TOPONESY: EXACT LIFETIME-SPACE VERIFICATION FOR GRAPH-TOPOLOGICAL QUERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Feature-space robustness verification can be conservative when candidate feature perturbations have no admissible input realization at the claimed distance. TopoNeSy studies this question for a fixed graph with independently perturbed edge-birth times, a one-dimensional preterminal filtration, and completion at time one.

Our main theoretical result is an exact image-and-lift theorem: every clipped ℓ_∞ input ball maps exactly to the ordered ℓ_∞ ball of the same radius around the global H_1 lifetime vector, and every target has a constructive input preimage at its distance from the nominal lifetime vector. This strengthens forward stability from an enclosure to exact realizability. Consequently, any frozen lifetime-only query has the same exact flip radius in input and lifetime space.

The reduction also permits exact elimination of unused lifetime coordinates. TopoNeSy instantiates the resulting verification interface with jointly trained affine predicates and Boolean DAGs under a no-bypass architecture, deployed with exact rational semantics. Controlled experiments validate reconstruction and computational reduction. On measured EEG connectivity graphs, a frozen full-H_1 DAG attains 76.11% balanced accuracy and 67.78% open-ball certified balanced accuracy at edge radius .01, averaged over three seeds. All 720 model/window cases have verified radius intervals of width at most 10^{-5} after refinement, with actual graph counterexamples at their upper endpoints. The guarantees concern the stated stored-edge domain, not raw EEG perturbations.

Code: https://anonymous.4open.science/r/toponesy-75FF/

1 INTRODUCTION

Prediction and explanation do not by themselves certify an answer. Verification in a feature space faces an additional question: does a candidate feature perturbation correspond to an admissible input at the claimed distance?

Our central question is *input realizability*: which topological targets can a declared input perturbation attain, at what minimum cost, and with what retained structural information? Marginal stability bounds alone provide an enclosure, not an input counterexample. We prove an exact image identity for global lifetimes under explicitly restricted filtration assumptions, together with minimum-distance input reconstruction. In this class, ordering is the only additional constraint needed: more elaborate graph verification cannot improve the exact flip-distance infimum. The consequence is an exact reduction of the verification problem.

The exact reduction removes the graph–feature realizability problem, but the remaining query must still admit practical verification. TopoNeSy jointly learns affine predicates and a hard Boolean DAG from output labels, with no direct evidence-to-label branch. The exported predictor admits a shared Boolean–linear encoding, so prediction and certification refer to the same deployed function. This symbolic architecture implements the verification interface; it is not a hypothesis of the image theorem, which applies to other frozen lifetime-only functions as well.

Our controlled graph-topological tasks, collectively called TopoLogic, supply known queries and explicit perturbation domains to isolate the reduction. An EEG connectivity study then tests prediction and certification of jointly learned queries on recording-condition labels not constructed from

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.

persistent homology (PH). The appendix gives the complete EEG protocol, predictive controls, certificate intervals and reconstructed graph counterexamples.

- an exact ordered-lifetime image and minimum-distance identity for a declared global H_1 filtration, with constructive input witnesses and an exact reduction of frozen query verification;
- a computable verification interface: exact elimination of unread coordinates and shared Boolean–linear verification of frozen rational predicates and Boolean DAGs;
- controlled reduction/scaling tests and a measured-graph application with nominal accuracy, tight graph-space intervals, reconstructed counterexamples and explicit computational costs.

2 EXACT IMAGE, MINIMUM-DISTANCE LIFT, AND VERIFICATION

Fix a graph structure $\bar{G}$ and write G_x for that structure equipped with edge-birth vector $x \in \mathcal{X}$. Certificates concern perturbations of these numerical coordinates, not arbitrary graph-edit distance. In the binary experiments the computational maps are

$$x \xrightarrow{L_{\bar{G}}} z \xrightarrow{P_\theta} b \xrightarrow{Q_\theta} \hat{y}, \qquad q_\theta = Q_\theta \circ P_\theta.$$

$L_{\bar{G}} : \mathcal{X} \to \mathcal{C}_\beta$ is the ordered lifetime map, $P_\theta : \mathcal{C}_\beta \to \{0,1\}^M$ the predicate map, and $Q_\theta : \{0,1\}^M \to \{0,1\}$ the Boolean rule function. We abbreviate $L_{\bar{G}}$ as L with the graph fixed. When raw measurements are first transformed into x, a certificate for x is not automatically a certificate in the raw measurement metric.

2.1 EXACT INPUT-TO-LIFETIME REDUCTION

Here we specialize the interface to a query on global H_1 lifetimes, without named persistent objects. Homology uses the field $\mathbb{F}_2$ with two elements. Let $\bar{G} = (V, E)$ be finite and simple, with $m_E = |E|$. Vertices are born at zero, observed edges at times ranging over the Cartesian domain $x \in [0,1]^{m_E}$, and the complex is completed to the full simplex at time one. The preterminal complex contains vertices and observed edges only, even if $\bar{G}$ contains a triangle; all positive H_1 intervals end at one. Write $\beta = m_E - |V| + c(\bar{G})$ for the cycle rank, where $c(\bar{G})$ is the number of connected components, counting isolated vertices. Define $L(x) \in [0,1]^\beta$ by sorting the *positive* H_1 lifetimes in nonincreasing order and appending zeros to length β. These zero coordinates are not off-diagonal persistence intervals. A nominal zero among the first β coordinates may become positive; extra coordinates beyond β, if used by an implementation, remain zero. The clique-expansion control enforces this filtration model using triangle-free graphs; the EEG study constructs a one-dimensional complex explicitly on K_{64}. Triangle-freeness is not needed by the proof. The ordered cycle-birth vector $B_i(x) = 1 - L_i(x)$ is an ℓ_∞ isometry of these lifetime coordinates; this change of variables is absorbed exactly into the already rounded affine export, without rounding again. EEG exports are evaluated in birth coordinates.

Let $\mathcal{C}_\beta = \{z \in [0,1]^\beta : z_1 \geq \cdots \geq z_\beta\}$ and let $\mathcal{B}(x,\epsilon) = [0,1]^{m_E} \cap \{y : \|y - x\|_\infty \leq \epsilon\}$. This closed box is used by the verifier; reported certificate balls are open. For $\beta = 0$, $\mathcal{C}_0$ contains the empty vector, whose norm is zero. A target z is *input-realizable* in a set A when some $y \in A$ satisfies $L(y) = z$.

Definition 1 (Formal certificate)**.** *For a frozen predictor $h : \mathcal{X} \to \mathcal{Y}$, a formal certificate at $x \in \mathcal{X}$ consists of a radius $R \in (0, +\infty]$ and a justification of $h(y) = h(x)$ for every $y \in \mathcal{X}$ with $\delta(x,y) < R$, with the domain, metric δ and deployed function specified. Here δ is the input ℓ_∞ metric. A symbolic trace records the predicate/rule computation but does not by itself justify invariance. In result tables, $R = 0$ is the sentinel for no positive certificate, not a formal certificate.*

Exclusion of all opposite answers in the exact image provides such a justification. Our implementations trust the Boolean/SMT encoding and solver; rational checking of LP conflicts and input witnesses does not independently verify the solver's UNSAT proof log.

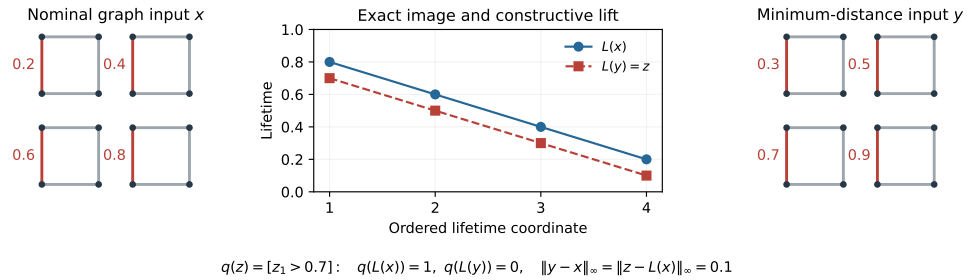

Figure 1: The central reduction, shown on four unfilled square components. Red edges carry the displayed birth times; gray edges are born at zero, vertices at zero, and completion occurs at one. A target lifetime vector has a nearest input preimage at the same ℓ_∞ distance. The displayed query flips at the boundary .1 of its invariant open ball. Values are exact rationals and the input is independently replayed. This analytic example illustrates reconstruction, not a learned result.

Theorem 1 (Exact ordered image and input distance)**.** *Under this filtration model, for every* $\epsilon \geq 0$*,*

$$L(\mathcal{B}(x,\epsilon)) = \{z \in \mathcal{C}_\beta : \|z - L(x)\|_\infty \leq \epsilon\}.$$

Moreover, every $z \in \mathcal{C}_\beta$ *has a constructive preimage satisfying*

$$\min_{y \in [0,1]^{m_E} : L(y) = z} \|y - x\|_\infty = \|z - L(x)\|_\infty.$$

Construction and proof sketch. Sort edges by nominal birth, refining ties arbitrarily, and mark cycle-closing positions $j_1 < \cdots < j_\beta$ in a greedy spanning-forest pass. Their births are $x_{j_i} = 1 - L_i(x)$. For a feasible target let $b_i = 1 - z_i$ and $l_j = \max(0, x_j - \epsilon)$, $u_j = \min(1, x_j + \epsilon)$ in this order. Set

$$y_j = \max\big(\{l_j\} \cup \{b_i : j_i \leq j\}\big).$$

Monotonicity of both interval endpoints gives $y_j \in [l_j, u_j]$ and $y_{j_i} = b_i$. The old edge order remains valid, so its forest decisions and cycle-closing positions remain valid, even with new ties. Hence $L(y) = z$. The forward inclusion follows from the standard monotonic cycle-rank bound, which makes each ordered lifetime 1-Lipschitz. Applying the construction at $\epsilon = \|z - L(x)\|_\infty$ proves the distance equality. The full proof, including terminal zeros and disconnected graphs, is in Appendix A.2. □

Figure 1 illustrates the distance-attaining construction. For any frozen Boolean query $q : \mathcal{C}_\beta \to \{0,1\}$, define $r_q^\star(x) = \inf\{\|y - x\|_\infty : y \in [0,1]^{m_E},\ q(L(y)) \neq q(L(x))\}$. The theorem gives the exact reduction

$$r_q^\star(x) = \inf_{z \in \mathcal{C}_\beta :\, q(z) \neq q(L(x))} \|z - L(x)\|_\infty, \qquad \inf \varnothing = +\infty.$$

For a disjunctive normal form (DNF) query, no single clause need hold throughout this ball; fixed object identities and fact assignments are not required either. An ordered-feature SAT assignment lifts to an input counterexample; UNSAT on the closed box certifies the domain-relative ball $\{y \in [0,1]^{m_E} : \|y - x\|_\infty < \epsilon\}$. Timeout yields only the lower bound already proved. Exactness of the reduction does not imply that a finite-budget solver reaches $r_q^\star$.

An immediate computational consequence is reuse across graph realizations: identical lifetime centers, cycle ranks, frozen queries and perturbation domains and metrics give the same reduced problem, directly by the displayed infimum. This shares verification, not labels or predictive generalization.

If the frozen query reads only indices $I = (i_1 < \cdots < i_s)$, its exact domain is smaller still: retain only their ordered coordinate intervals. Corollary 2 gives an $O(\beta)$ completion of any projected target at its minimum distance, followed by the input lift. The full cycle rank β and the query-read dimension s therefore play different computational roles. No reduction is claimed when all coordinates are read or when input side constraints are omitted.

