

A GPU-native, differentiable fiber simulator reveals a species-dependent selectivity penalty in reduced-order peripheral-nerve models

Max Habermusch

`max.habermusch@meduniwien.ac.at`

Medical University of Vienna <https://orcid.org/0000-0003-0712-4893>

David Lung

Medical University of Vienna <https://orcid.org/0009-0003-7748-2335>

Article

Keywords:

Posted Date: August 25th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10543706/v1>

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

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

A GPU-native, differentiable fiber simulator reveals a species-dependent selectivity penalty in reduced-order peripheral-nerve models

David Lung¹ and Max Haberbush^{1,2,*}

¹Center for Medical Physics and Biomedical Engineering, Medical University of Vienna, Vienna, Austria

²Ludwig Boltzmann Institute for Cardiovascular Research, Vienna, Austria

*Corresponding author: `max.haberbusch@meduniwien.ac.at`

Abstract

Computational models guide the design of selective peripheral nerve stimulation, but to stay tractable they almost universally reduce each fascicle to one or a few representative fibers — an approximation whose effect on predicted selectivity has gone untested, because testing it means simulating whole fiber populations under optimized extracellular fields. We built *jaxon*, a fully differentiable, GPU-native reimplementation of canonical peripheral-nerve fiber models in JAX, validated against NEURON (99.6 % of 943 configurations within 1 % threshold error) and roughly 820× faster at 100,000 fibers. Optimizing stimulation over a 12-contact cuff on 29 histology-segmented vagus nerves (18 swine, 11 human), we show that the standard single-fiber reduction overestimates deliverable selectivity for myelinated fibers — negligibly in swine (median 0.01) but substantially in human (median 0.14, up to 0.45), because human fascicles are large. Re-scoring candidate designs on the full

population, cheap in *jaxon*, recovers most of the loss, identifying population-scale validation as the missing step.

1 Introduction

Electrical stimulation of peripheral nerves is now a clinical therapy for drug-resistant epilepsy, treatment-resistant depression, rheumatoid arthritis, heart failure and chronic pain^{1–4}. Because peripheral nerves bundle motor, sensory and autonomic fibers into anatomically distinct fascicles whose number, caliber and somatotopic ordering vary substantially across species and individuals^{5,6}, the central therapeutic problem is not *whether* to stimulate but *which* fibers to engage — a problem of *selectivity*^{7–9}. Two design levers set selectivity: the spatial distribution of cathodic and anodic charge across the contacts of a multi-contact cuff, and the temporal shape of the stimulus waveform, both coupling nonlinearly to fiber threshold through the extracellular field and each fiber’s active membrane dynamics^{10–12}.

Because those levers are expensive to explore experimentally, device development increasingly relies on computational evaluation of candidate stimulation patterns. The dominant workflow couples a finite-element solve for the extracellular potential $V_e(\mathbf{x})$ in a realistic nerve geometry to compartmental NEURON¹³ simulations of canonical axon models — typically the McIntyre–Richardson–Grill (MRG) double-cable model (or the simpler Sweeney model) for myelinated fibers^{14,15} and the Sundt model for unmyelinated C-fibers¹⁶ — and scans parameters to map activation thresholds across fascicles. Open-source pipelines such as ASCENT¹⁷, NRV¹⁸, PyPNS¹⁹ and PyFibers²⁰ have made this workflow widely accessible, but they share a structural cost: every candidate design requires *thousands* of independent NEURON forward simulations, run serially over CPU cores, with no way to compute the gradient of a selectivity objective with respect to the design variables. Optimization is therefore confined to gradient-free methods — grid sweeps, finite-difference approximations²⁰ or evolutionary search⁷ — that scale poorly to the high-dimensional spaces opened by waveform shaping or many-contact cuffs.

A second economy compounds the first and is far less examined. To keep the fiber

count manageable, selectivity-optimization pipelines commonly sample only a few representative fibers per fascicle — in the limiting case a single centroid fiber, the reduction used, for example, in GPU-accelerated selectivity optimization²¹. Whether this reduction preserves the selectivity of a candidate design, or quietly biases it, has never been tested, for a simple reason: the test requires simulating the *entire* fiber population under the very fields the reduced model selects — exactly the computation the reduction exists to avoid. This is the question we set out to answer.

Recent progress in differentiable simulation offers a way in. Frameworks such as `jaxley`²² express Hodgkin–Huxley membrane dynamics and multi-compartment cable equations in JAX²³, in the same differentiable-physics spirit as tools like `JAX MD`²⁴, so that gradients of arbitrary functions of the membrane trajectory follow from reverse-mode automatic differentiation within a single GPU-vectorized forward pass, enabling gradient-based inference at large scale in the central nervous system. No comparable differentiable implementation has existed for peripheral-nerve fiber models: the MRG double cable requires a coupled (V_i, V_{px}) backward-Euler integration that is not available out of the box, the Sweeney, Sundt and Rattay models carry calibration-sensitive constants that must reproduce NEURON to benchmarking accuracy⁹, and selectivity optimization demands a batched forward solver that vectorizes across large fiber populations on a single device.

Here we build that solver — *jaxon*, a fully differentiable, GPU-native reimplementa- tion in JAX of four canonical axon models (the myelinated MRG double-cable and Sweeney models and the unmyelinated Sundt and Rattay C-fibers) — and use it to inter-rogate the reduced-order sampling assumption that population-scale simulation finally makes testable. Figure 1 lays out the whole approach end to end: from a histology- derived nerve cross-section and its per-contact extracellular fields, through the GPU- batched differentiable fiber solver validated against NEURON, to the selectivity objective whose gradient drives the per-contact stimulus optimization. *jaxon* reproduces NEURON across strength–duration, conduction-velocity and intracellular-trace benchmarks (99.6 % of 943 configurations within 1 %) while running roughly 820× faster than `PyFibers` at $N = 100,000$ fibers on a single A100, fast enough to optimize and re-score whole-nerve

populations. On a cross-species cohort of 29 histology-segmented vagus nerves, we find that stimulation configurations optimized on sparsely sampled fibers — the per-fascicle reduction used in GPU-accelerated selectivity work²¹ — lose substantial selectivity when deployed on the full fiber population. This *deployment penalty* scales with the number of fibers per target fascicle and therefore differs systematically between species: small in swine, large in human. It is a translational bias that reduced-order models cannot detect, because they never simulate the fibers on which the bias acts. The finding is established for the myelinated population at a single fiber diameter; the distinct and larger selectivity limit on the small unmyelinated autonomic C-fibers is a separate effect, shown on these same nerves by the companion golgi work²⁵ and left to future study. Enabling the analysis is a selectivity-optimization pipeline whose GPU-fast forward pass makes optimization tractable directly on the full fiber population rather than on per-fascicle centroids — to our knowledge the first full-population selectivity optimization on real cross-species histology geometry. The results reported here use finite-difference gradients over the contacts, as in `PyFibers`²⁰; the enabling capability is the fast full-population forward pass, not the gradient type. *jaxon* additionally exposes *exact* end-to-end gradients — verified against finite differences (Supplementary Fig. S1) and unavailable to NEURON-wrapped pipelines — for the waveform-shape optimization left to future work.

2 Results

Our central result is that the reduced-order fiber sampling used to keep selectivity models tractable systematically overestimates how much selectivity is actually delivered to the nerve, by an amount set by fascicular architecture. We first validate the simulator against NEURON and benchmark its throughput (Sections 2.1–2.2), then optimize selectivity on real cross-species anatomy (Section 2.3), where this central result emerges (Section 2.4). All raw data and analysis code are available as described in the Data Availability statement.




