Hard Limits on Sparse Bioelectric Control: Minimum Actuation Rank for Gap-Junction Network Steering

Chris Cable

March 17, 2026

Abstract

Bioelectric intervention is usually framed as a dose problem: stronger, longer, or broader stimulation is assumed to improve repair. We show that intervention dimensionality imposes a harder constraint. We formulate bioelectric tissue repair as finite-horizon optimal control on a gap-junction-coupled network and show that below a critical actuator count, no control signal of any amplitude can remove a damaged displacement on the chosen horizon. This actuation-rank threshold is a structural property of the spectral overlap between the damage pattern and the controllable subspace of the frozen linearized network. On an 8×8 lattice, the exact two-site result and greedy higher-rank placements show a geometry-dependent threshold at projected residual below 10^{-4}: 2 for a corner wound, 3 for a side patch, and 3 for a central lesion. On a 16×16 lattice under structured greedy placement, the same threshold phenomenon persists but the side patch becomes substantially harder, requiring 6 sites, while corner and central lesions remain low-rank (2). Below these thresholds, the unreachable residual is nonzero regardless of energy budget. Above them, each additional actuator site reduces residual by roughly one to two orders of magnitude. Random gap-junction dilution shifts the side-patch threshold upward, from stable rank 3 at 0–10% dilution to mostly 4–5 by 30% dilution. The controllable subspace is Laplacian-structured but not Laplacian-pure: dominant controllable A-modes overlap frozen Laplacian eigenvectors only moderately (0.46–0.57), showing that intrinsic voltage bistability materially reshapes network control directions. These results define a finite-horizon linearized baseline for intervention complexity in passive tissue and suggest a direct experimental program: compare predicted actuation-rank thresholds to observed minimum intervention to quantify endogenous morphogenetic competency.

1 Introduction

Bioelectric signaling is now understood to participate directly in growth, regeneration, and anatomical pattern control [3]. Experimental tools for perturbing and reading these states already exist: optogenetic actuators, ion-channel drugs, voltage-sensitive dyes, and applied electric fields all allow bioelectric manipulation at varying spatial scales. What remains missing is not the ability to perturb bioelectric state, but a quantitative theory telling an experimentalist where to intervene, how many independently addressable sites are required, and whether increasing stimulation strength can compensate for poor intervention geometry.

In practice, bioelectric intervention is still mostly interpreted through a dose lens. If a perturbation fails to redirect development or repair, the first adjustments are more current, longer exposure, higher drug concentration, or broader tissue coverage. That framing assumes that intervention strength is the main bottleneck. But in a distributed tissue network coupled by gap junctions, a more basic constraint may come first: whether the intervention has enough independent degrees of freedom to access the damaged state's relevant directions in voltage space at all.

The central quantity in this paper is the actuation-rank threshold: the smallest number of independently addressable intervention sites required to reduce the unreachable component of a damaged voltage displacement below a chosen tolerance on a fixed horizon. This is not an energy threshold. It is a dimensional threshold. Below it, no increase in control amplitude helps because part of the damage vector lies outside the controllable subspace generated by the actuator family. The motivating biological picture is that broad perturbations such as gap-junction blockade in planaria can alter regenerative polarity [5], but do so without a quantitative account of how many structured control directions are actually required.

Using a frozen linearization of a gap-junction-coupled tissue model, we show five results. First, sparse actuation generates a controllable subspace that is both small and ill-conditioned. Second, the dominant controllable directions are Laplacian-structured but not Laplacian-pure, indicating genuine mixing between network topology and intrinsic voltage dynamics. Third, realistic wound-like displacements can be exactly infeasible under low-rank actuation: on 8×8, every two-site actuator pair fails to exactly steer the side-patch wound to zero on the chosen horizon. Fourth, greedy actuation-rank thresholds depend on damage geometry. Fifth, the threshold shifts upward under both lattice scaling and gap-junction disorder. Conceptually, the work sits between structural controllability [4], minimum-energy network control [6], and nonlinear network control [2], but exploits the specific physics of bioelectric coupling rather than treating the tissue as a generic graph.

