DECISION-IDENTIFIABLE CONTROL OF PARTIALLY OBSERVED EPILEPSY NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Selecting a suppressive target in an epileptic network requires accounting for pathways that intracranial electrodes cannot observe. PULSE connects local probe responses to suppression benefits and uses the uncertainty in those benefits to decide where to measure next. In a stable nonnegative Hawkes network with calibrated probes, integrated observed responses and baseline activity determine stationary outcomes under local gain suppression, even when hidden structure is not uniquely identified. A target's baseline rate and two response summaries suffice; correcting for recurrent returns explains why the strongest evoked response need not identify the best target. In 24 network families whose members differ in hidden structure and transient response, all 2,880 intervention conditions share the benefits predicted by this formula, and predictions from finite measurements err by 2.38 percentage points of baseline observed burden on average across 3,072 interventions. The same summaries also direct measurement: in 96 benchmark networks, decision-focused acquisition reaches 2.35% selection regret at 2,560 trial blocks against 5.33% for uniform allocation, and at 5,120 blocks it keeps regret at 1.72% against 3.11% for an acquisition rule that reconstructs the responses more accurately. Independent post-action simulations of the intervened networks confirm the predicted benefits, and experiments on hidden coverage, probe calibration, and response windows delineate the required measurement conditions, including a 250-ms benchmark built on empirical human response timing. The information a suppression decision needs can thus be read from local probes, without reconstructing the hidden network.

1 INTRODUCTION

Focal treatment of drug-resistant epilepsy requires choosing where suppression will most reduce pathological activity. Intracranial electrodes sample only part of the network, while unrecorded populations can relay and return activity (Figure 1a). The sustained benefit of responsive neurostimulation (Nair et al., 2020) makes target selection consequential, and the hidden pathways must be accounted for through measurements at the recorded sites.

Local stimulation reveals how activity propagates (Keller et al., 2014), but a large response can include repeated returns to the stimulated population. Sustained suppression also interrupts these returns, so the strongest response need not identify the most beneficial target. Appendix C.5 gives a concrete ranking reversal with equal baseline rates. Localization, propagation, and response-prediction methods (Li et al., 2021; Proix et al., 2017; Brinyark et al., 2026) characterize complementary network properties. Establishing which response measurements determine sustained suppression benefits connects these observations to target selection.

PULSE shows that integrated probe responses and baseline activity suffice for this comparison. Observed responses accumulate the effects of propagation through recorded and hidden populations. In a stable nonnegative Hawkes network with calibrated local probes, networks sharing these measurements have identical observed stationary rates under local gain suppression, even when their hidden connections and transient responses differ. We call this property *decision identifiability*.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

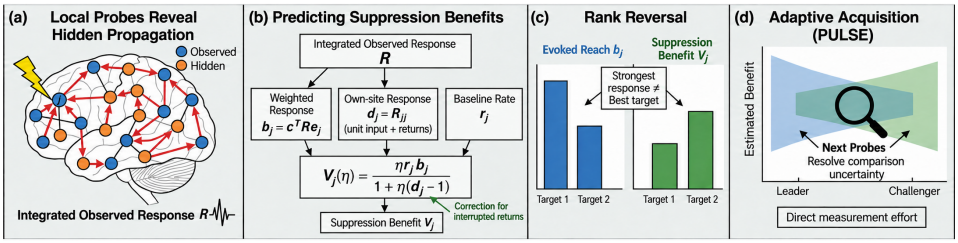


Figure 1: **From local probes to suppression decisions.** (a) A probe at population j propagates through observed (blue) and hidden (orange) populations into the integrated response R. (b) With baseline rate r_j, the summaries $b_j = c^\top Re_j$ and $d_j = R_{jj}$ determine suppression benefit. (c) Evoked reach and suppression benefit can rank targets differently. (d) Comparison uncertainty guides further probes.

For one target j, the calculation requires its baseline event rate r_j, total weighted probe response b_j, and own-site response d_j (Figure 1b). Reducing its event-generation gain by a fraction η lowers weighted observed event burden by

$$V_j(\eta) = \frac{\eta r_j b_j}{1 + \eta(d_j - 1)}, \qquad b_j = c^\top Re_j, \quad d_j = R_{jj}. \tag{1}$$

Here R is the integrated observed response, including the directly induced unit event, and c specifies the burden weights. The quantity $d_j - 1$ measures recurrent returns; the denominator accounts for their interruption by suppression. For fixed r_j and b_j, a larger own-site response therefore gives a smaller suppression benefit.

The formula also guides further measurements. Active stimulation can efficiently learn neural dynamics (Wagenmaker et al., 2024), but overall response accuracy need not settle the target comparison (Wilder et al., 2019; Elmachtoub & Grigas, 2022). PULSE translates uncertainty in the two response summaries into uncertainty about suppression benefits and directs further probes toward the current leader and its strongest challenger (Figure 1d). We make two contributions:

Predicting suppression benefits from local probes. We prove that calibrated probe responses and baseline rates determine local gain-suppression outcomes across compatible hidden networks (Theorem 2). The formulas account for recurrent propagation without reconstructing hidden connectivity; response-equivalent network families, finite-measurement convergence, and independent post-action simulations validate this connection (Section 6.1).

Allocating probes to resolve target comparisons. We develop a sequential acquisition rule using uncertainty in predicted suppression benefits. Across 96 networks, it reduces mean selection regret from 5.33% under uniform allocation to 2.35% at the same budget of 2,560 trial blocks, with matched comparisons separating value scoring from acquisition and response reconstruction from decision quality (Section 6.2).

Two further experiments distinguish incomplete spatial coverage from probe miscalibration (Section 6.3) and examine response windows and empirical human N1 timing (Section 6.4). Public CCEP analyses document the response structure and anatomical support of these measurements. Our code is available at https://anonymous.4open.science/r/PULSE0-20C7/README.md.

2 Related Work

Point-process structure learning. Hawkes inference links excitation kernels to temporal influence and supports flexible bases (Xu et al., 2016). Bayesian models infer latent *edges* among observed event processes (Linderman & Adams, 2014); unobserved node processes require additional treatment. Rank constraints can expose hidden structure (Huang et al., 2022). PULSE identifies the effective mean response from randomized perturbation contrasts, allowing heterogeneous hidden pathways without assuming that the observed marginal process has a closed Hawkes likelihood.

Hidden neural dynamics and epilepsy interventions. The decomposition of effective interactions into direct connections and hidden paths is established in neural network analysis (Brinkman et al.,

2018). PULSE uses this decomposition to establish when a measured probe response determines the benefit of sustained local gain suppression. The intervention value is identical across hidden networks sharing that response and baseline activity. Individualized seizure models and virtual-resection studies address related clinical decisions (Proix et al., 2017; Weiss et al., 2024), while neural fragility and stimulation-based dynamical models provide complementary localization and response information (Li et al., 2021; Brinyark et al., 2026). PULSE evaluates these decisions through the reduction in weighted observed event burden under the specified intervention.

Adaptive experimental design. Best-arm identification allocates measurements to distinguish the best option from its competitors (Audibert & Bubeck, 2010; Garivier & Kaufmann, 2016; Russo, 2020). PULSE applies this principle to suppression benefits computed from two correlated response summaries. Their nonlinear relationship to intervention value determines how response uncertainty enters a target comparison and how a further probe batch can reduce it. The finite-budget acquisition rule is evaluated empirically, and an interval-based decision result provides deterministic regret bounds from simultaneous measurement intervals.

3 Model and intervention target

Consider p observed event processes N_O and q hidden processes N_H. The complete stationary network has conditional intensity

$$\lambda(t) = \mu + \int_{(-\infty,t)} \Phi(t-u)\, dN(u), \qquad \Phi(t) = \begin{bmatrix} \Phi_{OO}(t) & \Phi_{OH}(t) \\ \Phi_{HO}(t) & \Phi_{HH}(t) \end{bmatrix}. \tag{2}$$

The receiving population indexes rows. Write $K = \int_0^\infty \Phi(t)dt$. We assume integrable nonnegative kernels and $\rho(K) < 1$ (Hawkes, 1971; Bacry & Muzy, 2016). Kernels can have different shapes, delays, and supports on every edge. The summary network may contain self-excitation, observed–hidden feedback, and hidden cycles. Nonnegativity describes the propagation of the selected pathological *event type*, a modeling choice about the event stream rather than about synaptic sign.

Assumption 1 (Calibrated, exchangeable probes)**.** *A probe at observed population j adds a known unit immigrant at that population, independently of the pre-probe network state. Pulse and sham trials have the same distribution of ongoing background activity. Probes do not directly recruit hidden populations, and the propagation mechanism remains unchanged across these trials.*

Known non-instantaneous or spatially mixed observed input kernels are also allowed when their transfer matrix is invertible; Appendix B gives the calibration formula. Calibrating direct recruitment brings a clinically applied electrical pulse within this model, and the experiments perturb the hidden-recruitment condition to quantify its effect.

The desired action is a local gain reduction. For $D = \mathrm{diag}(\alpha_1, \ldots, \alpha_p)$ with $0 \le \alpha_i \le 1$, multiply each observed intensity in Equation (2) by α_i, including its baseline and incoming excitation. Hidden intensities retain their mechanisms. A single target has $D = I - \eta e_j e_j^\top$, where $\eta \in [0,1]$ is an *event-generation suppression fraction*, not an electrical stimulation amplitude. Let $r = \mathbb{E}[dN_O(t)]/dt$ and r^D be the observed rates before and after the action. Given nonnegative, prespecified weights c, the decision is

$$j^\star(\eta) = \arg\max_{j \in \mathcal{J}} V_j(\eta), \qquad V_j(\eta) = c^\top (r - r^{D_j}), \tag{3}$$

where $\mathcal{J}$ is an eligible set of recorded targets. This formulation permits anatomical exclusions through $\mathcal{J}$ and prioritization of event streams through c.

4 Predicting suppression benefits from local probes

PULSE first acquires probe responses, then predicts suppression benefits from them, and finally uses the remaining uncertainty to guide further probes (Figure 2).

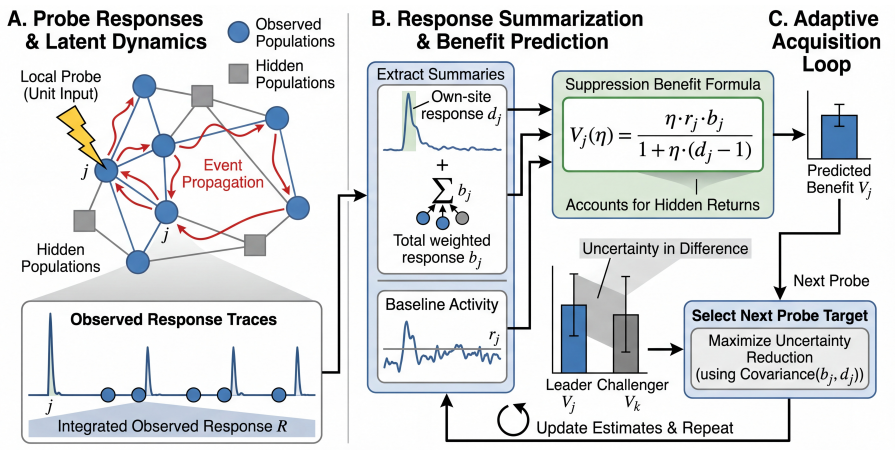


Figure 2: **PULSE method framework.** (A) Pulse/sham contrasts of calibrated probes yield the integrated observed response R. (B) The summaries b_j and d_j combine with baseline rate r_j to predict suppression benefit. (C) Response-summary covariance guides the next probe batch; the loop selects the largest estimated benefit at the total budget.

4.1 Randomized responses reveal hidden-pathway effects

For $s \geq 0$, let $\widehat{\Phi}(s) = \int_0^\infty e^{-st}\Phi(t)dt$ be the Laplace-transformed excitation kernel. Let $\mathcal{R}_{ij}(dt)$ denote the difference in expected observed event count following a unit probe at j and following sham, including the directly induced event at time zero. Its transform is $R(s) = \int_{[0,\infty)} e^{-st}\mathcal{R}(dt)$.

