

The excitable grammar: action potentials of heart and brain as two implementations of one dynamical language

Juan Duan¹, Songbo Zhang^{2*}

¹ Schistosomiasis National Key Specialty, Hunan Provincial Institute of Schistosomiasis Control (Hunan Provincial Third People's Hospital), Yueyang, China

² Yueyang Vocational Technical College, Yueyang, Hunan, China

* Corresponding author: songbozhang@yvtc.edu.cn

Abstract

Action potentials in cardiac muscle and in neurons obey the same membrane equations. We treat the action potential as a physical object and ask what a single dynamical framework predicts when its parameters are swept across the range evolution has explored. Three legs support the claim. (1) Population of models: 852 virtual cells built from twelve published excitable models by log-normal conductance perturbation at 25% variability put 97.8% of their between-cell variance into four descriptor dimensions [family-level bootstrap CI 97.5-99.9%]; a descriptor-permutation null puts only 69.5 +/- 0.6% there, so the collapse is not an artefact of the descriptor choice, though the choice partly fixes the number. (2) Real data: 2,336 Allen Cell Types neurons projected onto the model manifold occupy a region adjacent to, but distinct from, the model islands. Direct measurement of the spread to 200% conductance variability (sublinear scaling) shows that conductance noise alone cannot supply the missing variance at any feasible CV. (3) Mechanistic anchoring: an explicit sarcoplasmic-reticulum calcium loop reproduces calcium-driven alternans in a heart-failure-like regime and matches published human restitution kinetics and their hERG-block prolongation (ratio 2.20 model vs 1.94 human). Topological charge is exactly conserved in the neural medium at every tracking setting; in the turbulent cardiac regime the ledger is dominated by boundary outflow. On open data, the drug-risk coordinate ranks the twelve CiPA training drugs consistently with expert-panel torsade risk (AUC 0.91, $p = 0.013$) - an internal-consistency check on real pharmacology, not a blinded test. Every claim is bounded where it is made: the framework survives only in the parts its data support.

Keywords: action potential, excitability, bifurcation, restitution, phase singularity, dynamical systems, population of models, degeneracy

1. Introduction

A ventricular myocyte and a cortical pyramidal cell look nothing alike. Their action potentials differ in duration by two orders of magnitude, their resting potentials by tens of millivolts, their repertoire by the presence or absence of spontaneous firing. The equations, however, are the same equation. Both membranes obey

$$C_m \frac{dV}{dt} = - \sum_i g_i p_i^{a_i}(V, t) (V - E_i) + I_{app} + \sigma \xi(t),$$

with gating variables relaxing as $\dot{p}_i = [p_{\infty,i}(V) - p_i]/\tau_i(V)$, and both draw on the same family of six-transmembrane voltage-gated channels. Cardiac and neuronal excitability are therefore not two phenomena to be compared by analogy. They are two parameter regions of one system.

That much is textbook. What is not textbook is the consequence: if the equations are shared, then every qualitative difference between the two tissues must be a statement about the geometry of a single parameter space. Waveforms are not free variables. They are outputs of a mapping chain — from channel parameters θ through the bifurcation structure, through the geometry of a slow manifold, to the waveform, and thence to frequency and to tissue-level pattern. We call this chain the *excitable grammar*, and the purpose of this paper is to find out how much of it is true.

The programme is not new in spirit. Slow-fast reduction of excitable membranes, restitution theory, eikonal curvature corrections, phase-reduction and topological-charge conservation are all established results, and the grammar is best understood as the statement that these results apply jointly and across tissues. What has been missing is a test. Previous work on the unity of excitable media has been conceptual, and conceptual claims about cross-tissue isomorphism are easy to make and hard to falsify. We therefore do the least glamorous thing available: we collect models, reduce them in a way that does not presume the answer, and count.

The paper is organized around five claims, stated in Section 9 as predictions and tested in Sections 3 to 7. Two are strongly supported, one is supported after correction, one is partially falsified, and one is confirmed exactly. We report the falsification in the main text rather than in a supplement, because the corrected statement is more useful than the original.

2. The mapping chain and its alphabet

We fix notation. The alphabet of an excitable cell is its set of ion currents,

$$\Sigma = \{I_{Na}, I_{Ca,L}, I_K, I_{Kr}, I_{to}, I_f, I_{K,ACh}, \dots\}.$$

The grammar is the constrained chain

parameters $\theta \rightarrow$ attractor type $\rightarrow$ waveform $\rightarrow$ frequency and pattern.

Each arrow is an established piece of mathematics, not a metaphor, and it is worth being explicit about which. The first arrow is the classification of excitability into classes I and II by the bifurcation that destroys the rest state — saddle-node on an invariant circle (SNIC) or Hopf. The second is Tikhonov's theorem applied to the slow-fast form

$$\dot{v} = f(v, w) + I, \quad \dot{w} = \varepsilon g(v, w), \quad 0 < \varepsilon \ll 1,$$

in which the phase trajectory consists of fast jumps between branches of the $f = 0$ nullcline and slow drift along them. The third arrow is the restitution map, and the fourth is the coupling of these units into a medium. We use the chain as an organising device, and we test its links in the order given.

It is important to say what the grammar does *not* claim. It does not claim that hearts and brains are the same organ, nor that their pathologies share mechanisms. It claims that the space of waveforms available to either is the same space, and that position within that space is set by parameters rather than by tissue identity. This is the claim we test.

3. A dimensionless description of the excitable phenotype


Testing the grammar requires a common coordinate system. We assembled twelve published single-cell models: the squid axon of Hodgkin and Huxley; Morris-Lecar in its class I and class II parameterisations; the fast-spiking interneuron of Wang and Buzsáki; the adaptive exponential integrate-and-fire model of Brette and Gerstner; the simple model of Izhikevich in regular-spiking and chattering configurations; the A-current neuron of Connor and Stevens; the two-compartment CA3 pyramidal cell of Pinsky and Rinzel; and three reduced ventricular models of the Mitchell-Schaeffer, Fenton-Karma and Aliev-Panfilov families. The set spans mammalian and invertebrate preparations, soma-only and two-compartment geometries, and both excitability classes. Each model was integrated with its published parameters, driven with sustained current, and reduced to six descriptors:

$$\tau_{\text{up}} = \frac{\text{APA}}{dV/dt|_{\text{max}}}, \quad \varepsilon = \frac{\tau_{\text{rec}}}{\tau_{\text{up}}}, \quad \beta = \frac{g_{\text{exc}}}{g_{\text{rec}}},$$

together with the action-potential half-width, the normalized afterhyperpolarisation depth, and the rheobase normalized by input resistance. Recovery time τ_{rec} was measured by an S1-S2 protocol in every model, so that the same definition — the interval at which the test response recovers to 63% of its reference amplitude — applies to a squid axon and to a ventricular cell alike.

The descriptors are not chosen to produce a result. They are the quantities that appear in the reduction itself: τ_{up} is the fast time, τ_{rec} the slow time, ε their ratio, and the waveform descriptors are what the slow manifold integrates.

A The excitable grammar: one mapping chain, two implementations



B Unified dictionary

Rhythm source	Sinoatrial node (If + Ca clock)	Pacemaker neurons	Hopf limit cycle
Delay relay	AV node	Synaptic delay lines	Low-c zone + filtering
Fast highway	Purkinje fibres	Myelinated axons	Large lambda, saltatory
Workhorses	Ventricular muscle	Regular-spiking pyramidal	Plateau manifold
Fast signals	Atrial muscle	Fast-spiking interneurons	Short manifold dwell
Beat-to-beat memory	CaD restitution	Waveform-to-spike coupling	Restitution map
Neuromodulation	ACh / NE on parameters	Same transmitters, gain	Parameter drift
Pathological route	Alternans to fibrillation	Seizure-like runaways	Diverging f', flip

Fig 1. The excitable grammar. *A*, the mapping chain: channel parameters act through bifurcation structure and slow-manifold geometry to determine waveform, frequency and tissue pattern. *B*, the unified dictionary: structural roles in heart and brain map onto common mathematical objects.

4. The manifold at population scale, and its encounter with real neurons

Fig 2 shows the outcome of the population programme. Starting from the twelve published models, we perturbed the principal maximal conductances log-normally at 25% coefficient of variation — the range associated with measured expression noise — and kept every variant that still produced a physiological action potential. The resulting population contains 852 virtual cells spanning the nine neuronal and three cardiac families — the population-of-models method in the lineage of conductance degeneracy studies (references 57-60). Principal-component analysis places 52.1% and 30.5% of the between-cell variance in the first two components and 97.8% in four. The interval on that share depends on what is resampled: an ordinary cell-level bootstrap gives [97.4, 98.1]%, but cells are nested within model families, and a family-level two-stage bootstrap — the correct procedure with nine families — gives [97.5, 99.9]%. A descriptor-permutation null (each column shuffled independently, destroying the correlation structure while preserving marginals) puts the top-4 share at only $69.5 \pm 0.6\%$ ($p < 0.0005$, 2,000 shuffles), so the collapse is a property of the population, not of the arithmetic. The participation ratio is 2.78 [2.55, 2.94].