These results define a passive-tissue finite-horizon baseline. A real tissue with endogenous error-correction may require fewer external degrees of freedom than the frozen linearized calculation predicts, but that monotonic relation is not proved here. The gap between predicted intervention threshold and experimentally observed minimum intervention is therefore best viewed as a candidate quantitative measure of morphogenetic competency rather than as a formally bounded quantity.

2 Model

2.1 Controlled bioelectric dynamics

We model tissue as a gap-junction-coupled network of N cells with membrane-potential state $V(t) \in \mathbb{R}^N$. The controlled dynamics take the form

$$\frac{dV}{dt} = \frac{1}{\tau_V} F(V, G) + \frac{1}{\tau_V} B u(t),$$

where G is the gap-junction conductance field, $F(V, G)$ combines local voltage dynamics with diffusive coupling through the tissue network, B is the actuator matrix, and $u(t)$ is the control input.

The local voltage sector is the double-well force

$$f_{\text{local}}(V_i) = aV_i - bV_i^3,$$

with default parameters $a = 1$ and $b = 1$. Cells are coupled by gap-junction-mediated voltage differences over the tissue graph:

$$f_{\text{couple},i}(V, G) = \sum_{j \in \mathcal{N}(i)} g_{ij}(V_j - V_i).$$

The full sitewise drift is therefore

$$\frac{dV_i}{dt} = \frac{1}{\tau_V} \left[aV_i - bV_i^3 + \sum_{j \in \mathcal{N}(i)} g_{ij}(V_j - V_i) + h_i + c_i(t) \right],$$

where h_i collects static bias terms present in the damaged state and $c_i(t)$ is the projected site control.

The control cost is the standard finite-horizon quadratic energy

$$E[u] = \int_0^T \|u(t)\|^2 \, dt.$$

2.2 Linearization at the damaged operating point

We linearize around a damaged pre-intervention snapshot $(\bar{V}, \bar{G})$, with conductance frozen over the steering horizon. Writing $\delta V = V - V^*$ relative to the target reference state yields

$$\frac{d\delta V}{dt} = A\,\delta V + \widetilde{B}\,u(t),$$

with

$$A = \frac{1}{\tau_V}\,\frac{\partial F}{\partial V}\Big|_{(\bar{V},\bar{G})}, \qquad \widetilde{B} = \frac{B}{\tau_V}.$$

In components,

$$A_{ii} = \frac{1}{\tau_V}\left[a - 3b\bar{V}_i^2 - \sum_{j\in\mathcal{N}(i)} \bar{g}_{ij}\right], \qquad A_{ij} = \frac{1}{\tau_V}\bar{g}_{ij}$$

for neighbors $i \neq j$, and $A_{ij} = 0$ otherwise. Thus A inherits network structure without reducing to a pure graph Laplacian.

2.3 Finite-horizon steering problem

We study the problem

$$\min_u \int_0^T \|u(t)\|^2 \, dt$$

subject to

$$\frac{d\delta V}{dt} = A\,\delta V + \widetilde{B}\,u(t), \qquad \delta V(0) = \delta V_0, \qquad \delta V(T) = \delta V_{\text{target}}.$$

The endpoint mismatch is

$$\eta = \delta V_{\text{target}} - e^{AT}\delta V_0,$$

and the finite-horizon controllability Gramian is

$$W(T) = \int_0^T e^{A\tau}\widetilde{B}\widetilde{B}^T e^{A^T\tau}\, d\tau.$$

When the required direction lies in the reachable subspace of $W(T)$, the minimum-energy control is

$$u^*(t) = \widetilde{B}^T e^{A^T(T-t)} W(T)^{-1}\eta,$$

with minimum energy

$$E_{\min} = \eta^T W(T)^{-1}\eta.$$

2.4 Reachable and unreachable subspaces

The central obstruction arises when $W(T)$ is not invertible along the required endpoint direction. Eigendecomposing the Gramian partitions state space into a reachable subspace and an unreachable subspace. For a given damage pattern δV_0, we track:

- projected energy: the cost of steering the reachable component to zero,
- unreachable residual: the norm of the component that cannot be corrected by the chosen actuator family on the chosen horizon,
- actuation-rank threshold: the smallest number of actuator sites for which the unreachable residual falls below a chosen tolerance ϵ.