Theorem 1 (Response identification)**.** *Under Equation* (2) *and Assumption 1,*

$$R(s) = [I - \widehat{\Phi}_{\text{eff}}(s)]^{-1}, \tag{4}$$

$$\widehat{\Phi}_{\text{eff}}(s) = \widehat{\Phi}_{OO}(s) + \widehat{\Phi}_{OH}(s)[I - \widehat{\Phi}_{HH}(s)]^{-1}\widehat{\Phi}_{HO}(s). \tag{5}$$

Consequently $\widehat{\Phi}_{\text{eff}}(s) = I - R(s)^{-1}$ is identifiable from observed pulse/sham response contrasts. The identities hold without knowing the hidden dimension, and they allow unequal hidden-path lengths and edge-specific temporal kernels.

The cluster representation of a linear Hawkes process (Hawkes & Oakes, 1974) gives the full response as $\sum_{k\geq 0}\widehat{\Phi}(s)^k$. Taking its observed block and applying block elimination proves the result. Equation (5) sums paths whose internal populations are hidden. It is an effective *mean-response* operator: hidden innovations can remain correlated and history dependent after marginalization. Full proofs and a finite-window bound appear in Appendix B.

With n_j paired pulse/sham blocks and response horizon H, estimate each response column by

$$\widehat{R}_{:j,H}(s) = \frac{1}{n_j}\sum_{\ell=1}^{n_j}\int_{[0,H]} e^{-st}\left[dN_O^{+,j,\ell}(t) - dN_O^{-,j,\ell}(t)\right]. \tag{6}$$

Pairing here defines a measurement block; the two recordings are independent conditional on the shared protocol. For fixed H, this estimates the truncated response. Increasing H reduces truncation but also adds background variance. For displaying an effective kernel, we solve a nonnegative regularized linear inverse problem using $(I - \Phi_{\text{eff}})R = I$; target selection below avoids this matrix inversion entirely.

4.2 Response-equivalent networks have identical intervention values

Set $R = R(0)$, $K_{\text{eff}} = I - R^{-1}$ and $\Gamma = I - D$. Hidden populations can be eliminated from the stationary mean equations, giving $r = K_{\text{eff}}r + \beta$, where

$$\beta = \mu_O + K_{OH}(I - K_{HH})^{-1}\mu_H. \tag{7}$$

Scaling the observed intensities changes this equation to $r^D = DK_{\text{eff}}r^D + D\beta$. The hidden mechanism itself is not refitted after the action.

Theorem 2 (Intervention transport and decision equivalence)**.** *For every admissible D, the observed stationary rate change is*

$$r - r^D = (I - DK_{\text{eff}})^{-1}\Gamma r. \tag{8}$$

Thus any two complete networks satisfying the model assumptions and agreeing on (R, r) have the same observed rates under every such intervention, even if their hidden networks differ. For k targets represented by selection matrix E and suppression matrix $\Lambda = \operatorname{diag}(\eta_1, \ldots, \eta_k)$,

$$r - r^D = RE\Lambda[I + E^\top(R - I)E\Lambda]^{-1}E^\top r. \tag{9}$$

For one target this reduces exactly to Equation (1)*.*

The inverse exists because $0 \le DK_{\text{eff}} \le K_{\text{eff}}$ and K_{eff} is subcritical. Subtracting the baseline and intervention mean equations proves Equation (8); a low-rank inverse identity gives Equation (9). Appendix C supplies the details and constructs two networks with different hidden-node counts but identical (R, r).

For a fixed burden vector c, single-target evaluation requires only the baseline rate r_j, total weighted observed response $b_j = c^\top R e_j$, and own-site response $d_j = R_{jj}$ (Figure 2B). For fixed j, V_j is increasing and concave in η; target ordering can nevertheless change with the requested suppression level.

5 Allocating probes to resolve target comparisons

The acquisition loop in Figure 2C uses the suppression formula to choose where further measurements can clarify the target ranking. For a pulse/sham block at j, retain $X_j = c^\top \Delta N_O([0, H])$ and $Y_j = \Delta N_j([0, H])$. Their means estimate b_j and d_j, while their sample covariance supplies the uncertainty calculation. With $z_j = 1 + \eta(d_j - 1)$, the gradient of V_j with respect to these two response summaries is

$$g_j = \left(\frac{\eta r_j}{z_j},\ -\frac{\eta^2 r_j b_j}{z_j^2}\right)^\top, \qquad v_j = g_j^\top \operatorname{Cov}(X_j, Y_j) g_j. \tag{10}$$

The covariance term enters because a returning event contributes both to total burden and to the target's own count. The leading pulse-dependent variance of the plug-in value is v_j/n_j (Oehlert, 1992). Uncertainty in r_j, estimated from a separate passive recording, contributes a fixed term independent of pulse allocation. Physical constraints $b_j \ge c_j$ and $d_j \ge 1$ are imposed on the plug-in estimates.

PULSE starts with 40 trial blocks per target. At each step it identifies the current leader and the competing target with the largest upper uncertainty index $\widehat{V}_j + 2\sqrt{\widehat{v}_j/n_j}$. Among these two, it chooses the target maximizing

$$\frac{\widehat{v}_j}{n_j(n_j + B)}, \tag{11}$$

where $B = 40$ is the next batch size. Multiplying this expression by B gives the predicted reduction in the variance of the leader–challenger value difference. One target receives each batch, after which the moments and their covariance are updated. At the fixed total budget M, PULSE returns $\arg\max_{j \in \mathcal{J}} \widehat{V}_j$. Appendix F gives pseudocode and implementation details. Appendix F.2 complements this acquisition heuristic with deterministic value bounds from simultaneous measurement intervals.

6 Experiments

The experiments first predict benefits from probes, then allocate probes to the decision, and finally characterize the measurement conditions, including empirical human response timing.

Shared benchmark and evaluation. The main benchmark has 96 networks with eight observed populations and $q \in \{0, 4, 8, 12\}$ hidden populations, 24 networks per setting. Networks contain hidden common causes, recurrent pathways, and edge-specific Gamma kernels with mean delays of 25–220 ms; their integrated spectral radius is 0.76. Methods share a 2,400-second passive recording and target-specific pulse/sham pools. A trial block pairs an independent probe and sham window, each six seconds long. The action suppresses one observed intensity by $\eta = 0.8$, with $c = \mathbf{1}$. Selection regret is $(V_{j^\star} - V_{\hat{j}})/V_{j^\star}$, evaluated by independent complete-network stationary equations. Unless stated otherwise, 95% intervals use 5,000 bootstrap resamples of whole networks or network families. Appendix D specifies the protocols and baselines.

6.1 From local probes to suppression benefits

Different hidden networks, identical benefits. The experiment uses 24 independent families, each with four observed populations and one, two, or four hidden populations. Splitting a hidden population into recurrent clones preserves (R, r) while changing latent dimension and connectivity. Mean delays of 45, 90, and 180 ms also change the temporal responses. Figure 3a–e shows the first family: distinct hidden realizations and transients share the same target-specific benefit curves. Across 72 complete networks, four targets, and ten suppression levels.

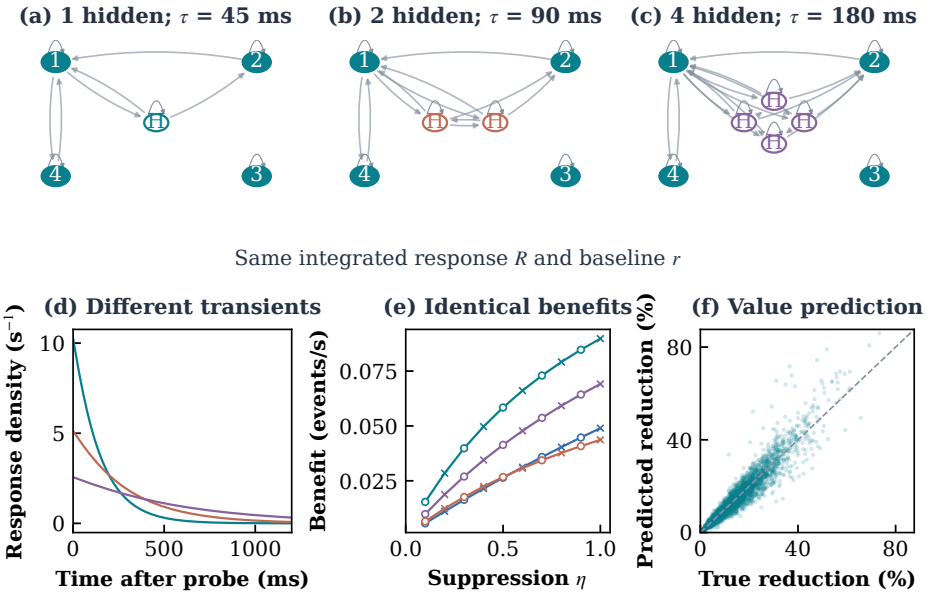


Figure 3: **Hidden structure can vary while suppression benefits remain identifiable.** (a–c) One of 24 response-equivalent families; filled nodes observed, outlined hidden. (d) Responses to target 1, excluding the instantaneous seed. (e) Four targets' benefits; marker and line styles denote one, two, and four hidden populations. (f) Finite-measurement predictions for 3,072 interventions (percentage points of baseline observed burden); dashed line is equality.

Predicting benefits from finite measurements. At 640 blocks per target, predictions using estimated responses and baseline rates achieve a mean absolute error of 2.38 percentage points of baseline observed burden across 3,072 interventions (Figure 3f). To isolate response-estimation error, we evaluate 40, 160, and 640 blocks per target in the 24 response-equivalent families with their known common baseline. Value NMAE, defined as mean absolute target-value error divided by the best target's benefit, decreases from 15.67% [13.24,18.24] to 9.07% [7.65,10.65] and 4.77% [3.89,5.78], respectively (Figure 4a). The three hidden variants are averaged within each family before resampling.

Checking sustained benefits against independent simulations. The population-response prediction is also tested against post-action event rates. Four networks, four targets per network, and four suppression levels give 64 conditions; each uses four independent 12,000-second event streams generated from the intervened model. Predictions use the population (R, r) before intervention, which isolates the transport identity from finite-sample response error, while measured reductions subtract the post-action rate from the population baseline. The mean absolute discrepancy is 0.0088 events/s across all conditions (Figure 4b). Appendix D gives the stream-generation protocol, and Appendix E specifies the family construction and convergence evaluation.

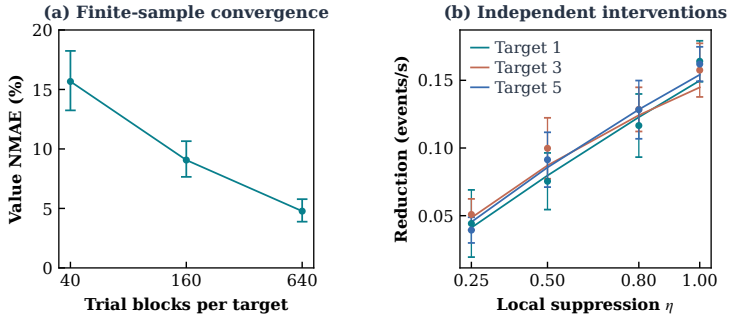


Figure 4: **Convergence and independent intervention checks.** (a) Value NMAE across the 24 response-equivalent families, averaged within family; intervals resample families. (b) Predicted population responses (lines) and independently simulated post-action rate reductions (points) for three targets in the first validation network; error bars are 1.96 standard errors over four 12,000-second event streams. Baseline rates are known in both panels.

6.2 Better target decisions with fewer measurements

Value correction before measurement allocation. Table 1 separates target-scoring inputs. With identical uniform allocation, evoked reach incurs 19.65% regret, baseline-weighted reach without return correction incurs 3.54%, and the PULSE value incurs 2.40%. The known-direct-edges reference has extra structural information but omits hidden paths, yielding 7.77% regret. Passive Hawkes fits use the common baseline recording.

