are the biological grouping units; 24 fitted conditions are not 24 independent cells.

Reproducibility and controlled updating

No model is retrained while preparing this manuscript. Hashes bind the selected complete model, predictions, protocols and figures to the accepted frozen package. A read-only reproduction entry point verifies those inputs and recomputes summary tables, figures and new fixed-waveform ablations. The 227 existing software tests are preserved as implementation checks, not biological validation. Future learning is versioned, bounded and reversible, with separate fitting, calibration and evaluation roles; autonomous longitudinal improvement has not been demonstrated.

Hardware and paired-sensor contextual modules

We retained the original Sato hardware panels and all 240 records. New diagnostics compare programmed-relative area error with full-scale normalization without changing any output calibration. The programmed occupancy is an input-state reference, not independent current metrology. In the real optical/electrode record, we calculated descriptive Pearson correlations for optical delays from minus 8 to plus 8 native frames, 5 mV voltage-bin medians/IQRs and 20-frame block-median trajectories. Positive delay pairs earlier electrode samples with later optical samples. No delay was selected to fit or evaluate IonGate. All original image pixels and electrode panels were retained; the added analyses characterize sensor-reference structure rather than increase the independent sample count.

Image panels were reconstructed from raw repository TIFF data rather than copied from published article figures. Cardiac displays use camera channel A, a left-half field crop and linear percentile-based pseudocolour scaling with record-specific limits; absolute brightness is not compared across cohorts. Neuronal displays use the recorded 0.59 micrometre pixel size and fixed field/ROI crops. The retained processing applies linear pseudocolour, without denoising, motion correction or spatial/temporal filtering; frame times and electrode time

stamps determine pairing. Source identifiers, crop/scale conventions and display settings are recorded in Data S2. Sham and neuronal source data are CC BY 4.0; TAC source data are CC0 [9,10,11].

AI-assisted preparation

OpenAI Codex assisted with manuscript drafting and editing, code preparation, figure assembly and consistency checks. The authors are responsible for verifying the scientific content, computational materials, interpretation and references, and for approving the final submission. No model was retrained and no new experimental data were generated during this journal adaptation.

Results

IonGate reduces typical and tail voltage errors against matched controls

IonGate reconstructed independently measured membrane voltage with setting-level RMSE median/mean/P95 of 0.252515/0.529503/2.565794 mV. Matched affine values were 0.285783/0.691783/3.613359 mV, and FIR values were 0.285681/0.685227/3.563487 mV. Voltage P95 was 28.99% lower than affine and 28.00% lower than FIR. Thus, the benefit extends beyond a lower typical error to the tail of the setting-level error distribution. Figure 2 presents all 16 current and voltage traces, paired errors, phase-specific residuals and fixed-setting close-ups; Table 2 reports every primary summary. The 16 settings belong to one acquisition session, so these are descriptive setting-level gains rather than independent-session effect estimates.

Dynamic branch memory improves the fixed-waveform functional readout

On the same final voltage waveforms, the full state-aware readout achieved a nominal functional discrepancy median/P95 of 0.105459%/0.443297%. Removing only the branch-

memory contribution increased these values to 0.349058%/1.345885%; retaining only the conductance integral yielded 0.380219%/4.136016%. The full readout had a smaller absolute discrepancy than the no-memory readout in 14/16 settings. All paired reversals are retained in Figure 3 and Supplementary Table V30_SAME_WAVE_MEMORY_ABLATION. This post-freeze intervention did not alter predictions, parameters, integration windows or reference definitions.

The matched ablation localizes a functional benefit to state-aware readout without changing voltage reconstruction. The endpoint branch state carries a history-dependent contribution that a conductance integral or a capacitance endpoint alone omits. Equation (2) specifies that contribution; the paired comparison tests its practical value against the common measured-voltage-derived reference. This is evidence for the branch-memory readout on these records, not an ablation of every dynamical feature in the observer.

Protected learning combines functional and waveform gains

The final functional discrepancy median was 49.47% below matched affine and 52.18% below FIR. Its P95 was 19.01% below affine and 1.66% below FIR. Together with voltage median and P95, all four prespecified point summaries favored IonGate over the current equally enhanced controls. These are summary-level advantages, not a claim of a win in every setting or against every historical control.

Relative to its protected predecessor, the new functional readout lowered median discrepancy from 0.119793% to 0.105459%, an 11.97% reduction, while preserving the voltage summaries and empirical reference-hit count. Functional P95 changed from 0.442318% to 0.443297%, an increase of 0.000979 percentage points. The complete version was selected for its combined improvement with this disclosed small tail change; the supplement retains

the stricter predecessor and unmet historical numerical targets. Figure 3 separates matched-model comparisons from predecessor and calibration attribution.

Complete reference containment entails a width and score trade off

All 16 nominal functional references were contained by the final intervals, compared with 15/16 for each matched control. Median normalized width was 1.364174%, versus 1.117981% for affine and 1.042015% for FIR. IonGate therefore achieved complete empirical containment with wider intervals. Its median proper score was 1.362887%, higher than affine at 1.159797% and FIR at 1.069416%; the strong controls remained more efficient on this score.

With byte-identical point predictions, residual-kernel calibration reduced median interval score by 8.14% relative to the conservative alternative, from 1.483720% to 1.362887%, with 16/16 hits in both cases. This is a calibration-layer improvement within IonGate, not interval-score superiority over the matched strong controls. Figure 4 displays all intervals, misses, widths and setting-level scores rather than only favorable summaries.

Known-drive hardware demonstrates why the reference definition matters

The retained real-hardware challenge contains 240 recordings across three SNR strata, with 80 records per stratum. Median full-scale area discrepancy ranges from 1.292% to 3.280% across the complete strata. Figure 5 retains native waveforms, all-record comparisons and the inverted-state control, while new modules show how low programmed occupancy amplifies programmed-relative errors. These are complementary measurement-chain diagnostics, not new V30 wins or independent analog-charge truth.

Biological application retains identifiable fits and visible failures

The preserved biological module yielded 20 identifiable activation fits among 24 conditions from 450 sweeps. Counts were 9/10, 5/5, 4/5 and 2/4 for cells 3, 6, 9 and 18. The four failed

or unidentifiable conditions remained in the table and figure, together with their fitted values and reasons. Enhanced controls also attained 20/24, so this endpoint demonstrates retained applicability rather than a unique identification advantage. Figure 6 and Table 3 connect the waveform method to a biologically interpretable downstream analysis without treating a fitted curve as independent biological voltage truth.

Real images and electrode recordings expose a complementary reference channel

The original multimodal panels are retained rather than replaced by a schematic. Figure 6 preserves the 24 displayed cardiac/neuronal frames as external context alongside the complete biological application. Figure 7 retains the independent electrode and optical record, the 18 event-frame montage and all original event analyses. New paired-sensor diagnostics use 2273 frames from the same neuron; zero-lag descriptive Pearson correlation between electrode voltage and optical signal is -0.741. Lag profiles, voltage-bin summaries and time-ordered trajectories make the measured correspondence visible without declaring another IonGate prediction benchmark. This separates a physically observable reference channel from a model-generated target; repeated frames do not provide additional independent cells.

Discussion

Circuit memory and bounded learning coordinate waveform and function

IonGate addresses two coupled inference objectives: recovering the membrane-voltage waveform and estimating a finite-window circuit functional. Its principal advantage is their coordination within one frozen model. The feedback-aware observer reconstructs voltage, the RC readout retains branch-state memory, and bounded learning improves the functional estimate while constraining changes to the same returned waveform. The lower voltage and functional median/P95 summaries against both matched controls support the complete

architecture; the fixed-waveform intervention identifies a distinct benefit of its memory-aware readout.

The method builds on established voltage-clamp circuit modeling [4,5,6]. Its contribution is not the RC identity alone, but the integration of that structure into a learnable voltage–function pipeline with explicit update limits and separable component comparisons. Three interventions address different questions: fixed-waveform ablation tests branch memory, the protected-anchor comparison tests bounded functional learning, and identical-point calibration tests interval construction. This organization makes the model contribution inspectable rather than assigning every performance gain exclusively to one component.

Relevance to cardiac functional assays

Functional ion-channel assays inform comparisons of variants, pharmacological conditions and cellular phenotypes [1,2,5]. Such comparisons depend on distinguishing biological changes from changes introduced by recording resistance and compensation. IonGate contributes at this measurement-inference layer: it improves voltage and nominal RC-functional estimates together, while exposing the assumptions and uncertainty costs associated with the derived quantity. This provides a reusable basis for more interpretable electrophysiological comparisons, rather than treating a visually plausible trace as sufficient evidence of accurate function.

The preserved biological analysis links the method to activation fitting and explicitly retains unidentifiable conditions. Figures 5–7 add known-drive hardware, published functional summaries and real optical/electrode signals as complementary measurement contexts. These modules explain what different reference channels can establish; they are not interchangeable replications of V30 accuracy. Improved clinical variant classification, drug-safety prediction

and biological ionic-charge recovery remain application hypotheses, not measured outcomes of this study.

Calibration exposes the containment–efficiency trade-off

IonGate contained all 16 available nominal references, whereas both matched controls missed one. This gain in empirical containment required wider intervals and produced worse proper scores. The residual-kernel comparison nevertheless shows an internal improvement: it lowered the median score by 8.14% while keeping point predictions identical and retaining every reference. Calibration therefore contributes independently of point estimation, but has not removed the strong controls' interval-efficiency advantage. Reporting containment, width and score together identifies that trade-off rather than equating a hit count with universal reliability.