Figure 1: Overview of the *jaxon* differentiable pipeline. **(a) Forward-and-optimize loop.** A per-contact stimulus $I_k(t)$ drives the superposed extracellular field $V_e = \sum_k I_k V_e^{(k)}$; the coupled-cable GPU solver integrates the membrane dynamics of every fiber in parallel and aggregates them into a fascicle-selectivity objective $\mathcal{L}$. Because the whole pipeline is built on *jaxley*/*JAX*, the gradient $\partial \mathcal{L} / \partial \theta$ with respect to the stimulus is available end-to-end by reverse-mode autodiff (dashed loop). The solver is validated against *PyFibers*/*NEURON* on thresholds and conduction velocity. **(b) Anatomy to model.** A histological vagus-nerve cross-section is segmented, area-preservingly circularized and embedded in the 12-contact cuff finite-element model, from which the per-contact extracellular lead fields $V_e^{(k)}$ (and the derived activating function) are computed. **(c) Fiber models and validation.** The four implemented fiber models (myelinated MRG and Sweeney; unmyelinated Sundt and Rattay) with their equivalent-circuit double-cable / single-cable schematics. **(d) Throughput and selectivity optimization on GPU.** Left, the *jaxon*-vs-*PyFibers* forward-pass speedup versus fiber count (Section 2.2); right, one example nerve with the optimized cuff pattern (cathode teal, anode orange, inactive white, over the cuff’s four angular columns $\times$ three axial rows) and fibers colored by recruitment (target fired / target silent / off-target fired / off-target silent), the peripheral target-fascicle cluster outlined.

2.1 Validation against NEURON

We first establish that *jaxon* reproduces NEURON across models and scales to whole-nerve populations — the accuracy and throughput the optimization analysis below relies on. For each axon model we drove *jaxon* and PyFibers-wrapped NEURON with identical stimuli and compared activation thresholds across diameters, eight waveforms (monophasic and biphasic cathodic- and anodic-first, sinusoidal, sawtooth, exponential, Gaussian) and pulse widths 0.02–1 ms — 943 configurations in total. 99.6 % **agree to within 1 % relative threshold error** (Table 1; *jaxon*–NEURON threshold correlation $r = 0.999998$, i.e. $R^2 > 0.99999$ as plotted in Figure 2c; mean signed bias +0.11 %). Only the Rattay model exceeds the criterion, in a single edge regime — biphasic cathodic-first pulses at the four smallest diameters (peak 2.24 % at $D = 0.3 \mu\text{m}$) — which we attribute to reduced bisection resolution near the strength–duration inflection; the intracellular traces agree to NEURON-level precision (Figure 2a).

Table 1: Strength–duration threshold agreement of *jaxon* vs. PyFibers-wrapped NEURON, aggregated across all tested diameters, waveforms and pulse widths per model. Diameters tested (μm): MRG = {5.7, 7.3, 8.7, 10.0, 11.5, 12.8, 14.0, 15.0, 16.0}; Sundt = {0.3, 0.5, 0.8, 1.0, 1.2}; Sweeney = {10}; Rattay = {0.3, 0.5, 0.8, 1.0, 1.5}. Sweeney is evaluated at $D = 10 \mu\text{m}$ only because Sweeney’s original parameterization¹⁵ was derived and validated at $D = 10 \mu\text{m}$ and the PyFibers reference implementation²⁰ follows the same convention; the 17 cells dropped from Sweeney’s 48-cell sweep at $D = 10$ are NEURON **extracellular**-mechanism singular-solve failures under anodic-leading pulses, not *jaxon* errors. Pulse widths tested: {0.02, 0.05, 0.1, 0.2, 0.5, 1.0} ms.

Model	configs	median err. [IQR]	95 th %ile	peak err.
MRG	432	0.04 % [0.02–0.09]	0.20 %	0.39 %
Sundt	240	0.34 % [0.28–0.39]	0.42 %	0.46 %
Sweeney	31	0.02 % [0.02–0.05]	0.06 %	0.06 %
Rattay	240	0.07 % [0.04–0.10]	0.15 %	2.24 %
total	943	0.08 % [0.03–0.20]	0.40 %	2.24 %

Seven of the eight waveforms agree within 1 % on *every* configuration; the lone exceptions are the four biphasic cathodic-first Rattay cases above (and even there 96.5 % pass). Conduction velocity, which probes the internodal and nodal time constants, agrees to a median < 0.1 % across all models and diameters (peak 0.24 % for Sundt at $D = 0.3 \mu\text{m}$). Per-waveform and per-diameter breakdowns are in Supplementary Tables S7–S8; the per-

waveform error distribution is Figure 2c. Beyond thresholds, *jaxon* reproduces canonical propagation phenomena — action-potential propagation, DC (depolarizing-hold) conduction block and AP collision across all four fiber models (Supplementary Fig. S2), and 20 kHz kilohertz-frequency block (Supplementary Fig. S3) — each matching NEURON within the scan resolution or the floating-point floor.

Figure 2 compactly summarizes the validation and the computational-efficacy result for all four models in a single panel set: representative extracellular AP and gating-variable traces with the stimulus overlaid (panel a, *jaxon* solid over NEURON dashed), the strength–duration curves at every tested diameter (panel b), the activation-threshold agreement broken down as a *jaxon*-vs-NEURON scatter and by model, waveform and diameter (panel c), CV-vs-diameter agreement (panel d), and the forward-pass throughput benchmark (panel e, detailed in Section 2.2). Across panel c the *jaxon*-vs-NEURON scatter lies on the unity line at $R^2 > 0.99999$, and every per-waveform and per-diameter error distribution sits below the 1 % acceptance criterion.

2.2 Forward-pass throughput

The principal performance motivation for the JAX implementation is the expectation that GPU vmap-batching will outperform CPU-parallel NEURON as the fiber count grows. We benchmarked all five model/parameterization variants (MRG with both the original and the interpolated parameter tables, Sweeney, Sundt and Rattay) at $N \in \{1, 10, 100, 1\,000, 10\,000, 100\,000\}$ fibers on a single NVIDIA A100 GPU, against **PyFibers** running 8-thread CPU-parallel on a matched Xeon host. Figure 2e reports the per-model wall time and the resulting *jaxon*-vs-PyFibers speedup; Table 2 summarizes the headline numbers.

At small N , GPU launch overhead dominates and **PyFibers** on CPU is faster, but the vmap-batched forward pass amortizes the per-step launch cost across the batch, so *jaxon* overtakes **PyFibers** at $N \approx 10$ for every model and the advantage grows over the small- N regime, reaching order 500–1400× at the largest N (Table 2; the MRG double-cable speedup peaks near $N = 10^4$ and rolls off slightly by $N = 10^5$, while the




Figure 2: Validation of *jaxon* against PyFibers-wrapped NEURON and computational-efficacy benchmark, all four fiber models. (PyFibers wraps NEURON; the reference is labelled “PyFibers” in panel annotations and “NEURON” in the text interchangeably.) (a) Membrane-potential and gating-variable response to a single extracellular cathodic stimulus (overlaid as the solid black step at resting potential), one column per model; *jaxon* (colored solid) over NEURON (dashed gray). (b) Strength-duration threshold (absolute value) vs. pulse width at every tested diameter (viridis = small→large diameter); *jaxon* solid line + circle markers, PyFibers dashed line + cross markers in a lighter tint of the same color. (c) Activation-threshold agreement: *jaxon*-vs-PyFibers scatter with R^2 and median relative error (top left), and relative-error distributions broken down per model, per stimulus waveform and per diameter bin. Median is the black bar; the 1% acceptance criterion is the dotted red line. (d) Conduction velocity vs. fiber diameter for the A-fiber (MRG, Sweeney) and C-fiber (Sundt, Rattay) models; same *jaxon*-solid / PyFibers-lighter-dashed convention. MRG and Sweeney differ in conduction velocity at matched diameter — an expected between-model biophysical difference, not a *jaxon*-vs-NEURON validation discrepancy (each model’s *jaxon* and NEURON curves coincide). (e) Forward-pass throughput on a single NVIDIA A100 GPU: absolute wall time (left) and *jaxon*-vs-PyFibers speedup (right) vs. fiber count N , log-log (Section 2.2).

lighter models stay monotonic). At $N = 100\,000$ the speedup reaches $544\times$ for MRG, $1\,377\times$ for Sweeney (smaller per-fiber state) and $\sim 1\,000\times$ for the C-fibers, a geometric mean of $\sim 820\times$ — taking what would be multiple CPU-hours per forward pass down to under twenty seconds. At the $\sim 1\,000$ -fiber scale of the cohort selectivity study the speedup is more modest but still decisive — $209\times$ (MRG) to $304\times$ (Sundt), the operating point that actually makes the thousands-of-evaluations optimization loop below tractable. This is a GPU-versus-eight-thread-CPU comparison against the field-standard `PyFibers` workflow; against many-core CPU or GPU-accelerated NEURON variants the multiplier would be smaller, but the order-of-magnitude advantage from `vmap`-batching the forward pass persists, and CPU NEURON is the established baseline a new solver must beat. This is what brings the population-scale optimize-and-re-score analysis below ($\sim 1\,000$ fibers per nerve, simulated thousands of times per optimization, across 29 anatomies and five sampling levels) within reach on a single GPU-day.