The metric contribution. Cycle-rank monotonicity supplies the coefficient-one bound for ordered lifetimes used in the forward enclosure; this is stronger than merely invoking a generic lifetime bound from diagram matching. The reverse inclusion and minimum-distance lift establish joint attainability and remove the graph variables from query search. They use standard greedy graphic-matroid structure (Henselman & Ghrist, 2016); no general verification speedup is asserted. After sorting and the forest pass, the lift costs $O(m_E + \beta)$; our measured implementation repeats that preprocessing per witness. The theorem permits any lifetime-only query. Rules give a compact query representation for the verification problem. Our implementations support threshold formulas and frozen affine-predicate Boolean circuits with threshold readout; DNF and cardinality rules are special cases used in controlled studies. The distance identity is anchored at x. In fact, the canonical choice of preimages preserves pairwise distances (Appendix F.1), although L itself is not injective. Forward stability alone does not give such an identity: for the 1-Lipschitz map $f(t) = \min(t, 1/2)$ on $[0, 1]$, the target $1/4$ lies at output distance $1/4$ from $f(1)$ but has input distance $3/4$ from 1. The extra conclusion here is a distance-attaining lift under the stated filtration and box assumptions, not a general inverse of persistence. Appendix F.5 states the finite-matroid nullity-profile construction underlying this order argument. The filtration assumption identifies that profile with the graph's H_1 lifetimes.

Proposition 10 extends the construction to boxes whose lower and upper endpoint sequences share one edge order. Uniform clipping supplies this condition. It cannot be replaced by arbitrary unequal edge boxes: in three internally disjoint paths, fix one path's maximum birth at $1/2$ and let the other two maxima vary in $[0, 1]$. The two global cycle births are the largest two of these three values. The ordered lifetime target $(3/4, 1/4)$ lies within the endpoint bounds $[0, 1] \times [0, 1/2]$, but would require births $(1/4, 3/4)$ and is impossible. Named local-cycle features are different statistics; the theorem concerns only the declared global lifetime map.

Scope and extensions. Input constraints generally change the exact image. Appendix F.6 retains rational polytopes using order cells and spanning-tree regions. A conserved-flow extension uses a weighted additive cost (Appendix F.8). Full H_0 has an analogous exact lift; joint H_0/H_1 requires additional order consistency and need not preserve minimum distances globally (Appendices F.1 and F.2). Each extension states its own domain and metric.

2.2 A LEARNED SYMBOLIC QUERY

Definition 2 (Strict no-bypass)**.** *Let $E_\theta : \mathcal{X} \to \mathcal{Z}$ be an evidence map, $P_\theta : \mathcal{Z} \to \{0,1\}^M$ a hard predicate map, $Q_\theta : \{0,1\}^M \to \{0,1\}^J$ a frozen Boolean rule map, and $\mathrm{Dec} : \{0,1\}^J \to \mathcal{Y}$ a fixed deterministic decoder. The deployed hard predictor is* strict no-bypass *if*

$$h_\theta(x) = \mathrm{Dec}(Q_\theta(P_\theta(E_\theta(x)))).$$

Here M counts predicates, J counts Boolean queries and $K = |\mathcal{Y}|$ counts classes. There is no other input to Dec*. Our binary experiments use $J = 1$, $E_\theta = L$ and the identity decoder. This definition specifies the evaluated computational interface.*

No-bypass is a structural restriction, not semantic interpretability: a predicate could encode an arbitrary classifier and feed a one-rule decoder. Interpretable evidence additionally requires a declared predicate vocabulary, bounded rule complexity and empirical semantic validation.

Deployed Boolean DAG. For $z = L(x)$, input bits are $p_j(z) = \mathbf{1}[a_j^\top z + b_j > 0]$. Each gate selects two predecessor bits and one of the 16 binary Boolean operations. Two acyclic layers may read the original predicates and earlier layers, sharing rather than duplicating their Boolean values. The entire DAG and its learned vote belong to Q_θ:

$$g_v = \mathrm{op}_v(u_v, w_v), \qquad Q_\theta(p) = \mathbf{1}[\alpha^\top(p, g) + \alpha_0 \geq 0].$$

The vote receives only hard predicate/gate bits, never continuous features or pre-threshold scores; direct predicate-to-vote connections remain allowed. This is an adaptation of differentiable logic gates (Petersen et al., 2022), not a new general circuit-learning principle.

Joint training and exact deployment. We fit standardization on training subjects, then jointly optimize predicates, wiring, gate operations and vote using output labels. Hard forward selections

use surrogate gradients; binary cross-entropy trains the vote score. Configuration and epoch selection evaluate the hard rational export on validation subjects. Standardization is absorbed into affine coefficients before rounding to six decimals; the resulting rational model is frozen for both evaluation and verification. Predicate equality is false, while vote equality is class one. Appendix B specifies the actual loss, surrogate derivatives, grids and decimal-input interpretation. Training the head jointly does not mean training the graph/PH frontend.

2.3 Verification of the frozen Boolean query

For $z_0 = L(x)$, a radius query asks whether

$$\exists z \in \mathcal{C}_\beta : \quad \|z - z_0\|_\infty \le r, \quad Q_\theta(P_\theta(z)) \neq Q_\theta(P_\theta(z_0)).$$

One shared variable per predicate/gate, truth-table constraints and the rational threshold vote encode the DAG without DNF expansion. UNSAT raises a safe lower bound; an exactly checked SAT point lifts by Theorem 1 to a graph counterexample at the same distance, yielding an upper bound. Radius bisection stops at the requested interval width or computational allowance, retaining established bounds on unknown/timeout. An exact decision procedure does not imply that a finite number of calls computes an exactly optimized infimum.

Our checked Boolean–linear refinement proposes opposite assignments, tests joint realizability by LP, and excludes assignments using checked rational LP conflicts or exact-SMT UNSAT results. Accepted points satisfy the deployed strict/nonstrict semantics and undergo graph reconstruction and replay. Exact SMT handles exceptional strict-boundary cases; Boolean/SMT UNSAT still trusts the backend. Appendix C.1 gives the encoding and conflict checks. Individual predicate bounds are a sound starting abstraction, not a joint realizability test. Dense predicates may read all 1,953 EEG coordinates; 32 predicate bits do not imply 32 continuous verification variables.

Let $b = P_\theta(L(x))$ and let ρ_j be the exact infimum of input distances flipping predicate j. The strongest independent-fact bound is

$$R_{\mathrm{Bool}} = \min_{v : Q_\theta(v) \neq Q_\theta(b)} \max_{j : v_j \neq b_j} \rho_j, \qquad \min \varnothing = +\infty, \quad \max \varnothing = 0.$$

Proposition 6 shows its invariance under Boolean rewriting on this fixed vocabulary. It need not respect joint input realizability of the predicates.

Proposition 1 (A gap beyond Boolean clause rewriting)**.** *Let $L(x) = (.8, .6, .4, .2)$ and $0 < \eta < .1$. On $\mathcal{C}_4$ set*

$$p_1 = [z_1 > .7], \quad p_2 = [z_2 > .5], \quad p_3 = [z_3 > .4 - \eta], \quad p_4 = [z_3 < .4 + \eta].$$

For $q = (p_1 \wedge p_3) \vee (p_2 \wedge p_4)$, the exact input flip radius is $r_q^\star(x) = .1$, attained, while the best fixed-literal bound R_{Bool} over all equivalent Boolean formulas on this vocabulary is η. Thus $r_q^\star / R_{\mathrm{Bool}}$ is unbounded as $\eta \downarrow 0$. The infima and Boolean vocabulary are the ones defined above.

Proof. The exact predicate flip infima are $(.1, .1, \eta, \eta)$. All four predicates are initially true; the Boolean assignment $(1, 1, 0, 0)$ changes the answer at weighted discrepancy η, and no smaller predicate discrepancy can change any bit. Hence $R_{\mathrm{Bool}} = \eta$. That assignment is not input-realizable, since $p_3 \vee p_4$ holds for every real z_3. Below distance .1, p_1, p_2 remain true, so q remains true. At $(.7, .5, .4, .2)$ both are false; the exact lift realizes distance .1. □

The gap concerns the independent-fact abstraction, not the optimality of a particular solver: ordered SMT also obtains the exact answer radius. The query is nonconstant, both clauses are indispensable, and no single conjunction of these literals represents it on $\mathcal{C}_4$. The example isolates coupled predicates; it is not a claim that this phenomenon occurs only for topological evidence. Indeed, adding the valid pair constraint $p_3 \vee p_4$ closes this example's gap. Appendix C.5 formalizes a dependency-aware Boolean control and separates its power from full joint verification.

3 Related work and positioning

Persistent homology and topological graph learning. Persistence stability bounds how much a persistence diagram can move when the underlying filtration changes (Cohen-Steiner et al., 2007;

Bauer & Lesnick, 2015). Topology layers and persistence-based graph models use topological representations or rationale filtrations (Brüel-Gabrielsson et al., 2020; Carrière et al., 2021; Immonen et al., 2023; Verma et al., 2024; Xin et al., 2025; Verma et al., 2025). Our question concerns invariance of a frozen logical answer, possibly through different proofs, rather than predictive representation quality alone.

Inverse persistence and ordered coordinates. Leygonie & Tillmann (2022), Theorem 2.2, describe persistence fibers as polyhedral complexes and establish a stratified fiber-bundle structure for filters on a fixed complex. Leygonie et al. (2022) develop barcode templates, their invariance under filtration-order equivalence (Proposition 4.5) and local differential coordinates. Their Proposition 4.26 includes local isometry for the full barcode tuple near an injective filter; our anchored distance identity concerns the full declared lifetime cube, including ties and zero coordinates, under a narrower filtration. Leygonie & Henselman-Petrusek (2024) algorithmically enumerate the polytopes forming a persistence fiber, including extensions to based chain complexes and restricted filtration families. Reconstructing an entire fiber, finding one preimage, and finding a preimage at minimum distance from a specified input are different problems. These results motivate, rather than exclude, inverse constructions. Our specialization supplies an anchored minimum-distance preimage over the entire ordered domain, including ties and zero lifetimes. Neither a description of the fiber nor local coordinates alone give this global distance identity: the additional step constructs a closest preimage for every target while preserving an admissible edge order. The inverse-matroid problem instead changes weights to make a specified basis, or bases constrained by a subset, optimal (Bérczi et al., 2026, Theorem 3.3 and Section 4). It does not specify the complete ordered persistence target. Minimum-cost verification of an uncertain matroid basis instead selects information queries establishing basis optimality (Diwan et al., 2026); that certificate is not an input-perturbation radius. The conserved-flow extension uses classical assignment (Kuhn, 1955); its contribution is the target reconstruction and the indexed/global information boundary, not the assignment algorithm. The global-target hardness proof uses classical 3-PARTITION (Garey & Johnson, 1979); the contribution is its application to this filtration, cost and target-information distinction.

Certified graph robustness. Margin-based certificates combine score differences with Lipschitz bounds (Hein & Andriushchenko, 2017; Tsuzuku et al., 2018). Convex relaxations and smoothing provide other routes to output invariance (Wong & Kolter, 2018; Cohen et al., 2019). Graph-specific certificates protect against declared binary-feature or edge-perturbation models (Zügner & Günnemann, 2019; Bojchevski & Günnemann, 2019). Our exact-image reduction is not a replacement for these threat models. Its certificate is an open ℓ_∞ ball in the edge-birth coordinates x on the fixed candidate graph. Arbitrary Hamming-budget edge edits and chemistry-level interventions are outside the formal claim unless separately encoded with verified constants.