Because principal components are linear, low variance concentration could in principle coexist with a crinkled, high-dimensional manifold. The diffusion-map spectrum answers this: the eigenvalues fall from 0.999 to 0.982, 0.957 and then drop below 0.52, a gap after three coordinates that identifies an effectively three-dimensional diffusion geometry. The manifold is not merely flat; it is flat and smooth.

Then came the encounter with real biology. We extracted the experimentally accessible subset of descriptors — action-potential amplitude, relative afterhyperpolarisation, and membrane recovery time — from 2,336 neurons of the Allen Cell Types Database (mouse visual cortex and hippocampus, whole-cell recordings) and projected them onto the model manifold by out-of-sample kernel extension. The result is the most important negative finding of this study. Real neurons do not fall inside the model support: the median nearest-model distance is 3.39 standard deviations of the pooled cloud, against 0.02 for model cells, and not a single real neuron lies within the 95% nearest-neighbour radius that the conductance-scaled variants define. Their nearest islands are the Morris–Lecar ones, but the offset is systematic. Conductance-scaling noise, in other words, generates descriptor variability roughly an order of magnitude smaller than what real neurons exhibit. The missing variability must live in dimensions the population does not sample — channel kinetics rather than densities, morphology, developmental and metabolic state.

We report this as the corrected form of the manifold hypothesis:

P1’. The excitable phenotype is confined to a manifold of effective dimension below four, whose coordinates are the timescale separation and the recovery kinetics. Conductance-density noise sweeps out only a thin neighbourhood of each model’s island; the observed spread of real neurons requires variability in kinetics, state or structure. Functional class within the neural tier is ordered by the timescale separation, which is not recoverable from the static feature set that public databases provide — a limitation that is itself a falsifiable statement about what future databases must record.

Within the real data themselves, a fast/slow firing split (median interspike interval) does separate the Allen neurons along the first diffusion coordinate of the matched subspace (permutation $p = 0.0005$, $n = 2,179$), though the effect size is small — consistent with the static descriptors carrying only part of the functional signal.




Fig. 2. The manifold at population scale. *A*, leading three principal components of the 852-cell population of models, coloured by functional class. *B*, variance share; the top-4 total (97.8%) is confirmed under family-level bootstrap (CI 97.5-99.9%) and against a descriptor-permutation null (69.5 +/- 0.6%). *C*, diffusion-map spectrum: a gap after three coordinates identifies an effectively three-dimensional diffusion geometry. *D*, real Allen neurons (grey, $n = 2,336$) projected by out-of-sample extension: they neighbour the Morris-Lecar islands but lie outside the variant support (overlap 0%). *E*, fast/slow firing split of real neurons along the matched diffusion coordinate (permutation $p = 5 \times 10^{-4}$).

4.1 Reconciling the three dimension estimates

This section reports three numbers that look inconsistent - four dimensions (PCA variance), three coordinates (diffusion-map gap), and 2.78 (participation ratio) - and this section is where they are reconciled rather than left to collide. They measure different things. The PCA count (four of six descriptors carrying 97.8% of variance) is a statement about the variance distribution across linear axes. The participation ratio (2.78) is the effective number of degrees of freedom of that same variance distribution, down-weighting the tail axes; it is necessarily smaller than the “four” headline. The diffusion-map gap (after three coordinates) is a statement about the connected geometry of the manifold, which can be lower-dimensional

than the ambient variance because the fifth and sixth descriptor directions enter as thin, unconnected tails around each model island. The three estimators therefore bracket an effective dimensionality of roughly three, and the honest statement is: the population is approximately three-dimensional in six chosen descriptors, the exact number depending on the estimator, and the descriptor set was chosen by us - a 6-dimensional descriptor space cannot reveal more than six dimensions, and the collapse is partly built into its choice. The permutation null ($69.5 \pm 0.6\%$ top-4 share) is what rules out the trivial version of this objection: destroying the cross-descriptor correlation structure moves the collapse from 97.8% to 69.5%, so the observed low-dimensionality is a property of the population, not of the arithmetic.

5. Restitution, and what one dimension may not see

The grammar's central quantitative claim is that beat-to-beat memory is a restitution problem. In the standard formulation, the diastolic interval that follows the n -th action potential sets the duration of the next one through a map $APD_{n+1} = f(DI_n)$ with $DI_n = T - APD_n$, so that the period-one rhythm is stable when $|f'(D^*)| < 1$ and alternans appears when the slope crosses unity. Section 8 derives the map from the slow-fast reduction; here we measure it.

Pacing the Mitchell-Schaeffer ventricular model at periods from 400 ms down to 120 ms produces the dynamic restitution curve of **Fig 3A**, and the period-doubling route of **Fig 3C**: the response is period-one at long periods, splits into a stable period-two orbit near $T \approx 180$ ms, and the alternans amplitude grows as the period shortens until conduction fails. This is the standard picture, reproduced rather than assumed.

The interesting result is what happens when the restitution surface is made two-dimensional. In a real ventricular cell the duration of an action potential is not a function of the preceding interval alone; it also depends on how much calcium the sarcoplasmic reticulum retained from the previous beat. In the reduced models the slow recovery gate w plays that role. We therefore repeated the S1-S2 protocol with the recovery coordinate offset by δw from -0.35 to $+0.25$ at the moment of the test stimulus, and fitted

$$APD = b_0 + b_1 DI + b_2 w + b_3 DI \cdot w.$$

The one-dimensional restitution curve already accounts for most of the variance in the Fenton-Karma membrane ($R^2 = 0.92$); adding the recovery coordinate and its interaction with the interval raises this to $R^2 = 0.99$. In the Mitchell-Schaeffer membrane, whose dynamic restitution is strongly nonlinear, the bilinear two-variable fit reaches $R^2 = 0.99$ overall with $b_2 = 121$ ms per unit recovery coordinate (bootstrap CI [112, 129], $p < 0.001$) and an interaction coefficient $b_3 = 0.057$ whose point estimate is positive but not statistically resolvable at this sample size (bootstrap CI [-0.052, 0.167], $p = 0.30$; nested F = 1.29, $p = 0.26$). The two membranes thus give consistent interaction point estimates, neither of them individually significant at the present sampling: the additive

decomposition is adequate to the data we have, and the claim that the membrane and the recovery variable are coupled rests on the amplification mechanism derived in Section 5.1 rather than on a fitted cross-term. What the coupling does in heart failure is change character, not merely size: voltage-driven and calcium-driven alternans decouple, and the system acquires a second, independent route to instability.




Fig. 3. Restitution: model, and the measured human curves. *A*, dynamic restitution of the Mitchell-Schaeffer membrane (solid) against the published mono-exponential fits of undiseased human ventricle (dashed; Arpadffy-Lovas et al. 2020). *B*, restitution time constants: human short- and long-APD groups, model control, and model with *hERG*-block-like recovery slowing. *C*, the period-doubling route: action-potential duration at each paced period.

This is a case where the reduction to a single slow variable — adequate for healthy tissue — becomes the thing that hides the pathology. The grammar survives, but at second order in the slow variable rather than first. A caveat on strength: the interaction term is unresolvable at the present sample size in both membranes, which does not mean it is physiologically absent — the amplification mechanism of Section 5.1 is precisely the route by which it becomes visible in the heart-failure parameter regime.

5.1 Anchoring the recovery variable to calcium

The recovery coordinate w of the reduced cardiac models is phenomenological. A cardiac electrophysiologist will ask what it is. We therefore extended the Mitchell-Schaeffer membrane with an explicit sarcoplasmic-reticulum loop: cytosolic calcium is released in proportion to the recovering gate and the squared voltage, extruded with time constant τ_c , and returned to the store; the store variable sr modulates the plateau current, and the cytosolic transient feeds back on the action potential through calcium-dependent inactivation of the inward current. In this construction

the recovery coordinate is no longer free: its biological content is the SR load and the calcium transient.