Table 2: Wall time per forward simulation (in seconds) and speedup of *jaxon* over `PyFibers` at representative fiber counts. *jaxon* runs on a single NVIDIA A100 GPU; `PyFibers` runs 8-thread CPU-parallel on a matched Xeon host. All timings exclude JAX compile time (5–30s one-time per model, amortized over each batched run). Geometric mean of the $N = 100\,000$ speedups across the five benchmarked models is $\sim 820\times$.

Model	$N = 1\,000$		$N = 10\,000$		$N = 100\,000$	
	pyf. (s)	speedup	pyf. (s)	speedup	pyf. (s)	speedup
MRG	49.3	$209\times$	525	$671\times$	5 172	$544\times$
MRG-interp.	58.0	$246\times$	537	$695\times$	5 021	$527\times$
Sweeney	62.6	$268\times$	603	$1\,322\times$	5 972	$1\,377\times$
Sundt	183	$304\times$	1 742	$985\times$	16 593	$1\,006\times$
Rattay	151	$259\times$	1 425	$868\times$	14 068	$929\times$

2.3 Selectivity optimization on histology-derived vagus nerves

Using *jaxon* — validated against NEURON and fast enough to optimize whole-nerve populations (Sections 2.1–2.2) — we ran the differentiable optimizer on a cohort of $n_{\text{cohort}} = 29$ (18 swine, 11 human; 116 target-cluster seeds) subject-specific FE-modeled vagus-nerve cross-sections drawn from the SPARC quantified morphology releases of swine²⁶ and human²⁷, with per-contact extracellular lead-fields computed in-house with

the golgi pipeline²⁵ (Section 4.3). The 12-contact cuff (three axial rows $\times$ four angular columns) gives the rectangular optimizer $K = 12$ amplitudes per nerve; the target is the peripheral fascicle cluster within one contiguous angular sector (Section 4.3). Figure 3a shows dense- and sparse-optimized activation maps and cuff patterns for a representative swine and human nerve (the sparse condition is defined in Section 2.4), fibers colored by activation status (target/off-target $\times$ fired/silent).

Dense optimization recovers focal, selective drive. The rectangular optimizer uses a probe-based smart initialization (Section 4.2): a small set of bipolar and single-column tripolar spatial patterns is scored at eight magnitudes in both polarities (80 forward passes per seed), and the best initializes a single gradient (Adam-FD) trajectory with the probe’s zero contacts frozen to preserve sparsity. Across the cohort (111 selectable-target seeds over 29 nerves; five degenerate-target seeds excluded, see below) this achieves mean SI = 0.90 (median 0.95), with 70 % of seeds reaching SI $\geq$ 0.90. Swine and human reach a near-identical median (0.946 vs. 0.964), but the human *mean* is lower (0.845 vs. 0.926) because a heavy tail of geometric-ceiling anatomies (below) pulls it down rather than a uniform downward shift. The optimizer converges to a focal, sparse drive: 103 of 111 seeds (93 %) use exactly three active contacts — a single-column tripolar configuration (cathode at the column’s z -middle contact, charge-balancing anodes at z -top and z -bottom); per-species figures are summarized in Table 3. Per-sample dense and sparse-transfer activation maps with their cuff patterns, for every nerve and seed, are provided as Supplementary Figures S34–S60. All results here use the fixed charge-balanced biphasic pulse with optimized per-contact amplitudes; waveform-shape optimization is left to future work (Section 3).

Emergent stimulation strategy. The optimized configurations are physically interpretable rather than arbitrary (Figure 3c). Circumferentially, the optimizer steers cathodic drive *toward* the target sector: the fraction of contacts that are cathodic is highest in the target-facing direction (0.34 swine, 0.40 human) and near zero to the side. Longitudinally it forms a center-surround current profile — a mid-cathode tripole flanked

Table 3: Per-species summary of the histology cohort. All metrics are over the 111 selectable-target seeds (target fraction ≤ 0.6 ; five degenerate-target seeds, all human, excluded — see text). The deployment penalty ($SI_{\text{dense}} - SI_{\text{transfer}}$, for the single-representative-fiber (centroid) reduction) and fibers per target fascicle are per nerve.

	count		dense SI		penalty	fibers/fasc.
	nerves	seeds	med. (mean)	≥ 0.90	med. (mean)	median
Swine	18	72	0.946 (0.926)	72 %	0.011 (0.013)	30
Human	11	39	0.964 (0.845)	67 %	0.136 (0.128)	151
All	29	111	0.948 (0.897)	70 %	—	—

Species penalty difference (centroid reduction): two-sided Mann–Whitney $p = 0.01$ ($n = 29$ nerves; robust at the subject level, $p = 0.011$ over 11 swine / 7 human subjects — see Section 4.5).

Strategy-averaged penalty vs. fibers per target fascicle: Spearman $r = +0.47$, $p = 0.01$ ($n = 29$ nerves; within-species not significant, see text). The fibers-per-fascicle counts (median 30 swine, 151 human) are sampled-fiber counts under area-weighted seeding of 1 000 fibers per nerve — i.e. proportional to fascicle area, not anatomical fiber counts.

by charge-balancing anodes — whose polarity is a species-dependent tendency (more often the classic mid-cathode form in swine, the inverse in human). The optimizer thus rediscovers, without being told to, the steered tripolar configurations that the selectivity literature arrives at by design.

Geometric-ceiling cases. A minority of human anatomies cannot be steered well. These are not primarily cases where off-target fascicles happen to sit as close to the contacts as the target — across the cohort the minimum fascicle-to-contact distance is only a weak predictor of achieved selectivity (Spearman $r \approx -0.13$, n.s.) — but cases where the target fascicle is large relative to the steerable field, so that sparing every off-target fiber while recruiting the whole target is geometrically impossible. These geometric-ceiling cases form the heavy lower tail of the human distribution and are *retained* in all statistics — they are genuine selectivity outcomes on selectable targets, not solver artifacts. Separately, we exclude five seeds (all human) whose target cluster spans more than 60 % of the nerve cross-section (target fraction > 0.6): such targets are degenerate by construction — no selective solution exists — so they are uninformative about achievable selectivity.