Table 1: **Target scoring and acquisition across 96 networks.** Pulse methods use 5,120 trial blocks and the same passive recordings. Evoked reach and uncorrected transport use uniform allocation. Regret intervals resample networks; known direct edges is an additional-information reference.

Method	Information	Regret (%) ↓	Exact best (%) ↑
Rate only	Passive	11.29 [8.18, 14.55]	51.0
Shared-kernel Hawkes	Passive	2.85 [1.66, 4.18]	74.0
Multi-kernel Hawkes	Passive	3.36 [1.87, 5.06]	70.8
Known direct edges	Known edges	7.77 [5.41, 10.32]	57.3
Evoked reach	Pulse	19.65 [15.33, 24.08]	36.5
Uncorrected transport	Pulse + rate	3.54 [2.06, 5.18]	72.9
PULSE uniform	Pulse + rate	2.40 [1.19, 3.84]	78.1
Response-focused	Pulse + rate	3.11 [1.63, 4.76]	75.0
PULSE w/o covariance	Pulse + rate	2.25 [1.25, 3.55]	75.5
PULSE adaptive	Pulse + rate	1.72 [0.95, 2.60]	78.1

Probing unresolved target comparisons. The four acquisition rules share the value formula, initialization, batch size, baseline estimates, and trial pools; only the allocation rule changes. At 2,560 blocks, PULSE achieves 2.35% mean regret versus 5.33% for uniform acquisition, a paired improvement of 2.98 percentage points [1.15,5.05] (Figure 5a). Uniform acquisition reaches 2.40% regret at 5,120 blocks: PULSE achieves comparable mean decision quality with half the measurements.

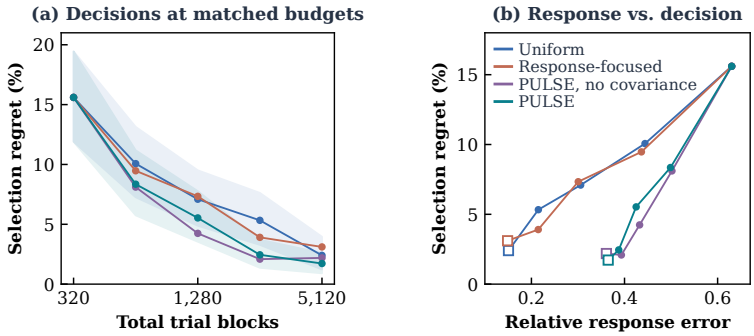


Figure 5: **Measure for the decision.** (a) Mean selection regret across the 96 matched networks; shaded bands are 95% network-bootstrap intervals for PULSE and uniform acquisition. (b) Mean relative response error versus regret at successive budgets.

Response accuracy and decision quality are distinct. At 5,120 blocks, response-focused acquisition has lower relative response error than PULSE, 0.149 versus 0.368, but higher regret, 3.11% versus 1.72% (Figure 5b). Both uniform and adaptive PULSE select the exact optimum in 78.1% of networks at this budget, with the latter incurring smaller loss on the remaining choices. The allocation advantage is largest at intermediate budgets; at 5,120 blocks the gap narrows to 0.68 percentage points in favour of PULSE. Appendix F.3 reports all budgets, a covariance ablation, and a fixed-network acquisition trajectory.

6.3 Hidden coverage and probe calibration

Prediction without specifying hidden-population count. The benchmark is stratified into 24 networks per hidden-population count, each with eight observed populations. At 5,120 blocks, PULSE's mean regret is 1.58%, 1.34%, 3.09%, and 0.87% for $q = 0, 4, 8, 12$, respectively (Figure 6a). Low regret persists even when hidden populations outnumber recorded ones.

Direct probe recruitment requires calibration. In 16 separate networks, each probe additionally recruits one random hidden population with probability $\epsilon \in \{0, 0.1, 0.3, 0.6\}$. With 800 blocks per target and known baseline rates, value NMAE rises from 6.4% at $\epsilon = 0$ to 15.6% at $\epsilon = 0.3$ and 21.8% at $\epsilon = 0.6$. Population responses still give 14.2% and 26.7% error at the latter levels (Figure 6b). Additional samples do not remove this input mismatch (Appendix H).

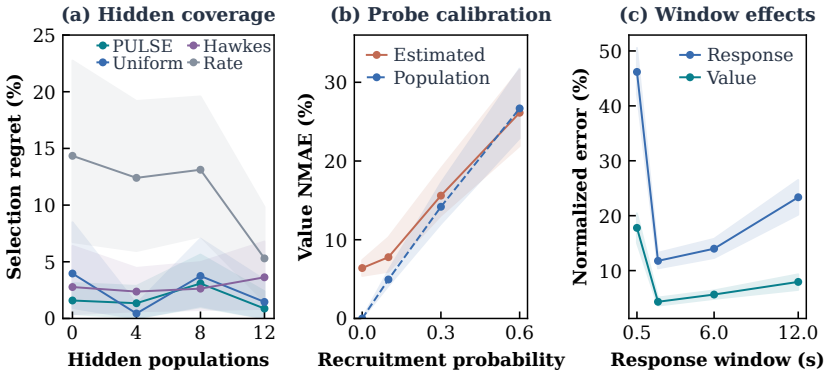


Figure 6: **Hidden coverage and measurement fidelity.** (a) Selection regret at 5,120 blocks (24 networks per setting). (b) Direct hidden recruitment biases values despite population responses. (c) Response-window effects at $\rho = 0.90$. Panels (b–c) assume known baseline rates; bands are 95% network-bootstrap intervals, with value NMAE normalized by the best target's benefit and response error by $\|R\|_F$.

6.4 Response windows and empirical timing

Response windows matched to recurrent propagation. Across five feedback strengths, six horizons, and eight replicates per condition, the lowest tested value NMAE shifts from 2.0% at 0.5 s for $\rho = 0.45$ to 7.8% at 2 s for $\rho = 0.90$ (Figure 7a). Later returns require longer windows, which in turn add background variance. In 12 separate networks with $\rho = 0.90$, using 1,600 blocks per target, NMAE is 15.2%, 4.3%, and 7.9% at 0.5, 2, and 12 s, respectively (Figure 6c). Appendices G and H give the full grid and protocols.

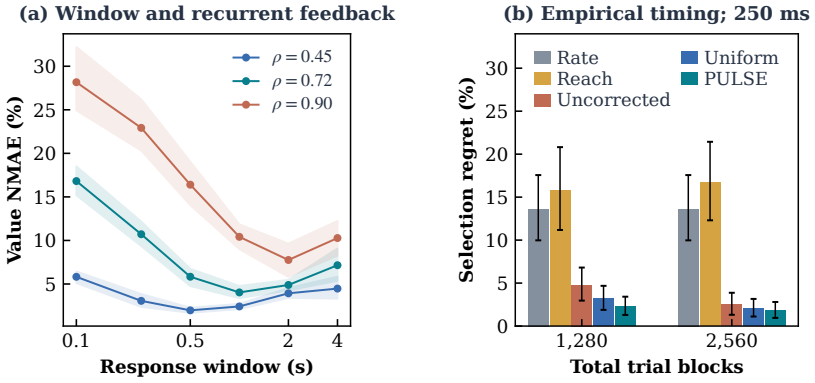


Figure 7: **Measurement conditions shape decision quality.** (a) Value error for three feedback strengths, with known baseline rates; bands resample eight independent replicates. (b) Selection regret at two budgets in 74 networks with empirical N1 timing and a 250-ms window; bars are network-bootstrap 95% intervals. Rate, reach, and uncorrected denote rate-only, evoked-reach, and uncorrected-transport scores.

Decisions on empirical time scales. The public UMC Utrecht RESPect release contains N1 measurements from 74 participants (Van Blooijs et al., 2023b;a). We use each participant's empirical latency distribution to specify propagation timing in this benchmark network, whose event streams are otherwise model-generated, and acquire responses over 250 ms. At 2,560 blocks, mean regret is 1.52% [0.82,2.28] for PULSE and 2.45% [1.45,3.52] for uniform allocation; uncorrected transport, evoked reach, and rate-only selection yield 2.52%, 16.68%, and 13.63%, respectively (Figure 7b). Appendix I gives the protocol; Appendix J documents 680 pathway-average waveforms from 73 participants and sampling coverage from 5,532 electrode coordinates in 74 participants.

7 Discussion and Conclusion

Local probe responses, once corrected for recurrent returns, determine which observed population benefits most from suppression, and measuring where the decision is uncertain matters more than measuring the entire response uniformly. This separates two goals that stimulation experiments often conflate: reconstructing network responses and choosing an intervention. The benchmarks favor the decision-focused objective even though it leaves parts of the response poorly estimated.

The measurement conditions delineate where the approach applies. Hidden pathways are absorbed into the observed response, whereas direct recruitment of unrecorded tissue enters as an input error that only calibration can address, and the response window must balance recurrent-signal truncation against background variance. Empirical N1 timing and anatomical coverage from public CCEP recordings ground these conditions in human data. Translating event-burden predictions toward clinical use further requires calibrated recruitment and an actuator-to-gain mapping to seizure outcomes. Joint-target values are evaluated in Appendix K; joint acquisition, signed interactions, and state-dependent propagation offer further directions.

AI USE STATEMENT

Generative AI was used to polish the language of this manuscript and to review its drafts.

REFERENCES

Jean-Yves Audibert and Sébastien Bubeck. Best arm identification in multi-armed bandits. In *COLT-23th Conference on learning theory-2010*, pp. 13–p, 2010.

Emmanuel Bacry and Jean-François Muzy. First-and second-order statistics characterization of hawkes processes and non-parametric estimation. *IEEE Transactions on Information Theory*, 62 (4):2184–2202, 2016.

Braden AW Brinkman, Fred Rieke, Eric Shea-Brown, and Michael A Buice. Predicting how and when hidden neurons skew measured synaptic interactions. *PLoS computational biology*, 14(10): e1006490, 2018.

Helen Brinyark, Caila Coyne, Nicholas Reyes, Benjamin Cox, Arie Nakhmani, and Rachel Smith. Dynamical network models predict spectral response to single-pulse electrical stimulation. *Journal of Neural Engineering*, 2026.

Bradley Efron. Bootstrap methods: another look at the jackknife. In *Breakthroughs in statistics: Methodology and distribution*, pp. 569–593. Springer, 1992.

Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management science*, 68(1): 9–26, 2022.

Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, pp. 998–1027. PMLR, 2016.

Alan G Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58 (1):83–90, 1971.

Alan G Hawkes and David Oakes. A cluster process representation of a self-exciting process. *Journal of applied probability*, 11(3):493–503, 1974.

Harold V Henderson and Shayle R Searle. On deriving the inverse of a sum of matrices. *SIAM review*, 23(1):53–60, 1981.

Biwei Huang, Charles Jia Han Low, Feng Xie, Clark Glymour, and Kun Zhang. Latent hierarchical causal structure discovery with rank constraints. *Advances in neural information processing systems*, 35:5549–5561, 2022.

Corey J Keller, Christopher J Honey, Pierre Mégevand, Laszlo Entz, Istvan Ulbert, and Ashesh D Mehta. Mapping human brain networks with cortico-cortical evoked potentials. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1653):20130528, 2014.

Adam Li, Chester Huynh, Zachary Fitzgerald, Iahn Cajigas, Damian Brusko, Jonathan Jagid, Angel O Claudio, Andres M Kanner, Jennifer Hopp, Stephanie Chen, et al. Neural fragility as an eeg marker of the seizure onset zone. *Nature neuroscience*, 24(10):1465–1474, 2021.

Scott Linderman and Ryan Adams. Discovering latent network structure in point process data. In *International conference on machine learning*, pp. 1413–1421. PMLR, 2014.

Dileep R Nair, Kenneth D Laxer, Peter B Weber, Anthony M Murro, Yong D Park, Gregory L Barkley, Brien J Smith, Ryder P Gwinn, Michael J Doherty, Katherine H Noe, et al. Nine-year prospective efficacy and safety of brain-responsive neurostimulation for focal epilepsy. *Neurology*, 95(9): e1244–e1256, 2020.