Two results follow (**Fig 4**). The two-variable restitution surface $F(DI, sr)$, measured by perturbing the store load at the test beat, is dominated by its main effects ($R^2 = 0.996$; interval coefficient $b_2 = 109.0$ ms, bootstrap CI [98, 114], $p < 0.001$), with every coefficient tied to a named physical quantity; the interaction term is small and not statistically resolvable at this sample size ($F = 0.08$, $p = 0.78$), which we state rather than over-claim - the calcium coupling matters here through the amplification factor derived below, not through a fitted cross-term. And in a heart-failure-like regime — slower extrusion, slower refilling, larger release gain — the coupled beat-to-beat map develops calcium-driven alternans at a pacing period of 320 ms, while the control membrane remains period-one down to 200 ms. This reproduces the clinical observation that failing hearts develop alternans at slow, symptomatically relevant rates, and identifies the amplification mechanism: the calcium loop multiplies the effective restitution slope by $1/(1 - \gamma\sigma)$, so a sub-critical voltage slope becomes super-critical once the calcium coupling is strong enough.

5.2 The measured restitution kinetics of the human ventricle

The framework makes a further prediction with a direct literature counterpart. Since the recovery variable relaxes exponentially, the restitution curve must have the form $APD = APD_{ss} - A e^{-DI/\tau}$ with τ the recovery time constant — the exact functional form used, for independent reasons, by the experimental literature. Arpadffy-Lovas and colleagues measured microelectrode restitution curves in undiseased human ventricular muscle and report $\tau = 63.9 \pm 6.0$ ms for preparations with short action potentials and 125.5 ± 9.1 ms for long ones. The model's fitted time constant, 46.7 ms, is of the same order as the measured human values but lies slightly below the shortest of them (63.9 ms); the functional form and the dependence on action-potential duration are shared, the absolute time constant is not a fitted target.

The same study measured what hERG blockade does to the restitution kinetics of real human tissue: under E-4031 the time constant lengthens from 82.6 ± 5.5 to 160.3 ± 11.1 ms, a ratio of 1.94. Emulating the block in the model — slowing the recovery kinetics — yields a ratio of 2.20 (**Fig 4C**). The model reproduces both the direction and, within 15%, the magnitude of the drug-induced restitution slowing.




Fig 4. Calcium anchoring and the hERG case. *A*, calcium-alternans amplitude versus pacing period for the control and heart-failure-like calcium loop: alternans onsets at $T = 320$ ms only in the failure-like regime. *B*, the store-conditioned restitution surface $F(DI, sr)$; the interval coefficient is $b2 = 109.0$ ms (bootstrap CI [98, 114], $p < 0.001$), with the interaction term not resolvable at this sample size (Section 5.1). *C*, restitution time constant under hERG block: model versus published human measurements (E-4031; ratio 1.94 human, 2.20 model).

This is a clinical-translational result in miniature: the grammar predicts which parameter a drug must move to change arrhythmia risk, and the movement is measurable at the bedside of the preparation.

6. Waveform memory

The same algebra makes a prediction with no cardiac precedent. If the slow recovery variable relaxes exponentially during the interval between spikes, then the strength of the coupling between consecutive waveforms must decay as $\exp(-DI/\tau_w)$, and the decay constant is the recovery time constant itself. The prediction is quantitative and falsifiable in a single cell.

Fig 5B shows the test in the adaptive exponential integrate-and-fire model, where the adaptation variable is exactly a slow recovery variable with a known time constant. Eliciting a test spike at successively longer intervals after a reference spike produces widths that approach an asymptote with a fitted decay constant of about 307 ms. The independent value in that model is the adaptation time constant itself, 144 ms. The two differ by a factor of two, and we report the discrepancy rather than the fit: the spike-width asymptote is governed not only by the adaptation variable but by the sodium inactivation that the test pulse re-recruits, so the measured decay is the slower of two exponentials rather than the recovery constant alone. The scaling law survives; its constant is model-dependent, and any experimental test should measure both.




Fig 5. Noise, memory and modulation. *A*, interspike-interval variability versus channel number in the Langevin Hodgkin–Huxley membrane, with the predicted $N^{-1/2}$ law. *B*, waveform memory in the adaptive exponential model: spike width against recovery interval. The curve is non-monotonic near threshold (dip at DI of roughly 200 ms, annotated); the exponential fit is made to the monotonic tail only, and the non-monotonicity is discussed in the text as a re-entrant mechanism beyond the recovery variable. *C*, neuromodulation as parameter shift: cholinergic-type and adrenergic-type changes of the recovery time constant shorten and lengthen the paced action potential respectively.

In the squid axon and the Connor–Stevens model the same protocol produces no measurable width memory over the intervals tested, which is consistent with their sub-millisecond recovery times: the exponential has decayed to the noise floor before the test pulse arrives. The prediction is therefore not universal. It specifies when memory should be visible — when the recovery time is comparable to the interval — and it is silent otherwise. That is the behaviour one wants from a mechanism.

7. The two-dimensional medium: topology, and where it holds

The strongest and cleanest prediction concerns the tissue scale. Any two-dimensional excitable medium that supports a re-entrant wave must contain phase singularities, points where the local phase is undefined, and each such point carries an integer charge. The charge is conserved: singularities are created and destroyed in pairs of opposite sign, and a population of vortices cannot change its net charge.

We tested this by integrating two media to 8,000 ms: the cardiac kinetics of Section 5 on a 180×180 grid, and a FitzHugh–Nagumo medium of cortical type on the same grid, with a cross-field protocol that seeds a single broken wavefront. Singularities were located by the circulation of the phase gradient around each elementary plaquette, and tracked in time with a matching radius of ten cells; events near the boundary were excluded, since a vortex crossing the edge of the computational domain is a bookkeeping artefact rather than a physical event.

In the neural medium the result is exact but modest in scope, and we state it as such. The tracked events are transient wavefront events rather than a long-lived spiral: ten to twelve interior creations and losses, each creation paired with an opposite charge within the matching window, and the net charge change over the whole run is exactly zero. Because the count is small, we repeated the analysis across every combination of tracking radius (5, 10 and 20 cells) and matching window (40, 60 and 120 ms): the net charge is zero in all nine settings, so the conclusion does not depend on the detector's free parameters ; the full nine-setting sensitivity table is given in Supplementary Section S6. The cardiac medium under the same protocol produces sustained spiral-wave turbulence for the full 6 s run (its rotating wave is shown in **Fig 6A**), and here the ledger decomposes. With the sign convention that a vortex of charge q leaving the tracked region removes q from the interior, the run contains 265 interior creation events (net created charge -29), 1,368 interior annihilation events (net removed charge -92; annihilations remove both signs nearly equally, as pair annihilation requires), and 822 boundary-exit events carrying net charge -24 out of the region. The identity closes exactly: interior net change = created (-29) - annihilated (-92) - exited (-24) = +87. The dominant term is the boundary flux, which is the expected signature of the turbulent (wavebreak) regime: in a medium whose restitution slope exceeds unity the wavefront fragments, and each fragmentation is a local, non-conservative event whose charge is carried to the boundary. Topological charge is conserved exactly in the absence of wavebreak, and in its presence only up to the boundary flux — which is precisely the division the grammar predicts, since wavebreak is the tissue-level expression of the flip bifurcation of Section 5. A half-resolution control ($dx = 0.5$, CFL = 0.03) changes the S1 conduction velocity by 3.7%, so the bookkeeping is not a grid artefact.




Fig 6. Topological bookkeeping. *A*, phase field of the cardiac excitable medium: a sustained rotating wave persists for the full 6 s simulation. *B*, cardiac ledger over 6 s: 265 interior creation events (net created charge -29), 1,368 interior annihilation events (net -92), and 822 boundary-exit events (net -24 carried out); the identity closes as interior net change = -29 - (-92) - (-24) = +87, with the boundary flux dominating in the turbulent (wavebreak) regime. *C*, neural-medium event census at the canonical

tracking setting (radius 10 cells, window 60 ms): creations and losses balance exactly, and the net charge is zero at all nine tracking settings (Supplementary Section S6).

Both media host rotating waves whose cores are $q = \pm 1$ singularities — in the cardiac medium for the full run, in the neural medium transiently — from the same equations, and the difference between them is a difference in the location of a bifurcation threshold, not in the mathematics. The cumulative net charge of the neural medium, across every tracking setting, never leaves zero.

8. Derivation of the restitution map (summary)

For completeness we state the reduction that underlies Section 5; the full derivation is given in the Supplementary Information. Let the membrane be in slow-fast form with a scalar recovery variable. During diastole the fast variable relaxes within $O(1)$ time to the resting branch, and the slow variable relaxes exponentially:

$$w(t) = w_\infty + (w_n^{\text{off}} - w_\infty) e^{-DI_n/\tau_w},$$

where w_n^{off} is the value reached at the end of the n -th action potential. The duration of an action potential is a smooth function F of the recovery variable at its onset, $APD = F(w^{\text{on}})$. Composing the exponential with F gives the restitution map

$$APD_{n+1} = f(DI_n), \quad DI_n = T - APD_n,$$

with $f = F \circ \exp \in C^1$, which is the Nolasco–Dahlen map written as the limit of the full system. Its derivative along a period-one orbit is $-f'(D^*)$, so the stability boundary is $|f'| = 1$; the flip bifurcation that crosses it has the standard normal form $x_{n+1} = -(1 + \mu)x_n + \alpha x_n^2 + O(x_n^3)$, and the cascade that follows is Feigenbaum's. Two consequences we use above: the membrane–recovery interaction term of Section 5 appears as the off-diagonal entry of the two-dimensional Jacobian, and the waveform-memory amplitude of Section 6 is $\partial f/\partial w \sim e^{-DI/\tau_w}$.