Figure 3: Selectivity optimization on the SPARC histology-derived vagus-nerve cohort (29 nerves: 18 swine, 11 human; 111 selectable-target seeds, degenerate targets excluded). “Sparse” in panels **a–b** denotes a configuration optimized on the single centroid fiber per fascicle — a standard reduced-order workflow — and evaluated on the full ~ 1000 -fiber population (transfer); panel **d** sweeps the sampling density. **(a)** Dense- versus sparse-optimized activation for a representative swine and human nerve, matched on target size ($\sim 28\%$ of fibers) and dense SI (~ 0.97); the swine deployment penalty is near zero (and can be slightly negative from sampling noise), consistent with the swine cohort median of 0.011; fibers colored by activation status (target/off-target $\times$ fired/silent), with the optimized 12-contact cuff pattern below (contacts colored by polarity — teal cathodic ($I_k < 0$), orange anodic ($I_k > 0$), white inactive — and sized by $|I_k|$). **(b)** Dense versus sparse performance per nerve (bars = mean $\pm$ SEM with per-nerve dots overlaid, shown for readability; all inference is on the per-nerve medians reported in the text and Table 3): selectivity index, on-target and off-target recruitment, and total stimulation current. Brackets give the two-sided per-nerve Wilcoxon p , Holm-corrected over the eight metric $\times$ species cells (the same values quoted in the text). **(c)** Emergent stimulation strategy: (left) fraction of contacts that are cathodic by target-relative direction — cathodes steer toward the target; (right) longitudinal mean current by axial row — a center-surround (tripolar) profile. **(d)** Sparse fiber-sampling: dense-optimized SI ceiling (dashed line $\pm$ bootstrap CI, 10 000 resamples) versus sparse-optimized transfer SI per strategy (centroid, 1/3/10 fibers per fascicle) for swine and human; right, the deployment penalty $SI_{\text{dense}} - SI_{\text{transfer}}$ versus the number of fibers per target fascicle (Spearman r).

2.4 Reduced-order fiber sampling overestimates deliverable selectivity unless the design is re-scored on the population

Prior GPU-accelerated selectivity work reduces each fascicle to a single centroid fiber to keep the population tractable²¹. Because *jaxon* simulates the full ~ 1000 -fiber population directly, we can ask what this approximation costs at deployment. For each nerve $\times$ seed we re-ran the identical optimization on four sparse fiber-sampling levels — one centroid fiber per fascicle, and 1, 3 or 10 randomly-sampled fibers per fascicle — by pure row-subsampling of the precomputed V_e lead-field with no FEM recomputation. Each density level uses a single random draw with a fixed RNG seed (set in the released code), so the sampling is reproducible; the 116 target seeds within the 29 nerves provide the spread, and the per-nerve density curves aggregate over a nerve’s seeds. The selectivity such a model *reports* is optimistic, because its amplitudes are tuned to the few fibers it can see; we therefore take each sparse-optimized amplitude vector and re-score it on the full population *without re-optimizing*, giving the transfer selectivity SI_{transfer} — the selectivity actually delivered. The deployment penalty is the drop from the dense ceiling, $SI_{\text{dense}} - SI_{\text{transfer}}$ (Figure 3d).

The penalty is negligible in swine but substantial and heavy-tailed in human. For the single-representative-fiber reduction that reduced-order pipelines actually use — one fiber at each fascicle centroid — the per-nerve penalty is a median 0.011 in swine (mean 0.013, maximum 0.05) versus 0.136 in human (mean 0.128; Mann–Whitney $U = 157$, $p = 0.01$, rank-biserial $r = 0.59$, 95% bootstrap CI (10 000 resamples, seed fixed in the released code) $[0.15, 0.93]$), against a dense ceiling of ~ 0.95 . The human distribution is not a simple right-skew but a *low-penalty cluster plus a heavy upper tail*: five of eleven human nerves lose less than 0.07 while the remainder form an upper tail reaching 0.45 (so the mean, 0.128, sits just below the median, 0.136), whereas no swine nerve loses more than 0.05. Sampling *more* fibers reduces the penalty: across the three random-density levels it falls significantly as density rises (one to ten random fibers per target fascicle, median $0.19 \rightarrow 0.03$; Friedman $\chi^2 = 11.4$, $p = 0.003$, $n = 11$ human nerves over the three levels, with the centroid reduction as a separate reference point; averaged over the

sparse strategies the human penalty is 0.083) — but it does not close on the largest fascicles, where three of eleven human nerves still lose more than 0.10 even at ten fibers (Supplementary Figure S4). Re-expressed as the fraction of each target fascicle sampled — the species-fair axis, since ten fibers is $\sim 33\%$ of a typical swine fascicle but only $\sim 7\%$ of a human one — the penalty concentrates below $\sim 5\%$ sampling and the species gap persists at matched low fraction (human ~ 0.10 vs swine ~ 0.00): a large fascicle needs a proportionally larger sample to be represented, and on the human vagus the standard single-fiber reduction rarely provides it.

What sets the species difference is fascicle *size*. The penalty is essentially a between-species step — negligible in swine, sizeable in human — tracking the number of fibers per target fascicle, which has median 30 in swine versus 151 in human. The strategy-averaged penalty correlates with fibers-per-fascicle across the cohort (Spearman $r = +0.47$, $p = 0.01$, $n = 29$ nerves), though within each species the trend is only directional at these sample sizes; we therefore read fibers-per-fascicle as a *between-species* driver rather than a graded within-species predictor. This explains the species gap: swine vagus nerves carry many small fascicles, so a handful of sampled fibers represents the target well and the optimized field transfers; human nerves carry a few large fascicles, so a field tuned to one sampled fiber does not cover the rest of that fascicle. This is the within-fascicle counterpart of an established geometric effect — cuff-electrode activation thresholds vary systematically with a fascicle’s perineurium thickness and cross-sectional area²⁸ — so within one large fascicle the spread of fiber positions produces a spread of thresholds that a single centroid fiber cannot represent. Counterintuitively, the more fascicular nerve (swine) is the more forgiving one.

The recruitment signature confirms the mechanism (Figure 3b). Relative to the dense optimum, the centroid-transfer configuration in human ($n = 11$ nerves) loses on-target recruitment (median 100% $\rightarrow$ 83%; $p = 0.023$) and selectivity ($p = 0.014$), whereas off-target recruitment and total delivered current are statistically unchanged ($p = 0.70$ for both) — all p -values here and in Figure 3b are two-sided per-nerve Wilcoxon signed-rank, Holm-corrected across the eight metric $\times$ species cells. In swine ($n = 18$) the

273 same dense-vs-transfer SI difference is Holm-significant ($p = 0.005$) but negligible in
 274 magnitude (~ 0.01 SI; on-target $p = 0.11$) — statistical significance without practical
 275 consequence, in contrast to the human loss. The penalty is thus a failure to recruit
 276 the large target fascicle, not off-target spill or reduced drive. It is, however, largely a
 277 cost of the *standard workflow* — optimizing on the reduced set and reading selectivity
 278 off it, without re-scoring on the population. Because *jaxon* re-scores a candidate on
 279 the full population cheaply, selecting per nerve the sparse-optimized design that scores
 280 best on the full population (optimize-then-validate) collapses the penalty to a median
 281 0.000 in swine and 0.016 in human. The honest conclusion is therefore narrow: centroid-
 282 only sampling *as standardly used*²¹ over-states deliverable selectivity on the large-fascicle
 283 human anatomies most relevant to clinical translation, but a *modeler* who brings the full
 284 population into the selectivity pipeline — at optimization or at least at validation, now
 285 computationally feasible — recovers most of it, leaving only a small residual (~ 0.02 in
 286 human; the optimize-then-validate residual above, 0.016, rounds into this range) that no
 287 sparse design fully closes. This is a statement about the modeling workflow, not a bedside
 288 prescription: it says a selectivity design tuned on centroid fibers should be re-scored on
 289 the full population before being trusted, not that a fixed stimulus amplitude be used
 290 clinically (see Discussion).