Scope and transferability

The primary evaluation is retrospective and limited to one hardware acquisition session. Compensation-setting holdouts separate prediction roles but do not create independent instruments, dates or cells. Repeated development on this session and partial inner out-of-fold selection limit generalization claims. The biological adapter is unchanged and its 20/24 identification count is shared by enhanced controls. Functional references depend on nominal component values and a specified zero initial branch state, rather than independent ionic-charge metrology.

The next transferability test should freeze the present architecture across independent acquisition sessions and varied electrical loads, with calibrated component uncertainty or an independently measured charge reference. Such evaluation would test calibration transport and biological functional utility without further selecting on the current 16 settings. The present evidence already demonstrates a concrete methodological gain: state-aware readout

and bounded learning can improve functional inference while retaining a more accurate voltage reconstruction than the current matched controls.

Conclusions

IonGate combines circuit-memory readout and waveform-protected learning to improve voltage-clamp functional inference. In one retrospective hardware-session evaluation, the same frozen model reduced median and P95 voltage RMSE and nominal RC-functional error relative to equally enhanced affine and FIR controls. Fixed-waveform ablation identified a contribution of branch memory, while identical-point comparison isolated calibration gains. All 16 nominal references were contained at a cost in interval width and proper score. The unchanged biological adapter retained 20/24 identifiable activation fits, matching enhanced controls. The demonstrated advance is a circuit-informed architecture that improves waveform and functional point estimates together and makes the uncertainty cost explicit.

Declarations

Ethics approval and consent to participate

This study was a secondary analysis of publicly available experimental recordings and images. No new participants were recruited, animal experiments were performed, or biological samples were collected. No new ethics approval was required for the present secondary analysis. Ethics approvals for the original data collection are described in the source publications [5,9–11,13,14].

Consent for publication

Not applicable. The manuscript contains no identifiable individual person’s details, images, or videos.

Availability of data and materials

The Lei electronic-model-cell recordings and NaV1.5 recordings analyzed in this study are publicly available from Figshare (https://doi.org/10.6084/m9.figshare.27193878). The source circuit model is available from Zenodo (https://doi.org/10.5281/zenodo.15200038), and the Sato known-drive recordings are available from Zenodo (https://doi.org/10.5281/zenodo.8288407). The cardiac and neuronal imaging datasets are cited in References 9–11 and 13–14. Processed figure data, numerical tables, and attribution records are supplied as Additional file 2 (DataS2.zip); the frozen inference/scoring archive, processed inputs, calibration records, and replay materials are supplied as Additional file 3 (DataS3.zip); supplementary methods, tables, and Figures S1–S20 are supplied as Additional file 1 (DataS1.pdf). The software project is IonGate V30 frozen replay; it was tested on Linux with Python 3.9.18, NumPy 2.0.2, pandas 2.2.3, and SciPy 1.13.1. A permanent public repository DOI for the software has not been assigned.

Competing interests

The authors declare that they have no competing interests.

Funding

No funding was received for this study.

Authors' contributions

YAX developed the model. TM collected the publicly available data and formatted the manuscript. PX and JL provided overall scientific guidance and supervision.

Acknowledgements

The authors acknowledge the publicly released recordings, images, and source models cited in this article. OpenAI Codex assisted with manuscript drafting and editing, code preparation, figure assembly, and consistency checks, as described in Methods. The authors are

responsible for verifying the scientific content, computational materials, interpretation, and references.

Additional files

Additional file 1. DataS1.pdf: supplementary methods, complete tables, and Figures S1–S20.

Additional file 2. DataS2.zip: machine-readable processed figure data, full numerical tables, and attribution records.

Additional file 3. DataS3.zip: frozen model, processed inputs, calibration records, and executable inference/scoring replay.

References

1. Ma JG, O'Neill MJ, Richardson E, Thomson KL, Ingles J, Muhammad A, et al. Multisite validation of a functional assay to adjudicate SCN5A Brugada syndrome-associated variants. Circ Genom Precis Med. 2024;17:e004569. https://doi.org/10.1161/circgen.124.004569.

2. Kemp JM, Whittaker DG, Venkateshappa R, Pang Z, Johal R, Sergeev V, et al. Electrophysiological characterization of the hERG R56Q LQTS variant and targeted rescue by the activator RPR260243. J Gen Physiol. 2021;153:e202112923. https://doi.org/10.1085/jgp.202112923.

3. Hamill OP, Marty A, Neher E, Sakmann B, Sigworth FJ. Improved patch-clamp techniques for high-resolution current recording from cells and cell-free membrane patches. Pflugers Arch. 1981;391:85–100. https://doi.org/10.1007/BF00656997.

4. Lei CL, Clerx M, Whittaker DG, Gavaghan DJ, de Boer TP, Mirams GR. Accounting for variability in ion current recordings using a mathematical model of artefacts in voltage-clamp experiments. Philos Trans A Math Phys Eng Sci. 2020;378:20190348. https://doi.org/10.1098/rsta.2019.0348.

5. Lei CL, Clark AP, Clerx M, Wei S, Bloothooft M, de Boer TP, et al. Resolving artifacts in voltage-clamp experiments with computational modeling: an application to fast sodium current recordings. Adv Sci (Weinh). 2025;12:e00691. https://doi.org/10.1002/advs.202500691.

6. Abrasheva VO, Kovalenko SG, Slotvitsky M, Romanova S, Aitova AA, Frolova S, et al. Human sodium current voltage-dependence at physiological temperature measured by coupling a patch-clamp experiment to a mathematical model. J Physiol. 2024;602:633–661. https://doi.org/10.1113/jp285162.

7. Huys QJM, Paninski L. Smoothing of, and parameter estimation from, noisy biophysical recordings. PLoS Comput Biol. 2009;5:e1000379. https://doi.org/10.1371/journal.pcbi.1000379.

8. Li S, Liu N, Yao L, Zhang X, Zhou D, Cai D. Determination of effective synaptic conductances using somatic voltage clamp. PLoS Comput Biol. 2019;15:e1006871. https://doi.org/10.1371/journal.pcbi.1006871.

9. Ou X. Dual voltage-calcium optical mapping in cardiac slice of Sham C57 murine. figshare. 2020. Version 1. https://doi.org/10.6084/m9.figshare.11936610.v1. CC BY 4.0.

10. Ou X. Dual voltage-calcium optical mapping in cardiac slice of TAC C57 murine. figshare. 2020. Version 2. https://doi.org/10.6084/m9.figshare.11931666.v2. CC0.

11. Rozsa M, Svoboda K. anm454597 Cell 0 of ‘Simultaneous Voltron (1.0) imaging and whole-cell patch-clamp recordings of somatosensory cortex layer 1 interneurons in vivo.’ Janelia Research Campus. 2021. Version 1. https://doi.org/10.25378/janelia.14143118.v1. CC BY 4.0.

12. Sato M, Hariyama M, Komiya M, Suzuki K, Tozawa Y, Yamamoto H, et al. Model-free idealization: adaptive integrated approach for idealization of ion-channel currents. Biophys J. 2023;122:3959–3975. https://doi.org/10.1016/j.bpj.2023.08.019.

13. He S, Kou K, O'Shea C, Chen T, Mu-u-min R, Dong R, et al. A dataset of dual calcium and voltage optical mapping in healthy and hypertrophied murine hearts. Sci Data. 2021;8:314. https://doi.org/10.1038/s41597-021-01085-5.

14. Xie ME, Adam Y, Fan LZ, Böhm UL, Kinsella I, Zhou D, et al. High-fidelity estimates of spikes and subthreshold waveforms from 1-photon voltage imaging in vivo. Cell Rep. 2021;35:108954. https://doi.org/10.1016/j.celrep.2021.108954.

15. Savitzky A, Golay MJE. Smoothing and differentiation of data by simplified least squares procedures. Anal Chem. 1964;36:1627–1639. https://doi.org/10.1021/ac60214a047.

16. Clerx M, Collins P, de Lange E, Volders PGA. Myokit: a simple interface to cardiac cellular electrophysiology. Prog Biophys Mol Biol. 2016;120:100–114. https://doi.org/10.1016/j.pbiomolbio.2015.12.008.

17. Hoerl AE, Kennard RW. Ridge regression: biased estimation for nonorthogonal problems. Technometrics. 1970;12:55–67. https://doi.org/10.1080/00401706.1970.10488634.

18. Lei J, G'Sell M, Rinaldo A, Tibshirani RJ, Wasserman L. Distribution-free predictive inference for regression. J Am Stat Assoc. 2018;113:1094–1111. https://doi.org/10.1080/01621459.2017.1307116.

19. Barber RF, Candès EJ, Ramdas A, Tibshirani RJ. Predictive inference with the jackknife+. Ann Stat. 2021;49:486–507. https://doi.org/10.1214/20-AOS1965.

Figure Legends

Figure 1. Circuit-memory inference with waveform-protected learning. Panels a–d show the model-cell circuit, state-aware readout, analytical gain surface and ratio compression; analytical panels are not experimental results. Panels e–m show the

group roles, separated prediction/reference paths, signed endpoint RC terms, measured command–membrane mismatch, native current/command and sensor/observer traces, bounded waveform changes, functional-learning comparison and identical-point calibration comparison. Role colors are pale blue fitting, gold calibration and teal test. The 16 base rotations each use 9 fitting, 3 calibration and 4 test settings; they do not represent 16 independent sessions. The waveform-protection line is 0.15 mV. All primary measurements come from the same electronic-model-cell session.