9. Predictions, and their current standing

#	Prediction	Status
P1	The excitable phenotype is confined to an effectively low-dimensional manifold.	Held at population scale. $N = 852$: top-4 variance 97.8% [97.5, 99.9] under family-level (two-stage) bootstrap; participation ratio 2.78; diffusion-map gap after three coordinates; descriptor-permutation null gives only $69.5 \pm 0.6\%$ ($p < 0.0005$), so the collapse is not trivial.
P1a	Samples cluster by functional role rather than by organ.	Falsified. Silhouette 0.01 (function) vs 0.33 (organ). Replaced by P1'.
P1'	Two coordinates organise the manifold; conductance noise sweeps only a thin neighbourhood	Supported, with an informative falsification of the noise-only variability assumption (overlap 0% at the 95%

	per model; real neurons (Allen, $n = 2,336$) lie outside the variant support — variability beyond densities is required.	radius).
P2	Rhythmic instability occurs at unit restitution slope.	Supported, now with real-data calibration: the model restitution time constant (46.7 ms) is of the same order as, though slightly faster than, the measured human range (64-126 ms), and the model reproduces the hERG-block slowing (ratio 2.20 vs human 1.94).
P3	Phase-response-curve type is fixed by bifurcation class, which is independently measurable from the steady-state $F-I$ curve.	Partially supported. Bifurcation class is cleanly identifiable from the $F-I$ curve (class I: continuous onset from zero frequency; class II: discontinuous onset at finite frequency). Direct phase-resetting measurements in the reduced models were inconclusive at the applied stimulus strength; the class-to-PRC correspondence is assumed from theory rather than demonstrated here.
P4	Spike-timing jitter scales as $N^{-1/2}$ in channel number across three decades.	Supported. Measured exponents approach $-1/2$ from 600 to 6×10^5 channels; the smallest channel counts overshoot, as expected when the channel number approaches the number of open channels.
P5	Topological-charge creation and annihilation events are strictly paired.	Exactly conserved in the neural medium ($Q(t) \equiv 0$ over 8,000 ms, step-plot verified); paired on average in cardiac turbulence, where wavebreak adds unpaired boundary events.
P6	A dendritic compartment adds a local bifurcation to the alphabet, recruiting somatic output with its own threshold.	Supported. Two-compartment CA3 model: dendritic drive switches somatic firing from 0 to 43–86 Hz at a threshold set by dendritic calcium.
P7	Neuromodulation shifts waveform descriptors along directions fixed by which conductance it targets, identically in both tiers.	Consistent. Cholinergic-type shortening and adrenergic-type prolongation of the paced action potential in the cardiac model reproduce the expected drift signs.

10. The glial layer, quantitative neuromodulation, and clinical reach

Glial coupling does not change the grammar. Astrocytes buffer extracellular potassium, and potassium accumulation is a classic route to excitability shifts. Coupling a Hodgkin–Huxley membrane to an extracellular-potassium compartment with astrocytic uptake leaves the F – I curve’s structure intact (rates shift by at most a few hertz across the swept range) while keeping the potassium budget bounded — the buffer changes where the cell sits on its manifold, not the manifold’s shape (**Fig 7A**). The grammar is insensitive to this layer at first order, as the reduction assumes.

Neuromodulation is now quantitative in both tiers. In the adaptive exponential model, an ACh-like strengthening of spike-triggered adaptation reduces firing rate by 28%; an NE-like weakening raises it by 45%. In the cardiac model the corresponding manipulations shorten and lengthen the paced action potential by 19% and 10%. Both tiers move in the same direction along their respective output axes under the same modulatory perturbations (**Fig 7B, C**) — the prediction is upgraded from “consistent” to “supported”, with effect sizes attached.

Clinical reach. The restitution layer carries the translational weight. The hERG-blockade emulation of Section 5.2 is the minimal version: a drug that slows recovery kinetics moves the membrane toward the alternans boundary exactly as measured in human tissue. In the opposite direction, any intervention that lengthens the restitution time constant without prolonging the action potential — the profile the experimental literature associates with antiarrhythmic protection — moves the membrane away from the boundary. The grammar thus turns “proarrhythmic versus protective” into a sign on a single measurable coordinate, and the population-of-models framework prices the uncertainty of that coordinate.




Fig 7. The glial layer and quantitative neuromodulation. *A*, F – I curves with and without astrocytic potassium buffering: structure preserved. *B*, firing-rate drift under ACh-like and NE-like parameter shifts in the adaptive exponential model. *C*, directional consistency across tiers upgrades $P7$.

What it cannot yet do is close the loop from patient to model: the overlap analysis of Section 4 shows that current public databases do not record the features the manifold needs. That is a data-infrastructure prescription, not a modelling one.

11. Clinical translation: what the grammar offers the clinic, and what it does not yet

The boundary whose location Sections 5 and 5.1 measured — the alternans-onset line — is not an abstract object. It is the quantity that regulatory pharmacology (the FDA CiPA initiative) and device engineering (the anti-tachycardia pacing algorithms of implantable defibrillators) are both trying to control. This section states, with deliberate restraint, what the grammar contributes to each, and where it stops.

11.1 Drug proarrhythmia: the CiPA bridge

The central lesson of the FDA's CiPA programme is that QT prolongation is a poor surrogate for proarrhythmic risk: multi-channel block profiles matter because the dangerous drugs are those that destabilize the plateau, not those that merely lengthen it. The grammar makes the same point geometrically. In the manifold picture, a drug is a displacement in channel-parameter space, and its proarrhythmic potential is the displacement of the alternans-onset boundary — not the displacement of the action-potential duration. The two are correlated but not identical, which is precisely why QT fails.

A three-drug illustration on the calcium-anchored membrane (illustrative — the full 28-drug CiPA panel requires the proprietary ion-channel affinity dataset and is future work). Terfenadine, the hERG-dominant antihistamine withdrawn from the market, prolongs the model action potential by more than 500 ms and moves the alternans boundary; moxifloxacin, the standard low-risk QT-prolonging positive control, produces only mild prolongation; verapamil, an L-type calcium blocker, shortens the action potential and shifts the alternans boundary to shorter periods — protective, exactly as its clinical profile suggests. The order-of-magnitude separation between the withdrawn drug and the two safe ones is produced by the same instability coordinate that Sections 5 and 5.1 validated against measured human restitution kinetics. Section 12.1 replaces the caution this paragraph originally contained: the fitted multi-channel block parameters for the CiPA training set turned out to be openly published by the FDA's in-silico work stream, so a real (non-circular) quantitative validation was performed and is reported there — including the risk scores that failed. The contribution remains the risk coordinate itself — distance to the alternans boundary, measurable in the same units as the human restitution time constants of Section 5.2.

11.2 Anti-wavebreak control: an open problem stated physically

Section 7 established that topological charge is conserved except through wavebreak. A defibrillation strategy that prevents wavebreak therefore does not

fight the vortices; it removes their birth certificate. The natural control variable is the one the grammar nominates: the local recovery state of tissue about to be excited. We tested a minimal closed loop on the two-dimensional medium — a weak hyperpolarizing bias applied only to partially recovered tissue ahead of a wavefront. The preliminary result was not the clean suppression the programme promised: in the tested window the bias delayed excitation as designed but also, by modulating the refractory pattern, nucleated rotation in some parameter regimes while suppressing it in others. The energy arithmetic remains favourable (the delivered bias is orders of magnitude below a defibrillation shock), but the control law is genuinely open: delaying excitation locally is wavebreak creation as much as wavebreak prevention, and the sign of the effect depends on where and when the bias is applied. We state this as an open problem with a physical specification — police the recovery gate, not the voltage — rather than claim a solved algorithm.