Neural-symbolic reasoning and verification. Neural-symbolic systems integrate learned perception with logical structure, for example through differentiable constraints, probabilistic logic programs, or Datalog provenance (Xu et al., 2018; Manhaeve et al., 2018; Huang et al., 2021; Green et al., 2007). Differentiable logic gate networks (Petersen et al., 2022) motivate our jointly trained hard DAG; learned affine predicates and wiring adapt that idea to continuous PH inputs. We do not claim a new general logic-learning architecture. We use standard Boolean semantics: fixing a sufficient set of literals can certify an answer, but need not characterize all input-realizable answers. Verification interfaces already address the neural–symbolic embedding gap (Daggitt et al., 2025); our contribution addresses the exact lifetime-space reduction within such an interface. Formal circuit discovery and NeSy robustness verification also protect structured computations (Hadad et al., 2026; Manginas et al., 2025). Neither structural factorization nor concept accuracy excludes reasoning shortcuts or residual information (Andolfi & Giunchiglia, 2025; Almudévar et al., 2026; Bortolotti et al., 2025). We therefore distinguish certificate soundness from semantic recovery; the latter is tested only in the controlled tasks with explicitly aligned predicate names. Incremental scopes, existential elimination and equivalence checks are standard satisfiability-modulo-theories (SMT) facilities (Bjørner et al., 2019). Our controls include them; verified pruning or gains over an unsimplified encoding alone would not establish a new verification mechanism. Deployment arithmetic can invalidate an otherwise sound mathematical certificate (Szász et al., 2025); our rational and rounded-input replays have explicitly different roles, rather than implying arbitrary floating-point soundness.

Specification-aware reduction and other topological guarantees. Specification-guided clipping and latent-space refinement already reduce verification domains (Zhou et al., 2025; Koller et al., 2026). Our selected-coordinate corollary instead gives exact elimination and an input witness under the ordered-box assumptions; it does not transfer their neural-network or accelerator guarantees. Minimum homological cuts relate individual persistent-feature robustness to arbitrary outliers (Roos Hoefgeest & Slot, 2026). That perturbation model and protected object differ from invariance of a frozen answer under continuous edge birth changes. Neither result is an interchangeable empirical baseline.

The exact-image theorem applies to any frozen lifetime-only query. The implemented verifier supports the specified Boolean–linear class. Our controls isolate representation and verification; they do not establish superiority over general graph classifiers.

4 EXPERIMENTS

Supplied queries isolate the exact reduction. We then evaluate one jointly learned full-H_1 Boolean DAG family on measured EEG graphs and certify the same three frozen models. Appendix E gives the controlled protocols; Appendix D gives the EEG protocol, predictive controls and certificate accounting.

For $\epsilon > 0$, open-ball balanced certified accuracy is

$$\mathrm{BCA}_{\mathrm{open}}(\epsilon) = \frac{1}{2}\sum_{c=0}^{1}\frac{1}{n_c}\sum_{i:y_i=c}\mathbf{1}[\hat{y}_i = y_i,\ \underline{r}_i \geq \epsilon], \qquad \mathrm{BCA}_{\mathrm{open}}(0) = \mathrm{BA}.$$

Here $\underline{r}_i$ is a proved lower radius and n_c counts class c. Incorrect predictions contribute zero even when stable. Each participant has six windows per class, so this equals subject-macro BCA; uncertainty resamples participants, not individual windows or model/window pairs. A fully certified binary constant has BCA .5 and recalls $(1, 0)$ or $(0, 1)$.

4.1 EXACT REDUCTION, COMPETING CLAUSES AND SCALING

We compare exact ordered-domain procedures with raw-input verification. Independent-coordinate and independent-fact relaxations are distinct controls: the former drops lifetime ordering from a specified formula, whereas the latter drops dependencies between its predicate truths. Neither relaxation's satisfying assignment is automatically an admissible input counterexample.

The information lost by independent facts. Figure 3 illustrates Proposition 1 at $\eta = .005$. The input answer radius is .1, whereas the strongest Boolean fixed-literal bound is .005: its opposite assignment violates the shared-coordinate constraint on p_3, p_4. Rational optimization and an independently replayed input witness check this analytic family at three values of η. One valid pair constraint closes the gap; the example motivates retaining dependencies, not a unique solver design.

Frozen nontrivial-query and scaling tests. The fixed manifest contains 690 cases and 4,140 method/budget arms. Competing queries have 40 generated and 40 constructed program/input identities, each on three graph families and two sizes. Exact checks require nonconstant queries, indispensable clauses and inequivalence to a single conjunction on the ordered domain. Constructed queries test timing, not separation from Boolean rewriting. Logic scaling varies 4–16 facts and 2/4/8 clauses; graph scaling fixes eight facts and four clauses. All methods receive .1/1 seconds, cap .125 and tolerance 10^{-4}. Appendix E.3 specifies the fixed workloads, timeouts across the tested regimes and paired-seed uncertainty. These are designed workloads; enumeration remains exponential in the number of facts.

All 4,140 arms return records without an execution failure; 1,561 searches stop at their time budget, and all 11,726 accepted SAT witnesses pass replay. In Table 1, enumeration and ordered SMT complete both competing strata, with lower bounds differing only within the reporting tolerance. Enumeration is about three times faster there, but this does not persist as logic grows: with sixteen facts and eight clauses, mean times are 317 versus 122 ms. At .1 seconds that enumerator times out on all ten cases, whereas ordered SMT times out on four. The result supports reducing graph

Table 1: Computational reduction, not separation from dependency-aware Boolean verification: 40 program/input identities and 240 graph realizations per row, one-second budget. Lower is the mean certified lower bound; time includes construction and replay. Full .1-second results and scaling are in Appendix E.3.

Cohort	Verifier	Lower	Time (ms)	Timeouts
Generated	Enumeration	.02146	51.1	0
	Ordered SMT	.02147	152.7	0
	Raw SMT	.01691	848.9	115
Proof-switch	Enumeration	.09990	44.7	0
	Ordered SMT	.09998	137.0	0
	Raw SMT	.09262	808.7	95

variables before verification, not a uniform preference for enumeration or larger exact radii than ordered SMT.

Removing unread coordinates. Corollary 2 reduces the lifetime dimension from β to the number s read by a query, without changing its exact verification problem. A frozen matrix varies $\beta = 4, 16, 64$, $s = 2$ or $s = \beta$, three seeds, three radii and three timing repeats. All 324 full/projected decisions agree and all 216 satisfiable cases pass independent graph replay. The nine size/seed identities yield a median paired time reduction of 55.2% for $s = 2$, including encoding, solving, completion and replay; it is 0.5% when every coordinate is read. These millisecond-scale measurements on one shared host support a scoped representation saving, not a general solver speedup or radius gain.

4.2 A LEARNED PREDICTOR ON MEASURED EEG GRAPHS

We classify eyes-open versus eyes-closed resting recordings from PhysioNet EEGMMIDB (PhysioNet, 2009); labels come from the recording protocol, not PH. After common-average referencing, 8–12 Hz magnitude-squared coherence between 64 channels gives $x_e = 1 - W_e$ on K_{64}: 2,016 edge births and 1,953 full ordered H_1 births. The complex remains one-dimensional before time one, exactly as in Theorem 1.

Subjects 1–50 fit the candidates, 51–65 select configuration/epoch, and refitting uses 1–65. Evaluation retains all 240 windows of subjects 66–85 (20 people, 12 windows each); no participant crosses the fitting/evaluation split. These are held-out development participants, not an untouched confirmation cohort. Subjects 86–109 remain unopened. All three frozen seeds (17, 29, 43) use 32 affine predicates and two layers of 16 Boolean nodes, followed by a vote over their 64 bits. The rational exports achieve mean subject-macro BA 76.11% (seed SD .96 percentage points). Appendix D gives preprocessing and same-feature MLP/linear/tree controls; no architectural superiority is claimed.

4.3 TIGHT GRAPH-SPACE CERTIFICATION AND COUNTEREXAMPLES

For every model/window pair, the theorem turns the checked feature interval into the same graph-space interval:

$$\underline{r}_i \le r^\star_{\mathrm{PH},i} = r^\star_{\mathrm{graph},i} \le \overline{r}_i.$$

The upper bound is attained by an explicitly reconstructed opposite-prediction graph x_i^{adv}: its recomputed PH equals the target and $\|x_i^{\mathrm{adv}} - x_i\|_\infty = \|L(x_i^{\mathrm{adv}}) - L(x_i)\|_\infty = \overline{r}_i$. This inverse direction supplies an input counterexample; forward stability alone would supply only the sufficient lower guarantee.

Table 2 summarizes all 720 intervals: each has width at most 10^{-5} and an exactly replayed graph witness at its upper endpoint. Figure 2 shows the certified-accuracy curves and radius distributions. Mean $\mathrm{BCA}_{\mathrm{open}}(.01)$ is 67.78%; at that radius the certified class-0/class-1 recalls are 79.17/55.00%, 80.00/58.33%, and 80.83/53.33% for the three seeds. The primary 900-second search allowance closes 710 intervals; ten further 3,600-second refinement attempts close the remainder. Replay and controller overhead are additional, so these allowances are not total wall-time caps. No model

Table 2: Full-H_1 EEG DAGs on the same 240 windows per seed. All accuracy columns are percentages. Every upper endpoint is a replayed graph counterexample; all cases, including nominal errors, remain. Bounds combine primary and additional refinement evidence.

Seed	BA (%)	$\mathrm{BCA}_{\mathrm{open}}(.01)$	$\mathrm{BCA}_{\mathrm{open}}(.02)$	Gap $\leq 10^{-5}$
17	75.00	67.08	59.58	240/240
29	76.67	69.17	63.33	240/240
43	76.67	67.08	58.75	240/240

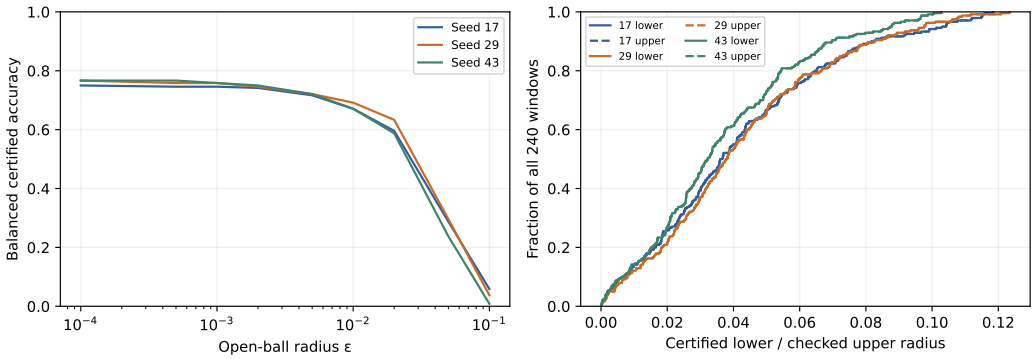

Figure 2: The three frozen full-H_1 EEG DAGs. Left: certified balanced accuracy for open edge balls. Right: empirical distributions of the proved lower and attained upper radii over all 240 windows per model, including errors. Close endpoint curves reflect tight intervals, not exact optimization. The cohort comprises 20 development participants.

is retrained or selected by certificate quality, and nominal errors remain in the denominator. Appendix D.7 reports full costs, and refinement accounting. Appendix D.5 bounds the robust BA from both sides: the mean lower and upper values coincide at .01, and their maximum gap over all radii is .277778 percentage points. Appendix D.6 shows one fully replayed graph counterexample. These are verified intervals of specified width, not claims that the exact optimum was attained.

5 Discussion and Limitations

The exact-image theorem and its computational interface remain distinct: exact geometric transfer does not ensure fast optimization, predictive superiority or semantic meaning of learned predicates. The guarantees start at independently perturbed stored edge births; perturbed coherence matrices need not be realizable by an EEG signal or estimator. They are not raw-signal or clinical certificates. Eyes-open/closed is also confounded with recording run, and the current development cohort does not establish independent confirmation. Named-edge controls are stronger predictors (Table 6). The learned affine predicates have no assigned or validated semantic interpretation. Constrained polytopes and conserved flows have their own assumptions; none is silently covered by the EEG independent-edge guarantee.