Figure 2. Independent membrane-voltage reconstruction across 16 compensation settings. Panel a shows all 16 native output-current traces and paired membrane-voltage traces; purple circles mark the independent sensor, teal IonGate, gold enhanced affine, blue enhanced FIR and dotted grey command voltage. The primary window is 45–80 ms inclusive. Panels b–d show all-setting RMSE, paired MAE and plateau RMSE. Panels e–g report phase-specific errors, FIR minus IonGate phase RMSE (positive favors IonGate), and error-energy fractions for onset, hold, offset and the remaining samples. Panels h,i show median absolute residuals and0-to-P95 descriptive ribbons across settings, not confidence intervals. Panel j gives four point-summary ratios to current matched affine/FIR. Panel k shows complete residual maps on a common symmetric mV scale. Panels l–n show native phase close-ups for fixed setting 00. There are16 settings from one session, not16 independent experiments.

Figure 3. Dynamic circuit memory and waveform-protected learning in the RC-functional readout. Panel a shows all 16 settings with cumulative full/no-memory/leak-only readouts of identical V30 waveforms and the full nominal functional of measured Vm as the common reference. Baseline is each waveform's Vm at 45 ms and initial branch current is zero at 45 ms. Panels b–g give signed terms, paired readout-error comparisons, full distributions, endpoint state correction and every memory-effect reversal. Panels h–j give complete dynamic/capacitor/state maps in pC; k–m compare observer and sensor-derived component trajectories for fixed setting00. Panels n–p add the bounded readout update against its protected anchor, residual-kernel versus conservative calibration with byte-identical points, and four matched point-summary ratios. Full-readout median/P95 are 0.105459%/0.443297%; no-memory 0.349058%/1.345885%; leak-only 0.380219%/4.136016%. Memory improves 14/16 settings. The reference is nominal circuit charge, not independently metered ionic charge; no component ablation is refitted.

Figure 4. Empirical reference containment and utility of actual asymmetric intervals at alpha = 0.2. Panel a shows all 16 nominal references and every method interval, including both matched-control misses. Panels b–g show width–error relations, all paired proper scores, all actual base-calibration residuals, reference hits, identical-point calibration alternatives and held-out residuals. Each outer fold contains 12 correlated calibration residual candidates across four base rotations; these are not12 independent experiments. Panels h–m show every interval width, signed containment margin, proper-score ECDF, prediction-width budget acceptance, accepted-set containment and accepted-set scores. Margin is100min(Qref−L,U−Qref)/max(|Qref|, 100 fC), not a symmetric-radius approximation. Both reporting budgets and relative-width summaries use prediction-normalized width; proper scores use the reference denominator. Empty accepted sets remain

unavailable. IonGate has16/16 hits with wider intervals and worse proper scores than the enhanced controls; neither a nominal independent-session guarantee nor 95% ionic-charge metrology is claimed.

Figure 5. Known-drive hardware challenge and denominator diagnostics. Panels a–l show native traces, edge responses, kinetics/noise maps, all 240 recordings, calibrated output/state areas and inverted-state negative control. Panels m–o show programmed occupancy versus programmed-relative discrepancy, full-scale versus programmed-relative normalization, and median (bars)/P95 (dark dots) full-scale errors for all 80 records per SNR stratum. These diagnostics do not change calibration or refit IonGate. State labels are programmed inputs, not independently measured analog current; this is a complementary measurement-chain challenge, not an additional final-model accuracy benchmark.

Figure 6. Biological activation application and complementary recorded signals. Panels a–c show 24 Sham, TAC and Voltron frames with scale bars. Images are external cohorts from the published optical resources [13,14], not images of the four HEK cells and not model-generated pictures. Panels d–g include all 450 sweeps, all 24 activation conditions, four unidentifiable cases and equal 20/24 identification for the enhanced biological controls. In e, error bars show fit-reported SE only for identifiable fits; unidentifiable uncertainties remain in the complete table. Panels h,i show all available paired SCN5A peak-density/recovery summaries (recovery axes are symlog), j,k recorded WT/R56Q hERG control/drug currents, l their native AP-segment integrals, m recorded tissue fluorescence, n recorded NaV compensation traces and o cell-held-out model-defined activation audit loss. These distinct endpoints are not pooled into V30 voltage or charge accuracy. The biological adapter is unchanged, the audit loss is not independent measured-voltage RMSE. Data attribution: cardiac source pixels, Xianhong Ou [9,10] (Sham CC BY 4.0; TAC CC0); neuronal source pixels, Marton Rozsa and Karel Svoboda [11] (CC BY 4.0). Displays are cropped and linearly false-coloured from the raw data; brightness limits are record-specific and are not used for cross-cohort intensity comparisons.

Figure 7. Paired optical/electrode recordings and sensor diagnostics. Panels a–j show the field, cell detail, synchronized signals, 3-by-6 event frames, current command and event/repeat analyses. Panels k–m use the same 2273-frame, one-neuron record: lag-dependent descriptive Pearson correlation, optical medians/interquartile ranges within5 mV electrode-voltage bins, and a time-colored trajectory of 20-frame block medians. Positive lag means the optical sequence follows the electrode sequence; no optimum is used to tune IonGate. Repeated frames and events are not independent neurons; the bin ribbons are descriptive IQRs, not confidence intervals. These independently recorded sensor pairs illustrate reference observability, not final-model accuracy on optical data. Source images and electrode recordings are attributed to [14]. Raw-data attribution: Marton Rozsa and Karel Svoboda, Janelia Research Campus, dataset 10.25378/janelia.14143118.v1 [11], CC BY 4.0. Crops, mean-frame display and pseudocolour were produced for this study; the accompanying paper [14] is cited for scientific context, not as the source of copied article figures.

Tables

Table 1. Recordings, reference types and evaluation units

Data	Extent	Reference	Endpoint
Electronic model cell	16 settings / 1 session	Independent second-amplifier Vm	Voltage RMSE; nominal RC functional
NaV1.5 application	450 sweeps / 4 cells / 24 conditions	Model-derived activation fits	Identifiability and retained failures
External resources	Separate source-specific units	Programmed states / published features / optical pairing	Main Figures 5–7 contextual modules; not pooled V30 accuracy

590 Nominal RC charge uses measured Vm and nominal circuit parameters. No independently metered ionic charge is available.

591 Table 2. Final frozen model and equally enhanced controls

Metric	IonGate	Affine	FIR
Voltage RMSE median (mV)	0.252515	0.285783	0.285681
Voltage RMSE mean (mV)	0.529503	0.691783	0.685227
Voltage RMSE P95 (mV)	2.565794	3.613359	3.563487
Nominal Q error median (%)	0.105459	0.208699	0.220536
Nominal Q error mean (%)	0.193961	0.214569	0.230084
Nominal Q error P95 (%)	0.443297	0.547341	0.450802
Reference hits /16	16/16	15/16	15/16
Interval score median (%)	1.362887	1.159797	1.069416
Interval score mean (%)	1.377353	1.266840	1.194838
Interval score P95 (%)	2.091487	1.840088	1.790461
Median relative width (%)	1.364174	1.117981	1.042015

592 Alpha0.2;16 settings from one session. Lower error/score is better. Complete containment is accompanied by wider intervals

593 and worse scores. The enhancement search budget is matched; total historical development cost is not asserted equal.

594 Table 3. Biological identification with all failures retained

Cell	Identifiable	Total conditions	Failed/unidentifiable
cell3	9	10	1
cell6	5	5	0
cell9	4	5	1
cell18	2	4	2

Cell	Identifiable	Total conditions	Failed/unidentifiable
Total	20	24	4

450applied sweeps,4cells. Biological module unchanged; enhanced controls also20/24. Not an independent voltage-accuracy or clinical classification endpoint.