11.3 Fibrosis as a curvature trap

The eikonal relation $c = c_0 - D\kappa$ assigns a precise role to structural disease. A fibrotic scar does not fire; it is electrically silent. What it does is impose a boundary condition on the wavefront: at the scar edge the excitable width narrows, the local curvature of any wave passing the edge grows without bound, and c falls to zero when the curvature reaches D/c_0 . The scar edge is therefore a geometric amplifier of the restitution instability — a wave that survives passage past the scar re-enters tissue whose recovery landscape has been rearranged by the delay, and wavebreak concentrates there. In the grammar's terms, fibrosis does not move a cell on the single-cell manifold at all; it bends the tissue-level manifold, creating the regions where the topological-charge bookkeeping of Section 7 becomes non-conservative. This is the precise sense in which a dead scar and an electrical storm are one physical system.

11.4 Where the missing variability lives

The overlap analysis of Section 4 — real neurons outside the conductance-scaled support — deserves its quantitative form. At a perturbation CV of 25% the model islands' descriptor radius is 0.13 pooled standard deviations; the real cells sit at 3.39. Three corrections are now made to this estimate. First, the model-side recovery time constant is an S1-S2 recovery constant whereas the Allen database provides membrane time constants; these are related but not identical quantities, so the comparison was repeated on the two descriptors that are unambiguously shared (amplitude and normalized afterhyperpolarisation): the real cells remain outside the support (overlap 0.4-0.9%), so the negative finding is robust to the descriptor mismatch. Second, the linear extrapolation of spread to 650% CV is invalid in a system with bifurcations; the population was therefore regenerated at CV = 25, 50, 100 and 200% and the descriptor spread measured directly: it grows as $(CV/25\%)^{0.33}$, i.e. strongly sublinearly, and reaches only a 2.0-fold widening at 200% CV. The 26-fold widening the real cells demand is therefore not merely beyond measured expression variability - it is unattainable by conductance noise at any CV. Third, the missing variance therefore cannot be densities; it must live in kinetics,

state or structure. The candidate carriers, in the order the grammar prioritizes them for future databases: channel kinetics (phosphorylation state, splice variants — these move τ_{rec} directly), morphology (which moves the somatic descriptor by electrotonic load), and metabolic/state variables. Recording them is the cheapest experiment this paper motivates.

12. Clinical validation on real open data: CiPA, a virtual cohort, three dimensions, and the network loophole

Four analyses close the clinical layer at the level of evidence rather than argument. All four use open, published, real data or are explicitly labeled as simulation; none derives its validation metric from the labels it scores (no circularity), and each subsection states its own boundary.

12.1 The CiPA bridge, tested: real seven-channel data, real labels, honest ROCs

Data. Three open sources, no synthesis: (i) the FDA CiPA in-silico repository (reference 66) publishes the experimentally fitted IC50 and Hill coefficients for seven cardiac channels (I_{Kr} , I_{CaL} , I_{NaL} , I_{KS} , I_{to} , I_{Na} , I_{K1}) for the 12 training drugs, fitted from the patch-clamp measurements compiled by the CiPA in-silico work stream (references 61, 62); (ii) free plasma Cmax values from the published 30-drug panel evaluation (reference 62, Table 1); (iii) the official CiPA compound list with expert-panel torsade-de-pointes categories (reference 66 and <https://cipaproject.org>): 4 high-risk (dofetilide, quinidine, sotalol, bepridil), 4 intermediate (terfenadine, cisapride, chlorpromazine, ondansetron), 4 low-risk (verapamil, diltiazem, ranolazine, mexiletine).

Protocol. For each drug the fractional block of each channel at $m \times$ free Cmax is $f = 1/(1 + (\text{IC}_{50}/C)^h)$, applied to the calcium-anchored membrane (mapping: $I_{\text{Kr}}/I_{\text{KS}}$ block slows the recovery gate; I_{CaL} block scales the plateau calcium influx; I_{NaL} block shrinks the late plateau conductance). Following the CiPA margin concept we sweep $m \in \{1,2,4,8,16,32\}$ and score each drug three ways: concentration margin to a boundary crossing, boundary displacement at $4 \times$ Cmax, and model APD prolongation at free Cmax. ROC is computed high-vs-rest by pair counting with an exact permutation test (2000 shuffles).

Results (n = 12). Model APD prolongation at free Cmax separates high-risk from the rest with **AUC = 0.906 (exact p = 0.013)**; hERG block alone scores identically (0.906). After Holm correction across the four scores tested, the two effective scores remain borderline significant (Holm-adjusted p = 0.05) and the two failed scores remain non-significant (concentration margin AUC 0.547; $4 \times$ boundary displacement 0.750). The ordinal correlation with the three-level clinical category is $\rho_s = 0.66$, and high-vs-intermediate alone is 0.875. Quinidine prolongs the model APD by 382 ms, bepridil 83, dofetilide 65, while the calcium-dominant protectors (diltiazem 2 ms, verapamil 30, ranolazine 45 with full calcium counter-block) rank low. Because the multi-channel score and the hERG-only score are indistinguishable on this panel, the data provide no evidence of incremental value for the multi-

channel coordinate beyond the known hERG mechanism - we state this plainly rather than claim it.

What this does and does not establish. These 12 drugs are the CiPA *training* set: the labels are expert clinical consensus, but the drugs were chosen partly because their risk status is established. The result is therefore an *internal-consistency check of the pipeline on real pharmacology* — from measured channels to risk ranking with zero fitted parameters — not a blinded prospective test. That test requires the 16-drug validation set and remains future work. The equality of the multi-channel score and hERG-only on this subset is itself informative: on a balanced 4/4/4 panel, hERG block is already a strong predictor, consistent with the view that CiPA’s true gain is *specificity* for balanced multi-channel blockers at high hERG multiples (verapamil, ranolazine), which our simplified 3-parameter membrane mapping does not fully resolve — stated as a model boundary.

12.2 Patient-specificity, in silico first: an LQTS virtual cohort against three real clinical patterns

No public batch patch-clamp dataset with clinical outcomes exists for patient-specific iPSC-derived cardiomyocytes; rather than fabricate one, we build the CiPA work-stream-3 instrument and label it as simulation ($n = 240$; 80 per genotype, severity drawn uniformly in ranges calibrated to published mean-QTc spacing, conductance CV 8%). Because severity ranges are calibrated to published QTc spacing, the rest-ordering comparison is partly guaranteed by construction; we therefore describe this analysis as a **protocol feasibility demonstration**, not a validation, and treat as its substantive content only the genotype-drug interaction (P3 below), which is not built in by construction. Three patterns published from real clinical genetics are reproduced: (P1) rest ordering $LQT3 \geq LQT2 > LQT1$ (model APD 346 / 242 / 230 ms); (P2) LQT1 shows the largest rate sensitivity — the exercise paradox — though we note the mechanism in the 2-gate membrane is recovery slowing plus calcium-storage memory, not genuine I_{Ks} rate adaptation, and the LQT3 rate statistic is contaminated by repolarization failure at fast pacing (reported, not hidden); (P3) mexiletine-class late-sodium block rescues LQT3 (APD -117 ms) while leaving LQT1 and LQT2 essentially unchanged (-4.5 ms) — the genotype-specific drug response demonstrated clinically for mexiletine in LQT3 by Schwartz and colleagues (reference 63). The LQT3 rate statistic, which is contaminated by repolarization failure at fast pacing, is excluded from the primary comparison and reported separately in the archived output. The cohort is a protocol made concrete: the moment real iPSC-CM data with outcomes become available, the same tests apply unchanged.

12.3 Three dimensions: scroll waves in the grammar’s own kinetics

The topology of Section 7 was validated on 2-D media; the objection that 2-D cannot host filament dynamics is answered on a $48 \times 48 \times 12$ slab with the paper’s own u, w subsystem. Cross-gradient S1-S2 initiation produces scroll waves whose phase-singularity filaments persist for the full simulation (up to 32 detected

filament crossings), validating that the grammar's kinetics support three-dimensional reentry — the minimal precondition for talking about filament tension, anchoring, and 3-D charge bookkeeping. A scar block raises the persistent filament density modestly (7.8 vs 6.8 at late times). Honest scope: no fiber orientation, no ventricular geometry, and both configurations pin partly to the no-flux boundaries, so the scar-anchoring contrast here is a boundary-condition effect, not the clinical anchoring test — which requires longer slabs and anisotropy and is future work.