Gary W Oehlert. A note on the delta method. *The American Statistician*, 46(1):27–29, 1992.

Timothée Proix, Fabrice Bartolomei, Maxime Guye, and Viktor K Jirsa. Individual brain structure and modelling predict seizure propagation. *Brain*, 140(3):641–654, 2017.

Daniel Russo. Simple bayesian algorithms for best-arm identification. *Operations Research*, 68(6): 1625–1647, 2020.

D Van Blooijs, MA van den Boom, JF van der Aar, GJM Huiskamp, G Castegnaro, M Demuru, WJEM Zweiphenning, P van Eijsden, KJ Miller, FSS Leijten, et al. Ccep ecog dataset across age 4-51. *DOI: doi*, 10, 2023a.

Dorien Van Blooijs, Max A Van Den Boom, Jaap F Van Der Aar, Geertjan M Huiskamp, Giulio Castegnaro, Matteo Demuru, Willemiek JEM Zweiphenning, Pieter Van Eijsden, Kai J Miller, Frans SS Leijten, et al. Developmental trajectory of transmission speed in the human brain. *Nature Neuroscience*, 26(4):537–541, 2023b.

Andrew Wagenmaker, Lu Mi, Marton Rozsa, Matthew S Bull, Karel Svoboda, Kayvon Daie, Matthew D Golub, and Kevin Jamieson. Active learning of neural population dynamics using two-photon holographic optogenetics. *Advances in neural information processing systems*, 37: 31659–31687, 2024.

Shennan Aibel Weiss, Michael R Sperling, Jerome Engel Jr, Anli Liu, Itzhak Fried, Chengyuan Wu, Werner Doyle, Charles Mikell III, Sima Mofakham, Noriko Salamon, et al. Simulated resections and responsive neurostimulator placement can optimize postoperative seizure outcomes when guided by fast ripple networks. *Brain Communications*, 6(5):fcae367, 2024.

Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 1658–1665, 2019.

Hongteng Xu, Mehrdad Farajtabar, and Hongyuan Zha. Learning granger causality for hawkes processes. In *International conference on machine learning*, pp. 1717–1726. PMLR, 2016.

A NOTATION AND THE IDENTIFIED OBJECT

Symbol	Meaning
O, H	Observed and hidden population index sets
$\Phi_{AB}(t)$	Excitation from populations in B into populations in A
K_{AB}	Integrated excitation $\int_0^\infty \Phi_{AB}(t)dt$
$R(s)$	Laplace transform of the observed unit-probe response
$R = R(0)$	Integrated observed response, including the unit seed
K_{eff}	Integrated effective operator $I - R^{-1}$
r, r^D	Observed mean rates before and after gain intervention
β	Effective baseline vector in Equation (7)
$b_j = c^\top R e_j$	Scalar burden-response moment for target j
$d_j = R_{jj}$	Scalar self-return moment for target j
D, Γ	Observed retained gain and suppression: $\Gamma = I - D$
η, B, H	Suppression fraction, acquisition batch size, response horizon

Table 2: Notation for the response operator, intervention, and decision statistics.

The operator Φ_{eff} summarizes pathways that begin and end in the observed set and have only hidden internal populations. It is distinct from the physical observed-edge kernel Φ_{OO}. Its inverse resolvent additionally accounts for repeated passages through observed populations. Identifying Φ_{eff} does not identify the number, labels, or individual edges of hidden populations. The equivalence statement covers all models sharing the observed *response and baseline means*.

B RESPONSE IDENTIFICATION

B.1 PROOF OF THEOREM 1

Write $F(s) = \widehat{\Phi}(s)$. Since $0 \le F(s) \le K$ entrywise for $s \ge 0$ and $\rho(K) < 1$, the Neumann series converges:

$$Q(s) = \sum_{k=0}^{\infty} F(s)^k = (I - F(s))^{-1}. \tag{12}$$

The kth term is the expected transform of the kth offspring generation from one immigrant. The zeroth term is the immigrant itself. Nonnegativity and integrability justify exchanging expectations, sums, and integrals. Adding an independent immigrant to a linear Hawkes process superposes its descendant cluster on the background process. Therefore, subtracting matched sham expectations leaves exactly the immigrant cluster response. Observing only O gives $R(s) = Q_{OO}(s)$.

For clarity, solve the response equations for an arbitrary observed input column u:

$$x_O = u + F_{OO}x_O + F_{OH}x_H, \tag{13}$$

$$x_H = F_{HO}x_O + F_{HH}x_H. \tag{14}$$

The principal block F_{HH} is subcritical, so $x_H = (I - F_{HH})^{-1}F_{HO}x_O$. Substitution yields

$$[I - F_{OO} - F_{OH}(I - F_{HH})^{-1}F_{HO}]x_O = u. \tag{15}$$

This proves both the response identity and the Schur expression in Theorem 1. Each power of F_{HH} represents an additional hidden step, so the argument allows arbitrary unequal path lengths and recurrent hidden paths, with unrelated temporal shapes across entries of F and no assumption of covariance rank faithfulness or a fixed number of latent sources.

B.2 KNOWN INPUT CALIBRATION

If the physical probe generates a known observed input-transfer matrix $G(s)$, the measured contrast is $M(s) = R(s)G(s)$. A square invertible $G(s)$ gives $R(s) = M(s)G(s)^{-1}$; a full-row-rank input design gives the corresponding right-inverse solution. Thus direct spread within the observed set is permitted when it is known and spans the relevant observed directions. Unknown direct

hidden recruitment adds a different term to $M(s)$ and cannot generally be removed by this observed calibration. This is the mechanism tested in Appendix H.

The direct immigrant at time zero must be included in the response convention. If artifact blanking removes its recorded signature, the known unit contribution is added analytically. Subtracting the identity a second time, or excluding this contribution when computing R_{jj}, changes the return correction and gives the wrong intervention value.

B.3 Finite response windows

Let $W(t) = \sum_{k\geq 1} \Phi^{*k}(t)$, where $*$ is matrix convolution. With the identity atom included, $Q(s) = I + \int_0^\infty e^{-st} W(t)dt$. The finite-window bias of the observed response estimate is

$$R(s) - R_H(s) = \int_H^\infty e^{-st} W_{OO}(t)dt \geq 0. \tag{16}$$

This is an entrywise statement. If there exists $a > 0$ such that

$$K_a = \int_0^\infty e^{at}\Phi(t)dt, \qquad \rho(K_a) < 1, \tag{17}$$

then

$$0 \leq R(s) - R_H(s) \leq e^{-(a+s)H} \left[(I - K_a)^{-1} - I\right]_{OO}. \tag{18}$$

Indeed, $e^{-(a+s)t} \leq e^{-(a+s)H}$ for $t \geq H$, and the exponentially weighted integral of the convolution series is $\sum_{k\geq 1} K_a^k$. For Gamma kernels the exponential moment exists below the smallest edge-specific inverse scale; subcriticality at a sufficiently small positive a follows by continuity. The bound establishes vanishing bias; at a fixed measurement budget, longer windows also add variance through the subtracted background counts.

For fixed H, independent trial blocks with finite second moments give the usual law of large numbers and multivariate central limit theorem for Equation (6). In the limit $H \to \infty$, the estimand approaches $R(s)$. Continuity of inversion at the nonsingular population $R(s)$ then identifies the effective kernel. A joint statistical limit must balance increasing trial counts against the growing variance of longer windows; the fixed-window experiments examine this balance empirically.

B.4 Nonnegative operator estimation

For the operator plots, given $\widehat{R}(s)$ we estimate each row a_i of $\widehat{\Phi}_{\text{eff}}(s)$ by

$$\arg\min_{a_i \geq 0} \left\| a_i \widehat{R}(s) - [\widehat{R}(s) - I]_{i:} \right\|_2^2 + n^{-1/2}\|a_i\|_2^2. \tag{19}$$

This is a nonnegative least-squares problem with an augmented design matrix. To keep fitted systems stable, any resulting matrix whose spectral radius exceeds 0.98 is projected onto the spectral-norm ball by uniform rescaling to 0.98. This is a standard finite-sample stability constraint applied identically to every diagnostic fit; the intervention-value estimator uses the two raw response projections with physical lower bounds and does not use this inversion or spectral projection.

C Intervention transport

C.1 Effective mean equations and stability

The complete mean rates $\bar{r} = (r, r_H)$ satisfy $\bar{r} = \mu + K\bar{r}$. Eliminating the hidden block gives

$$r_H = (I - K_{HH})^{-1}(\mu_H + K_{HO} r), \tag{20}$$

$$r = \underbrace{[K_{OO} + K_{OH}(I - K_{HH})^{-1}K_{HO}]}_{K_{\text{eff}}} r + \underbrace{[\mu_O + K_{OH}(I - K_{HH})^{-1}\mu_H]}_{\beta}. \tag{21}$$

All entries of K_{eff} and β are nonnegative. Block inversion shows $(I - K_{\text{eff}})^{-1} = R$, which is nonnegative and includes the identity. This implies $\rho(K_{\text{eff}}) < 1$. To see this directly, if a nonnegative Perron vector $v \neq 0$ had eigenvalue $\rho > 1$, then $Rv = v/(1-\rho)$ would contradict nonnegativity; $\rho = 1$ would contradict invertibility. Consequently every DK_{eff} with $0 \leq D \leq I$ is also subcritical.

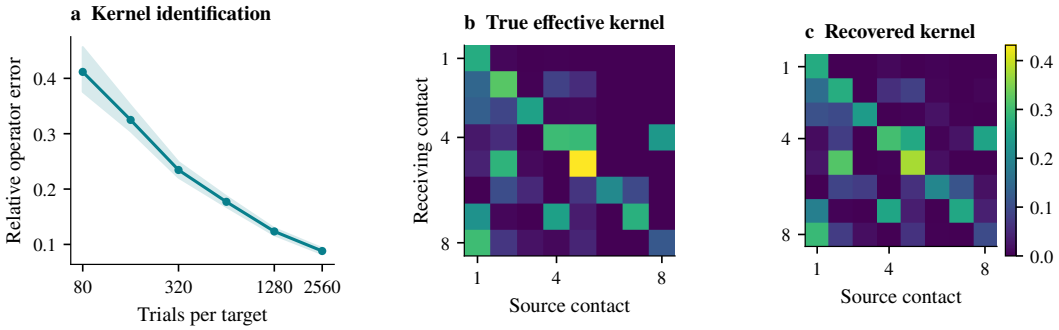


Figure 8: **Recovering effective dynamics with heterogeneous kernels.** (a) Relative Frobenius error, averaged over four Laplace frequencies and 12 independent networks; shading is a network-bootstrap 95% interval. (b,c) Population and estimated zero-frequency kernels for the prespecified first diagnostic network, with eight observed and eight hidden populations. The estimator uses 2,560 pulse/sham blocks per target.

C.2 Proof of Theorem 2

The intervention scales rows in the observed intensity equations. The hidden equations keep the same coefficients, so eliminating hidden means after the action gives

$$r^D = DK_{\text{eff}} r^D + D\beta. \tag{22}$$

Using $\beta = (I - K_{\text{eff}})r$ and subtracting yields

$$(I - DK_{\text{eff}})(r - r^D) = (I - DK_{\text{eff}})r - D(I - K_{\text{eff}})r \tag{23}$$

$$= (I - D)r = \Gamma r. \tag{24}$$

The inverse exists by the preceding stability argument, establishing Equation (8). This equation contains only (K_{eff}, r), which is equivalent to (R, r); hence it proves response-equivalent intervention values.

For multiple targets, write $\Gamma = E\Lambda E^\top$. Then

$$I - DK_{\text{eff}} = I - K_{\text{eff}} + E\Lambda E^\top K_{\text{eff}}. \tag{25}$$

The push-through inverse identity (Henderson & Searle, 1981) gives

$$(I - DK_{\text{eff}})^{-1} E\Lambda = RE\Lambda[I + E^\top K_{\text{eff}} RE\Lambda]^{-1} \tag{26}$$

$$= RE\Lambda[I + E^\top (R - I)E\Lambda]^{-1}. \tag{27}$$

Multiplication by $E^\top r$ gives Equation (9). The expression permits complete suppression, since it does not require Λ^{-1} or D^{-1}.