Fig. 1 | Physics-constrained functional inference

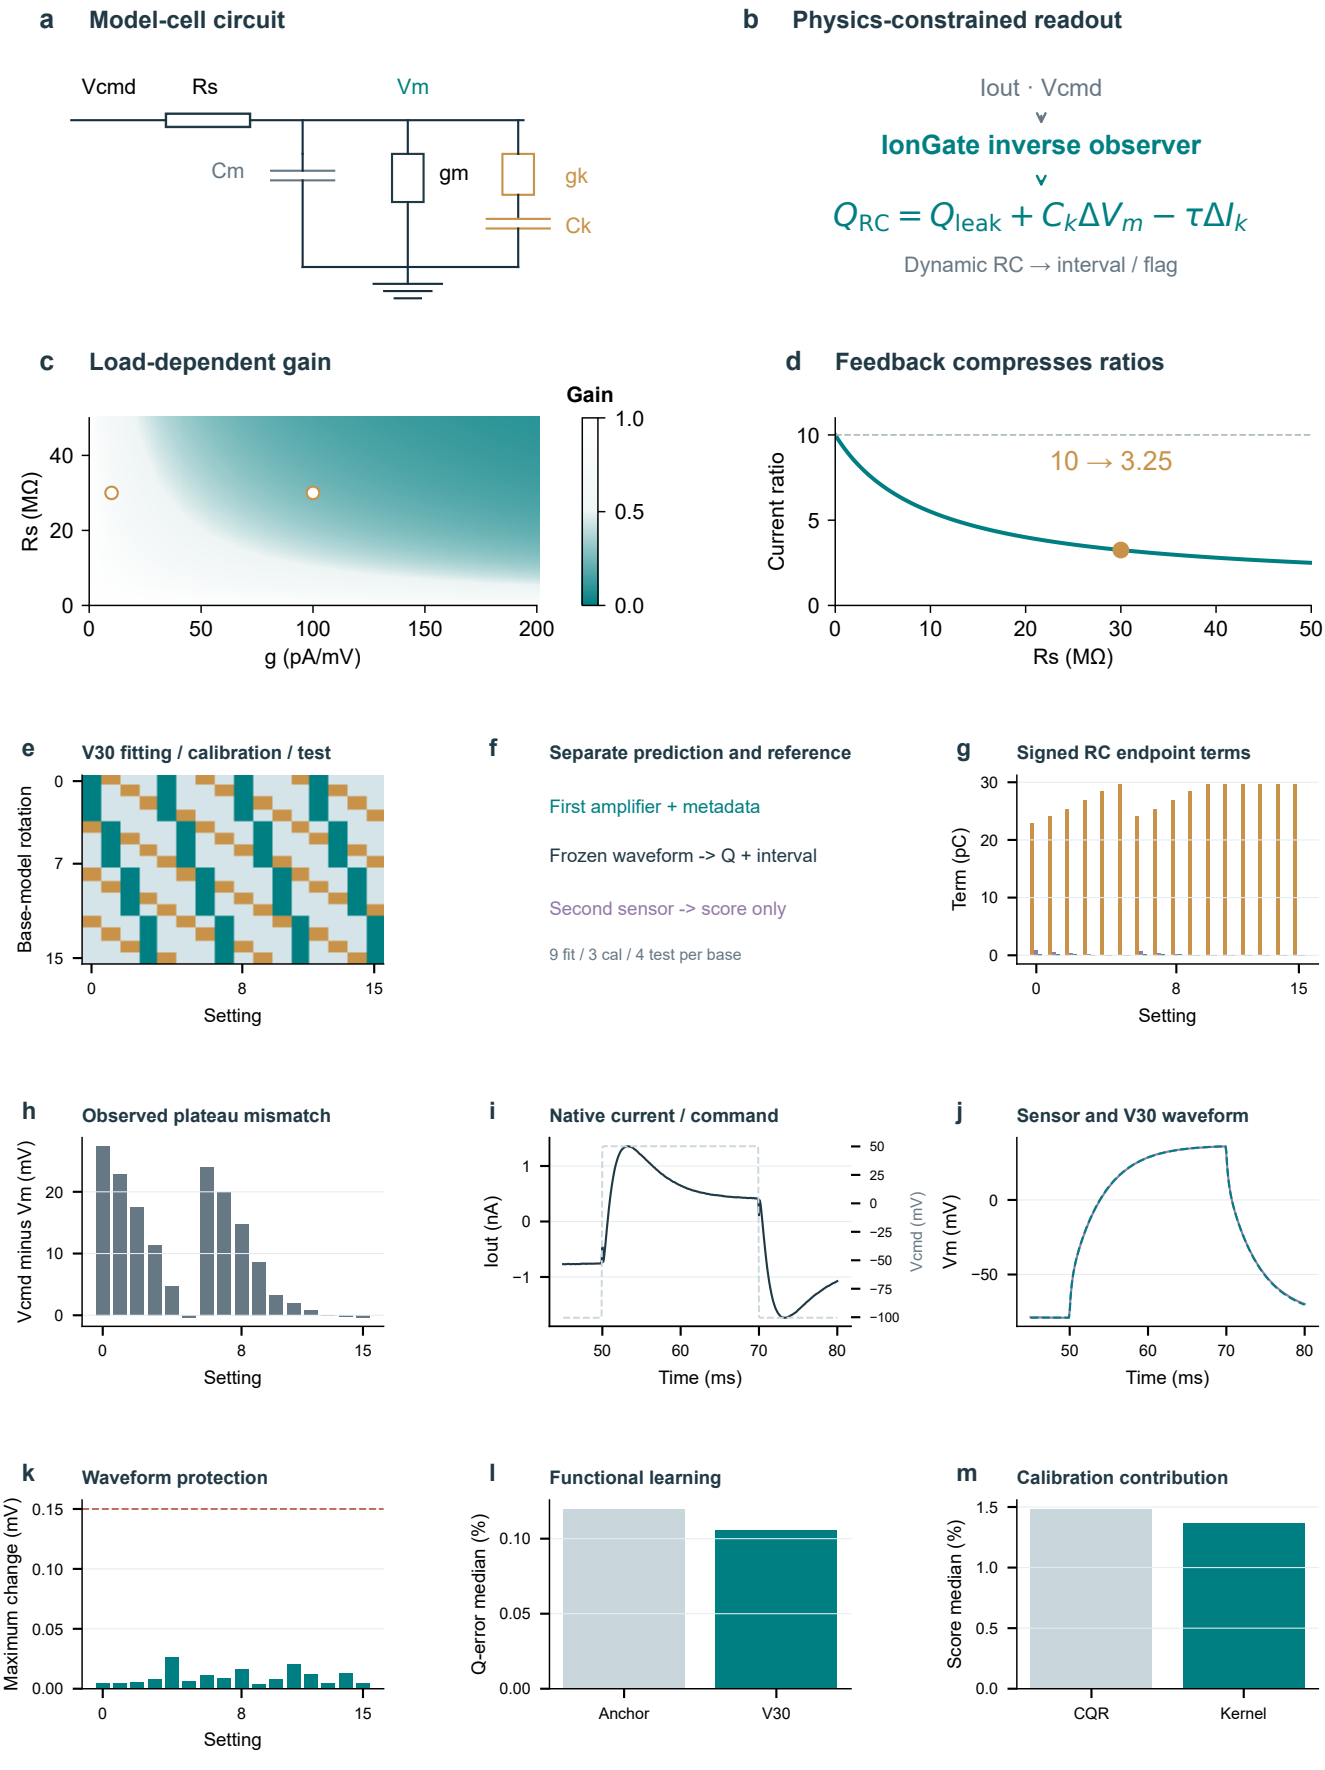

Fig. 2 | Independent Vm recovery with frozen V30

Real first-amplifier inputs | independent second sensor | equally enhanced affine / FIR