Reproducibility and ethics. Full proofs, training/export details, complete certificate accounting and replayed witnesses accompany the paper. Floating LP results are proposals; accepted LP conflicts are checked rationally, while Boolean/SMT UNSAT trusts the backend. EEG uses public de-identified recordings without new acquisition. Statistical summaries preserve participant dependence. Appendix G.2 describes research and writing assistance.

References

Antonio Almudévar, José Miguel Hernández Lobato, and Alfonso Ortega. There was never a bottleneck in concept bottleneck models. In *International Conference on Learning Representations*,

2026. URL https://iclr.cc/virtual/2026/poster/10010544.

Luca Andolfi and Eleonora Giunchiglia. Right for the right reasons: Avoiding reasoning shortcuts via prototypical neurosymbolic ai. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://papers.nips.cc/paper_files/paper/2025/hash/e6bb63d00f597c380c012bd99c68b0c1-Abstract-Conference.html.

Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *Journal of Computational Geometry*, 6(2):162–191, 2015. doi: 10.20382/jocg.v6i2a9.

Kristóf Bérczi, Mirabel Mendoza-Cadena, and José A. Soto. Subset-constrained inverse matroid optimization, 2026. URL https://arxiv.org/abs/2507.00930v3. Version 3, 17 July 2026.

Nikolaj Bjørner, Leonardo de Moura, Lev Nachmanson, and Christoph Wintersteiger. Programming z3. Microsoft Research tutorial, 2019. URL https://z3prover.github.io/papers/programmingz3.html.

Aleksandar Bojchevski and Stephan Günnemann. Certifiable robustness to graph perturbations. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Samuele Bortolotti, Emanuele Marconato, Paolo Morettin, Andrea Passerini, and Stefano Teso. Shortcuts and identifiability in concept-based models from a neuro-symbolic lens. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/fcdccd419c4dc471fa3b73ec97b53789-Abstract-Conference.html.

Rickard Brüel-Gabrielsson, Bradley J. Nelson, Anjan Dwaraknath, and Primoz Skraba. A topology layer for machine learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pp. 1553–1563, 2020.

Mathieu Carrière, Frédéric Chazal, Marc Glisse, Yuichi Ike, Hariprasad Kannan, and Yuhei Umeda. Optimizing persistent homology based functions. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 1294–1303. PMLR, 2021.

Jeremy M. Cohen, Elan Rosenfeld, and J. Zico Kolter. Certified adversarial robustness via randomized smoothing. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1310–1320. PMLR, 2019.

David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37:103–120, 2007.

Matthew L. Daggitt, Wen Kokke, Robert Atkey, Ekaterina Komendantskaya, Natalia Slusarz, and Luca Arnaboldi. Vehicle: Bridging the embedding gap in the verification of neuro-symbolic programs. In *10th International Conference on Formal Structures for Computation and Deduction*, volume 337 of *Leibniz International Proceedings in Informatics*, 2025. doi: 10.4230/LIPIcs.FSCD.2025.2. Invited contribution.

Haya Diwan, Lisa Hellerstein, Nicole Megow, and Jens Schlöter. Optimal verification of a minimum-weight basis in an uncertainty matroid. In *43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026)*, volume 364 of *LIPIcs*, pp. 32:1–32:20, 2026. doi: 10.4230/LIPIcs.STACS.2026.32.

Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

Todd J. Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pp. 31–40, 2007. doi: 10.1145/1265530.1265535.

Itamar Hadad, Guy Katz, and Shahaf Bassan. Formal mechanistic interpretability: Automated circuit discovery with provable guarantees. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/b5afe13494c825089b1e3944fdaba212-Abstract-Conference.html.

Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier against adversarial manipulation. In *Advances in Neural Information Processing Systems*, 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/e077e1a544eec4f0307cf5c3c721d944-Abstract.html.

Gregory Henselman and Robert Ghrist. Matroid filtrations and computational persistent homology. *arXiv preprint arXiv:1606.00199*, 2016. URL https://arxiv.org/abs/1606.00199.

Jiani Huang, Ziyang Li, Binghong Chen, Karan Samel, Mayur Naik, Le Song, and Xujie Si. Scallop: From probabilistic deductive databases to scalable differentiable reasoning. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Johanna Immonen, Amauri Souza, and Vikas Garg. Going beyond persistent homology using persistent homology. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Lukas Koller, Tobias Ladner, and Matthias Althoff. Out of the shadows: Exploring a latent space for neural network verification. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/14bb27f680bee45d83bc769738e7f9b5-Abstract-Conference.html.

Harold W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2):83–97, 1955. doi: 10.1002/nav.3800020109.

Jacob Leygonie and Gregory Henselman-Petrusek. Algorithmic reconstruction of the fiber of persistent homology on cell complexes. *Journal of Applied and Computational Topology*, 8: 2015–2049, 2024. doi: 10.1007/s41468-024-00165-w. URL https://doi.org/10.1007/s41468-024-00165-w.

Jacob Leygonie and Ulrike Tillmann. The fiber of persistent homology for simplicial complexes. *Journal of Pure and Applied Algebra*, 226(12):107099, 2022. doi: 10.1016/j.jpaa.2022.107099. URL https://doi.org/10.1016/j.jpaa.2022.107099. Accepted manuscript: arXiv:2104.01372v3.

Jacob Leygonie, Steve Oudot, and Ulrike Tillmann. A framework for differential calculus on persistence barcodes. *Foundations of Computational Mathematics*, 22:1069–1131, 2022. doi: 10.1007/s10208-021-09522-y.

Vasileios Manginas, Nikolaos Manginas, Edward Stevinson, Sherwin Varghese, Nikos Katzouris, Georgios Paliouras, and Alessio Lomuscio. A scalable approach to probabilistic neuro-symbolic robustness verification, 2025. URL https://arxiv.org/abs/2502.03274. 19th Conference on Neurosymbolic Learning and Reasoning.

Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. DeepProbLog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems*, volume 31, pp. 3753–3763, 2018.

Subrat Prasad Panda, Blaise Genest, and Arvind Easwaran. Approximation-free differentiable oblique decision trees. *Journal of Machine Learning Research*, 27(125):1–31, 2026. URL https://jmlr.org/papers/v27/25-2047.html.

Felix Petersen, Christian Borgelt, Hilde Kuehne, and Oliver Deussen. Deep differentiable logic gate networks. In *Advances in Neural Information Processing Systems*, 2022. URL https://arxiv.org/abs/2210.08277.

PhysioNet. EEG Motor Movement/Imagery Dataset, version 1.0.0, 2009. URL https://physionet.org/content/eegmmidb/1.0.0/.

Pepijn Roos Hoefgeest and Lucas Slot. Robustness of persistent topological features and minimum homological cuts. In *42nd International Symposium on Computational Geometry (SoCG 2026)*, volume 367 of *LIPIcs*, pp. 87:1–87:15, 2026. doi: 10.4230/LIPIcs.SoCG.2026.87.

Attila Szász, Balázs Bánhelyi, and Márk Jelasity. No soundness in the real world: On the challenges of the verification of deployed neural networks. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 58088–58105, 2025. URL https://proceedings.mlr.press/v267/szasz25a.html.

Yusuke Tsuzuku, Issei Sato, and Masashi Sugiyama. Lipschitz-margin training: Scalable certification of perturbation invariance for deep neural networks. In *Advances in Neural Information Processing Systems*, 2018. URL https://papers.nips.cc/paper_files/paper/2018/hash/485843481a7edacbfce101ecb1e4d2a8-Abstract.html.

Yogesh Verma, Amauri H. Souza, and Vikas Garg. Topological neural networks go persistent, equivariant, and continuous. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 49388–49407. PMLR, 2024.

Yogesh Verma, Amauri H Souza, and Vikas K Garg. Positional encoding meets persistent homology on graphs. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267, pp. 61323–61343. PMLR, 2025. URL https://proceedings.mlr.press/v267/verma25b.html.

Eric Wong and J. Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5286–5295. PMLR, 2018.

Cheng Xin, Fan Xu, Xin Ding, Jie Gao, and Jiaxin Ding. Toping: Topologically interpretable graph learning via persistent rationale filtration. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267, pp. 68847–68869, 2025.

Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. A semantic loss function for deep learning with symbolic knowledge. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5502–5511. PMLR, 2018.

Duo Zhou, Jorge Chavez, Hesun Chen, Grani A. Hanasusanto, and Huan Zhang. Clip-and-verify: Linear constraint-driven domain clipping for accelerating neural network verification. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/ffa977364ab7046c803da0e04dbb2832-Abstract-Conference.html.

Daniel Zügner and Stephan Günnemann. Certifiable robustness and robust training for graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 246–256, 2019. doi: 10.1145/3292500.3330905.

APPENDIX ROADMAP

Appendix A proves the exact lifetime image and query-radius reduction. Appendix B specifies the learned DAG and rational deployment, and Appendix C gives its verification procedure and the limits of independent-predicate abstractions. Appendix D contains the EEG protocol, predictive comparisons, certificates, counterexamples and computation costs. Appendix E details the controlled reduction and scaling experiments. Appendix F develops isometric sections, full and joint H_0/H_1 profiles, co-ordered boxes, finite matroids, constrained input polytopes and conserved flows under their respective assumptions. Appendix G gives reproducibility details and the research-assistance statement.

A EXACT LIFETIME IMAGE AND QUERY-RADIUS REDUCTION

A.1 ASSUMPTIONS, NOTATION AND INFERENCE ARTIFACTS

Theorem 1 concerns the declared graph filtration and input domain. Its query-radius consequence holds for any frozen function of the lifetime vector. The rational predicate/DAG representation specifies the evaluated implementation. Every certificate concerns the same deployed predictor whose nominal prediction is reported. Table 3 separates the geometric, query and arithmetic assumptions; Table 4 collects the notation.

Table 3: Three distinct levels of assumptions: the image identity, its query consequence, and the evaluated implementation. None is asserted necessary for all possible certificates.

Assumption	Role	Boundary of the claim
Image theorem (real-valued mathematical statement)		
Fixed finite graph, independently variable edge-birth coordinates	Cartesian input box	No arbitrary graph edits or conserved-flow perturbations
Zero-born vertices; preterminal dimension at most one	Cycle-rank increments identify births	Clique expansion with early triangles is different
Completion at one	All positive H_1 deaths equal one	Variable deaths require another encoding
Uniform clipped ℓ_∞ ball	Co-ordered interval endpoints	Arbitrary unequal boxes can violate the image identity
Query-radius consequence		
Frozen lifetime-only query	Answer depends only on the exact image	An extra input branch requires separate verification
Arithmetic and implementation		
Declared rational computation	Solver and replay refer to one function	Not a certificate for arbitrary floating-point execution

Terminology. A formal certificate has the meaning in Definition 1. A symbolic trace records the predicate and rule computation, without implying invariance. A certified radius is a sufficient lower bound; the exact flip radius is an infimum over admissible answer-changing inputs. A satisfying Boolean or feature assignment is a solver witness. It becomes an input counterexample only after reconstruction and verification of input admissibility, feature consistency, distance and the opposite deployed answer. A DNF clause is a conjunction. An input-realizable feature or predicate assignment has a preimage in the specified input set. These terms are not synonyms. A frozen predictor fixes preprocessing, weights, rules and the decoder; a fixed training algorithm can yield different predictors.

For $\beta = 0$, the lifetime vector is empty and its norm is zero. An infinite certificate means invariance on the entire declared domain. Metric and decision-status definitions appear in Appendix D.4. The input is the stored edge-birth vector on a fixed graph; a perturbed EEG edge vector need not have a signal-space realization.