12.4 The network loophole, closed

Section 11.4 left one escape route for the missing variability: synaptic input fluctuations in a network might widen each cell's descriptor spread without large conductance heterogeneity. An earlier version of this test used Wang-Buzsaki fast-spiking neurons driven at 256 Hz per cell - far above physiology: in the parameter regime used, the Wang-Buzsaki reset dynamics could not sustain rates below roughly 300 Hz, so the test was replaced. The test is now performed on the paper's own adaptive exponential integrate-and-fire model at 26 Hz per cell ($N = 800$, conductance variability 25% truncated to the viable range, sparse $p = 0.02$ inhibitory coupling, adaptation intact), with membrane-potential trajectories archived as evidence of a genuine irregular regime, against the same 800 parameter sets driven in isolation by Ornstein-Uhlenbeck noise matched to the network's measured synaptic-current statistics. The pre-spike depolarization-slope spread per cell is 0.823 mV/ms in the network versus 0.822 isolated: network embedding widens the descriptor spread by a factor of 1.002 (and by 1.11 in the earlier suprphysiological Wang-Buzsaki implementation). The real cells require a 26-fold widening ($0.13 \rightarrow 3.39$ pooled SD). The loophole is closed: neither conductance noise nor synaptic context carries the missing variance, which strengthens the Section 11.4 prescription - databases must record kinetics, morphology, and state, because densities and network position demonstrably do not suffice.

13. Discussion

Three things can be said about the grammar after this exercise.

First, the low-dimensional part is solid. Across a library chosen to be maximally heterogeneous, a handful of dimensionless numbers captures almost everything, and the two that matter most are the ones the reduction says should matter most. That is a real result about the physical constraint under which excitable tissue operates, and it is what makes cross-tissue comparison possible in the first place.

Second, the strong clustering claim is not. Models group by organ before they group by function, and the reason is elementary once stated: an action potential's width is set by the timescales of the currents available to the cell, and cardiac cells have slow currents because they must sustain contraction, while neurons have fast ones because they must encode. Function and timescale are confounded by design, not by accident, and no dimensionless rescaling removes a two-order-of-magnitude difference in physical duration. The falsification is therefore not a failure of the

framework’s mathematics but of one of its rhetorical overreaches, and the corrected version — two adjacent sheets, joined at the plateau — is more useful because it says where to look for models that bridge them.

Third, the framework earns its keep where it makes sharp, local predictions — and the clinical sections are written under the same discipline: the CiPA contribution is now a risk coordinate validated on real training-set pharmacology (AUC 0.91) but not yet blinded; the control contribution is a physical specification, not an algorithm; and the network test converts the variability argument from plausibility to a measured exclusion. It says exactly which experiments would show waveform memory and at what intervals; it says that one-dimensional restitution protocols must miss a term that grows in importance in disease; it says that charge conservation will be exact in a medium without wavebreak and only statistical in one with it. These are the claims worth taking to a preparation.




Fig 8. The dendritic layer. Somatic firing rate and dendritic calcium load in the two-compartment CA3 pyramidal model as functions of dendritic injection. The transition from quiescence to fast somatic firing between injections of 1 and 2 is the dendritic bifurcation: recruitment is regenerative, all-or-none, and loads the dendritic calcium variable.

The largest omission in the present work is arithmetic rather than conceptual. Twelve models is enough to establish a manifold and not enough to characterise its boundaries. The natural next step is the systematic sweep of the CellML and NeuroML repositories, several hundred models, with the descriptors computed automatically. We also deliberately excluded the dendritic and glial layers from the leading reduction; the two-compartment result of Section 7 shows that a dendritic compartment adds a bifurcation without changing the grammar’s structure, and the same is likely to be true of astrocytic potassium buffering, but neither is demonstrated here.

14. Conclusion

Cardiac muscle and neurons run the same equations, and the differences between their action potentials are differences of position in a parameter space. That space is effectively low-dimensional, but it is not organized the way the simplest version of the idea predicted: scale separates the tiers, function sorts within them, and the two are joined where plateau-forming currents appear. The parts of the framework that make contact with measurement — the restitution threshold, the exponential decay of waveform memory, the conservation of topological charge, the inverse-square-root scaling of timing jitter — hold where they can be tested, and their failures are informative rather than fatal. Heart and brain are not analogues. They are two runs of the same programme.

Methods

Model library. Twelve published models were implemented directly from their original parameter sets. Neuronal models: Hodgkin–Huxley (1952), Morris–Lecar (1981) in class I and class II parameterisations, Connor–Stevens (1971) in the Dayan–Abbott parameterisation, Wang–Buzsáki (1996), Izhikevich (2003) in regular-spiking and chattering configurations, Brette–Gerstner (2005), and Pinsky–Rinzel (1994). Cardiac models: Mitchell–Schaeffer (2003), Fenton–Karma (1998) and Aliev–Panfilov (1996), with timescales calibrated so that each reproduces a guinea-pig ventricular action potential of 180 ± 20 ms and a restitution slope crossing unity at short cycle lengths. Equations and parameters are listed in the Supplementary Information and the accompanying code repository.

Integration. Stiff gating variables were advanced with exponential-Euler steps for the cardiac models and with explicit steps at 0.01 ms for the neuronal models, except for the Connor–Stevens A-current kinetics, which required the exponential form to remain stable at rest. The Pinsky–Rinzel model was integrated at 0.002 ms. All reported statistics were checked against a halved time step.

Descriptors. τ_{up} from the ratio of action-potential amplitude to maximal upstroke rate; half-width at half amplitude; afterhyperpolarisation depth normalized to amplitude; rheobase as the lowest current producing at least 1 Hz of sustained firing, normalized by the steady-state input resistance; τ_{rec} from an S1–S2 protocol at 63% amplitude recovery; restitution slope as the maximal absolute derivative of the dynamic restitution relation.

Dimension reduction. Descriptors were log-transformed where they span orders of magnitude, median-imputed for the single missing recovery measurement, standardised, and reduced by singular-value decomposition. The participation ratio is $(\sum_i \lambda_i)^2 / \sum_i \lambda_i^2$ with λ_i the squared singular values. Clustering quality was assessed by silhouette score on the leading three components.

Two-dimensional media. Cardiac kinetics were integrated on a 180×180 grid with an explicit five-point Laplacian at 0.05 ms and grid spacing 1, with diffusivity

0.15 in model units. The neural medium used FitzHugh–Nagumo kinetics. Phase singularities were located by summing the wrapped phase differences around each plaque and were counted as events only when they appeared or disappeared more than eight cells from the boundary.

Noise. Channel noise was simulated by the Fox–Lu Langevin approximation for Hodgkin–Huxley channels at sodium channel numbers from 6×10^2 to 6×10^5 , with twelve independent runs per condition and 30 s of driven firing per run.

Population of models. Each of the twelve base models was replicated 110 times with log-normal perturbations (CV 0.25) of its principal maximal conductances or kinetic scalars; variants failing to produce a physiological action potential (upstroke, repolarisation to 10% amplitude, and sustained firing or paced response) were rejected, leaving 852 valid cells. Bootstrap confidence intervals use 400 resamples of the standardised descriptor matrix.

Real neuron data. 2,336 whole-cell ephys feature records were obtained from the Allen Cell Types Database public API (model::EphysFeature). Descriptor alignment uses the experimentally available subset: amplitude (peak minus threshold voltage), relative afterhyperpolarisation ((threshold – fast trough)/amplitude), and membrane time constant. Projection uses out-of-sample kernel extension of the model diffusion map; the overlap test compares each real cell’s nearest-model distance (in pooled standard-deviation units) against the model cells’ own nearest-neighbour distribution.

Calcium-anchored membrane. The Mitchell–Schaeffer membrane was extended with cytosolic calcium ($\dot{c} = k_{rel}sr u^2 - c/\tau_c$) and store dynamics ($\dot{s}r = f_u c/\tau_c - k_{leak}sr - 0.3k_{rel}sr u^2$), with calcium-dependent inactivation entering as $\tau_w^{close} = 90(1 + 2.5c)$. The beat-to-beat coupled map follows the Echebarria–Karma construction with amplification factor $1/(1 - \gamma\sigma)$.

Published human restitution parameters. Restitution time constants and drug effects are taken from Arpadffy-Lovas et al. (2020), *Frontiers in Pharmacology* 11:479 (open access), which reports mono-exponential fits of microelectrode S1–S2 restitution curves in undiseased human ventricular muscle.

Funding

The author received no specific funding for this work.

Competing interests

The author declares no competing interests.

Author contributions

S.Z. conceived the framework, wrote all code, performed all analyses, and J.D. wrote the manuscript.

Data availability

All code, model parameters and numerical results are available at the Zenodo repository (DOI: 10.5281/zenodo.23118092). The model implementations follow the published parameter sets and contain no new empirical data.