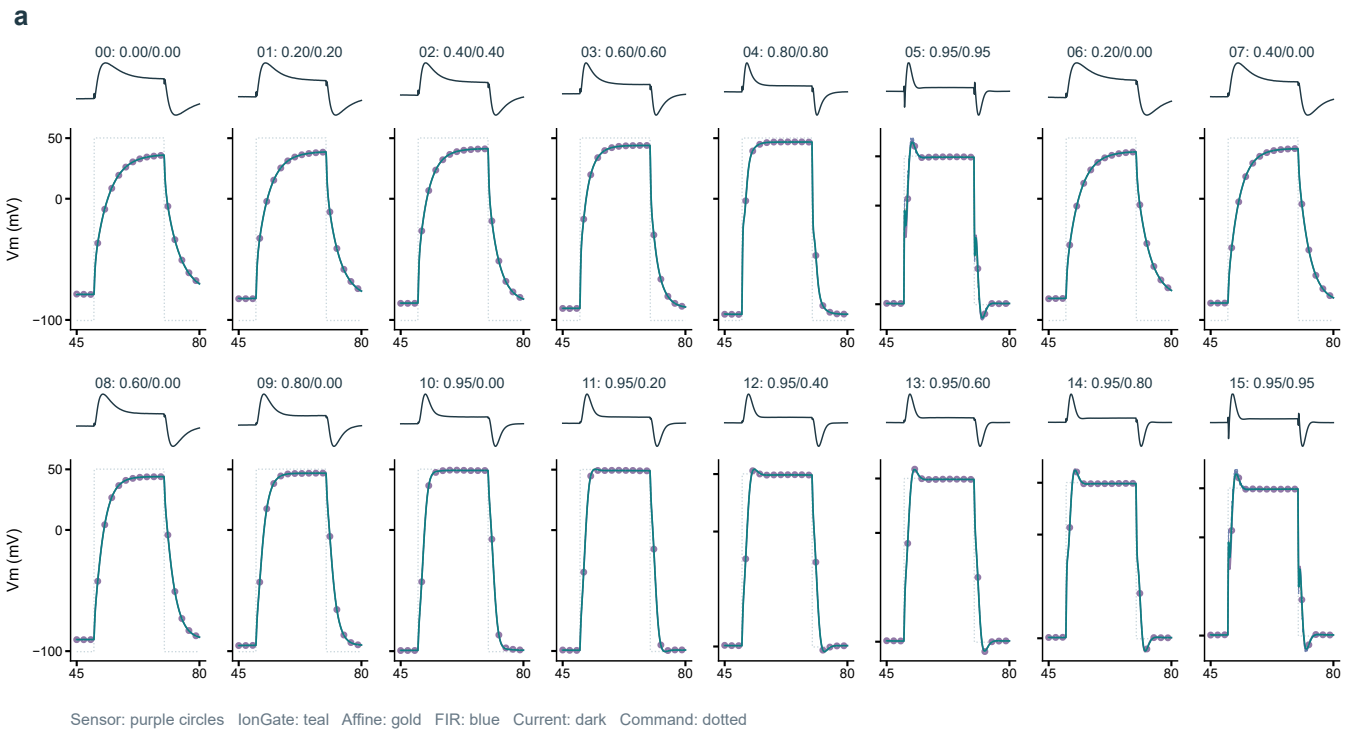

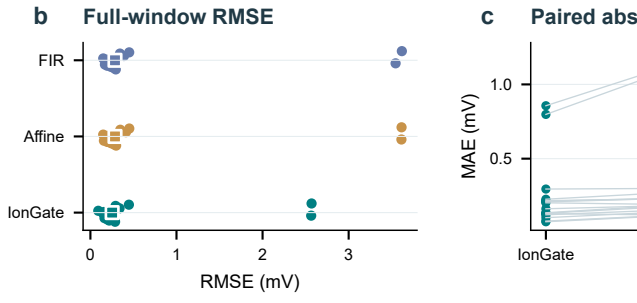

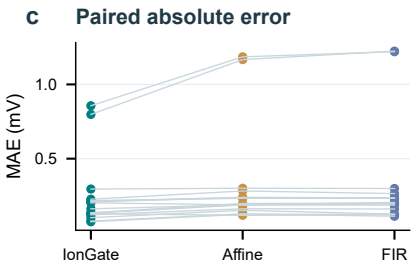

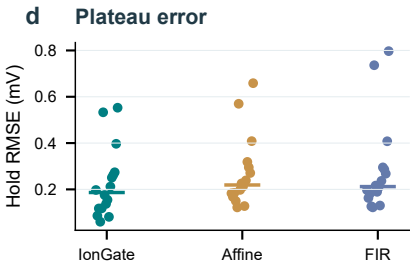



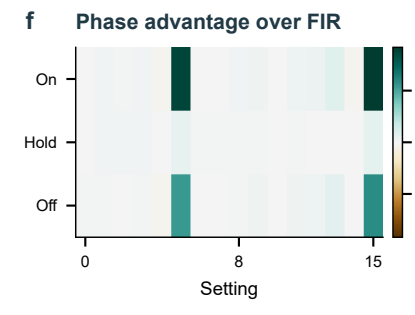

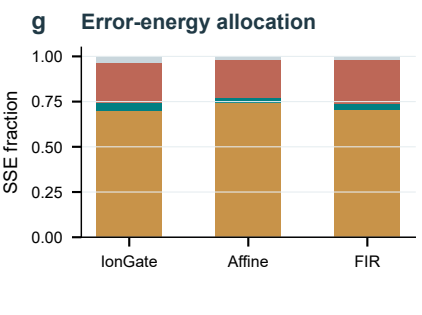

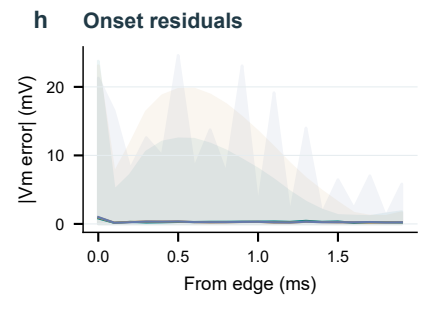

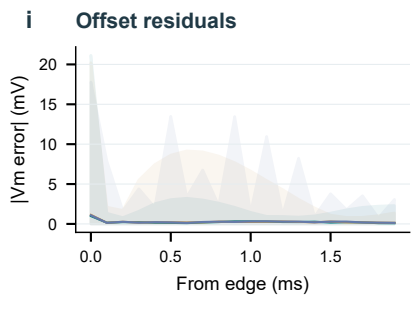

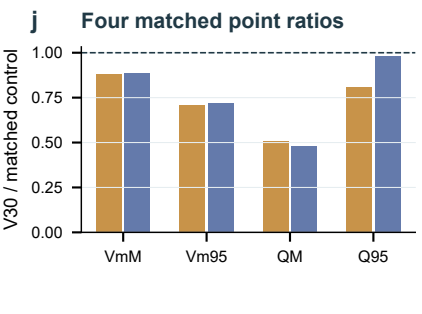

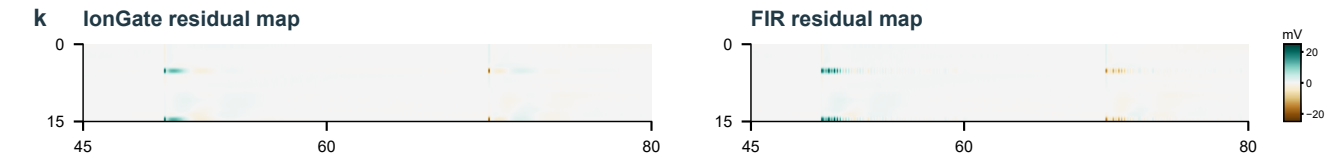

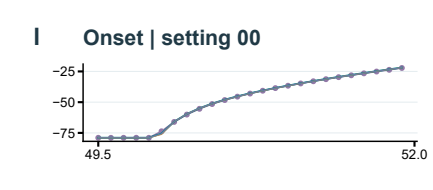

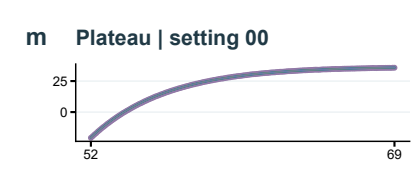

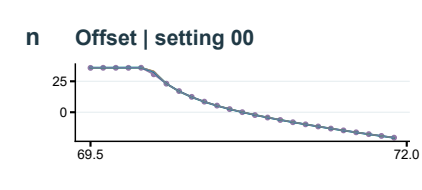

Fig. 3 | Dynamic RC memory and protected learning

Identical V30 waveforms | one full nominal reference | no refit or target redefinition

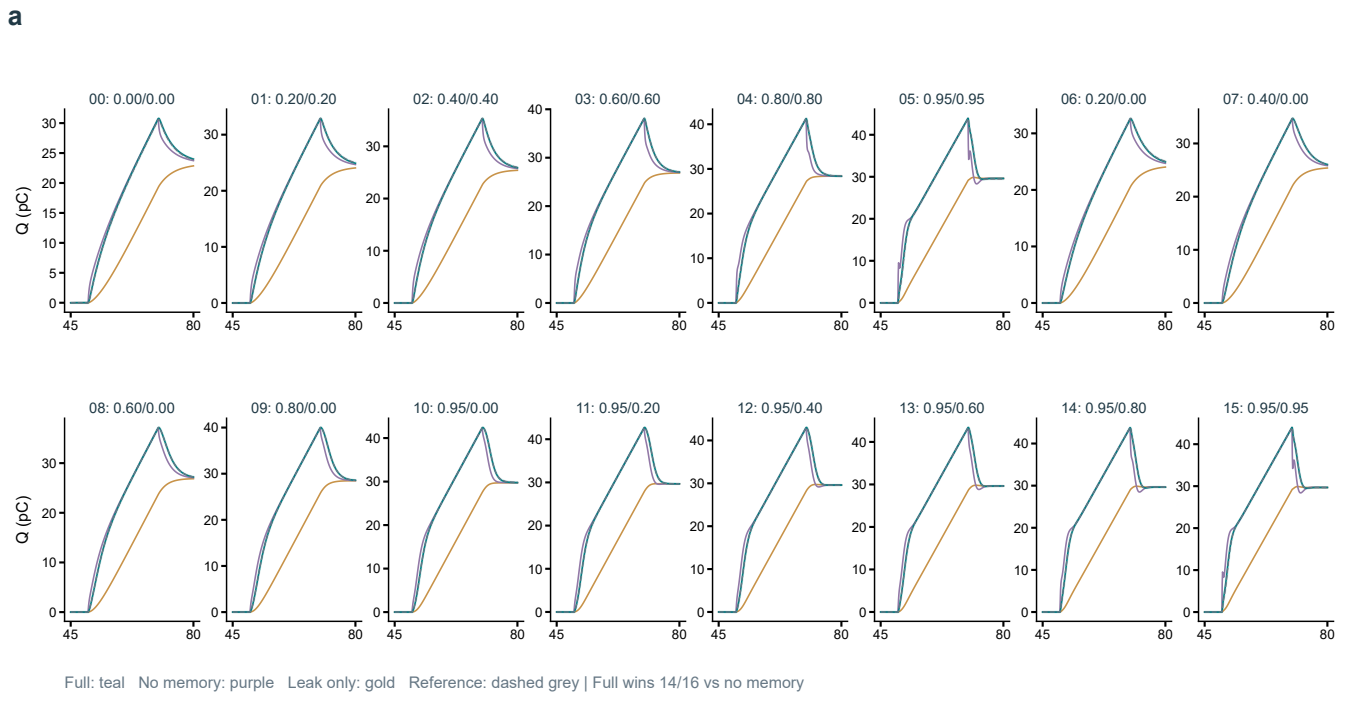

Full: teal No memory: purple Leak only: gold Reference: dashed grey | Full wins 14/16 vs no memory