Table 4: Notation for the central result and evaluated controls.

Symbol	Definition
$\bar{G}$, G_x, x, δ	Fixed structure, weighted graph, edge-birth vector and input ℓ_∞ metric
m_E, β	Observed edge count and cycle rank
$L(x)$, $\mathcal{C}_\beta$	Zero-padded ordered global lifetimes and their ordered domain
d, M, J	Declared learning dimension, predicate count, query count; $d = \beta$ only in experiments using the full lifetime vector
I, s, π_I	Increasing list of query-read coordinate indices, its length, and the corresponding projection of $L(x)$
P_θ, Q_θ, q_θ	Frozen predicates, Boolean rule function, and their composition $q_\theta = Q_\theta \circ P_\theta$
t_j, t_{death}	Predicate threshold and shared death parameter, respectively
ρ_j	Exact input flip-distance infimum of predicate j
$r^\star_q$, $r^\star_{q,G}$	Exact answer flip-distance infimum; the latter specifies the graph
R_{Bool}	Optimal independent-fact, fixed-literal bound
R_2	Radius-local bound retaining individual and pairwise predicate constraints
$R_i \in [0, +\infty]$	Reported lower bound: zero means no positive certificate

A.2 Exact global image and minimum-distance preimage

Global H_1 filtration contract. Let $G = (V, E)$ be a fixed finite simple graph, with $n = |V|$ vertices and $m_E = |E|$ observed edges. All vertices are born at zero; observed edge births $x_e \in [0, 1]$ are the perturbed coordinates. Before one, the complex consists of these vertices and the available observed edges. At one, it is completed to the full simplex. A two-skeleton is sufficient for the H_1 computation. Thus every positive H_1 interval dies at one. Define $\beta = m_E - n + c(G)$, where c includes isolated vertices. Sort the positive lifetimes and append zeros to length β; optionally append further structural zeros to a fixed length $P \geq \beta$. The resulting statistic is $L_1(x) \geq \cdots \geq L_P(x)$. Zero entries are not intervals in the off-diagonal diagram. We use coefficients in $\mathbb{F}_2$. Triangle-freeness is a sufficient restriction for the evaluated clique-expansion implementation; the proof uses the stronger explicit specification that the preterminal complex has dimension at most one. A graph with triangles treated as an unfilled one-complex obeys the same proof; the EEG study constructs an explicit one-dimensional complex on K_{64}.

For every $a \in \mathbb{R}$, define $E_{<a}(x) = \{e : x_e < a\}$ and $E_{\leq a}(x) = \{e : x_e \leq a\}$. Write $b(E') = |E'| - n + c(V, E')$ for the cycle rank of the spanning subgraph. This formula and the connection to persistence computations are standard; see, for example, Henselman & Ghrist (2016) for a broader matroid viewpoint.

Proposition 2 (Exact global lifetime threshold facts)**.** *Under the declared contract, for $1 \leq k \leq P$ and threshold t,*

$$L_k(x) > t \iff b(E_{<1-t}(x)) \geq k \qquad (0 \leq t \leq 1),$$
$$L_k(x) \geq t \iff b(E_{\leq 1-t}(x)) \geq k \qquad (0 < t \leq 1).$$

Every coordinate L_k is 1-Lipschitz in the observed-edge ℓ_∞ metric.

Proof. Adding an edge between components leaves the cycle rank unchanged; adding one within a component increases it by one. No two-simplex kills a cycle before terminal completion. Hence the rank counts exactly the H_1 births already encountered, including simultaneous births, and subtracting birth from one gives the two equivalences. If k exceeds the full graph cycle rank, both right-hand sides are false and the left sides are false by the stated threshold ranges. The excluded inclusive threshold zero would instead make padded coordinates true. For stability, if $\|x - y\|_\infty \leq \delta$, the threshold edge sets satisfy $E_{\leq a-\delta}(x) \subseteq E_{\leq a}(y) \subseteq E_{\leq a+\delta}(x)$. Cycle rank is monotone under adding edges. Taking the generalized inverse of these count inequalities bounds each ordered birth displacement by δ. The same bound holds for its lifetime. Only structural coordinates beyond β remain identically zero; nominal zeros within cycle rank can change. □

An exact image without a graph SMT encoding. The uniform perturbation box admits a stronger reduction than the count compiler above. With $\beta = m_E - n + c(G)$, let $\mathcal{C}_\beta = \{z \in [0, 1]^\beta : z_1 \geq \cdots \geq z_\beta\}$. Only coordinates beyond β are structurally padded; a nominal zero among the first β coordinates is allowed to become positive.

Full proof of Theorem 1. The forward inclusion follows from Proposition 2. For the reverse inclusion, index observed edges in nondecreasing nominal birth order, using an arbitrary fixed refinement of ties. A greedy spanning-forest pass marks the cycle-closing positions $j_1 < \cdots < j_\beta$; their births x_{j_i} are nondecreasing and $L_i(x) = 1 - x_{j_i}$. For a proposed z, set $b_i = 1 - z_i$, so b_i is nondecreasing and lies in $[l_{j_i}, u_{j_i}]$, where $l_j = \max(0, x_j - \epsilon)$ and $u_j = \min(1, x_j + \epsilon)$. Assign every edge in this order the birth

$$y_j = \max\big(\{l_j\} \cup \{b_i : j_i \le j\}\big).$$

Both the lower endpoints and the prescribed birth prefix maxima are nondecreasing. At $j = j_i$, the assignment gives $y_{j_i} = b_i$. For an earlier prescribed birth b_i, monotonic upper endpoints give $b_i \le u_{j_i} \le u_j$; hence $l_j \le y_j \le u_j$ everywhere. The original edge order remains a valid nondecreasing order after reassignment, and its graph-theoretic forest decisions do not depend on the numerical births. It therefore has the same cycle-closing positions and the proposed births. Other refinements of new ties have the same rank increment in each tied block, so the persistence multiset is unchanged. This also covers terminal-zero lifetimes, disconnected graphs and forests. Applying the construction at $\epsilon = \|z - L(x)\|_\infty$, together with the forward Lipschitz bound, establishes the minimum-distance identity. □

Corollary 1 (Exact reduction of an answer radius)**.** *For any frozen Boolean query function $q : \mathcal{C}_\beta \to \{0, 1\}$ depending only on these lifetimes, its infimal input flip distance is*

$$\inf_{\substack{z \in \mathcal{C}_\beta \\ q(z) \neq q(L(x))}} \|z - L(x)\|_\infty,$$

with empty infimum $+\infty$. An ordered-lifetime SAT assignment lifts to an input counterexample, regardless of which proof supports the nominal answer.

Proof. Each input flip gives a lifetime vector at no greater distance. Conversely, the theorem lifts every candidate lifetime vector at exactly its distance. Taking the two infima gives equality, without assuming the infimum is attained or ignoring strict predicate boundaries. □

Thus relational graph constraints cannot improve the ideal ordered-interval answer radius in this contract. A finite-budget graph SMT solver can still return a worse lower bound by timing out. Preprocessing uses one edge sort and a spanning-forest pass; after its closing positions are known, reconstruction is linear in $m_E + \beta$. This is a witness-construction bound, not a bound on solving an arbitrary query formula. The proof uses standard greedy graphic-matroid structure under the stated common-terminal-death filtration. Uniform clipping guarantees the monotone interval endpoints used by the proof; some nonuniform boxes also satisfy that sufficient condition. Arbitrary unequal coordinate budgets and named overlapping local-cycle features do not automatically have this image. The implementation currently repeats sorting and the forest pass for each witness; measured costs include these operations.

A.3 Eliminating coordinates not read by a frozen query

The exact image also permits removal of unused lifetime coordinates. This is a constructive corollary of the ordered-box theorem, not a new general dimensionality-reduction or inverse-persistence principle. Let $c = L(x) \in \mathcal{C}_\beta$, let $I = (i_1 < \cdots < i_s)$ be selected coordinate indices, and write $\pi_I z = (z_{i_1}, \ldots, z_{i_s})$. Query support here is any declared set containing every coordinate read by the frozen formula; finding a smallest semantic support is not required.

Corollary 2 (Exact query-coordinate projection)**.** *For every $\epsilon \ge 0$ under Theorem 1,*

$$\pi_I L(\mathcal{B}(x, \epsilon)) = \{v \in \mathcal{C}_s : |v_a - c_{i_a}| \le \epsilon \ (a = 1, \ldots, s)\}.$$

Every $v \in \mathcal{C}_s$ has an input preimage under $\pi_I L$ at minimum distance $r = \|v - \pi_I c\|_\infty$. For any frozen $q = \bar{q} \circ \pi_I$,

$$r_q^\star(x) = \inf_{v \in \mathcal{C}_s : \bar{q}(v) \neq \bar{q}(\pi_I c)} \|v - \pi_I c\|_\infty.$$

Proof. Projection preserves order and coordinate bounds, giving one inclusion. For the converse, let $r = \|v - \pi_I c\|_\infty$ and form $a_j = \max(0, c_j - r)$, replacing a_{i_b} by v_b. Define $z_j = \max_{k \ge j} a_k$. This vector is decreasing and in $[0, 1]^\beta$. For $j = i_b$, all later selected values are at most v_b, while

$c_k - r \le c_{i_b} - r \le v_b$ for $k \ge i_b$; hence $z_{i_b} = v_b$. Every term in the suffix maximum is at most $c_j + r$: for an unselected term use $c_k - r \le c_j + r$, and for a selected term use $v_b \le c_{i_b} + r \le c_j + r$. Also $z_j \ge \max(0, c_j - r)$. Thus $\|z - c\|_\infty = r$, the lower bound following from the selected coordinates. Theorem 1 lifts this completion at distance r. Projection stability supplies the matching lower bound for every possible input preimage. Grouping opposite answers by v proves the final identity without assuming attainment of that infimum. For empty I, use $z = c$: the query is constant and its flip infimum is $+\infty$. Completion costs $O(\beta)$ after the indices are known. □

This removes variables without enlarging the exact radius. In particular, $s = \beta$ gives no variable reduction. The same argument is not valid for an arbitrary coupled input domain: if $L = (1 - a, a)$, $a \in [1/5, 2/5]$, the marginally feasible target $(4/5, 2/5)$ violates the joint constraint. Under side constraints one must project the full input–feature relation, not silently replace it by these ordered coordinate intervals.

A.4 Beyond the uniform fixed-death contract

Three statements must be distinguished. For an input set A and feature map F, an enclosure E satisfies $F(A) \subseteq E$; an exact image satisfies $E = F(A)$, requiring both inclusions. A minimum-distance lift additionally attains the distance from a specified nominal input. Theorem 1 establishes all three under its hypotheses; none is inferred for a different filtration from successful sampled replays.

A variable-death family. Consider k triangles sharing one vertex and otherwise disjoint, with $k \ge 2$. All vertices and two edges of each triangle are born at zero. Its remaining edge has birth $b_i \in [0, 1]$. For a shared $t_{\mathrm{death}} \in [0, 1]$ and fixed nonnegative offsets o_i, the triangle is filled at $\max\{b_i, \min(1, t_{\mathrm{death}} + o_i)\}$. Put $u = (b_1, \dots, b_k, t_{\mathrm{death}})$. Its H_1 lifetime, with zeros retained, is

$$\ell_i(u) = \max\{0, \min(1, t_{\mathrm{death}} + o_i) - b_i\}.$$

The vector $F(u)$ sorts these k values in nonincreasing order. Here deaths can precede one and depend on a common input coordinate. The perturbation is a clipped uniform ball in u, not an independent perturbation of all simplex filtration values.