291 **3 Discussion**

292 Using population-scale differentiable simulation, we have shown that the reduced-order
 293 fiber sampling standard in GPU-accelerated selectivity modeling — reducing each fas-
 294 cicle to one or a few representative fibers and reading selectivity off that reduced set
 295 — overestimates how much selectivity is actually delivered to the nerve when used that
 296 way. The error is not a fixed offset: this *deployment penalty* is governed by fascicular
 297 architecture, the number of fibers per target fascicle, and is therefore negligible in swine
 298 (median ~ 0.01) but sizeable and heavy-tailed in human (median 0.14 for the standard
 299 single-representative-fiber reduction, reaching 0.45 on the largest-fascicle nerves) — a

300 between-species translational bias that reduced-order models cannot detect because they
 301 never simulate the fibers the bias acts on. Mechanistically the penalty is a failure to
 302 recruit the large human target fascicle from a field tuned to a few sampled fibers, not
 303 off-target spill or reduced drive; it extends to the within-fascicle scale the established de-
 304 pendence of cuff-electrode thresholds on fascicle perineurium and cross-sectional area²⁸,
 305 a spread a single centroid fiber cannot capture. Crucially, it is also largely *recover-*
 306 *able*: because the full-population forward pass is cheap in *jaxon*, re-scoring or selecting
 307 a sparse-optimized design on the population removes most of it (residual human ~ 0.02),
 308 so the actionable message is that population-scale validation — not a different optimizer
 309 — is what reduced-order pipelines lack. All of this was measurable only because *jaxon*
 310 simulates the full ~ 1000 -fiber population at NEURON-level fidelity (99.6 % of 943 con-
 311 figurations under the 1 % criterion) and a geometric-mean $\sim 820\times$ throughput gain (peak
 312 $1377\times$) on a single A100; throughput is the enabling capability, since prior NEURON
 313 pipelines are slow enough that full-population optimization is impractical, forcing the
 314 per-fascicle reduction in the first place (Sections 2.3–2.4).

315 *jaxon* occupies a distinct position in the peripheral-nerve modeling ecosystem. The
 316 closest NEURON-based comparator is PyFibers²⁰, which wraps NEURON axon mod-
 317 els in Python with random-search or finite-difference-gradient optimization; *jaxon* re-
 318 moves the Python–NEURON overhead that blocks vmap-batched GPU forward passes
 319 and adds end-to-end gradients PyFibers lacks, without loss of fidelity (Section 2.1). The
 320 wider open-source ecosystem — ASCENT, NRV¹⁸ and PyPNS¹⁹ — shares the same
 321 gradient-free, CPU-NEURON basis. The one GPU-accelerated exception, AxonML²¹, is
 322 a fast learned surrogate of the fiber response; *jaxon* instead integrates the *exact* biophys-
 323 ical equations on the GPU, retaining NEURON-level fidelity and exposing exact gradi-
 324 ents without a learned approximation, and it is AxonML’s centroid-per-fascicle sampling
 325 whose deployment penalty we quantify (Section 2.4). Other GPU neuron simulators —
 326 NeuroGPU²⁹, Brian2GeNN³⁰, CoreNEURON³¹ and Arbor³² — accelerate the forward pass
 327 but none models *extracellular* field stimulation, the regime relevant to cuff electrodes,
 328 and none exposes end-to-end gradients, leaving CPU NEURON as the established base-

line a new solver must beat. Within the JAX ecosystem, `jaxley`²² supplies the multi-compartment Hodgkin–Huxley machinery but ships only canonical CNS channels; `jaxon` adds the peripheral specialization — the coupled (V_i, V_{px}) MRG solver, the calibrated Sweeney, Sundt and Rattay models, the NEURON-validation harness and the selectivity pipeline. It is thus complementary to, rather than a replacement for, this ecosystem: the subject-specific finite-element fields used here are produced by our companion open platform, `golgi`²⁵, which builds the anatomy, meshes and solves the extracellular field, while `jaxon` contributes the differentiable, GPU-native fiber solver and the population-scale sampling-penalty analysis and drops in as an additional backend.

Several limitations bound the scope of the finding. It is established *in silico*, with one cuff and one finite-element pipeline, on 29 nerves (11 human). We therefore frame fascicular architecture as the *driver* the between-species penalty difference scales with (cohort Spearman $r = +0.47$; the within-species trend is directionally consistent but under-powered), not as a manipulated cause. A diameter spot-check confirms the species ordering persists at 5.7, 7.3 and 10 μm (strategy-averaged human penalty 0.083, 0.066 and 0.054 versus swine ~ 0.01 throughout), so the direction is not an artifact of the single diameter, though we expect the magnitude — not the direction — to depend on fiber-type mix, cuff geometry and the conductivity model. Within-fascicle heterogeneity in fiber diameter we tested directly: adding realistic per-fiber variability ($\text{CV} = 0.15$) leaves the species-dependent penalty intact (swine small, human large; $p = 0.01$) and, if anything, deepens the heavy human tail (worst nerve $0.45 \rightarrow 0.66$) while lowering the achievable selectivity ceiling ($0.95 \rightarrow 0.88$ in human), so the single-diameter model is a conservative rather than optimistic idealization. That single fiber type at a single diameter, with identical channel densities and resting state across fibers, isolates the sampling effect we set out to quantify but is itself an optimistic setting for selectivity — with identical fibers a per-fascicle centroid is maximally representative, so the reported penalty is a conservative lower bound, and inter-fiber variability is known to broaden activation distributions and lower achievable selectivity⁷. A distinct and, on the vagus, larger limit on deliverable selectivity comes from the realistic *mix* of fiber types: the companion

platform²⁵ shows on these same nerves that moving from a single-diameter idealization to a morphometry-based A/B/C population collapses the selectivity of the small, high-threshold cardiac fibers, whose recruitment forces currents that also engage off-target fibers. That fiber-type axis is orthogonal to the sampling axis studied here and compounds it, and *jaxon's* throughput makes the multi-type sweeps needed to characterize their interaction tractable.

The remaining caveats concern geometry, operating regime and scoring. The histology results use a single 12-contact cuff (three axial rows $\times$ four angular columns) at a fixed radius; larger contact counts, conformable cuffs and multi-cuff arrays are increasingly common clinically but were out of scope, and *jaxon's* per-contact amplitude parameterization extends to arbitrary layouts with no solver change. The per-contact lead fields are computed with a simplified cylindrical-cuff finite-element model rather than a device-faithful encapsulation/electrode-impedance solve; because *jaxon* accepts V_e as a precomputed $[K, n_{\text{fiber}}, n_{\text{comp}}]$ tensor, a higher-fidelity model can be substituted without touching the optimizer. The deployment penalty is defined as under-recruitment of the target fascicle at the *optimized fixed amplitude*, whereas clinical neuromodulation titrates amplitude to a physiological endpoint; since the reduced-order design fixes a suboptimal *spatial* pattern, it sits on a strictly lower selectivity-versus-off-target trade-off curve, and a clinician raising overall amplitude could recover some lost target recruitment only by engaging more off-target fibers, never reaching the dense pattern's trade-off — the deficit is in the pattern, not the operating amplitude. The penalty should therefore be read as a bias in *model-predicted deliverable selectivity* that propagates into device and waveform design, actionable today for modelers building selectivity pipelines rather than at the bedside. Finally, recruitment is scored by a sigmoid-smoothed activation proxy (Section 4.2) rather than full spike detection; we verified it directly at population scale by re-scoring the optimized configuration in **PyFibers**-wrapped NEURON over the entire ~ 1000 -fiber population on eight cohort nerves (four human, four swine), where the proxy agrees with NEURON spike detection on 99.9% of fibers (minimum 99.6%) and the population selectivity index matches NEURON to within 0.005 SI on every nerve

— an order of magnitude below the human deployment penalty, so the penalty is not a proxy artifact. The proxy is nonetheless approximate in pathological regimes (failed propagation, partial recruitment, branching or blocked nodes) that were absent from the straight-axon MRG simulations here but would warrant revisiting for branched-fiber or blocked-conduction applications such as DC or kHz conduction block³³.