C.3 Single-target interpretation and dose dependence

Setting $E = e_j$ and $\Lambda = \eta$ gives

$$r - r^{D_j} = \frac{\eta r_j}{1 + \eta(R_{jj} - 1)} Re_j. \tag{28}$$

The burden formula follows by multiplying by $c^\top$. Define $A = r_j c^\top Re_j \geq 0$ and $B = R_{jj} - 1 \geq 0$. Then

$$V_j'(\eta) = \frac{A}{(1 + \eta B)^2}, \qquad V_j''(\eta) = -\frac{2AB}{(1 + \eta B)^3}. \tag{29}$$

The gain-reduction effect is monotone with diminishing returns. Comparing two targets involves two different denominators, so the target with the larger marginal effect at $\eta = 0$ need not be optimal at a larger suppression fraction.

C.4 Different latent structures with the same action values

Let

$$A = \begin{bmatrix} .15 & .12 \\ .05 & .20 \end{bmatrix}, \quad B = \begin{bmatrix} .40 \\ .25 \end{bmatrix}, \quad C = \begin{bmatrix} .25 & .08 \end{bmatrix}, \quad h = .20. \tag{30}$$

One network has the integrated matrix

$$K^{(1)} = \begin{bmatrix} A & B \\ C & h \end{bmatrix} \tag{31}$$

and one hidden population with baseline .04. A second network has

$$K^{(2)} = \begin{bmatrix} A & B/2 & B/2 \\ C & h & 0 \\ C & 0 & h \end{bmatrix} \tag{32}$$

and two independent hidden populations, each with baseline .04. Both have observed baselines $(.02, .03)^\top$. Their effective kernels and baseline means agree because the two half-strength outgoing hidden contributions add to the original contribution. Duplicating the temporal kernel of every corresponding edge extends this equality to all $s \geq 0$. The two networks therefore have identical observed intervention means under every admissible D. The equivalence concerns intervention means; higher-order observed fluctuations can differ. The supplied numerical check verifies the response/rate equality to machine precision.

C.5 Exact ranking-reversal example

The ranking distinction shown in Figure 1c has a concrete realization in the three-population network

$$K = \begin{bmatrix} .75 & 0 & 0 \\ 0 & .23 & 0 \\ .10 & .69 & .20 \end{bmatrix}, \qquad r = c = \mathbf{1}, \qquad \mu = (I - K)\mathbf{1} = (.25, .77, .01)^\top. \tag{33}$$

The first target has a larger integrated evoked reach because of its strong self-return. From $R = (I - K)^{-1}$, the first two targets have $(b_1, b_2) = (4.5, 2.4188)$ and $(d_1, d_2) = (4, 1.2987)$, rounded to four decimal places. At $\eta = .8$, Equation (1) gives $(V_1, V_2) = (1.0588, 1.5618)$ events/s: the second target has the larger suppression value despite its smaller evoked reach. These values follow directly from the specified network and equal baseline rates.

D Complete experimental specification

D.1 Network generation

Every matrix entry is oriented as receiving population by source population. For $n = p+q$ populations, initialize off-diagonal edges independently with probability 0.10 and draw their integrated weights uniformly from $[0.04, 0.28]$. Draw self-excitation uniformly from $[0.06, 0.38]$. For each hidden population, choose two distinct observed children and add weights drawn from $[0.12, 0.42]$. With probability 0.8, add an observed-to-hidden input of weight in $[0.12, 0.38]$. With probability 0.6, add a connection from the next hidden population in a circular indexing to the current one, with weight in $[0.08, 0.24]$. Add 0.25 to one randomly selected observed matrix entry. Finally, rescale the complete matrix to its specified spectral radius.

Each edge kernel is $\Phi_{ij}(t) = K_{ij}\,\mathrm{GammaPDF}(t; k_{ij}, \tau_{ij}/k_{ij})$, with k_{ij} sampled uniformly from $\{1, 2, 3\}$ and τ_{ij} sampled log-uniformly between 0.025 and 0.22 seconds. Background rates are sampled uniformly from $[0.008, 0.025]$ events/s; observed rates receive an additional independent factor from $[0.65, 2.5]$. Main-network spectral radius is 0.76; the near-critical diagnostic uses 0.90. Network seeds are fixed before any method is evaluated.

Main-network seeds are $13000 + 100q + u$, for $q \in \{0, 4, 8, 12\}$ and $u = 0, \ldots, 23$. Seed 12000 is reserved for the implementation pilot and lies outside the reported benchmark ranges. Additional networks use seeds 22000–22011 for operator recovery, 24000–24015 for direct hidden recruitment, 25000–25011 for response windows, and 26000–26003 for direct intervention validation.

D.2 CONTINUOUS-TIME EVENT GENERATION AND OBSERVATION

We generate independent background immigrants for every type on $[-12, H]$ in each trial. Each event of source type j generates a Poisson number of offspring with mean $\sum_i K_{ij}$; offspring types are drawn proportionally to K_{ij} and their delays from the corresponding Gamma density. Generation continues until no descendants remain within the observation horizon. This is the Poisson cluster construction. The process is stable, and no generation guard is triggered in the reported experiments.

Treatment trials add one known observed immigrant at time zero. Sham trials use independent background events from the same distribution. Their observed counts are subtracted to form the measurement. The estimator receives only observed trial counts; event ancestry, hidden counts, and the underlying excitation matrix are held out. The same finite burn-in is used for both arms; matching their background distributions makes the expected response contrast independent of their shared startup transient. The separate 2,400-second passive record uses the same burn-in.

For each main network, 2,560 trial blocks per candidate are generated into a reusable evaluation pool. A method with total budget M can inspect only the first n_j blocks it selects at each target, with $\sum_j n_j = M$. Uniform acquisition uses $n_j = M/8$. Adaptive acquisition cannot request beyond the stored per-target pool; if both preferred candidates are exhausted, it uses the available target with the largest variance-reduction index. This finite-pool convention is implemented identically for all budget evaluations. Budgets are $320, 640, 1280, 2560, 5120$ blocks.

D.3 PASSIVE BASELINES AND ABLATIONS

Observed-only Hawkes baselines use 50-ms binned counts from the passive recording. For each temporal basis, an exponentially filtered history is formed from previous bins only. The one-basis model has decay time 0.15 s; the three-basis model has decay times $(0.05, 0.15, 0.5)$ s. For each target, a nonnegative Poisson likelihood is optimized with an intercept and all observed source histories. The objective adds 0.3 times the sum of non-intercept coefficients. Up to 25,000 evenly spaced observations enter each fit. To ensure identically stabilized fits, integrated coefficient matrices whose spectral radius exceeds 0.97 are projected onto the spectral-norm ball by uniform rescaling to 0.97. These fitted models are evaluated by the same gain-intervention formula, using their fitted observed-only resolvents and the common empirical baseline rate vector.

Optimization uses L-BFGS-B with at most 600 iterations and relative objective tolerance 10^{-10}. The intercept is bounded below by 10^{-7} and each history coefficient lies in $[0, 2]$. All 1,536 target-level fits across the two baselines and 96 networks report convergence; their per-fit status and iteration counts are included in the source package.

The known-direct-edges reference instead inserts the physical K_{OO} in that observed-only calculation. It omits the hidden-pathway term, isolating the contribution of hidden-pathway information. The rate-only baseline chooses $\arg\max_j \widehat{r}_j$. Evoked reach chooses $\arg\max_j c^\top \widehat{R} e_j$. Uncorrected transport chooses $\arg\max_j \widehat{r}_j c^\top \widehat{R} e_j$ without the return denominator. Both response ablations use the uniform per-target trial allocation at each matched total budget. PULSE uses physically constrained moment estimates $\widehat{b}_j = \max(c_j, \overline{X}_j)$ and $\widehat{d}_j = \max(1, \overline{Y}_j)$.

D.4 GROUND TRUTH, METRICS, AND REPLICATION UNITS

For each intervention, the ground-truth complete matrix is obtained by multiplying row j of K and entry j of μ by $1 - \eta$. Its stationary rate is solved directly from this modified complete model. This calculation does not use the estimated R or the effective kernel. The normalized selection regret is

$$\mathcal{L} = \frac{\max_j V_j - V_{\widehat{j}}}{\max_j V_j}. \tag{34}$$

All candidates have a positive baseline and all burden weights equal one, so the denominator is nonzero. Exact-target accuracy uses the maximizing complete-model target without a near-tie tolerance. For dose-prediction plots, each value is divided by $c^\top r$; errors are therefore percentage points of baseline event burden. For stress tests, mean absolute value error across targets is divided by the network's maximum true target value, and operator error is the relative Frobenius norm.

Confidence bands over network experiments resample whole network seeds 5,000 times (Efron, 1992). For the operator plot, errors are first averaged over the four Laplace frequencies within each network, then resampled over the 12 networks. The paired acquisition comparison resamples the per-network difference between adaptive and uniform regret. Individual intervention conditions within the same network are not treated as independent samples for inferential claims.

D.5 Independent validation of intervention means

The four intervention-validation networks each use targets $0, 2, 4, 6$ in zero-based code indexing and suppression levels $0.25, 0.5, 0.8, 1.0$, giving 64 distinct actions. For every action, four independent 12,000-second event streams are generated from the modified complete model. The measured event-rate reduction is the population baseline observed rate minus the empirical post-action observed rate. This check predicts from population R, which isolates the transport identity from pulse estimation noise. Figure 4b displays targets $0, 2, 4$ from the first validation seed; the aggregate discrepancy includes all 64 actions.

Separately, 50 random networks with $q = 0, \dots, 9$ and $\rho(K) \in [0.55, 0.942]$ verify the single-target identity at four suppression levels for all eight targets, and the two- and three-target identities. Maximum errors against direct complete-model solutions are 1.53×10^{-15} and 1.67×10^{-16}, respectively. These checks cover matrix orientation and full suppression; the analytic proofs establish the general results.

E Constructing and evaluating response-equivalent networks

For family $f \in \{0, \dots, 23\}$, seed $51000 + f$ generates a four-observed/one-hidden network with spectral radius $0.68 + 0.02(f \bmod 4)$. Write its integrated matrix as $K = \left[\begin{smallmatrix} A & B \\ C & h \end{smallmatrix}\right]$. Replace the hidden population by $k \in \{1, 2, 4\}$ clones:

$$K^{(k)} = \begin{bmatrix} A & \frac{1}{k} B \mathbf{1}_k^\top \\ \mathbf{1}_k C & h S_k \end{bmatrix}, \qquad S_k = 0.4 I + 0.6 P_k, \qquad \mu^{(k)} = \begin{bmatrix} \mu_O \\ \mu_H \mathbf{1}_k \end{bmatrix}, \tag{35}$$

where P_k is a cyclic permutation. Since $S_k \mathbf{1}_k = \mathbf{1}_k$, eliminating the clones yields the same $A + B(1-h)^{-1}C$ and the same effective immigrant term. Hence both $R(0)$ and r agree. The subspace with identical clone rates recovers the original network; the remaining hidden modes have magnitude at most $h < 1$, retaining stability. Each variant uses an exponential kernel on every nonzero edge, with mean 45, 90, or 180 ms respectively. This diagnostic varies the hidden realization and the temporal response while preserving the integrated decision information, whereas the main benchmark uses heterogeneous edge kernels.

For all four targets and $\eta \in \{0.1, 0.2, \dots, 1.0\}$, we independently solve the complete post-action stationary system and compare it with the shared response formula. Maximum absolute discrepancies are 3.56×10^{-15} for R, 2.23×10^{-16} for baseline rates, and 5.97×10^{-16} events/s for intervention values. For the finite-sample experiment, each target supplies 640 pulse/sham blocks with a six-second horizon; prefixes of 40, 160, and 640 fix the common baseline at its population value to isolate response estimation. Trial seeds are $61000 + 10f + v$, $v \in \{0, 1, 2\}$. Value NMAE is mean absolute target-value error divided by the maximum true target value. The three variants are averaged within family before bootstrap, so clones are not treated as independent replicates.