Why the ordered marginal box is insufficient here. For $k = 2$, zero offsets, $u_0 = (1/5, 2/5, 1/2)$ and radius $1/20$, the ordered marginal lifetime bounds are $z_1 \in [1/5, 2/5]$, $z_2 \in [0, 1/5]$. The target $(2/5, 0)$ is inside these bounds but is unattainable: the larger lifetime can be $2/5$ only with $t_{\mathrm{death}} = 11/20$ and $b_1 = 3/20$, forcing the second lifetime to be at least $1/10$. More generally the query $q(z) = \neg[z_1 > 3/8] \vee [z_2 > 1/40]$ is invariant throughout this input ball. If $z_1 > 3/8$, then $t_{\mathrm{death}} > 21/40$ and the other birth is at most $9/20$, giving $z_2 > 3/40$. The marginal ordered box nevertheless contains false answers. This counterexample refutes an unrestricted transfer of the fixed-death image identity.

B Learning and rational deployment of the Boolean DAG

B.1 Architecture and coordinate convention.

The head is inspired by differentiable Boolean-gate learning (Petersen et al., 2022); learned affine input predicates and learned wiring are our implementation choices, not a claim of reproducing that method. Write $h = B(x) = \mathbf{1} - L(x)$ for the complete ordered H_1 birth profile. This coordinate change preserves ℓ_∞ distance. For the EEG graph, $h \in [0, 1]^{1953}$. The deployed head has P affine predicate bits

$$p_j(h) = \mathbf{1}\{a_j^T h + b_j > 0\}, \qquad v^{(0)} = p. \tag{1}$$

Layer ℓ selects two entries of $v^{(\ell)}$ for each of its W gates, applies a selected binary Boolean operation, and appends the W outputs to form $v^{(\ell+1)}$. Thus every layer can read the original predicates and all earlier layers, but not another node of the same layer. The selected inputs may coincide. Operation $o \in \{0, \dots, 15\}$ has truth-table entries $T_{o,k} = \lfloor o/2^k \rfloor \bmod 2$, $k = 0, 1, 2, 3$, with input

index $k = 2a + b$. All 16 operations, including constants and projections, are available. With two gate layers, the final decision is

$$f(h) = \alpha^T v^{(2)} + \alpha_0, \qquad q(h) = \mathbf{1}\{f(h) \geq 0\}. \tag{2}$$

The entire threshold vote belongs to the learned Boolean query; the external binary decoder is the identity. The vote receives only hard predicate and gate bits. Direct predicate-to-vote connections are allowed; no continuous feature or pre-threshold predicate score bypasses this interface. This restriction does not assign human meanings to the learned bits or establish that the selected gates are necessary for prediction.

B.2 Hard forward values and the actual surrogate derivatives.

Let sg denote stop-gradient and define

$$\mathrm{ST}(s, h) = \mathrm{sg}(h) + s - \mathrm{sg}(s). \tag{3}$$

For standardized input $\widetilde{h}_i = (h_i - \mu_i)/s_i$, training uses margins $m_j = \widetilde{a}_j^T \widetilde{h} + \widetilde{b}_j$ and

$$\widehat{p}_j = \mathrm{ST}(\sigma(2m_j), \mathbf{1}\{m_j > 0\}), \qquad \frac{\partial \widehat{p}_j}{\partial m_j} = 2\sigma(2m_j)(1 - \sigma(2m_j)), \tag{4}$$

where σ is the logistic function. The factor two is fixed. For either a wiring-logit vector ξ or an operation-logit vector η, the selector is

$$w(\xi) = \mathrm{ST}\big(\mathrm{softmax}(\xi), e_{\arg\max \xi}\big), \qquad J_{\mathrm{softmax}}(\xi) = \mathrm{diag}(\rho) - \rho\rho^T, \quad \rho = \mathrm{softmax}(\xi). \tag{5}$$

Both softmax temperatures are one, without annealing. For one gate, the two selected inputs and the selected table are

$$a = \sum_k v_k w_k(\xi^{(1)}), \quad b = \sum_k v_k w_k(\xi^{(2)}), \quad \tau = \sum_{o=0}^{15} w_o(\eta) T_o. \tag{6}$$

Its output is the multilinear extension

$$g = (1-a)(1-b)\tau_0 + (1-a)b\tau_1 + a(1-b)\tau_2 + ab\tau_3. \tag{7}$$

The forward value is exactly the selected Boolean truth-table entry. Backward propagation differentiates (6)–(7), with (5) as the selector Jacobian. Explicitly,

$$\frac{\partial g}{\partial a} = (1-b)(\tau_2 - \tau_0) + b(\tau_3 - \tau_1), \qquad \frac{\partial g}{\partial b} = (1-a)(\tau_1 - \tau_0) + a(\tau_3 - \tau_2), \tag{8}$$

$$\frac{\partial a}{\partial \xi_k^{(1)}} = \rho_k \left(v_k - \sum_j \rho_j v_j \right), \qquad \frac{\partial a}{\partial v_k} = \mathbf{1}\{k = \arg\max \xi^{(1)}\}. \tag{9}$$

The corresponding formula holds for b. If $\pi = ((1-a)(1-b), (1-a)b, a(1-b), ab)$ and $\zeta = \mathrm{softmax}(\eta)$, then

$$\frac{\partial g}{\partial \eta_o} = \zeta_o \left(\pi^T T_o - \sum_r \zeta_r \pi^T T_r \right). \tag{10}$$

These are local surrogate derivatives evaluated with the hard forward operands, not the Jacobian of an entirely soft forward network.

B.3 Objective and optimization.

For a minibatch $\mathcal{B}$, the code applies unweighted two-class cross entropy to logits $[0, f_i]$:

$$\mathcal{L}_{\mathcal{B}} = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \big[\log(1 + \exp f_i) - y_i f_i\big], \qquad \frac{\partial \mathcal{L}_{\mathcal{B}}}{\partial f_i} = \frac{\sigma(f_i) - y_i}{|\mathcal{B}|}. \tag{11}$$

The vote logit is differentiable in its weights; its final decision threshold is not inserted into the training loss. Predicates, wiring, operations, and vote are optimized jointly from output labels. There is no fact supervision, separately fitted rule program, or gradient through the fixed PH extraction. AdamW uses learning rate .003, weight decay 10^{-4}, batches of 128, and gradient-norm clipping at 10. Training uses float32 neural arithmetic, at most 200 epochs, and shuffled batches.

B.4 SELECTION AND REFIT.

Scaling and initial fitting use subjects 1–50. The candidate architectures are $(P, W, \text{layers}) = (32, 16, 2)$ and $(64, 32, 2)$. Every five epochs, and at the final allowed epoch, selection evaluates the current rational hard export on subjects 51–65 using mean participant balanced accuracy. Equal scores retain the earlier checkpoint; improvement must exceed 10^{-12}. Training stops after 35 epochs without improvement. Candidate ties select the smaller architecture. The selected architecture is refit from fresh initialization with the same seed, on subjects 1–65, for the selected epoch count, with scaling recomputed on those subjects. All model selections and refits were frozen before evaluating subjects 66–85. The full-H_1 models for seeds 17, 29, and 43 all select 32 predicates and two 16-node layers; their refit counts are 45, 35, and 60 epochs. They are three separate models, not an ensemble. Their complete 1953-coordinate inputs remain active: each coordinate has a nonzero coefficient in at least one exported predicate. Thus 32 predicates do not reduce the continuous verification dimension to 32.

B.5 ONE RATIONAL DEPLOYED FUNCTION.

Scaling is absorbed before coefficient rounding:

$$a_{ji}^{\text{fold}} = \widetilde{a}_{ji}/s_i, \qquad b_j^{\text{fold}} = \widetilde{b}_j - \sum_i a_{ji}^{\text{fold}} \mu_i. \tag{12}$$

The folded predicate weights/biases and vote coefficients are rounded to six decimal places and exported along with hard argmax wiring/operation indices. These serialized decimal coefficients define the deployed function. Stored feature numbers are interpreted as exact rationals from their decimal strings; the features themselves are not rounded to six places. Positive common-denominator scaling gives exact integer sign tests. The strict predicate convention assigns equality to false; the vote assigns equality to class one. Validation and test use the decimal integer evaluator. Verification uses equivalent Fraction or arbitrary rational common-denominator evaluation, with Fraction replay retained for the integer-accelerated solver. These are different implementations of the same rational function, not a claim that a floating-point network and its rounded export agree at every input.

To express this frozen model in lifetime coordinates, substitute $h = \mathbf{1} - L$ exactly after export: $a_j^L = -a_j^B$ and $b_j^L = b_j^B + \sum_i a_{ji}^B$. No additional rounding is applied; this preserves every existing prediction and certificate.

The encoding, checked LP conflicts and remaining Boolean/SMT trust boundary are specified in Appendix C.1.

C VERIFICATION OF THE FROZEN PREDICTOR

C.1 DAG VERIFICATION BY CHECKED BOOLEAN–LINEAR REFINEMENT

Proposition 3 (Exact encoding of a frozen DAG)**.** *Let the feature domain be specified by rational linear constraints. For each affine predicate introduce a bit $p_i \leftrightarrow (a_i^T h + b_i > 0)$, encode every gate by its truth table on its selected predecessor bits, and encode the vote with the deployed ≥ 0 convention. The resulting formula with output different from $q(h_0)$ is satisfiable exactly when the domain contains an opposite-label feature vector. DAG nodes and their shared uses require no DNF expansion.*

Proof. A feature vector uniquely determines the predicate bits. Induction in topological order uniquely determines every gate, and the vote determines the output. Conversely, the same induction forces any satisfying assignment to equal execution of the frozen classifier. Rational scaling of the vote preserves its inequality. Strict predicate boundaries remain strict in both directions. □

Proposition 4 (Individual predicate extrema on an ordered box)**.** *For sorted bounds $0 \leq \ell \leq u \leq 1$, let $X = \{x : \ell \leq x \leq u,\ x_1 \leq \cdots \leq x_d\}$. Put $S_k = \sum_{j=k+1}^{d} a_j$, $L(t) = \#\{j : u_j < t\}$ and $U(t) = \#\{j : \ell_j < t\}$. Then*

$$\min_{x \in X} a^T x = \int_0^1 \min_{L(t) \leq k \leq U(t)} S_k \, dt, \qquad \max_{x \in X} a^T x = \int_0^1 \max_{L(t) \leq k \leq U(t)} S_k \, dt.$$

Proof. Write $x_j = \int_0^1 \mathbf{1}\{t \le x_j\}\, dt$. For almost every t, the indicators form a suffix whose switch index lies between $L(t)$ and $U(t)$, proving the respective bounds. They are attainable: choose the smallest minimizing switch at each t. As both interval endpoints increase, this switch cannot move left; otherwise the newly preferred smaller index and the former index were both available before, contradicting minimization with the same tie rule. Its suffix indicators are therefore decreasing in t and define sorted coordinates in $[\ell, u]$ attaining the integral. The maximizing case is identical. Endpoint equalities occur on a measure-zero set. □

The integrands change only at coordinate bounds. Suffix sums and moving min/max queues give exact rational evaluation without an LP per predicate. This is an exact extremum for each predicate separately; it does not imply that arbitrary combinations of individually feasible predicate signs are jointly realizable.