Beyond these bounds, *jaxon's* GPU throughput and end-to-end gradients position it for several next steps that are intractable for gradient-free pipelines: fiber-type (diameter-based) selectivity via waveform-shape gradients; population-averaged designs that marginalize over nerve-to-nerve variability in a single batched pass; inverse modeling of cuff-recorded compound action potentials to infer recruitment^{34,35}; and closed-loop stimulation policies trained by backpropagation through the forward model. Released open-source, *jaxon* slots into existing pipelines as a gradient-enabled, GPU-native replacement for the NEURON forward solver.

References

- [1] Howland, R. H. Vagus nerve stimulation. *Current Behavioral Neuroscience Reports* **1**, 64–73 (2014).
- [2] Koopman, F. A., Chavan, S. S., Miljko, S. *et al.* Vagus nerve stimulation inhibits cytokine production and attenuates disease severity in rheumatoid arthritis. *Proceedings of the National Academy of Sciences* **113**, 8284–8289 (2016).
- [3] Bonaz, B., Sinniger, V. & Pellissier, S. The vagus nerve in the neuro-immune axis: Implications in the pathology of the gastrointestinal tract. *Frontiers in Immunology* **8**, 1452 (2017).
- [4] De Ferrari, G. M. *et al.* Chronic vagus nerve stimulation: a new and promising therapeutic approach for chronic heart failure. *European Heart Journal* **32**, 847–855 (2011).
- [5] Settell, M. L. *et al.* Functional vagotomy in the cervical vagus nerve of the domes-

tic pig: implications for the study of vagus nerve stimulation. *Journal of Neural Engineering* **17**, 026022 (2020).

[6] Pelot, N. A. *et al.* Human vagus nerve fascicular anatomy and its implications for targeted cardiac stimulation. *Frontiers in Neuroscience* **17**, 1169187 (2023).

[7] Aristovich, K. *et al.* Model-based geometrical optimisation and *in vivo* validation of a spatially selective multielectrode cuff array for vagus nerve neuromodulation. *Journal of Neuroscience Methods* **352**, 109079 (2021).

[8] Thompson, N. *et al.* Spatially selective stimulation of the pig vagus nerve to modulate target effect versus side effect. *Journal of Neural Engineering* **20**, 016029 (2023).

[9] Musselman, E. D., Pelot, N. A. & Grill, W. M. Validated computational models predict vagus nerve stimulation thresholds in preclinical animals and humans. *Journal of Neural Engineering* **20**, 036030 (2023).

[10] Rattay, F. Analysis of models for extracellular fiber stimulation. *IEEE Transactions on Biomedical Engineering* **36**, 676–682 (1989).

[11] Grill, W. M. & Mortimer, J. T. Stimulus waveforms for selective neural stimulation. *IEEE Engineering in Medicine and Biology Magazine* **14**, 375–385 (1995).

[12] Wongsarnpigoon, A. & Grill, W. M. Energy-efficient waveform shapes for neural stimulation revealed with a genetic algorithm. *Journal of Neural Engineering* **7**, 046009 (2010).

[13] Hines, M. L. & Carnevale, N. T. The NEURON simulation environment. *Neural Computation* **9**, 1179–1209 (1997).

[14] McIntyre, C. C., Richardson, A. G. & Grill, W. M. Modeling the excitability of mammalian nerve fibers: Influence of afterpotentials on the recovery cycle. *Journal of Neurophysiology* **87**, 995–1006 (2002).

- [15] Sweeney, J. D., Mortimer, J. T. & Durand, D. Modeling of mammalian myelinated nerve for functional neuromuscular stimulation. *IEEE 9th Annual Conference of the Engineering in Medicine and Biology Society* 1577–1578 (1987).
- [16] Sundt, D., Gamper, N. & Jaffe, D. B. Spike propagation through the dorsal root ganglia in an unmyelinated sensory neuron: A modeling study. *Journal of Neurophysiology* **114**, 3140–3153 (2015).
- [17] Musselman, E. D., Cariello, J. E., Grill, W. M. & Pelot, N. A. ASCENT (automated simulations to characterize electrical nerve thresholds): a pipeline for sample-specific computational modeling of electrical stimulation of peripheral nerves. *PLOS Computational Biology* **17**, e1009285 (2021).
- [18] Couppey, T. *et al.* NRV: An open framework for in silico evaluation of peripheral nerve electrical stimulation strategies. *PLOS Computational Biology* **20**, e1011826 (2024).
- [19] Lubba, C. H. *et al.* PyPNS: Multiscale simulation of a peripheral nerve in Python. *Neuroinformatics* **17**, 63–81 (2019).
- [20] Marshall, D. P., Farah, E. S., Musselman, E. D., Pelot, N. A. & Grill, W. M. PyFibers: An open-source NEURON-Python package to simulate responses of model nerve fibers to electrical stimulation. *PLOS Computational Biology* **21**, e1013764 (2025).
- [21] Hussain, M. A., Grill, W. M. & Pelot, N. A. Highly efficient modeling and optimization of neural fiber responses to electrical stimulation. *Nature Communications* **15**, 7597 (2024).
- [22] Deistler, M. *et al.* Jaxley: differentiable simulation enables large-scale training of detailed biophysical models of neural dynamics. *Nature Methods* **22**, 2649–2657 (2025).
- [23] Bradbury, J. *et al.* JAX: composable transformations of Python+NumPy programs (2018). URL <http://github.com/google/jax>.

- [24] Schoenholz, S. S. & Cubuk, E. D. JAX, M.D.: A framework for differentiable physics. In *Advances in Neural Information Processing Systems*, vol. 33 (2020).
- [25] Lung, D. *et al.* golgi: an open-source graphical platform for image-to-recruitment modeling of peripheral nerve stimulation. *bioRxiv* (2026). Companion platform paper (preprint).
- [26] Pelot, N. A., Goldhagen, G. B., Cariello, J. E. & Grill, W. M. Quantified morphology of the pig vagus nerve (2020). Quantified vagus nerve morphology across species.
- [27] Pelot, N. A. *et al.* Quantified morphology of the human vagus nerve with anti-claudin-1 (2021). Quantified vagus nerve morphology across species.
- [28] Davis, C. J., Musselman, E. D., Grill, W. M. & Pelot, N. A. Fibers in smaller fascicles have lower activation thresholds with cuff electrodes due to thinner perineurium and smaller cross-sectional area. *Journal of Neural Engineering* **20**, 026032 (2023).
- [29] Ben-Shalom, R. *et al.* NeuroGPU: Accelerating multi-compartment, biophysically detailed neuron simulations on GPUs. *Journal of Neuroscience Methods* **366**, 109400 (2022).
- [30] Stimberg, M., Goodman, D. F. M. & Nowotny, T. Brian2GeNN: accelerating spiking neural network simulations with graphics hardware. *Scientific Reports* **10**, 410 (2020).
- [31] Kumbhar, P. *et al.* CoreNEURON: An optimized compute engine for the NEURON simulator. *Frontiers in Neuroinformatics* **13** (2019).
- [32] Akar, N. A. *et al.* Arbor — a morphologically-detailed neural network simulation library for contemporary high-performance computing architectures. In *27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*, 274–282 (2019).
- [33] Kilgore, K. L. & Bhadra, N. Reversible nerve conduction block using kilohertz

frequency alternating current. *Neuromodulation: Technology at the Neural Interface* **17**, 242–255 (2014).

[34] Vissamsetti, M., Pelot, N. A. & Grill, W. M. Simulation insights on the compound action potential in multifascicular nerves. *PLOS Computational Biology* (2025).

[35] Eiber, C. D. *et al.* Computational modelling of nerve stimulation and recording with peripheral visceral neural interfaces. *Journal of Neural Engineering* **18**, 066020 (2021).

[36] Lung, D., Jia, Y., Moro, A., Fachino, M. & Haberbush, M. golgi: open-source software for automated nerve model generation and recruitment simulation. *bioRxiv* (2026). Companion software note (preprint).