The family-averaged finite-sample NMAE decreases from 15.67% [13.24,18.24] at 40 blocks per target to 4.77% [3.89,5.78] at 640 (Figure 4a). Across the 64 independent intervention conditions described in Appendix D, the mean absolute discrepancy is 0.0088 events/s (Figure 4b).

F ACQUISITION AND UNCERTAINTY

Algorithm 1 PULSE decision-focused acquisition

Require: Eligible targets $\mathcal{J}$, baseline rates $\widehat{r}$, weights c, suppression η, total budget M, batch size $B = 40$

```
Collect 40 pulse/sham blocks at each target
while fewer than M blocks have been used do
    Estimate b̂_j, d̂_j, V̂_j, and v̂_j for each j
    a ← arg max_j V̂_j
    b ← arg max_{j≠a} {V̂_j + 2√(v̂_j/n_j)}
    j ← arg max_{j∈{a,b}} v̂_j/[n_j(n_j + B)]
    Collect up to B further blocks at j, respecting the remaining budget
    Update the two response moments and their covariance
end while
return arg max_j V̂_j
```

F.1 WHY THE ACQUISITION CRITERION USES CORRELATED PROJECTIONS

For one target, each trial produces a vector $Z = (X, Y)$, where X is total weighted response and Y is self-response. Differentiating the ratio $f(b, d) = \eta r_j b/[1 + \eta(d - 1)]$ yields Equation (10). Under a multivariate central limit theorem,

$$\mathrm{Var}(\widehat{V}_j) \approx \frac{g_j^\top \Sigma_j g_j}{n_j} + \left(\frac{\eta b_j}{1 + \eta(d_j - 1)}\right)^2 \mathrm{Var}(\widehat{r}_j). \tag{36}$$

The second term is independent of pulse allocation when the passive recording is fixed. For two independent target-specific probe streams, their pulse-dependent comparison variance is $v_a/n_a + v_b/n_b$. A batch of B observations at j decreases this expression by $Bv_j/[n_j(n_j + B)]$, yielding the greedy rule. With fixed known variances and a two-target budget, the continuous allocation optimum satisfies $n_a/n_b = \sqrt{v_a/v_b}$. PULSE estimates these quantities online and also updates which two targets matter.

The online sufficient summaries for a fixed target are the counts of trials, the sums of X and Y, and their three distinct second moments. Full response matrices are unnecessary for single-target selection. The released benchmark nevertheless retains observed count vectors to support alternative burden weights and the operator-recovery diagnostics.

F.2 PROOF OF THE INTERVAL DECISION RESULT

Value intervals. Suppose simultaneous intervals contain (r_j, b_j, d_j) for every target, with physically valid endpoints. Then

$$L_j = \frac{\eta \underline{r}_j \underline{b}_j}{1 + \eta(\overline{d}_j - 1)}, \qquad U_j = \frac{\eta \overline{r}_j \overline{b}_j}{1 + \eta(\underline{d}_j - 1)} \tag{37}$$

contain the true value. For $\widehat{j} = \arg\max_j L_j$, regret is at most $\max_j U_j - L_{\widehat{j}}$ on that coverage event. A strict inequality $L_{\widehat{j}} > \max_{k \neq \widehat{j}} U_k$ certifies the optimal target under the same condition. The bound inherits the simultaneous coverage event of the input intervals.

Assume one event contains every r_j, b_j, and d_j within the supplied intervals. The value function increases in r_j and b_j and decreases in d_j when the other arguments are fixed. Its denominator is positive on the physically valid interval. Substituting opposite endpoints proves $L_j \leq V_j \leq U_j$ simultaneously. For $\widehat{j} = \arg\max_j L_j$,

$$V_{j^\star} - V_{\widehat{j}} \leq \max_j U_j - L_{\widehat{j}}. \tag{38}$$

If $L_{\widehat{j}} > \max_{k \neq \widehat{j}} U_k$, the selected target is strictly optimal on the coverage event. These conclusions are deterministic; a statistical guarantee inherits the coverage properties of the input intervals.

A Bonferroni construction of the input intervals satisfies this coverage condition, with physical endpoints truncated at $r_j \geq 0$, $b_j \geq c_j$, and $d_j \geq 1$. The bound is stated as a deterministic certification criterion for the selected target.

F.3 Matched acquisition objectives and covariance ablation

All four acquisition methods use the same 96 main-network trial pools, an initial 40 blocks at every target, batches of 40, and a per-target cap of 2,560 blocks. Uniform acquisition adds a batch to a least-sampled target. Response-focused acquisition chooses the target maximizing $\mathrm{tr}(\widehat{\Sigma}_j)/[n_j(n_j+40)]$, where $\widehat{\Sigma}_j$ is the covariance of the entire observed pulse-minus-sham response vector. This is the one-step variance-reduction objective for estimating all response columns under squared error. It serves as a matched objective comparator.

The no-covariance ablation keeps PULSE's value formula, physical clipping, competing-target rule, and trial pools, but replaces the projected variance by $g_{j1}^2\widehat{\mathrm{Var}}(X_j) + g_{j2}^2\widehat{\mathrm{Var}}(Y_j)$. The full rule also includes $2g_{j1}g_{j2}\widehat{\mathrm{Cov}}(X_j, Y_j)$. The ablation changes both the comparison index and allocation variance.

Table 3: Mean selection regret (%) across the 96 matched networks.

Acquisition	320	640	1,280	2,560	5,120
Uniform	15.60	10.07	7.09	5.33	2.40
Response-focused	15.60	9.47	7.34	3.91	3.11
PULSE without covariance	15.60	9.25	6.18	3.55	2.25
PULSE	15.60	8.05	4.15	2.35	1.72

At 5,120 blocks, relative response errors $\|\widehat{R} - R\|_F/\|R\|_F$ are 0.151, 0.149, 0.358, and 0.368, respectively. Unsampled or lightly sampled low-value columns can dominate this global metric without changing the preferred target. The paired PULSE-minus-uniform regret difference at 2,560 blocks is -2.980 percentage points $[-5.05, -1.15]$; at 5,120 it is -0.680 $[-1.252, -0.118]$. The covariance term reduces regret at every tested budget; its contribution is strongest at intermediate budgets, and at 5,120 blocks the full rule attains the lowest mean regret (1.72%).

The displayed trajectory is fixed to seed 13800, the first main network with eight hidden populations. We record cumulative allocation and all target estimates after every batch. Figure 9b shows the two highest population-value targets, fixed across the trajectory independently of the observed estimates, with standard-error bands for conditional pulse uncertainty.

G Feedback strength and response-window sensitivity

Eight independently generated eight-observed/eight-hidden networks, seeds 33000–33007, are rescaled to spectral radii 0.45, 0.60, 0.72, 0.82, and 0.90. Connectivity proportions, immigrant rates, and heterogeneous Gamma timing are retained within each replicate. For every spectral radius, horizons are 0.1, 0.25, 0.5, 1, 2, and 4 seconds, with 320 independent pulse/sham blocks per target and the known baseline rate. Thus the grid contains 240 network-condition replicates, with eight independent networks per cell; conditions within a replicate are paired.

Figure 10 distinguishes truncation of recurrent cascades from background-count variance. At $\rho = 0.45$, mean value NMAE is 5.8% at 0.1 s, 2.0% at 0.5 s, and 4.5% at 4 s. At $\rho = 0.90$, these become 28.2%, 16.4%, and 10.3%, with the lowest tested error of 7.8% at 2 s. Selection regret need not track absolute value error: common bias can preserve the winning target, while small errors near a tie can change it. The grid identifies the best tested horizon for each feedback strength.

H Spatial coverage and probe-calibration diagnostics

Figure 6 in the main paper compares incomplete spatial coverage, direct hidden recruitment, and a near-critical response-window diagnostic; this appendix provides the complete diagnostic protocols.

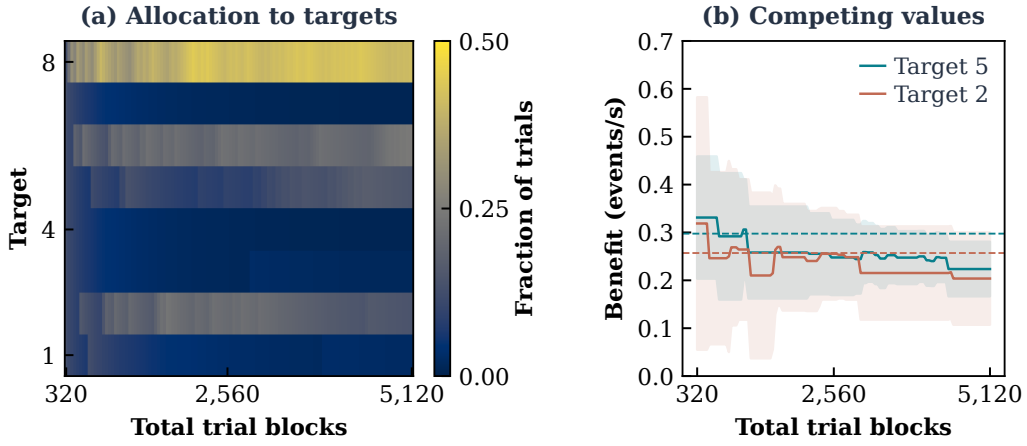
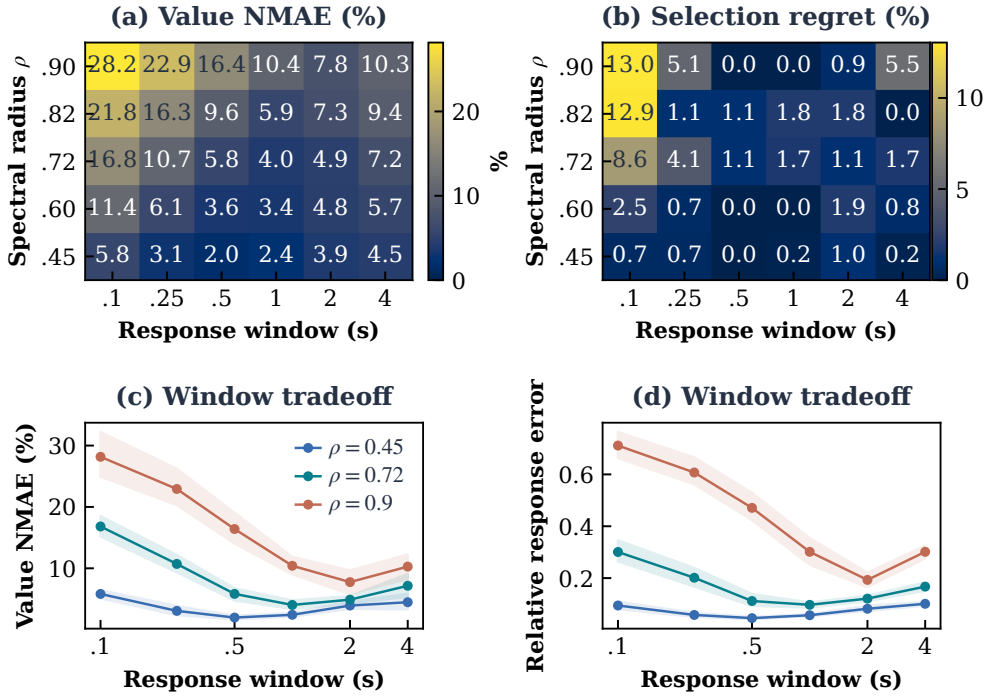


Figure 9: **A fixed-network acquisition trajectory.** (a) Cumulative allocation in the first eight-hidden-population benchmark network, seed 13800. (b) Estimates for its two highest population-value targets; shading shows ±2 plug-in standard errors from pulse uncertainty and dashed lines show true values.



8 independent replicates per condition; 320 blocks per target; baseline rates known.

Figure 10: **Window length must match recurrent propagation.** (a–b) Mean value error and regret across eight independent replicates per cell. (c–d) Value and whole-response error for three representative spectral radii; shaded bands resample the eight replicates. Short-window truncation and long-window background variance yield different error profiles across feedback strengths.