The frozen DAG is encoded with shared predicate values and an exact integer-scaled weighted vote. At radius ρ, the feature domain has bounds $\ell_j = \max(0, h_j - \rho)$ and $u_j = \min(1, h_j + \rho)$ and order constraints $Ch' \le 0$. For a Boolean assignment v, put $s_i = -1$ when $v_i = 1$ and $s_i = 1$ otherwise. A phase-I LP introduces $t \ge 0$ and relaxes the assigned predicate constraints to

$$s_i(a_i^T h' + b_i) - t \le 0.$$

For nonnegative order multipliers λ and predicate multipliers μ with $\sum_i \mu_i \le 1$, set

$$r = C^T\lambda + \sum_i \mu_i s_i a_i, \qquad \gamma = \sum_i \mu_i s_i b_i + \sum_j \min\{r_j \ell_j, r_j u_j\}.$$

Every feasible phase-I point has $t \ge \gamma$: this follows by adding the nonpositive constraints to the objective and minimizing the resulting affine expression over $h' \in [\ell, u]$ and $t \ge 0$. Thus $\gamma > 0$, recomputed in rational arithmetic, excludes the assignment. Only predicates with nonzero μ_i are needed in the resulting Boolean conflict. Floating LP multipliers and points are proposals, not trusted bounds. Relaxing > 0 to its closure is used only for exclusion; every accepted witness satisfies the original strict/nonstrict deployed semantics.

Boolean UNSAT after these checked conflicts proves the radius safe. The extrema in Proposition 4 strengthen the initial fixed-predicate set. The independent Boolean control stops at this abstraction: a satisfying assignment there supplies an *abstract* upper threshold, never a graph counterexample. Joint refinement additionally checks continuous realizability. An interior LP proposal pushes both true and false predicate inequalities away from zero to avoid numerical boundary misclassification; failure of this proposal proves nothing. Exceptional boundary cases use exact rational SMT under an externally enforced deadline. All returned unknowns retain only previously established lower bounds.

The search stores rational lower/upper endpoints, evidence, model/domain hashes, and actual opposite-label feature points. The main EEG implementation uses ascending births $B = 1 - L$ in these formulas; this exact isometry preserves the radius query. Replay recomputes exact dual residuals, Boolean coverage and witness distances. A tolerance-complete interval of width 10^{-5} is distinguished from an exactly optimized infimum. Strict predicates can make the infimum unattained.

The trust boundary includes the Boolean/SMT solver: replay checks rational conflicts and reruns Boolean coverage, but does not independently check a SAT/SMT UNSAT proof certificate. Numerical LP status alone never excludes an assignment.

C.2 Verification checks and trust assumptions

Table 5 separates checked witness properties from solver-trusted exclusion. Unknown and timeout supply no additional lower bound.

C.3 Sound refinement and bound composition

Sound composition of lower bounds. Let h, x, $\mathcal{X}$ and δ be the same frozen predictor, nominal input, admissible domain and metric for two certified lower bounds R_1, R_2. The domain-relative open balls of radii R_1, R_2 are nested. Consequently

$$R_{\text{safe}} = \max\{R_1, R_2\}$$

Table 5: Verification checks and their trust assumptions.

Component	Checked property	Remaining trust
Exact image and lift	Domain and distance equality	Mathematical assumptions and proof
Rational evaluation	Stored-constant function	Translation of the deployed model
Input-witness replay	Distance, exact $\mathbb{F}_2$ homology, opposite answer	PH and reconstruction implementations; SAT checks only
SMT exclusion	No opposite answer in the encoded box	Solver and encoding; no independently checked UNSAT proof log
Hashes/manifests	Source, model and input identity	Do not establish semantic correctness

is also a valid lower bound for h at x, including an infinite bound. This elementary composition prevents a completed sound certificate from being discarded when a subsequent search times out; it does not make that earlier computation free. Bounds for distinct predictors or perturbation contracts cannot be composed in this way. Our artifact interface checks their common identity.

What the independent-fact relaxation discards. Use the frozen predicate vector $P_\theta : \mathcal{C}_\beta \to \{0,1\}^M$, the Boolean rule function Q_θ, and $b = P_\theta(L(x))$. For $r > 0$ let

$$\mathcal{A}_r = \{P_\theta(L(y)) : y \in [0,1]^{m_E},\ \|y - x\|_\infty < r\},$$
$$\mathcal{B}_r = \{v \in \{0,1\}^M : v_j = b_j \text{ whenever } \rho_j \geq r\}.$$

where ρ_j is the exact infimum of input distances changing predicate j. Then $\mathcal{A}_r \subseteq \mathcal{B}_r$. A gap at radius r occurs precisely when Q_θ is constant on $\mathcal{A}_r$ but not on $\mathcal{B}_r$. This characterization concerns the same vocabulary and answer, not a change in the learned model or the Boolean spelling of its rules.

Proposition 5 (Sound radius-local exclusion)**.** *Let $\mathcal{C}_r$ initially equal $\mathcal{B}_r$. At each iteration choose an unexcluded $v \in \mathcal{C}_r$ with $Q_\theta(v) \neq Q_\theta(b)$ and check its realizability in the exact image of the closed input ball of radius r. On UNSAT, choose a conjunction of literals satisfied by v whose inconsistency with that image has been proved, and remove all assignments satisfying it. On SAT, lift and replay the opposite answer; on unknown, stop without increasing the accepted bound. Then $\mathcal{A}_r \subseteq \mathcal{C}_r$ after every removal. If no opposite-answer assignment remains in $\mathcal{C}_r$, r is a sound open-ball certificate. If the checks are exact and terminate, at most 2^M exclusions are possible at one fixed radius before the procedure either certifies the open ball or returns an input-realizable opposite answer at distance at most r.*

Proof. The initial inclusion follows from the definition of each head flip infimum. An excluded conjunction has no realization even in the closed image and therefore cannot remove an assignment from $\mathcal{A}_r$. Induction proves the inclusion. Excluding all opposite answers establishes invariance. At an UNSAT iteration, the removed set contains the selected, previously unexcluded v. Thus $\mathcal{C}_r$ strictly decreases in cardinality. There can be at most 2^M such iterations at fixed r. A feasible candidate is lifted by Theorem 1, and both its lifetimes and opposite answer are replayed. A returned boundary counterexample does not contradict invariance of the open ball. Timeouts give no new conclusion. □

The implementation obtains inconsistent literal subsets from solver UNSAT cores and uses them only at the radius where they were checked. Relations estimated from training correlations are never accepted as certificates. This is a standard abstraction-refinement strategy instantiated with the exact image, not a claim of a new general SAT algorithm. It has exponential worst-case Boolean complexity and cannot improve on an ideal exact ordered solver's radius. Its possible benefit is finite-budget cost. Replacing exact ρ_j by sufficient lower bounds only enlarges the initial Boolean set, so the inclusion and progress proof still apply. The exact gap characterization above, however, uses the true flip infima.

C.4 Independent-predicate verification

Let x_0 be the nominal input, $b = P_\theta(L(x_0))$ its predicate vector, and ρ_j the infimum of input distances that change predicate j. Theorem 1 allows the same infimum to be computed in the ordered

lifetime domain. Write $Q = Q_\theta : \{0,1\}^M \to \{0,1\}$ for the frozen Boolean rule function. Under the central contract the deployed answer is $Q_\theta(P_\theta(L(x)))$. For a Boolean assignment v, define

$$D_\rho(b,v) = \max_{\{j : v_j \neq b_j\}} \rho_j, \qquad \max \varnothing = 0, \quad R_{\mathrm{Bool}} = \min_{\{v : Q(v) \neq Q(b)\}} D_\rho(b,v), \qquad \min \varnothing = +\infty.$$

D_ρ is a weighted discrepancy used in a Boolean relaxation, not the distance to a graph realizing v; some assignments have no such graph.

Proposition 6 (Representation-invariant fixed-literal radius)**.** *The answer is invariant on the input ball $\delta(x, x_0) < R_{\mathrm{Bool}}$. Moreover R_{Bool} is the maximum of $\min_{j \in S} \rho_j$ over all index sets S for which the partial assignment $v_j = b_j$ for $j \in S$ logically implies $Q(v) = Q(b)$ on the full Boolean cube. Consequently $R_{\mathrm{Bool}} \leq r_q^\star(x_0)$. Equivalent Boolean formulas on the same predicate vocabulary give the same R_{Bool}. This does not assert invariance under changing the vocabulary.*

Proof. An input at distance s can change only facts whose $\rho_j \leq s$. Its truth vector v therefore has $D_\rho(b,v) \leq s$; if its answer differs, $s \geq R_{\mathrm{Bool}}$. For the optimality statement, any sufficient partial assignment of radius r must disagree with each opposite vector on some index of radius at least r, so $r \leq R_{\mathrm{Bool}}$. Conversely fix all indices with $\rho_j \geq R_{\mathrm{Bool}}$. An opposite vector consistent with these indices would have discrepancy strictly smaller than R_{Bool}, a contradiction. The same argument covers $R_{\mathrm{Bool}} = +\infty$: fix all infinite-radius indices. If there are no opposite vectors, the empty partial assignment suffices. Input soundness gives the upper bound by the exact flip radius. The definition depends on the truth function, not its formula. □

This is an optimal bound within an independent-fact relaxation, not a new exact-image theorem. A Boolean SAT query asks for an opposite answer while fixing every fact with $\rho_j \geq r$; binary search over the finitely many head radii computes the control without listing all prime implicants. Negated literals remain linked complements. The original input verifier can still be tighter by excluding jointly unattainable fact assignments. For the consensus diagnostic $q = (a \wedge c) \vee (b \wedge \neg c)$ with $a = [z_1 > .7]$, $b = [z_2 > .5]$, $c = [z_3 > .405]$ and nominal $(.8, .6, .4, .2)$, the implied consensus clause $a \wedge b$ obtains $R_{\mathrm{Bool}} = .1$. That diagnostic establishes no advantage over this stronger control. The distinct four-predicate family in Proposition 1 does: p_3, p_4 cannot both be false, and this relation is not a Boolean tautology before restricting the attainable facts. To check that the query is not a single conjunction on the ordered domain, note that each of the four literals takes both truth values across true inputs: use $(.65, .6, .4, .2)$ for $p_1 = 0$ and $(.8, .48, .4, .2)$ for $p_2 = 0$, then vary z_3 below $.4 - \eta$ or above $.4 + \eta$ with $z_1 = .8, z_2 = .6$. Thus no signed literal is true on every true input, while false inputs exist. Each clause can uniquely support the answer, proving indispensability.

C.5 Pairwise dependencies and the remaining joint constraint

Figure 3 illustrates how one coupled-coordinate constraint rules out the opposite assignment allowed by independent facts. Fix the same predicate map P, Boolean query Q, ordered domain $\mathcal{C}_\beta$ and nominal lifetime vector $c = L(x)$ as in the main text. For a nonempty index set I with $|I| \leq 2$ and a Boolean vector v, define

$$d_I(v_I) = \inf\{\|z - c\|_\infty : z \in \mathcal{C}_\beta,\ P_I(z) = v_I\}, \qquad \inf \varnothing = +\infty,$$

where P_I and v_I restrict to indices in I. The radius-local pairwise bound is

$$R_2 = \min_{v : Q(v) \neq Q(P(c))} \ \max_{1 \leq |I| \leq 2} d_I(v_I).$$

This is a relaxation of joint realizability, not a graph counterexample generator. A global pair control is weaker: it only forbids pairs with $d_I(v_I) = +\infty$ and otherwise uses individual flip distances. Implications, mutually exclusive facts and covering pairs are all instances of these forbidden assignments. A pair can be possible globally but require more perturbation locally than either of its individual facts.