[37] Schild, J. H. *et al.* A- and c-type rat nodose sensory neurons: model interpretations of dynamic discharge characteristics. *Journal of Neurophysiology* **71**, 2338–2358 (1994).

[38] Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR)* (2015). ArXiv:1412.6980.

[39] Bucksot, J. E. *et al.* Validation of a parameterized, open-source model of nerve stimulation. *Journal of Neural Engineering* **18**, 042001 (2021).

Data availability

The quantified vagus-nerve morphology underlying the histology cohort is openly available from the SPARC Portal: swine²⁶ (dataset doi:10.26275/MAQ2-EII4) and human^{6,27} (dataset doi:10.26275/OFJA-GHOZ). The finite-element field solver and the pre-computed per-contact lead fields used for the cohort are provided by the golgi pipeline^{25,36}. The input and intermediate data required to reproduce every figure and statistic in this paper are deposited in a companion Zenodo data archive (DOI:10.5281/zenodo.21708097).

Code availability

The `jaxon` source code, together with the reproduction code and analysis scripts for every figure and statistic in this paper, is developed openly on GitHub (<https://github.com/CellularSyntax/jaxon>) under the MIT license. The exact software version used for the results reported here is archived as a frozen release on Zenodo (DOI:10.5281/zenodo.21710149).

Acknowledgements

The authors thank the Medical University of Vienna for access to the HPC facility on which the NVIDIA A100 benchmarks and the randomized selectivity sweep were run. This work was supported by the Medical University of Vienna Focus Grant M (project VAGUSPEC).

Author contributions

D.L. and M.H. conceived and designed the study. D.L. implemented the *jaxon* simulator, the coupled solver and the selectivity-optimization pipeline, performed the NEURON validation and the cohort selectivity experiments, and analysed the results. M.H. supervised the project and acquired funding. D.L. and M.H. interpreted the results and wrote the manuscript. Both authors reviewed and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information accompanies this paper.

Correspondence and requests for materials should be addressed to M.H.

4 Methods

4.1 Axon models

jaxon implements four canonical peripheral nerve fiber models: the myelinated McIntyre–Richardson–Grill (MRG) double-cable and Sweeney single-cable axons, and the unmyelinated Sundt and Rattay C-fibers^{10,14–16}; the unmyelinated models follow the biophysical tradition of detailed vagal/nodose C-fiber membrane models³⁷. Each carries its published channel set and diameter-dependent geometry (resting potentials -80 mV for MRG/Sweeney, -60 mV Sundt, -70 mV Rattay); the MRG model is a coupled (V_i, V_{px}) double cable, the others single cables. The governing equations and all constants are given in the Supplementary Material; all four share a common compartment-level interface (`FiberStatics`, `batch.integrate.m.max`), so the library extends to further models without solver changes. Figure 1 summarizes the forward-and-optimize pipeline, FE construction, validation and cohort optimization.

Coupled solver and batched forward pass. The MRG double cable is integrated by backward Euler; each step is a 2×2 block-tridiagonal system in the interleaved intracellular/periaxonal unknowns, solved directly by a JAX-compatible block-Thomas sweep (Supplementary Material). The forward pass for one fiber is wrapped in `jax.lax.scan` over time steps and returns, for optimization, only the per-compartment maximum of the Na^+ activation gate $\mathbf{m}_{\text{Na}}$ — a differentiable recruitment proxy — at ≥ 2 orders of magnitude lower memory than the full trace. This is one of two independent memory savings: the shortened *output* (proxy scalar instead of the full voltage trace), and, with finite-difference gradients, the absence of a reverse-mode *tape* over the time integration (below); the two combine to let the whole population fit on one GPU. For N fibers under the same electrode, the per-fiber static parameters are stacked and a single `jax.vmap` call compiles to one N -way XLA kernel, eliminating the per-fiber Python overhead that dominates NEURON pipelines and driving the throughput gain (Section 2.2). We integrate at $\Delta t = 5 \mu\text{s}$.

4.2 Selectivity optimization

Loss function. For each fiber f we form a differentiable activation proxy $a_f \in [0, 1]$ from the per-compartment maximum over time of the Na^+ activation gate, $\mu_c = \max_t \mathbf{m}_{\text{Na}}(t)|_c$. At each of the two terminal nodes μ is passed through a sigmoid, $\sigma((\mu - \frac{1}{2})/\tau)$ with temperature $\tau = 0.1$, and a_f is the smaller of the two ends — so $a_f \rightarrow 1$ when an action potential propagates to both ends and $a_f \rightarrow 0$ otherwise. The sigmoid smoothing keeps the gradient informative away from threshold; its binary limit ($\tau \rightarrow 0$) is $a_f = \min(\mu_{\text{first}}, \mu_{\text{last}})$, whose gradient vanishes except in a narrow window at threshold. A fiber is counted *recruited* when $a_f > 0.5$ (equivalently, in the binary limit, when the max Na^+ activation gate exceeds $\frac{1}{2}$ at both terminal nodes). We minimize the weighted-quality (WQ) loss

$$L_{\text{WQ}} = \sum_f w_f [y_f(1 - a_f) + (1 - y_f)a_f] + \lambda \cdot \frac{1}{K} \sum_k I_k^2, \quad (1)$$

where $y_f \in \{0, 1\}$ is the target/off-target indicator (constructed per nerve by the peripheral-cluster selector of Section 4.3) and the w_f are class-balancing weights (each of the target and off-target populations carries total weight $\frac{1}{2}$, so neither dominates the imbalanced split). The second term is a small amplitude-energy regularizer that discourages the trivial “fire-everything” optimum; we set $\lambda = 10^{-3}$ for all results reported here, small enough to leave the single-diameter selectivity outcomes unchanged.

Selectivity index. The primary outcome, the (hard) *selectivity index* (SI), is the difference between the recruited fraction of the target population and the recruited fraction of the off-target population,

$$\text{SI} = \frac{1}{|\mathcal{T}|} \sum_{f \in \mathcal{T}} \mathbf{1}[a_f > 0.5] - \frac{1}{|\mathcal{O}|} \sum_{f \in \mathcal{O}} \mathbf{1}[a_f > 0.5], \quad (2)$$

where $\mathcal{T}$ and $\mathcal{O}$ are the target and off-target fiber sets ($y_f = 1$ and $y_f = 0$). SI is bounded on $[-1, 1]$: SI = 1 is perfect selectivity (every target fiber recruited, every off-target fiber spared), SI = 0 is no discrimination (equal recruited fractions), and negative SI is reversed selectivity (more off-target than target recruitment); the “hard SI” used

for the 0.95 early-stop and reported throughout applies the binary $a_f > 0.5$ recruitment rule of Eq. (2), whereas the soft WQ loss of Eq. (1) uses the smooth proxy a_f directly so the optimizer sees an informative gradient. A change of ΔSI corresponds, at fixed off-target recruitment, to a ΔSI change in the fraction of the target fascicle recruited (e.g. the human deployment penalty of ~ 0.14 is predominantly a ~ 14 -percentage-point loss of on-target recruitment; Section 2.4).