The coverage comparison uses the 96 main networks, with 24 networks at each hidden-population count; the remaining diagnostics use separate networks.

H.1 Direct recruitment of hidden populations

For each probe, independently add an immigrant to a uniformly selected hidden population with probability $\epsilon \in \{0, 0.1, 0.3, 0.6\}$. Use 800 trial blocks per candidate and exact baseline rates to isolate pulse contamination. Let $Q = (I - K)^{-1}$ be the complete integrated resolvent. The expected observed response in this modified protocol is

$$R_{\text{contaminated}} = R + \epsilon\, Q_{OH} \frac{\mathbf{1}_q}{q} \mathbf{1}_p^\top. \tag{39}$$

The second term is not the response to an observed-only input. Inserting it into Equation (1) yields the dashed population-bias curve. Mean normalized value errors from estimated responses are $6.4\%, 7.8\%, 15.6\%, 21.8\%$; their population counterparts are $0.0\%, 4.9\%, 14.2\%, 26.7\%$. This diagnostic identifies a causal calibration requirement that additional samples alone cannot fix.

H.2 Response horizon near criticality

For each of 12 new networks with spectral radius 0.90, acquire 1,600 blocks per target at horizons $0.5, 2, 6, 12$ seconds. Exact baseline rates again isolate response estimation. Mean relative response errors are $38.5\%, 11.8\%, 14.0\%, 23.4\%$; mean normalized action-value errors are $15.2\%, 4.3\%, 5.6\%, 7.9\%$. The nonmonotone pattern is consistent with the bias–variance decomposition of the pulse/sham contrast. These networks are generated independently of the six-second main benchmark.

I Computational benchmark with human-derived timing priors

Figure 7b evaluates whether the selection procedure remains useful with a short response window and measured physiological timing diversity. Each of the 74 participants supplies an empirical distribution of deduplicated N1 latencies. For participant number $u \in \{1, \ldots, 74\}$, construct an eight-observed/eight-hidden network with seed $31000 + u$, using the connectivity and baseline distributions in Appendix D. Draw each Gamma kernel's mean independently with replacement from that participant's empirical N1 distribution, retaining shapes in $\{1, 2, 3\}$ and integrated spectral radius 0.76. This uses N1 timing as a phenomenological prior on propagation times. Connectivity, event streams, and suppressive intervention outcomes are generated by the model.

The measurement horizon is 250 ms, with 1,280 stored pulse/sham blocks per target and an independent 2,400-second passive recording. Budgets are 1,280 and 2,560 blocks. Both PULSE variants, evoked reach, uncorrected transport, and rate-only selection use the same observed data. Initial allocations and batch sizes are 40, and the same finite-pool rule as the main benchmark applies. The action suppresses one observed intensity by $\eta = 0.8$. Ground-truth regret is computed from the complete post-action stationary model. The 250-ms horizon provides a conservative test: it is shorter than the recurrent propagation time of most networks, so this benchmark measures decision quality under truncation as well as under empirical timing. All 74 specified networks are included.

At 1,280 blocks, adaptive PULSE has 2.18% mean regret $[1.25, 3.20]$, compared with 3.35% $[2.10, 4.78]$ for uniform PULSE. At 2,560 blocks the respective means are 1.52% $[0.82, 2.28]$ and 2.45% $[1.45, 3.52]$; exact-target accuracies are 83.8% and 76.4%. Uncorrected transport reaches 2.52% $[1.31, 3.87]$, evoked reach 16.68% $[12.30, 21.44]$, and rate-only selection 13.63% $[9.97, 17.57]$. These values assess model decisions under empirical timing constraints.

J Public human responses and anatomical sampling

J.1 Dataset provenance and analysis units

The source is *CCEP ECoG dataset across age 4–51*, OpenNeuro ds004080, version 1.2.4, DOI `10.18112/openneuro.ds004080.v1.2.4`. The associated study is Van Blooijs et al.

(2023b). The recordings were collected by the UMC Utrecht RESPect investigators, not by the present secondary-analysis workflow. We use the release's `ccepData_V2.mat` (15,255,812 bytes) and `ccepAverages.mat` (155,745,595 bytes). Their SHA-256 digests are matched to the release's Git-annex identifiers; the source package stores both digests and the download URLs. The pinned dataset commit is `c4fd7418883e33b024292468eb14da1649f51aae`.

The N1 derivative contains 40,168 finite detections across 116 runs and 74 participants. Sample indices are one-based and are converted to seconds using each run's own time vector. We retain positive latencies in channels marked good and exclude the directly stimulated contacts; no finite entry is removed by the latter checks in this release. Stimulation pairs are canonicalized without regard to electrode ordering. Repeated measurements of the same participant, stimulated pair, and receiving contact are collapsed by their median latency, giving 37,371 distinct connections. Cohort summaries give equal weight to participants. This avoids allowing a densely sampled participant to dominate the latency distribution.

The waveform derivative averages connections within anatomical pathway and direction. The original processing code appends participants in the order of `ccepData`, within age-indexed cells, including all-NaN entries for absent pathway coverage. We recover the participant mapping from this documented order and assert every cell's dimensions against the participant list. Of $74 \times 16 = 1,184$ possible participant/pathway/direction combinations, 504 are absent and 680 are retained. These span 73 participants; all retained waveforms enter evaluation. Inclusion rules are independent of model reconstruction error. The released averages, rather than newly processed raw recordings, are the data analyzed in the waveform experiment.

J.2 Upstream processing and measurement scope

The original analysis groups single-pulse trials by stimulation pair, requires at least five trials for a condition, subtracts the median prestimulus baseline, and detects N1 peaks in good ECoG channels. Its amplitude threshold is $3.4 \max(\sigma_{\mathrm{pre}}, 50\,\mu\mathrm{V})$, with a search range ending at 100 ms. Its pathway averages exclude stimulation pairs separated by more than 18 mm and receiving contacts within 13 mm of either stimulated contact, and align responses on a 2,048-Hz time grid. We inherit those published derivatives and their selection of detected responses. The waveform endpoint therefore characterizes the released evoked responses, with the original preprocessing documented in the cited dataset's companion code.

Our additional preprocessing subtracts the median voltage over $[-200, -20]$ ms, then linearly interpolates to a 2-ms grid from 15 to 249 ms. Starting at 15 ms avoids the immediate stimulation artifact. We retain signed voltages because this experiment characterizes the released waveforms, whereas the transport theorem operates on nonnegative event intensities; the two analyses address the response representation and the decision identities respectively. A single response template and a shared elementary Hawkes kernel impose different restrictions: different propagation paths can yield different responses even from a common kernel.

J.3 Participant-held-out waveform completion

Let $y_{s\ell} \in \mathbb{R}^{118}$ be the released, baseline-corrected voltage response for participant s and pathway/direction ℓ. Define $u_{s\ell} = y_{s\ell}/\|y_{s\ell}\|_2$ for basis learning. For a held-out participant s, form the uncentered training covariance

$$C_{-s} = \frac{1}{S-1} \sum_{a \neq s} \frac{1}{L_a} \sum_{\ell=1}^{L_a} u_{a\ell} u_{a\ell}^{\top}, \qquad S = 73. \tag{40}$$

The k leading eigenvectors, $k \in \{1, 4\}$, form $B_{-s,k}$. Each training participant receives equal weight regardless of pathway coverage. Neither the held-out participant nor any of that participant's waveforms enters the covariance.

Five fixed temporal folds assign each contiguous 20-ms block to a fold, cycling through fold labels. For visible samples O_f and withheld samples M_f, solve

$$\widehat{a}_{s\ell f} = \arg\min_{a} \|y_{s\ell,O_f} - B_{-s,k,O_f} a\|_2^2 + 0.01\|a\|_2^2, \qquad \widehat{y}_{s\ell,M_f} = B_{-s,k,M_f} \widehat{a}_{s\ell f}. \tag{41}$$

The ridge coefficient is identical across ranks and folds. Every time sample is predicted once while withheld. The participant error is

$$E_{s,k} = \frac{1}{L_s} \sum_{\ell=1}^{L_s} \frac{\|y_{s\ell} - \widehat{y}_{s\ell}\|_2^2}{\|y_{s\ell}\|_2^2}. \tag{42}$$

This is temporal completion within a new participant, using visible portions of that participant's responses. Participant-bootstrap intervals in the expanded comparison use 5,000 resamples with seed 7357; the paired comparison resamples $E_{s,4} - E_{s,1}$. Figure 11a–c displays every retained waveform, with a fixed ordering independent of prediction error. Appendix J.5 adds interpolation and longer missing intervals.

Across all 680 waveforms, four bases give participant-averaged normalized squared error 0.071 [0.058,0.085], versus 0.345 [0.298,0.394] for one, with paired difference -0.274 $[-0.318, -0.231]$. Linear interpolation gives 0.089 [0.073,0.107]. Multiple bases capture temporal diversity that interpolation alone misses when the full waveform must be reconstructed, while local smoothness remains effective for very short interior gaps.

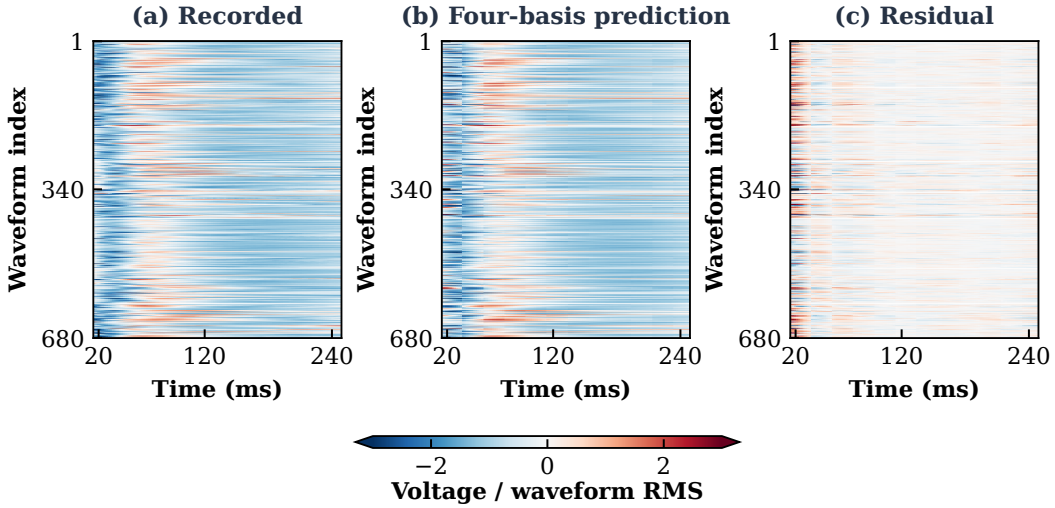


Figure 11: **Held-out human response atlas.** All 680 pathway-average waveforms from 73 participants, sorted by pathway, direction, and participant: (a) recorded responses, (b) out-of-fold four-basis predictions, and (c) prediction-minus-recording residuals. Each waveform is normalized by its recorded RMS; panels share ± 3 RMS limits, clipping below 0.35% of entries. White separators mark pathways.

J.4 Anatomical rendering and spatial sampling

We use the 74 released `electrodes.tsv` files at the pinned dataset revision. The coordinate-system JSON specifies fsaverage coordinates in millimetres, surface-projected and transformed using FreeSurfer. These are associated with the MNI305 reference, not MNI152. Rendering uses the matching fsaverage5 pial surfaces distributed by Nilearn and actual PyVista mesh rendering, so contacts stay on their released coordinate template. This provides a shared anatomical frame across participants.

Nonfinite electrode positions are excluded from anatomical displays. Contact names are canonicalized by case and integer numeric suffix; among finite contacts this mapping is one-to-one within each participant. Duplicate placeholder names with unavailable positions do not enter the join. All 37,371 N1 receiver connections match a finite coordinate. There are 3,652 left-hemisphere and 1,880 right-hemisphere finite contact entries. For each template vertex, coverage is the percentage of all

74 participants with at least one same-hemisphere contact within a 10-mm Euclidean radius. It is a sampling-coverage quantity.