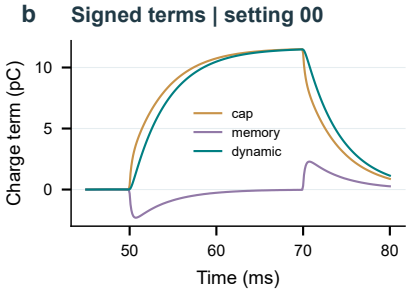

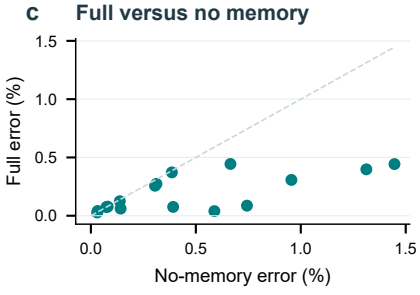



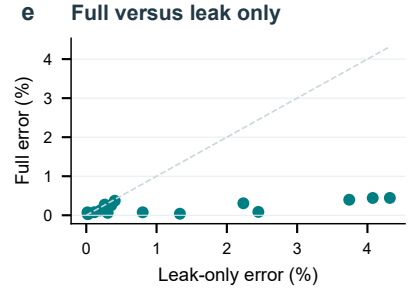

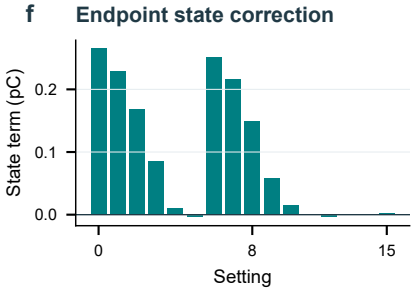

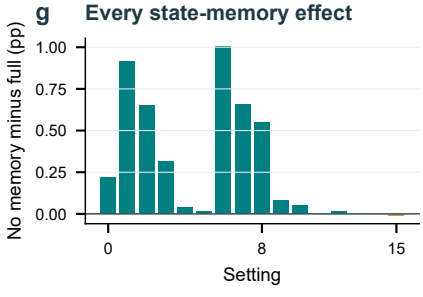

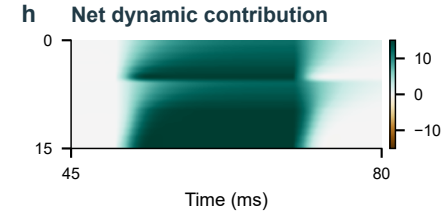

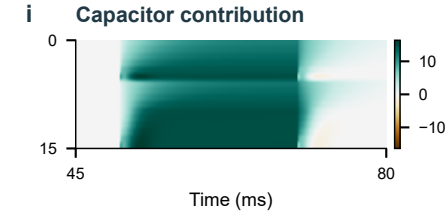

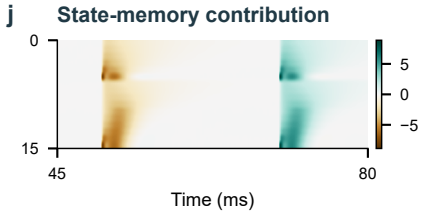

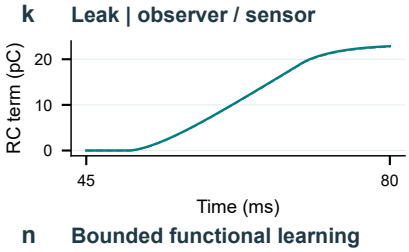

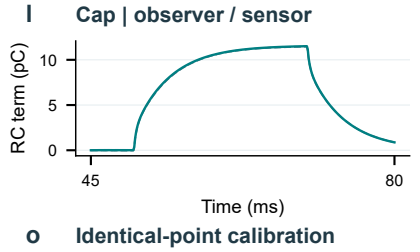

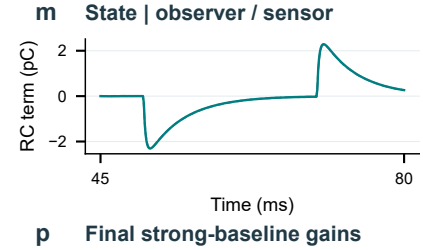

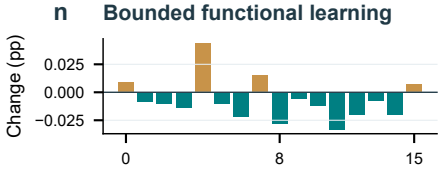

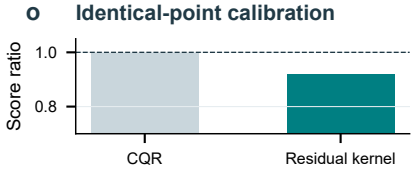

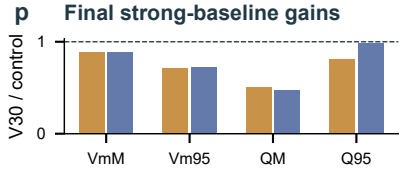

Fig. 4 | V30 intervals, utility and empirical containment

Actual asymmetric intervals | alpha = 0.2 | all settings and all matched-control misses retained