Amplitude optimization. The design variable is the per-contact amplitude vector $\mathbf{I} \in \mathbb{R}^K$ ($K = 12$); each contact scales a fixed, charge-balanced asymmetric-biphasic rectangular pulse (0.5 ms cathodic phase followed by a 2.0 ms anodic recharge at one-quarter amplitude, so net charge is zero — the clinical convention for chronic implants). All selectivity results use Adam³⁸ with one-sided finite-difference gradients: each step needs only $K+1$ forward passes (one per contact plus the baseline) and avoids back-propagating through the long integration, so it runs on the full ~ 1000 -fiber population — the regime prior NEURON pipelines could not afford. The one-sided step ($\epsilon = 0.010$ mA, 0.5 % of the ± 2.0 mA amplitude clip) was chosen small enough to resolve the local gradient yet large enough to avoid cancellation noise in the sigmoid recruitment proxy. That the loss landscape is smooth enough for a finite-difference gradient to steer the optimizer is established independently in Supplementary Fig. S1, where a *central* finite-difference estimate agrees with the exact reverse-mode autodiff gradient to a relative L_2 error of 3.7×10^{-3} ; the coarser one-sided step used in the cohort loop is a deliberately cheaper approximation of the same gradient, sufficient for FD-Adam. The amplitude clip, learning rate and finite-difference step are listed in Supplementary Table S6; the loop runs for up to 90 iterations and early-stops when the hard SI reaches 0.95 (a convergence floor: an SI within 0.05 of the perfect-selectivity ceiling, the resolution below which additional iterations no longer change which fibers fire) or when the loss and hard SI both plateau. On the histology cohort the trajectory is warm-started by a discrete probe of physically meaningful spatial patterns — one bipolar and four single-column tripolar (cathode at each column’s z -middle contact with charge-balancing anodes) — each scored at eight magnitudes in both polarities (80 forward passes per seed); the best initializes the tra-

jectory, with the probe’s zero contacts frozen to preserve sparsity. (*jaxon* also exposes exact autodiff gradients, used for waveform-shape optimization; left to future work.)

4.3 Subject-specific SPARC histology FEM cohort

For the cohort-scale selectivity study (Section 2.3) we use subject-specific finite-element (FE) nerve models built from quantified vagus-nerve morphology of swine²⁶ and human²⁷, both released through the SPARC Portal. The raw datasets provide, per nerve, calibrated histological cross-sections with manually segmented epineurium and per-fascicle endoneurium boundaries; the human dataset additionally includes anti-claudin-1 staining that resolves the perineurium thickness directly⁶. This study is a secondary computational analysis of these previously published, de-identified, openly available histology datasets; it involved no new experiments on humans or animals and no new tissue collection. Ethical oversight of the original tissue procurement is described by the source datasets^{6,26,27}; no additional ethics approval was required for the present reuse.

An automated pipeline (*golgi*²⁵) converts each segmented cross-section into a watertight 3-D multi-domain nerve-and-cuff FE model — nerve embedded in a concentric 12-contact cuff (3 axial rows $\times$ 4 angular columns), saline film and muscle bath — and solves the linear electro-quasistatic problem $\nabla \cdot (\sigma \nabla V) = 0$ in FEniCSx with literature conductivities and a zero-thickness perineurium contact impedance, giving a per-contact extracellular lead field $V_e^{(k)}$ (full construction, conductivities and the optimizer defaults in the Supplementary Material). The segmented anatomy, 3-D model and computed fields for every cohort nerve are shown in Supplementary Figs S5–S33.

Fiber sampling, target selection and cohort gating. 1 000 straight axial fibers per nerve are seeded by area-weighted rejection sampling inside the fascicles and the lead fields $V_e^{(k)}$ sampled at every MRG compartment, giving a per-fiber tensor $V_e^{(k)}_{f,n}$ that the optimizer superposes linearly with the amplitudes. Unless otherwise stated, all selectivity results use MRG fibers of a single diameter, $5.7 \mu\text{m}$ — the smallest of the three diameters in the spot-check of Section 3 — so the strategy-averaged human penalty reported in

640 the Results (Section 2.4) and the $5.7\ \mu\text{m}$ spot-check value coincide by construction. The
 641 target/off-target labels y_f in Eq. (1) are set per nerve by a peripheral-cluster selector that
 642 matches what the cuff can steer: a 90° sector (the cuff’s angular resolution) is slid around
 643 the perimeter, and four maximally-separated sectors carrying ≥ 50 peripheral fibers each
 644 become the four target “seeds” optimized per nerve, with the enclosed fascicles labeled
 645 target. Of 36 SPARC bundles processed, 29 nerves (18 swine, 11 human; 116 seeds)
 646 entered the study; the rest were rejected for having < 3 fascicles or insufficient FE-mesh
 647 z -resolution between contact rows (full criteria in the Supplementary Material). These 29
 648 nerves are axial cross-sections from 11 swine and 7 human subjects (four human subjects
 649 each contribute two sections); we treat the nerve cross-section as the unit of replication
 650 throughout and, to guard against pseudoreplication, confirm the species effect at the
 651 subject level in Section 4.5.

652 4.4 Validation protocol

653 For each model we drove *jaxon* and **PyFibers**-wrapped NEURON with identical stimuli,
 654 at a matched fixed time step ($\Delta t = 5\ \mu\text{s}$ — the field-standard MRG step, e.g. the Ax-
 655 onML default — so that the comparison isolates implementation agreement rather than
 656 discretization; NEURON integrated with its default backward-Euler method), and com-
 657 pared three quantities against NEURON: the strength–duration threshold (terminal node,
 658 $\text{PW} \in [0.02, 1.0]\ \text{ms}$), the conduction velocity over the published diameter range, and the
 659 full center-node intracellular trace; agreement to $< 1\%$ on all three was required for ac-
 660 ceptance, following the established practice of validating open-source nerve-stimulation
 661 models against NEURON³⁹. All results were produced on one NVIDIA A100 (40 GB)
 662 with **jax** 0.6.2 / **jaxley** 0.13.0 (CUDA 12.4) and **PyFibers** 0.8.5 (NEURON 8.2); the full
 663 environment and reproduction code are available as described in the Data Availability
 664 statement.

4.5 Statistical analysis

Selectivity comparisons on the histology cohort compare two matched conditions evaluated on the *same* nerve $\times$ target-cluster seed (e.g. dense-optimized vs. centroid-transfer), so we use the paired, two-sided Wilcoxon signed-rank test, which makes no Gaussian assumption on the heavy-tailed, ceiling-bounded SI distributions. The “pairing” refers only to this matching of conditions; the independent unit of replication is always the *nerve*: per-seed values are averaged within each nerve before any cohort-level test, so that nerves rather than nested seeds enter the statistics, and every reported p -value is computed on $n = 29$ nerves (18 swine, 11 human) unless stated otherwise. The deployment penalty (dense-optimized SI minus sparse-optimized transfer SI) is tested per species and sampling strategy with two-sided Wilcoxon signed-rank tests, Holm-corrected across the eight species $\times$ strategy cells; the species difference in penalty is tested with the Mann–Whitney U test, and its anatomical driver — the number of fibers per target fascicle — with Spearman rank correlation. Seeds whose target cluster is degenerate (target fraction > 0.6 — spanning most of the nerve, so no selective solution exists by construction; 5 of 116, all human) are excluded from all analyses; low-SI outcomes on selectable targets are retained as genuine geometric-ceiling cases rather than solver failures. The species penalty difference is insensitive to this cutoff: it is essentially unchanged for exclusion thresholds of 0.5, 0.6 and 0.7 (Mann–Whitney $p = 0.010$, 0.010 and 0.011 respectively). We report medians with inter-quartile ranges alongside means throughout, because the SI distributions are bounded at 1.0 and skewed toward low values (mean below median) by the geometric-ceiling subset, so the median is the more representative central estimate. The derived human deployment-penalty distribution is not a simple right-skew: it is a low-penalty cluster plus a heavy high-penalty tail (so the mean, 0.128, sits just below the median, 0.136), which the median again summarizes more faithfully than the mean. Because the 11 human nerves derive from 7 subjects (and the 18 swine from 11), we additionally confirm the species penalty difference at the subject level, aggregating each subject’s sections to a single value: the difference is essentially unchanged (swine median 0.008, $n = 11$ subjects; human median 0.099, $n = 7$ subjects; Mann–Whitney $p = 0.011$,

694 vs. $p = 0.010$ at the nerve level), so it is not an artifact of counting two sections from the
695 same subject as independent. Analyses use `scipy.stats` (SciPy 1.15.3, Python 3.11).

Supplementary Files

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

- [nrsoftwarepolicy.pdf](#)
- [nrreportingsummary.pdf](#)
- [supplement.pdf](#)