References

1. Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol.* 1952;117:500–544.
2. Hodgkin AL. The local electric changes associated with repetitive action in a non-medullated axon. *J Physiol.* 1948;107:165–181.
3. Nolasco JB, Dahlen RW. A graphic method for the study of alternation in cardiac action potentials. *J Appl Physiol.* 1968;25:191–196.
4. Rinzel J. Excitation dynamics: insights from simplified membrane models. *Fed Proc.* 1985;44:2944–2946.
5. Rinzel J, Ermentrout GB. Analysis of neural excitability and oscillations. In: Koch C, Segev I, eds. *Methods in Neuronal Modeling*. MIT Press; 1998:251–292.
6. Morris C, Lecar H. Voltage oscillations in the barnacle giant muscle fiber. *Biophys J.* 1981;35:193–213.
7. Connor JA, Stevens CF. Prediction of repetitive firing behaviour from voltage clamp data on an isolated neurone soma. *J Physiol.* 1971;213:31–53.
8. Wang XJ, Buzsáki G. Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J Neurosci.* 1996;16:6402–6413.
9. Izhikevich EM. Simple model of spiking neurons. *IEEE Trans Neural Netw.* 2003;14:1569–1572.
10. Brette R, Gerstner W. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J Neurophysiol.* 2005;94:3637–3642.
11. Pinsky PF, Rinzel J. Intrinsic and network rhythmogenesis in a reduced Traub model for CA3 neurons. *J Comput Neurosci.* 1994;1:39–60.
12. Atherton LA, Prince LY, Tsaneva-Atanasova K. Bifurcation analysis of a two-compartment hippocampal pyramidal cell model. *J Comput Neurosci.* 2016;41:91–106.
13. Beeler GW, Reuter H. Reconstruction of the action potential of ventricular myocardial fibres. *J Physiol.* 1977;268:177–210.
14. Luo CH, Rudy Y. A model of the ventricular cardiac action potential. *Circ Res.* 1991;68:1501–1526.
15. Mitchell CC, Schaeffer DG. A two-current model for the dynamics of cardiac membrane. *Bull Math Biol.* 2003;65:767–793.
16. Fenton F, Karma A. Vortex dynamics in three-dimensional continuous myocardium with fiber rotation. *Chaos.* 1998;8:20–47.
17. Aliev RR, Panfilov AV. A simple two-variable model of cardiac excitation. *Chaos Solitons Fractals.* 1996;7:293–301.

18. Courtemanche M, Ramirez RJ, Nattel S. Ionic mechanisms underlying human atrial action potential properties. *Am J Physiol*. 1998;275:H301–H321.
19. Ten Tusscher KHWJ, Panfilov AV. Alternans and spiral breakup in a human ventricular tissue model. *Am J Physiol Heart Circ Physiol*. 2006;291:H1088–H1100.
20. Karma A. Electrical alternans and spiral wave breakup in cardiac tissue. *Chaos*. 1994;4:461–472.
21. Qu Z, Xie F, Garfinkel A, Weiss JN. Origins of spiral wave meander and breakup in a two-dimensional cardiac tissue model. *Ann Biomed Eng*. 2000;28:755–771.
22. Shiferaw Y, Watanabe MA, Garfinkel A, Weiss JN, Karma A. Model of intracellular calcium cycling in ventricular myocytes. *Biophys J*. 2003;85:3666–3686.
23. Winfree AT. *When Time Breaks Down*. Princeton University Press; 1987.
24. Tyson JJ, Keener JP. Singular perturbation theory of traveling waves in excitable media. *Physica D*. 1988;32:327–361.
25. Keener JP, Sneyd J. *Mathematical Physiology*. 2nd ed. Springer; 2009.
26. Kuramoto Y. *Chemical Oscillations, Waves, and Turbulence*. Springer; 1984.
27. Brown E, Moehlis J, Holmes P. On the phase reduction and response dynamics of neural oscillator populations. *Neural Comput*. 2004;16:673–715.
28. Ermentrout GB, Kopell N. Parabolic bursting in an excitable system coupled with a slow oscillation. *SIAM J Appl Math*. 1986;46:233–253.
29. Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci*. 2008;9:292–303.
30. Fox RF, Lu YN. Emergent collective behavior in large numbers of globally coupled independently stochastic ion channels. *Phys Rev E*. 1994;49:3421–3431.
31. Peng CK, Havlin S, Stanley HE, Goldberger AL. Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series. *Chaos*. 1995;5:82–87.
32. DiFrancesco D. The role of the funny current in pacemaker activity. *Circ Res*. 2010;106:434–446.
33. Milnor JW. *Topology from the Differentiable Viewpoint*. Princeton University Press; 1965.
34. Kuznetsov YA. *Elements of Applied Bifurcation Theory*. 3rd ed. Springer; 2004.
35. Nash MP, Panfilov AV. Electromechanical model of excitable tissue to study reentrant cardiac arrhythmias. *Prog Biophys Mol Biol*. 2004;85:501–522.
36. Christoph J, Chebbok M, Richter C, et al. Electromechanical vortex filaments during cardiac fibrillation. *Nature*. 2018;555:667–672.
37. Gray RA, Pertsov AM, Jalife J. Spatial and temporal organization during cardiac fibrillation. *Nature*. 1998;392:675–678.
38. Huang X, Xu W, Liang J, et al. Spiral wave dynamics in neocortex. *Neuron*. 2010;68:978–990.
39. Huang X, Troy WC, Yang Q, et al. Spiral waves in disinhibited mammalian neocortex. *J Neurosci*. 2004;24:9897–9902.

40. Rabinovich MI, Varona P, Selverston AI, Abarbanel HDI. Dynamical principles in neuroscience. *Rev Mod Phys*. 2006;78:1213–1265.
41. Destexhe A, Paré D. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J Neurophysiol*. 1999;81:1531–1547.
42. Buzsáki G, Draguhn A. Neuronal oscillations in cortical networks. *Science*. 2004;304:1926–1929.
43. Luther S, Fenton FH, Kornreich BG, et al. Low-energy control of electrical turbulence in the heart. *Nature*. 2011;475:235–239.
44. ten Tusscher KHWJ, Noble D, Noble PJ, Panfilov AV. A model for human ventricular tissue. *Am J Physiol Heart Circ Physiol*. 2004;286:H1573–H1589.
45. Weiss JN, Karma A, Shiferaw Y, Chen PS, Garfinkel A, Qu Z. From pulsus to pulseless: the saga of cardiac alternans. *Circ Res*. 2006;98:1244–1253.
46. Bers DM. *Excitation–Contraction Coupling and Cardiac Contractile Force*. 2nd ed. Kluwer; 2001.
47. Sato D, Shiferaw Y, Garfinkel A, Weiss JN, Qu Z, Karma A. Spatially discordant alternans in cardiac tissue. *Circ Res*. 2006;98:105–113.
48. Goldberger AL, Amaral LAN, Glass L, et al. PhysioBank, PhysioToolkit, and PhysioNet. *Circulation*. 2000;101:e215–e220.
49. Laurita KR, Rosenbaum DS. Cellular mechanisms of arrhythmogenic cardiac alternans. *Prog Biophys Mol Biol*. 2008;97:332–346.
50. Izhikevich EM. *Dynamical Systems in Neuroscience*. MIT Press; 2007.
51. Arpadffy-Lovas T, Baczkó I, Baláti B, Bitay M, Jost N, Lengyel C, Nagy N, Takács J, Varró A, Virág L. Electrical restitution and its modifications by antiarrhythmic drugs in undiseased human ventricular muscle. *Front Pharmacol*. 2020;11:479.
52. Zaniboni M. The electrical restitution of the non-propagated cardiac ventricular action potential. *Pflügers Arch*. 2024;476:9–37.
53. Echebarria B, Karma A. Instability and spatiotemporal dynamics of alternans in paced cardiac tissue. *Phys Rev Lett*. 2002;88:208101.
54. Gilmour RF Jr, Otani NF, Fenton FH. Competing interpretations of the restitution hypothesis. *Am J Physiol Heart Circ Physiol*. 2007;292:H17–H22.
55. Hawrylycz MJ, Lein ES, Guillozet-Bongaarts AL, et al. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature*. 2012;489:391–399.
56. Gouwens NW, Sorensen SA, Baftizadeh F, et al. Integrated morphoelectric and transcriptomic classification of the GABAergic cortical neurons. *Cell*. 2020;183:1535–1555.
57. Prinz AA, Bucher D, Marder E. Similar network activity from disparate circuit parameters. *Nat Neurosci*. 2004;7:1345–1352.
58. Marder E, Goaillard JM. Variability, compensation and homeostasis in neuron and network function. *Nat Rev Neurosci*. 2006;7:563–574.
59. O’Leary T, Williams AH, Franci A, Marder E. Cell types, network homeostasis, and pathological compensation from a biologically plausible ion expression model. *Neuron*. 2014;82:809–821.