In the present numerical results we use $\epsilon = 10^{-4}$ as an operational analysis tolerance, not as a biologically privileged constant.

3 Results

3.1 The controllable subspace is small and ill-conditioned

On the default 16×16 damaged lattice with the default sparse two-site actuator, the state dimension is 256, but the finite-horizon controllability Gramian on horizon $T = 4.0$ has reachable rank only 15 at tolerance 10^{-12}. Thus, less than 6% of the local voltage state space is materially reachable under this actuator family.

The reachable directions are also extremely anisotropic. The same baseline yields a reachable-subspace condition number of 2.01717×10^{11}, with largest Gramian eigenvalue 0.307589 and smallest positive eigenvalue 1.52486×10^{-12}. Sparse actuation therefore creates two constraints at once: low rank and strong energetic disparity even within the reachable sector.

3.2 The controllable directions are Laplacian-structured but not Laplacian-pure

The dominant controllable A-modes have only moderate maximum overlaps with Laplacian eigenvectors, around 0.46–0.57. The controllable subspace clearly retains network structure, but it is not pure Laplacian structure. The local bistable voltage response materially reshapes the modal basis. In practical terms, a stimulation strategy derived from network geometry alone will miss part of the actual frozen control structure.

3.3 Two-site exact steering is infeasible for realistic wounds

On the 8×8 lattice, we exhaustively scanned all 2016 two-site pairs for a side-patch depolarization on the low-x boundary. The result is exact: the number of exactly reachable pairs at tolerance 10^{-12} is zero.

This is a structural statement about reachability, not a numerical failure. For every two-site Gramian in the exhaustive scan, the endpoint mismatch retains a nonzero projection onto the Gramian null space. The best projected pair $(8, 41)$ still leaves residual 1.37294×10^{-4}, and the median projected residual across all pairs is 2.74410×10^{-4}.

Damage geometry	Threshold rank
Corner wound	2
Side patch	3
Central lesion	3

Table 1: Greedy actuation-rank thresholds on the 8×8 lattice. The $k = 2$ result is exact because all pairs were exhaustively scanned; the $k \geq 3$ values are greedy upper bounds.

3.4 Greedy actuation-rank thresholds depend on damage geometry

Using greedy residual-minimizing placement on 8×8 at tolerance 10^{-4} yields:

The best corner-wound pair already achieves residual 2.90001×10^{-6}, well below threshold. By contrast, the best side-patch pair remains at 1.37294×10^{-4} and the best central-lesion pair remains at 1.89178×10^{-4}, so both require an additional site in the greedy construction.

3.5 Above threshold, residual drops rapidly with added sites

Across all three damage geometries, increasing actuator count from 2 to 4 reduces residual by roughly one to two orders of magnitude. For the corner wound, residual falls from 4.14099×10^{-6} to 3.44058×10^{-7}. For the side patch, it falls from 1.56095×10^{-4} to 2.00112×10^{-5}. For the central lesion, it falls from 2.40414×10^{-4} to 2.80263×10^{-5}.

The sharp transition is therefore not in stimulation amplitude. It is in actuation rank. Below threshold, residual is pinned away from zero by unreachable directions. Above threshold, a small increase in control dimensionality produces large gains in repair fidelity.

3.6 Thresholds shift with scale and disorder

For the 16×16 scaling check, using a structured greedy candidate pool built from damage support and high-weight controllable modes yields:

Damage geometry	Greedy threshold rank
Corner wound	2
Side patch	6
Central lesion	2

Table 2: Structured greedy scaling check on the 16×16 lattice. These are not exhaustive optimum thresholds.