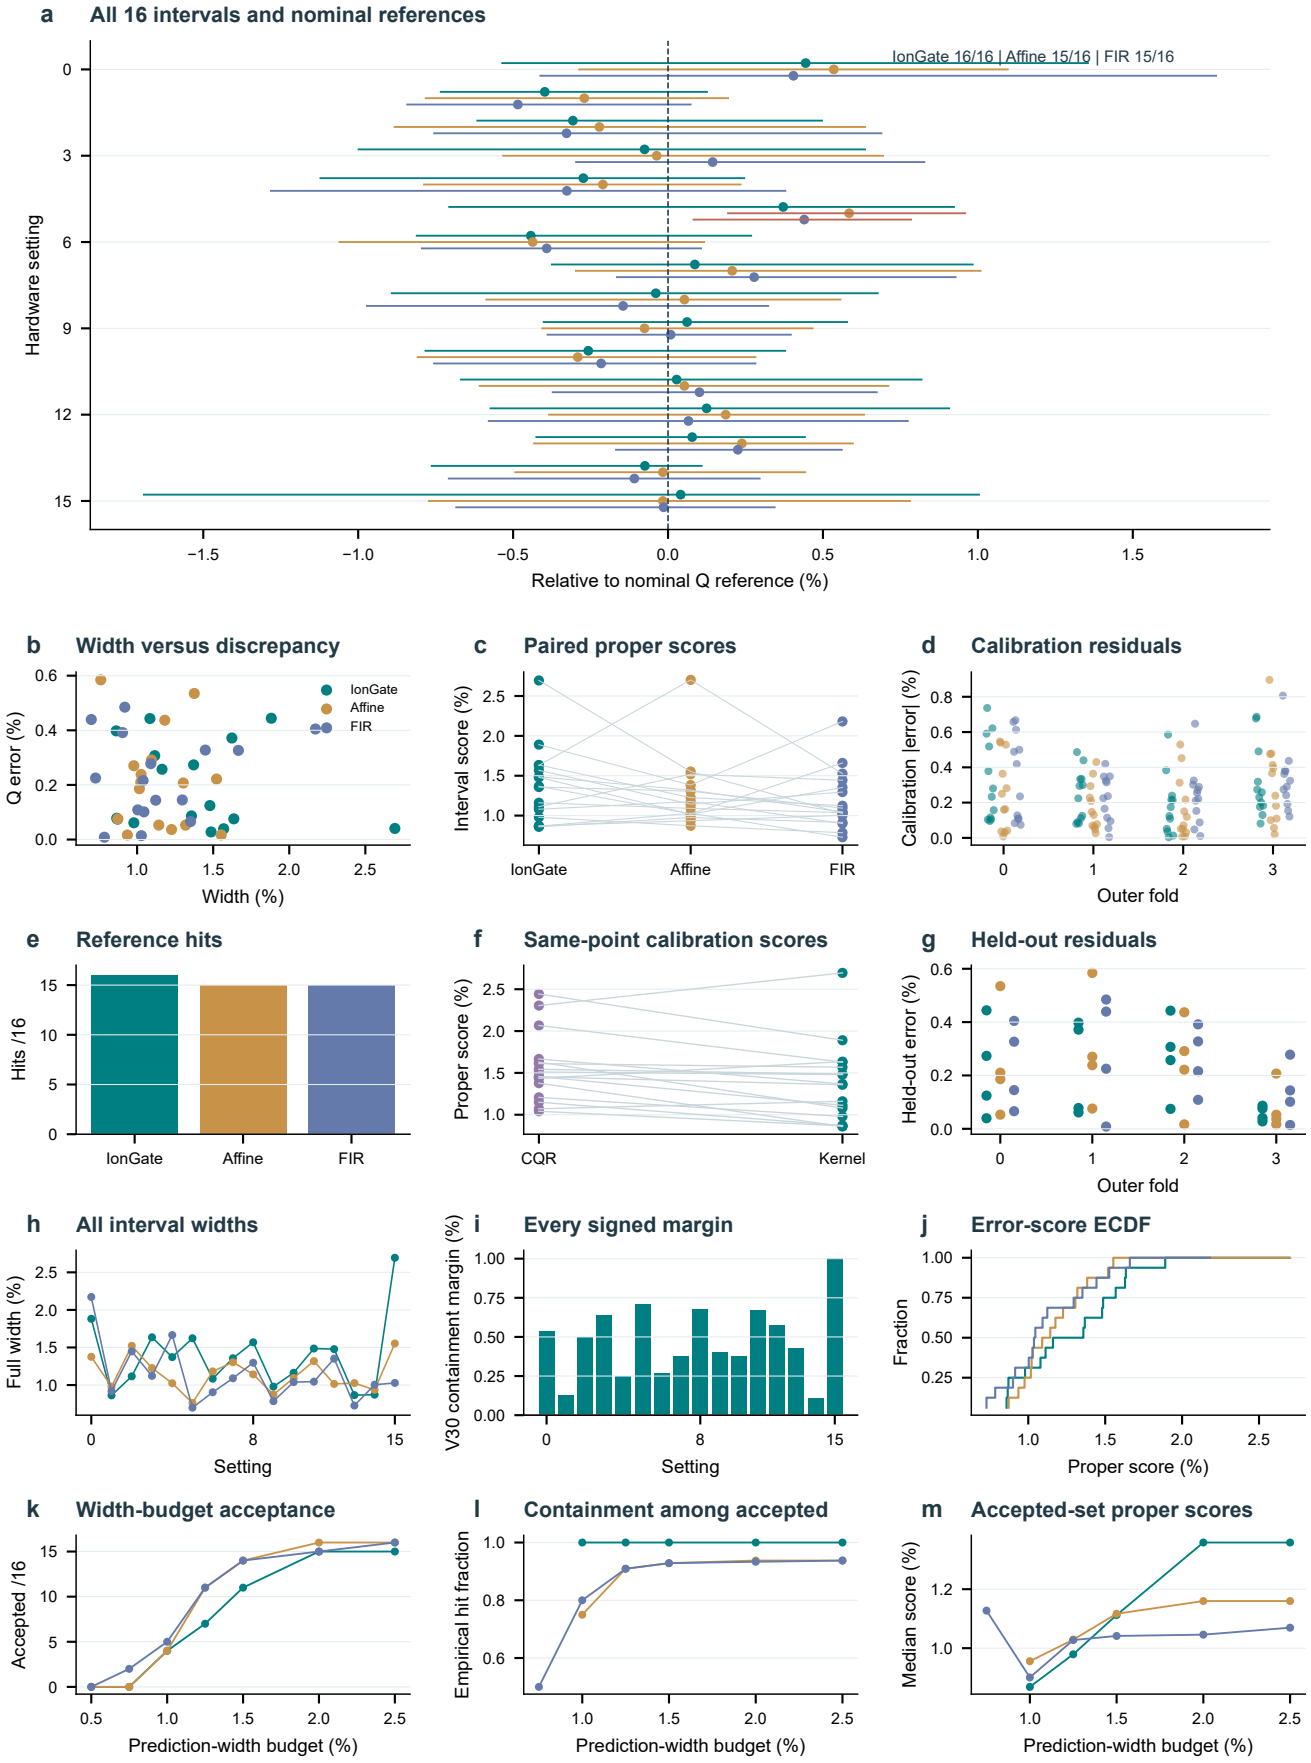

One session. Correlated cross-fit residuals. No independent ionic-charge or nominal coverage guarantee.

Fig. 5 | Native hardware noise and switching challenge



Fig. 6 | Real images and the full biological application

24 actual frames | external imaging cohorts

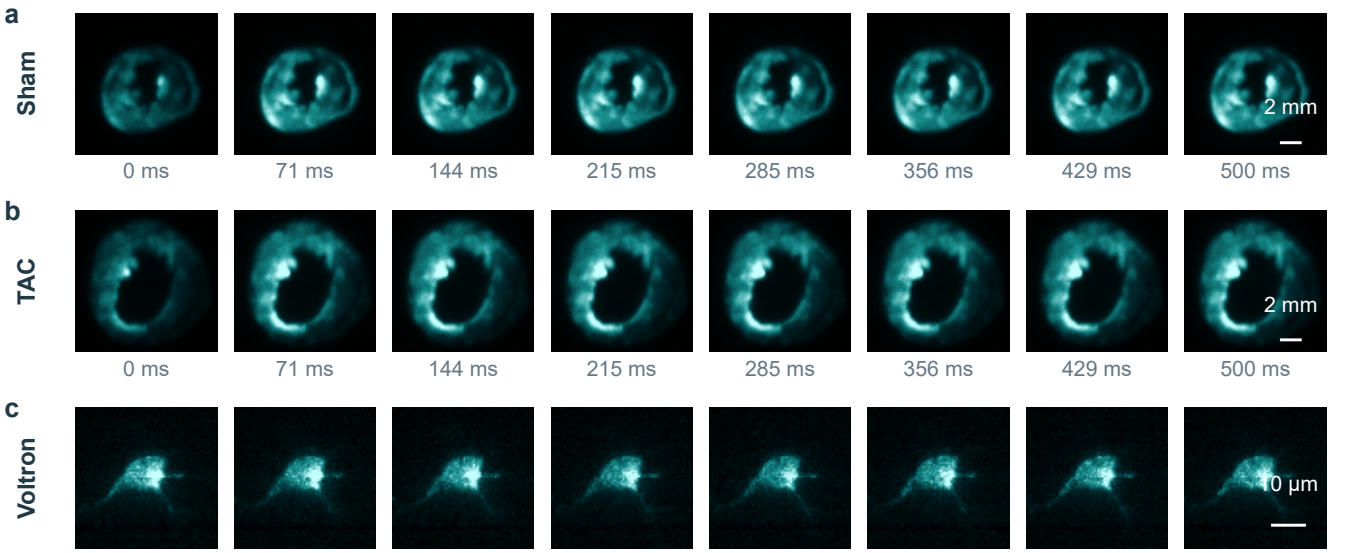

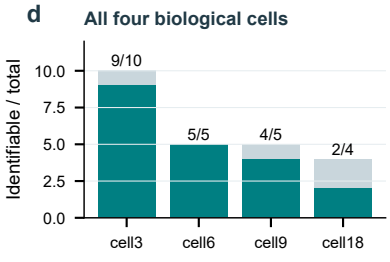

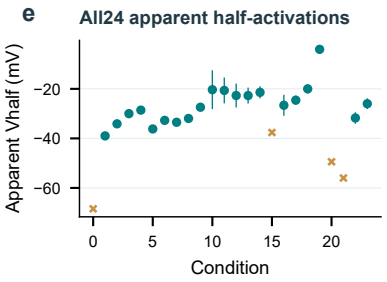

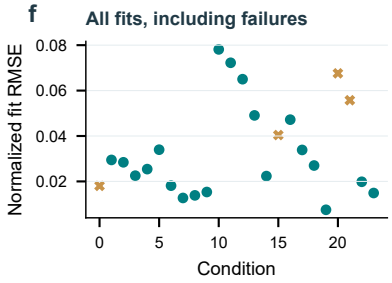

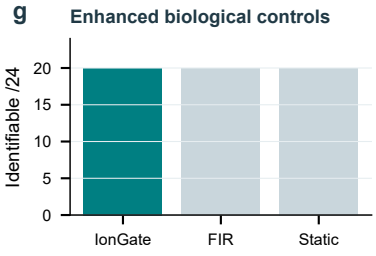

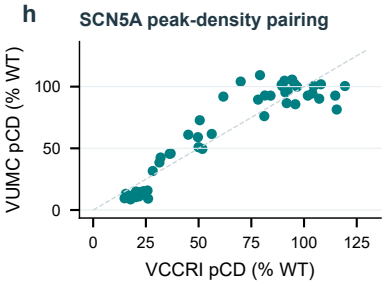

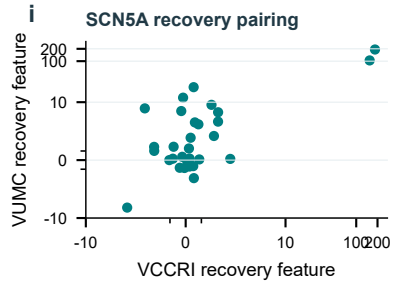

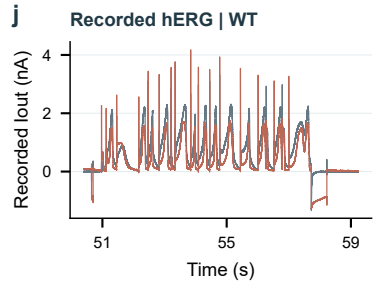

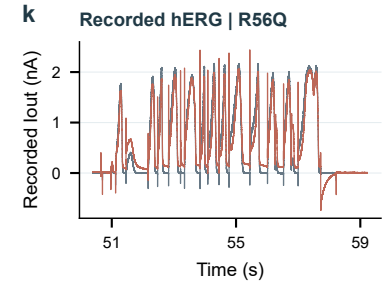

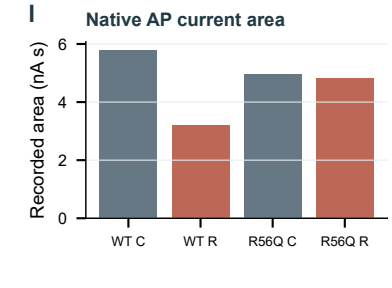

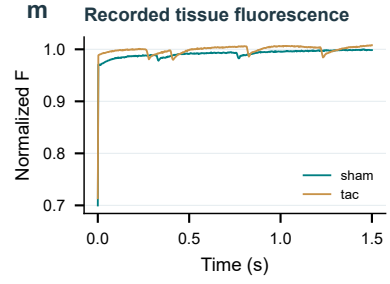

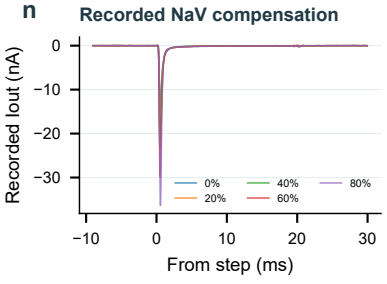

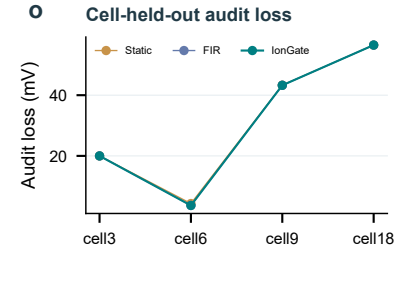

Images are external context, not these HEK cells. Biological adapter unchanged; all failures retained.