The single-participant example is the first participant, UMCU01; its displayed input is selected by the largest number of released receiver connections, independently of response magnitude, latency, or model error. OC6–OC7 supplies 38 receivers. Its stimulation marker is the midpoint of those two recorded contact coordinates. Receiver colors encode the published N1 latencies, ranging from 10.25 to 25.88 ms in this example. Straight segments represent input-to-recording associations. The pathway-average voltage derivative cannot be mapped to individual contacts, so the anatomical and waveform panels retain their respective measurement units.

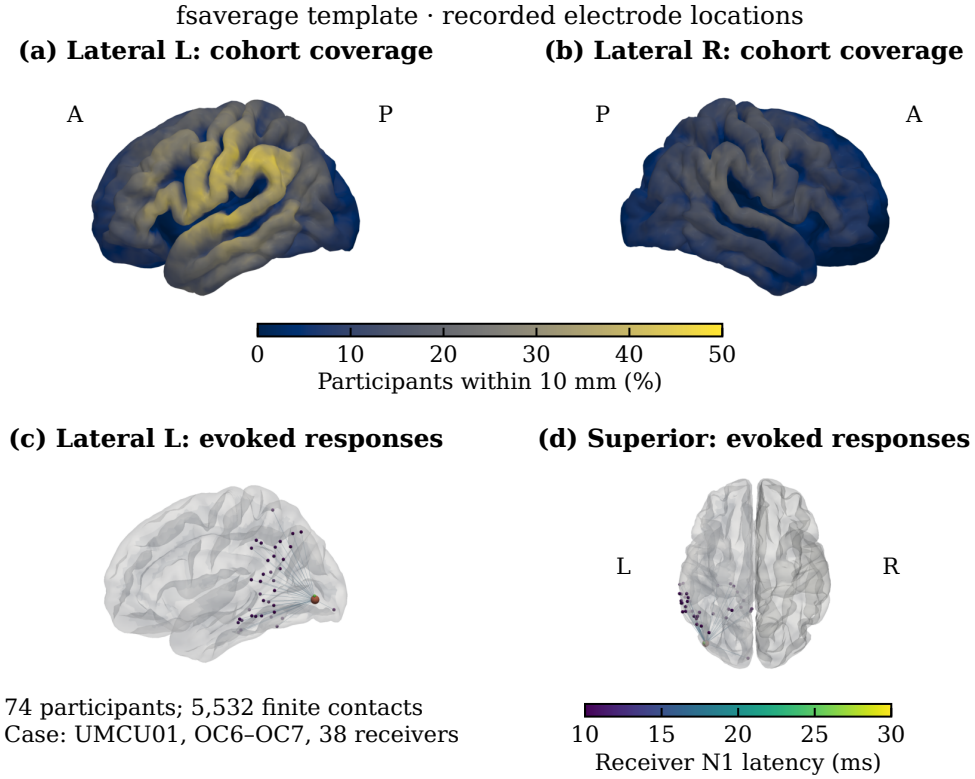


Figure 12: **Anatomical coverage and evoked-response geometry.** Lateral views show the left and right template hemispheres. Superior orientation displays anatomical left on the image left. A/P denote anterior/posterior. All nodes retain their released fsaverage coordinates; no EEG inverse solution or cortical voltage interpolation is applied.

J.5 Clinical completion over longer missing intervals

We extend the participant-held-out protocol with linear interpolation and a duration-by-position sweep. Interpolation uses only visible samples, with the nearest visible endpoint held constant outside their range. It has no training parameters. The one- and four-basis models keep the same participant-excluded covariance and ridge coefficient 0.01. For each waveform, one contiguous interval is withheld: durations are 10, 20, 40, and 80 ms; start times range from 15 to 215 ms in steps of 20 ms, retaining intervals ending no later than 250 ms. These rules produce 40 valid duration/start cases. Gray heatmap cells correspond to unavailable windows, not zero error.

For a missing interval M, the normalized gap error is

$$E_M = \frac{|M|^{-1} \sum_{t \in M} (\widehat{y}_t - y_t)^2}{118^{-1} \sum_{t=1}^{118} y_t^2}. \tag{43}$$

This uses whole-waveform energy in the denominator, avoiding an unstable ratio when a late gap has little signal. Errors are first averaged over waveforms within participant and then over eligible gap

starts; participant bootstrap uses seed 7357. In the full five-fold completion experiment, each sample is predicted once and the original full-waveform normalized error applies.

Table 4: Mean normalized gap error over 73 participants.

Model	10 ms	20 ms	40 ms	80 ms
One basis	0.396	0.391	0.391	0.405
Four bases	0.106	0.153	0.283	0.325
Linear interpolation	0.041	0.083	0.217	0.332

The early response is the most difficult region to withhold because its rapid deflection cannot be reconstructed from local smoothness alone. Four bases improve substantially over one throughout the duration sweep; interpolation remains a strong baseline, particularly for short interior gaps. These observations establish response flexibility and measurement sensitivity.

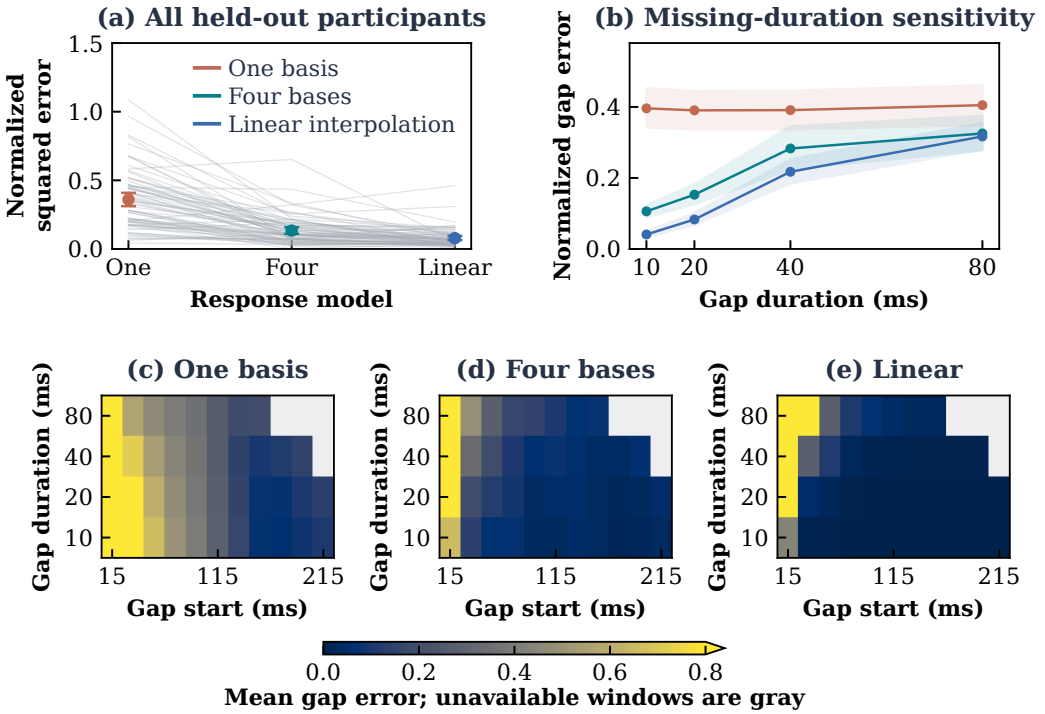


Figure 13: **Participant and temporal generalization.** (a) Each gray line pairs the three models for one participant; colored markers are means with 95% participant-bootstrap intervals. (b) Mean gap error versus missing duration, with participant-bootstrap intervals. (c–e) Duration-by-start-time error maps, using the same scale; values above 0.8 saturate the color scale. Late windows extending beyond the evaluated waveform are absent.

K JOINT SUPPRESSION AND OVERLAPPING BENEFITS

For each of the 96 main networks, we evaluate every unordered pair of distinct observed targets: $\binom{8}{2} = 28$ pairs and 2,688 conditions in total. Each target has suppression fraction 0.4. Equation (9) uses the population response and baseline rates, and its predictions are checked against an independent complete-network solve with both observed rows and immigrant intensities scaled. Maximum absolute rate discrepancy is 8.89×10^{-16} events/s.

Define the overlapping benefit of targets a, b as $V_a(0.4) + V_b(0.4) - V_{a,b}(0.4, 0.4)$. It is nonnegative up to floating-point precision in every tested pair. This is consistent with suppressing propagation paths that partly overlap in a nonnegative event network. Selecting the two highest isolated values

disagrees with the best joint pair in 22 networks; its relative pair-value regret is 0.77% [0.39,1.22]. The experiment evaluates the multiple-action formula directly at the population operator.

The fixed example is seed 13800. We show all pair values and overlaps at fraction 0.4, and a 41×41 gain grid for that example's best pair at those fractions. Both axes range from zero to complete event-generation suppression in increments of 0.025. Gain is a dimensionless model parameter.

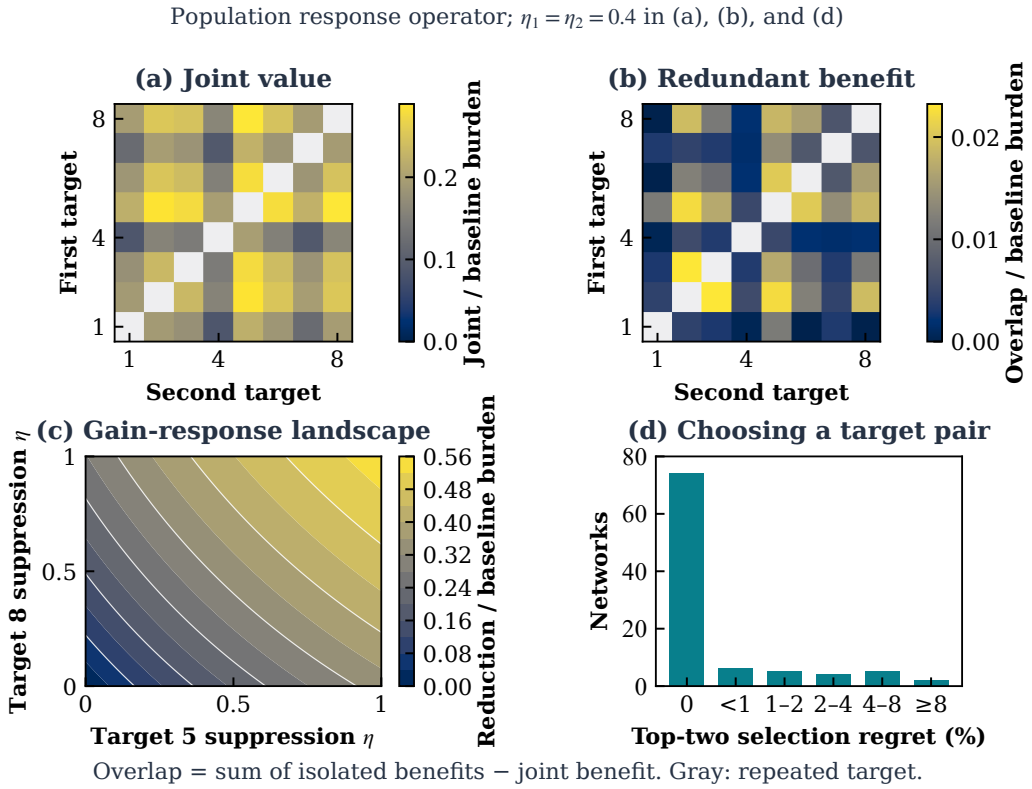


Figure 14: **Joint values need not follow isolated rankings.** (a–b) Pair-value and overlapping-benefit matrices for the fixed example; gray diagonal entries denote excluded repeated targets. (c) Gain-response landscape of its best pair at $\eta_1 = \eta_2 = 0.4$. Panels (a–c) normalize reductions by the baseline burden. (d) Regret from choosing the two best isolated targets, across all 96 networks; the zero-regret bin includes 74 networks. The exact formula supports joint evaluation when recurrent pathways overlap.