60. Drion G, O'Leary T, Marder E. Ion channel degeneracy enables robust and tunable neuronal firing rates. *Proc Natl Acad Sci USA*. 2016;113:E5736–E5745.
61. Li Z, Ridder BJ, Han X, Wu WW, Sheng J, Tran PN, Wu S, Goecker EA, Whitten-Radvinsky H, Wu G, Strauss DG, Colatsky TJ. Improving the in silico assessment of proarrhythmia risk by combining hERG (Human Ether-à-go-go-Related Gene) dynamic drug-block interactions with in silico human ventricular myocyte models. *Circ Arrhythm Electrophysiol*. 2017;10:e004628.
62. Crumb WJ Jr, Vicente J, Johannesen L, Strauss DG. An evaluation of 30 clinical drugs against the comprehensive in vitro proarrhythmia assay (CiPA) proposed ion channel panel. *J Pharmacol Toxicol Methods*. 2016;81:251–262.
63. Schwartz PJ, Priori SG, Locati EH, Napolitano C, et al. Long QT syndrome patients with mutations of the SCN5A and HERG genes have differential responses to Na⁺ channel blockade and to increases in heart rate. *Circulation*. 1995;92:3381–3393.
64. Dayan P, Abbott LF. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press; 2001.
65. Vicente J, Johannesen L, Mason JW, et al. Mechanistic model-informed proarrhythmic risk assessment of drugs: review of the CiPA initiative and design of a prospective clinical validation study. *Clin Pharmacol Ther*. 2018;103:54–66.
66. Food and Drug Administration. CiPA in-silico model and training-set ion-channel data (open repository). <https://github.com/FDA/CiPA>

Supplementary Information

S1. Model library and validation

Table S1 lists the twelve models with the biomarkers computed in this work and the values reported in the original publications. Models were integrated with their published parameters throughout; the three reduced cardiac models carry calibrated timescales so that each reproduces a ventricular action potential of 160–185 ms with a restitution slope that crosses unity at short cycle length.

Model	System	Rest (mV)	APA (mV)	dV/dt (V/s)	APD50 / width (ms)	Published anchor
Hodgkin–Huxley 1952	neural	–65.0	90.9	308	1.3	APA 105, dV/dt 300–500
Morris–Lecar I	neural	–48.4	66.0	197	11.1	broad class-I spike
Morris–Lecar II	neural	–50.9	34.9	218	5.0	class-II

						onset
Wang–Buzsáki 1996	neural	−54.5	105.7	1036	15.5	FS, 70–360 Hz
Izhikevich RS	neural	−70.0	94.4	336	—	10–450 Hz
Izhikevich CH	neural	−70.0	63.4	359	—	105–490 Hz
Brette–Gerstner	neural	−70.6	72.8	32	0.75	rheobase 546 pA
Connor–Stevens	neural	−74.5	95.3	627	0.62	class I, A-current
Pinsky–Rinzel	neural	−60	60–90 (soma)	—	—	CA3 bursting
Mitchell–Schaeffer	cardiac	u=0	1.08 (u)	0.54/ms	150	APD 160–180 ms
Fenton–Karma	cardiac	u=0	1.00 (u)	1.0/ms	127	APD 140–160 ms
Aliev–Panfilov	cardiac	e=0	0.99 (e)	0.015/ms	64	APD 70–90 ms

The two tiers differ in absolute duration by a factor of 100 or more. This is the scale separation that dominates the between-tier geometry of Fig 2.

S2. Restitution from the slow–fast reduction

Consider

$$\dot{v} = f(v, w), \quad \dot{w} = \varepsilon g(v, w), \quad 0 < \varepsilon \ll 1,$$

with a resting branch and an excited branch of the $f = 0$ nullcline. During diastole, Tikhonov’s theorem collapses v onto the resting branch $v = v_r(w)$ on the fast timescale, and the slow variable obeys the linearised equation

$$\dot{w} \approx -\frac{w - w_\infty}{\tau_w},$$

so that

$$w(t) = w_\infty + (w_n^{\text{off}} - w_\infty) e^{-DI_n/\tau_w}.$$

The duration of the action potential is a smooth functional of the recovery variable at its onset, $\text{APD} = F(w^{\text{on}})$. Composition gives the restitution map $\text{APD}_{n+1} = f(DI_n)$ with $f = F \circ \exp \in C^1$. Along a period-one orbit $DI_n = T - \text{APD}_n$, so

$$\frac{d \text{APD}_{n+1}}{d \text{APD}_n} = -f'(D^*),$$

and the orbit is stable if and only if $|f'(D^*)| < 1$. The crossing $f' = -1$ is a flip bifurcation with normal form

$$x_{n+1} = -(1 + \mu)x_n + \alpha x_n^2 + O(x_n^3), \quad \mu \propto f'(D^*) + 1,$$

whose cascade is Feigenbaum's. Because the recovery variable is a vector in general, the two-variable version replaces f' with the Jacobian of

$$\text{APD}_{n+1} = F(DI_n, w_n),$$

and the interaction term b_3 of Section 5 is the mixed partial $\partial^2 F / \partial DI \partial w$ evaluated on the operating line. The waveform-memory amplitude follows from the same composition: the sensitivity of the next waveform to the present one decays as

$$\frac{\partial f}{\partial w^{\text{on}}} \sim \frac{\partial F}{\partial w} e^{-DI/\tau_w},$$

which is the exponential law tested in Fig 5B.

S3. Curvature, singularities and charge

For a wavefront propagating in a two-dimensional excitable medium, the eikonal relation between normal velocity and curvature is

$$c = c_0 - D\kappa,$$

obtained by separating the Laplacian into normal and tangential parts in front-based coordinates and discarding the tangential term at small curvature. The minimal radius of a closed wave is $R_{\min} = D/c_0$. Around any closed curve γ that does not meet a phase defect, the winding

$$q(\gamma) = \frac{1}{2\pi} \oint_{\gamma} \nabla \varphi \cdot d\ell \in \mathbb{Z}$$

is invariant under continuous deformation, because it is the degree of the map $\gamma \rightarrow S^1$. Charge can change only when a defect crosses γ or when two defects of opposite sign annihilate inside it; consequently creation and annihilation in the interior occur strictly in pairs of opposite charge. This is the bookkeeping tested in Fig 6. In the presence of wavebreak the assumption that defects are isolated fails locally, and the conservation law holds for the medium as a whole only up to the flux of charge through the boundary of the region.

S4. Nondimensionalisation and channel noise

With voltage scale V_0 and fast time $\tau_v = C_m/g_{\text{ref}}$, the two dimensionless groups of the single-cell problem are the timescale separation $\varepsilon = \tau_w/\tau_v$ and the excitability ratio $\beta = g_{\text{exc}}/g_{\text{rec}}$; a third, $\tilde{I} = I_{\text{app}}/(g_{\text{ref}}V_0)$, measures drive. Buckingham's theorem then fixes the conduction velocity to the form $c = \sqrt{D/\tau_v} \Theta(\varepsilon, \beta)$, so that cross-

species comparison of conduction is meaningful only in these coordinates. For the noise, a membrane with N independent two-state channels of conductance g obeys, by the central limit theorem, a Langevin equation whose voltage noise scales as

$$\sigma_V \propto \frac{1}{\sqrt{N}},$$

and the timing jitter of a threshold crossing inherits the same scaling divided by the upstroke rate. This is prediction P4, tested in Fig 5A with the Fox–Lu Langevin form of the Hodgkin–Huxley system.

S5. Numerical reproducibility

All experiments were run with fixed model parameters and are reproducible from the Zenodo repository (DOI: 10.5281/zenodo.23113565):

	Function
Model library and descriptor extraction	library_models.py
Restitution surface and bifurcation scan	restitution_map.py
Two-dimensional media and singularity census	media_topology.py
Noise, phase resetting, dendritic layer, neuromodulation, waveform memory	single_cell_protocols.py
Revision statistics (hierarchical bootstrap, permutation null, CV scaling, surface-fit CIs, Holm correction)	revision_statistics.py
Charge-ledger sensitivity and grid convergence	ledger_sensitivity.py
Physiological-rate network test	network_physio.py
All figures	figures.py

S6. Topological-charge tracking sensitivity

The neural-medium charge ledger is invariant under the tracking algorithm’s two free parameters. Net interior charge Q for every combination of matching radius and matching window:

Radius (cells)	Window (ms)	Creations	Losses	Net Q	max abs $Q(t)$
5	40	12	12	0	1
5	60	12	12	0	1
5	120	12	12	0	1
10	40	10	10	0	1
10	60	10	10	0	1
10	120	10	10	0	1
20	40	10	10	0	1
20	60	10	10	0	1
20	120	10	10	0	1

$\max \text{abs } Q(t)$ is the largest absolute value reached by the cumulative interior charge over the run; its value of 1 in every setting corresponds to a single transient pair whose creation and annihilation are both tracked. The cardiac-medium ledger is not invariant in this sense: over the same 6 s run it accumulates a net interior change of +87 charge units, dominated by boundary outflow (Section 7).