The side patch becomes substantially harder on the larger lattice, remaining above threshold at $k = 4$ and crossing below 10^{-4} only at $k = 6$.

Random gap-junction dilution pushes the threshold upward in the same direction. For the 8×8 side patch, the greedy threshold rank is stable at 3 for dilution fractions 0.0 and 0.1, mixed between 3 and 4 at 0.2, and mostly 4–5 by 0.3 across five seeds. With only five seeds per dilution level, we interpret this as a robustness check rather than as a calibrated biological scaling law.

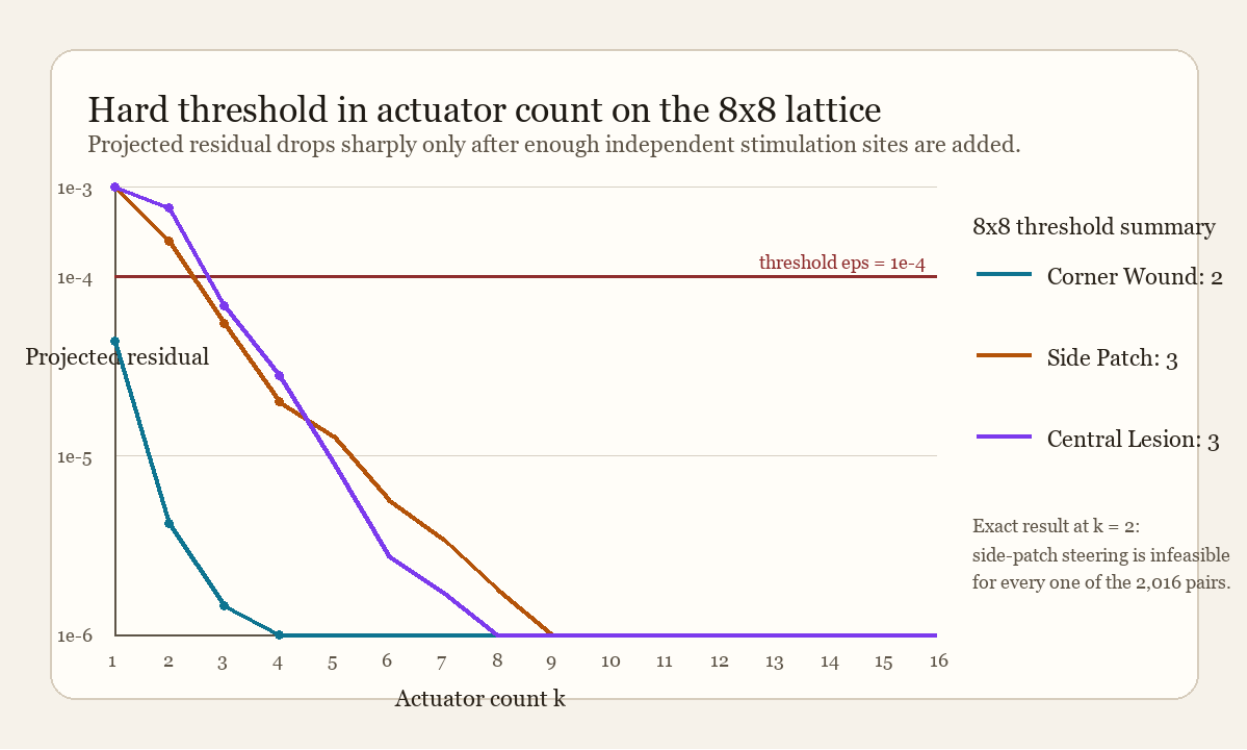

Figure 1: Projected residual versus actuator count on the 8×8 lattice for the corner wound, side patch, and central lesion. The horizontal line marks the working residual threshold $\epsilon = 10^{-4}$. The exact two-site result shows that the side patch remains infeasible at $k = 2$, while the greedy trajectories for the three geometries cross the threshold at different actuator counts.