Fig. 7 | Seeing voltage and recording it independently

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

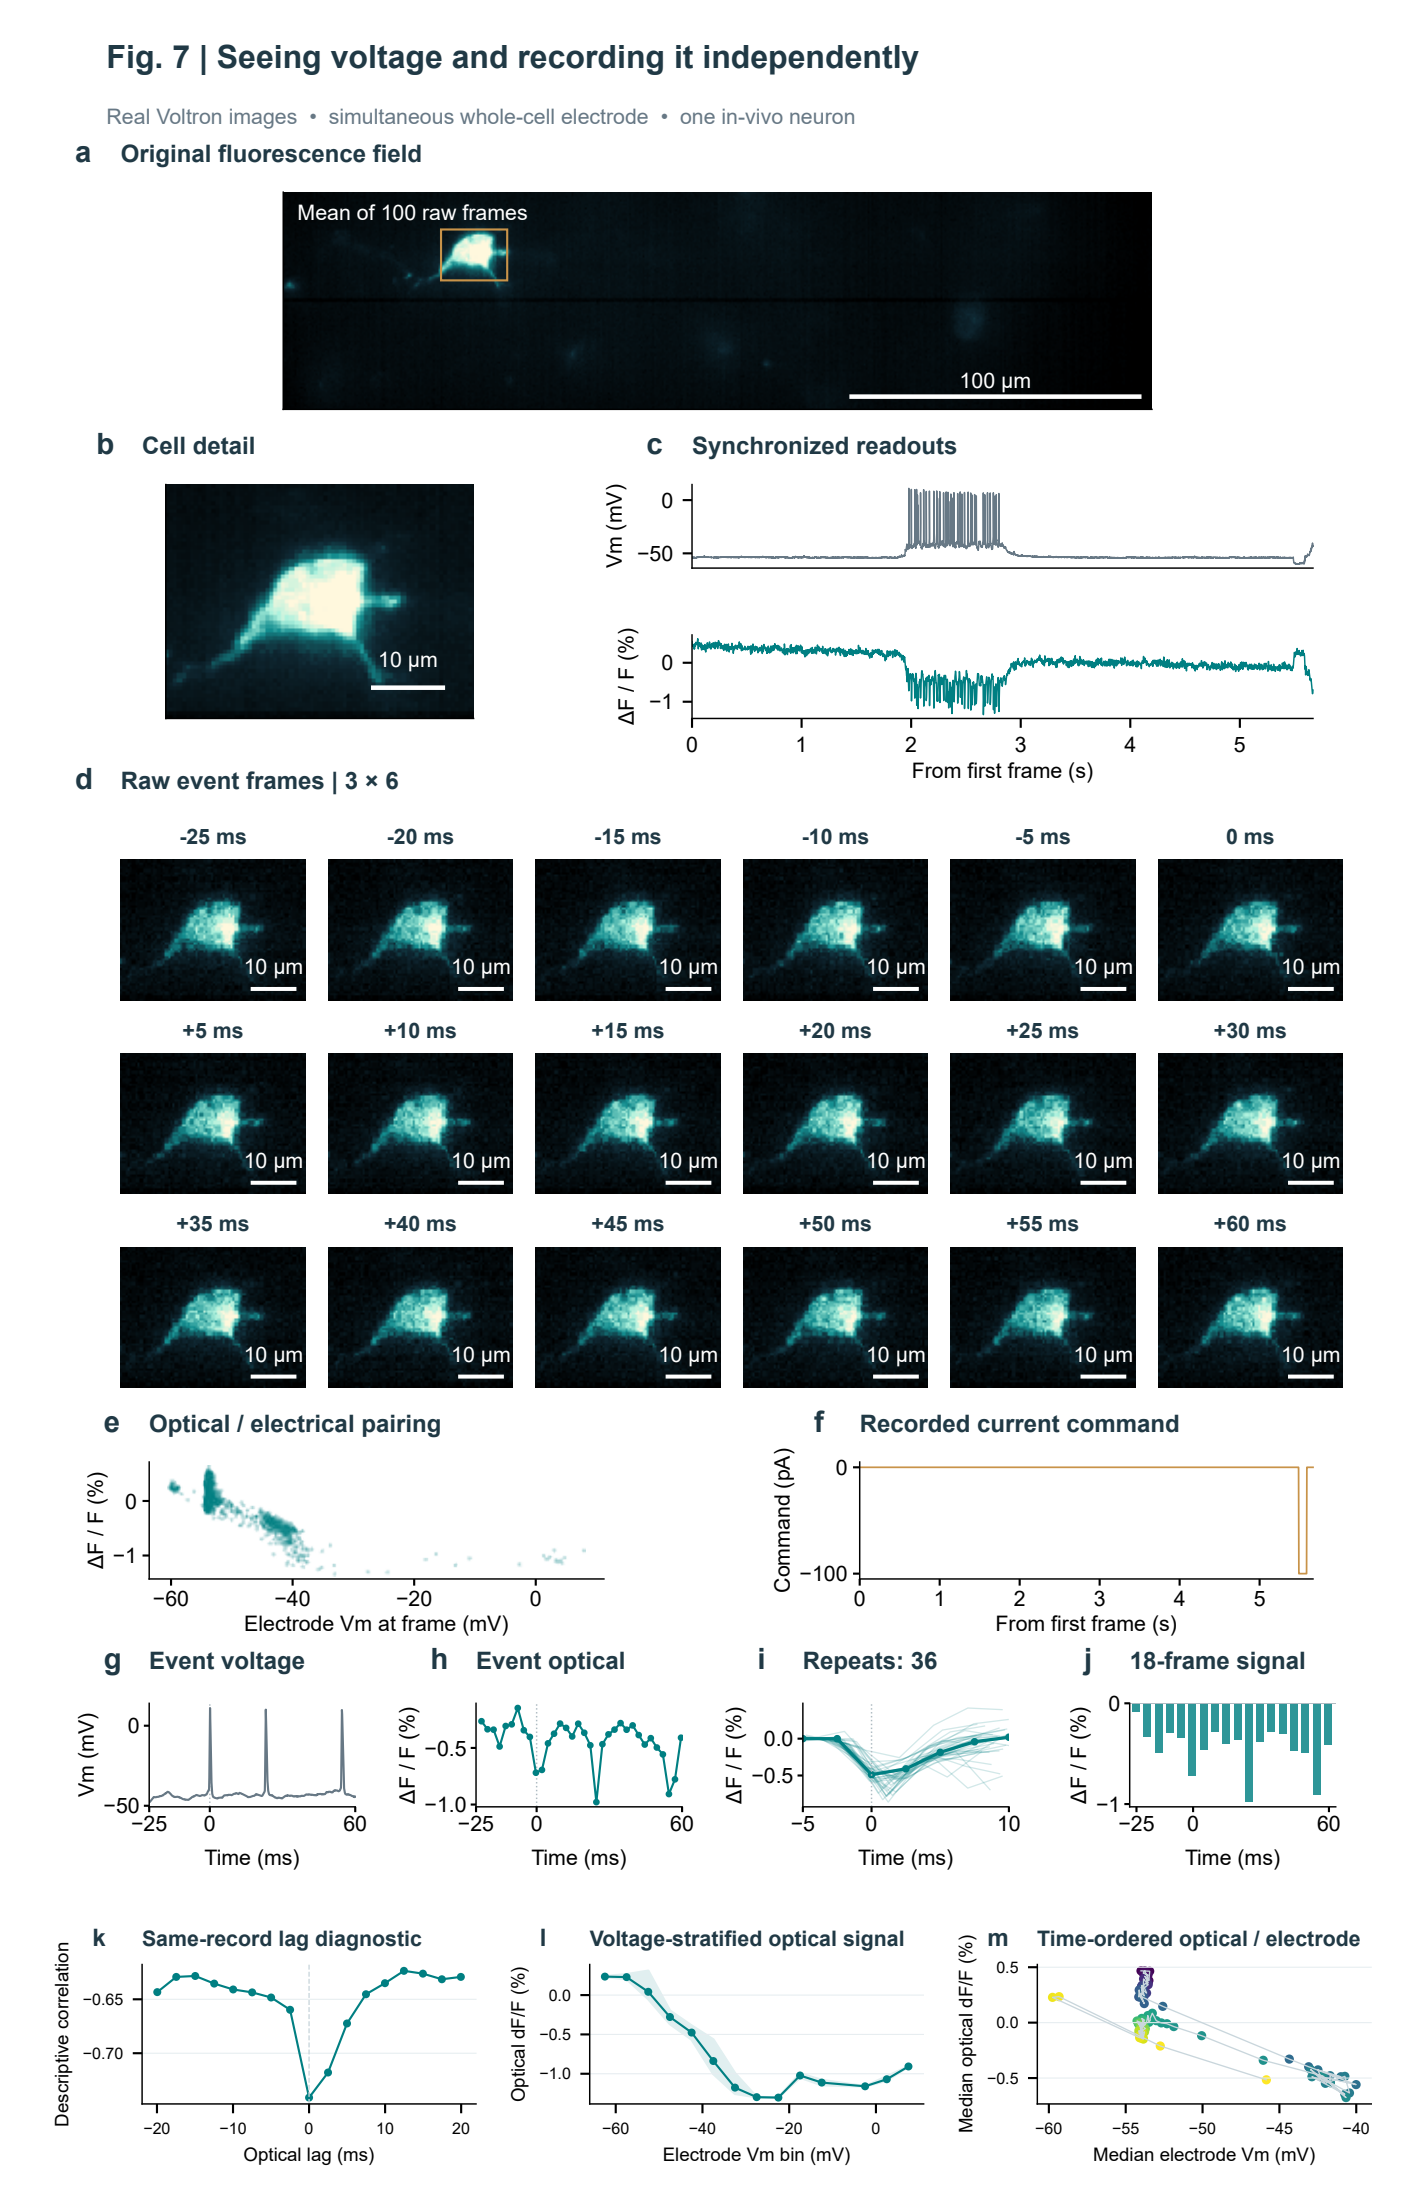

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Additionalfile2DataS2.zip
- Additionalfile3DataS3.zip
- Additionalfile1DataS1.pdf