4 Discussion

Everything computed here is a passive-tissue baseline. The frozen linearized model assumes that conductance remains fixed over the steering horizon and that repair is assessed in the local finite-horizon linearized system. That approximation may overestimate or underestimate the true intervention complexity once conductance adaptation and nonlinear basin geometry are restored. We therefore do not claim a formal upper bound.

Even without a formal bound, the framing creates a useful experimental quantity. If a real tissue repairs under fewer intervention sites than the frozen linearized prediction, then endogenous tissue dynamics are carrying part of the control burden. To make that comparison meaningful, the experimental protocol would need a count-reduction design analogous to the computational rank search: hold the damage geometry fixed, vary the number and placement of independently controlled stimulation sites, and record the smallest intervention that reliably restores pattern.

For intervention design, the first engineering question is not how much current or drug to apply. It is how many independent control directions the modality provides. The present framework converts the clinical question "will this intervention repair the tissue?" into the control-theoretic specification "does this modality's actuation rank exceed the threshold required for this damage pattern on this tissue's gap-junction network?"

5 Methods

The simulations use square lattices of size 8×8 and 16×16 with nearest-neighbor coupling. Default parameters are $\tau_V = 1$, base conductance 0.35, timestep $\Delta t = 0.02$, and 200 steps, corresponding to horizon $T = 4.0$. The local voltage sector is the double-well force $aV_i - bV_i^3$ with default $a = 1$ and $b = 1$, and the network coupling is diffusive through the gap-junction graph.

Two-site placement on 8×8 is evaluated exhaustively over all 2016 pairs. Higher-rank sparse placement uses greedy residual-first selection over candidate nodes. Accordingly, all reported $k \geq 3$ thresholds are greedy upper bounds rather than certified optima. The 16×16 scaling check uses a structured candidate pool built from damage support and high-weight controllable modes, so it is a scaling validation rather than an exhaustive optimum.

All numerical artifacts cited in the draft are saved in the associated public research repository.[1] The present formulation complements forward-oriented bioelectric modeling efforts [1, 7] by solving an inverse intervention problem.

References

[1] Javier Cervera, Alexis Pietak, Michael Levin, and Salvador Mafe. Bioelectrical coupling in multicellular domains regulated by gap junctions: a conceptual approach. *Progress in Biophysics and Molecular Biology*, 144:149–163, 2019.

[2] Sean P. Cornelius, William L. Kath, and Adilson E. Motter. Realistic control of network dynamics. *Nature Communications*, 4:1942, 2013.

[3] Michael Levin. Bioelectric signaling: Reprogrammable circuits underlying embryogenesis, regeneration, and cancer. *Cell*, 184(8):1971–1989, 2021.

[4] Yang-Yu Liu, Jean-Jacques Slotine, and Albert-László Barabási. Controllability of complex networks. *Nature*, 473(7346):167–173, 2011.

[5] N. J. Oviedo, J. Morokuma, P. Walentek, I. P. Kema, M. B. Gu, J.-M. Ahn, J. S. Hwang, T. Gojobori, and Michael Levin. Long-range neural and gap junction protein-mediated cues control polarity during planarian regeneration. *Developmental Biology*, 339(1):188–199, 2010.

[6] Fabio Pasqualetti, Sandro Zampieri, and Francesco Bullo. Controllability metrics, limitations and algorithms for complex networks. *IEEE Transactions on Control of Network Systems*, 1(1):40–52, 2014.

[7] Alexis Pietak and Michael Levin. Bioelectric gene and reaction networks: computational modelling of genetic, biochemical and bioelectrical dynamics in pattern regulation. *Journal of The Royal Society Interface*, 16(154):20190095, 2019.

[1] https://github.com/cablechris/hard-limits-sparse-bioelectric-control