Proposition 7 (Pairwise bound and its coordinate-threshold scope)**.** *With exact individual distances, $R_{\mathrm{Bool}} \leq R_2 \leq r_q^\star(x)$. If every predicate tests one coordinate against a threshold strictly between zero and one, with either orientation and fixed strictness, then $R_2 = r_q^\star(x)$ for every Boolean query. In this coordinate case, complete global pair exclusions together with individual distances already suffice. The conclusion does not extend to general affine predicates.*

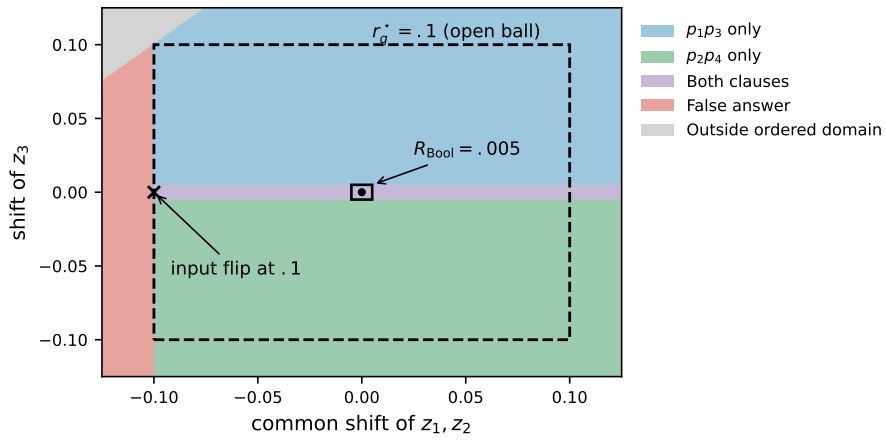

Figure 3: Analytic coupled-predicate example, not a selected fitted model. The coordinates shift z_1, z_2 together horizontally and z_3 vertically. Colours identify active clauses, false answers and violations of ordering. The dashed square outlines the input certificate's open ball; its marked boundary point $(.7, .5, .4, .2)$ changes the answer at distance .1. The small square marks $R_{\text{Bool}} = .005$; a dependency-aware Boolean control also reaches .1. The full four-coordinate claim follows from Proposition 1 and the minimum-distance lift.

Proof. Singleton sets I give the independent-fact discrepancy, proving the first inequality. Every input-realizable opposite assignment has distance at least each $d_I(v_I)$; taking infima proves the second inequality by the exact-image theorem. For coordinate tests, a fixed assignment induces intervals intersected with the ordering cone. Infeasibility arises from a lower endpoint exceeding an upper endpoint on the same coordinate or an earlier ordered coordinate, including incompatible strict endpoints. Each offending endpoint is supplied by one fact or by the fixed domain. Thus an infeasible assignment has an infeasible restriction to at most two facts. For a feasible assignment, Proposition 8 gives its distance as the maximum displacement of its interval endpoints. An unchanged fact contributes zero; a changed fact contributes its individual threshold distance $|c_{k_j} - t_j|$. Consequently singleton distances determine the distance of each feasible cell, while pair exclusions eliminate precisely the empty cells. Minimizing over opposite assignments proves equality. □

This proposition identifies a sufficient grammar for completeness of a standard pairwise relaxation. For the main analytic example, the single exclusion $\neg(\neg p_3 \wedge \neg p_4)$ recovers radius .1. For a contrasting affine example on $\mathcal{C}_2$, use $p_1 = [z_1 > 3/5]$, $p_2 = [z_2 > 2/5]$, $p_3 = [z_1 + z_2 < 1]$. Every pair can be true, but the triple cannot. At $c = (11/20, 7/20)$, all pairwise distances to the all-true assignment are $1/20$ as infima, while the query $p_1 \wedge p_2 \wedge p_3$ is constant false. This is a consistency counterexample, not a predictive benchmark.

D EEG: PROTOCOL, PREDICTIVE COMPARISONS AND EXACT CERTIFICATES

D.1 DATA AND FIXED GRAPH CONSTRUCTION

We use PhysioNet EEGMMIDB 1.0.0 (PhysioNet, 2009), resting runs R01/R02 (eyes open/closed), with protocol labels 0/1. The 85 included participants supply 170 recordings, without exclusions. Six nonoverlapping ten-second windows from the first 60 seconds of each recording give 1,020 balanced windows (510/class). Sampling is 160 Hz. Common-average referencing precedes magnitude-squared Welch coherence, with two-second periodic-Hann segments, 50% overlap and segment demeaning; values are averaged over 8–12 Hz. Each of the 2,016 channel pairs becomes an edge with birth $1 - W_e$. The graph is not clique-expanded before terminal time one. A deterministic Kruskal pass extracts the 1,953 sorted cycle births; $L = 1 - B$ gives the nonincreasing lifetime vector of Theorem 1.

Participants 1–50 fit each candidate and 51–65 select configuration/epoch. Refits on 1–65 precede evaluation on all 240 windows from 66–85. Train-only standardization is refitted on the applicable fitting subset. Each person has six windows/class, so window BA equals subject-macro BA here. Subjects 86–109 were not downloaded or opened for these results. The evaluation participants belong to a cohort used for representation development; participant-disjoint fitting and validation therefore do not constitute an untouched confirmation study. Eyes-open/closed is confounded with the corresponding recording run.

D.2 Predictive controls on fixed representations

Table 6 reports the non-PH graph and spectral controls. Table 7 reports the same classical/MLP controls and the verifiable heads on full H_1 and H_0. Together they compare representations and classifiers under one participant protocol. The same participant split is used throughout. Logistic

Table 6: Non-PH representation controls: nominal balanced accuracy (%) on the same participants; neural uncertainty is seed SD, not a confidence interval. Named edges and spectral inputs are predictive controls, not inputs covered by the full-H_1 transfer theorem.

Representation	Dim.	Logistic	SVM	Boosting	MLP
Named edges	2016	88.33	80.83	87.08	85.28 ± 0.64
Sorted all edges	2016	66.25	67.08	73.33	76.81 ± 0.87
Spectral control	64	82.08	81.25	81.25	85.28 ± 1.46

Table 7: Subject-macro balanced accuracy (%) on the same 20 development participants. Classical entries use one selected fit, including the converged primal SVM. Neural entries are means $\pm$ sample SD over seeds 17/29/43; NeSy evaluation uses the hard rational export. NeSy heads use the same candidate grids for both representations.

Classifier	H_1 (1953)	H_0 (63)
Logistic regression	71.67	67.08
Linear SVM	76.25	67.08
Gradient boosting	74.58	67.92
MLP	75.42 ± 1.10	70.83 ± 0.72
Oblique tree	65.00 ± 13.87	70.42 ± 4.35
Boolean DAG	76.11 ± 0.96	72.50 ± 1.91
Min/max rules	69.31 ± 0.48	70.69 ± 0.87

regression and linear SVM select $C \in \{.01, .1, 1, 10\}$; reported SVM fits use a converged primal optimization. Histogram gradient boosting selects 7/15/31 maximum leaf nodes. MLPs have two ReLU hidden layers, width 64/128 and dropout .1. Neural fits use AdamW, learning rate .003, weight decay 10^{-4}, at most 200 epochs, and patience 35. Oblique trees use depths 4/6, adapted from a differentiable classification construction (Panda et al., 2026); min/max heads have 8/16 clauses of four affine conditions. Both are classification adaptations.

The spectral control uses 64 channelwise log relative alpha-power features: 8–12 Hz power divided by 4–40 Hz power, floored at 10^{-15} before taking the natural logarithm. It does not use PH. Named edges retain channel-pair identity; sorted-edge and PH inputs discard it to different degrees. The stronger named-edge results limit any claim of predictive superiority of global PH. DAG–MLP full-H_1 BA differs by $+.69$ percentage points, with exploratory paired participant-bootstrap 95% interval $[-1.67, 2.92]$ (10,000 resamples after averaging seeds). It establishes neither superiority nor a prespecified noninferiority result. Neural table error bars are sample SD over three seeds, not participant CIs. The affine predicates are learned without a prescribed semantic vocabulary.

D.3 Graph-space certificates across prediction heads

Table 8 compares certificate intervals for three learned heads on the same full-H_1 input domain. All 240 evaluation windows from participants 66–85 are included for each of the three frozen seeds (17,

Table 8: Full-H_1 controls on all 240 evaluation windows of participants 66–85, pooled over seeds 17, 29 and 43: 720 model/window cases per head. All nominal errors are retained. Lower/upper medians are separate summaries; gap counts are per case. Every upper endpoint has an actual graph witness.

Head	Correct	Median lower	Median upper	Gap $\leq 10^{-5}$
Tree	468/720	0.003050	0.003050	720/720
DAG	548/720	0.035524	0.035531	720/720
Min/max	499/720	0.003068	0.003076	720/720

29 and 43), giving 720 model/window cases per head and 2,160 cases overall. No model is refitted for certification, and no nominal error is excluded. The full-H_1 theorem applies unchanged to each frozen head. Each lower endpoint guarantees open-ball invariance; each upper is the distance to an actually replayed opposite-prediction graph. The tree uses leafwise LP; min/max is compiled into an equivalent Boolean query and verified with ordered SAT/LP refinement. All 2,160 intervals have gap at most 10^{-5}, a tolerance guarantee rather than an exact rational optimum. The DAG row uses the same 720 certificates as the all-window evaluation below.

For these frozen models, the DAG has substantially larger certified radii: its lower bound exceeds the tree's upper bound in 690/720 paired cases, and the min/max upper bound in 702/720. At radius .01, certified balanced accuracy for open balls is 67.78% for DAG, 15.42% for tree and 9.44% for min/max. The DAG gains are 52.36 and 58.33 percentage points, with exploratory paired participant-bootstrap 95% intervals $[44.86, 60.28]$ and $[49.72, 67.64]$, respectively (5,000 resamples, retaining all windows and seeds of each sampled participant). These comparisons involve 20 participants, not 720 independent observations; they support a robustness advantage of the fitted DAGs on this development cohort, without isolating architecture from model capacity or training.

D.4 All-window certificate accounting

Metrics and decision status. For a fixed labelled evaluation set, n_c counts inputs with true class c. At $\epsilon > 0$, define $I_i(\epsilon) = \mathbf{1}[\hat{y}_i = y_i,\ \underline{r}_i \geq \epsilon]$. Certified class recall is $n_c^{-1} \sum_{i:y_i=c} I_i(\epsilon)$ and $\mathrm{BCA}_{\mathrm{open}}$ averages these recalls over the declared classes. At radius zero we report nominal balanced accuracy. An absent true class makes its recall and BCA undefined. A stable error contributes zero; a fully certified binary constant has BCA $1/2$ when both classes occur.

The exact flip-radius infimum, a proved lower bound, an attained witness distance and a search cap are distinct. An infimum r certifies the open ball of radius r, even when a flip occurs on its boundary. Unknown and timeout retain only established bounds. A zero lower bound supplies no positive certificate; an infinite bound requires invariance on the entire declared domain. Finite-budget search reports the interval, tolerance and measured cost rather than claiming an exact optimum from successful execution. All three full-H_1 DAG exports remain frozen; all 240 evaluation windows/model are certified without choosing seeds by certificate quality. Warm upper bounds search exact affine breakpoints on segments toward four nearby already-open evaluation profiles with the opposite model prediction. This uses predicted answers, not true labels, and proves only an upper bound. Each accepted target is lifted, its PH recomputed, and its exact rational prediction and distance replayed. Lower bounds come from proved predicate invariance, Boolean coverage and checked joint refinement; unknowns never become UNSAT. No profile is dropped for being misclassified or expensive.

Table 2 reports the all-window results. All 720 final intervals have gap at most 10^{-5}; their median, 90th percentile and maximum gaps are $7.06 \cdot 10^{-6}$, $9.35 \cdot 10^{-6}$ and $9.98 \cdot 10^{-6}$. Of the 720 nominal model/window predictions, 548 are correct. At radius .01, per-seed BCA participant-bootstrap 95% intervals are [58.33,76.25], [60.42,78.33], [58.75,76.25]%, using 2,000 participant resamples. Each resample retains all windows of a participant; cross-seed comparisons must also preserve pairing. These are 20 people and three models, not 720 independent statistical trials.

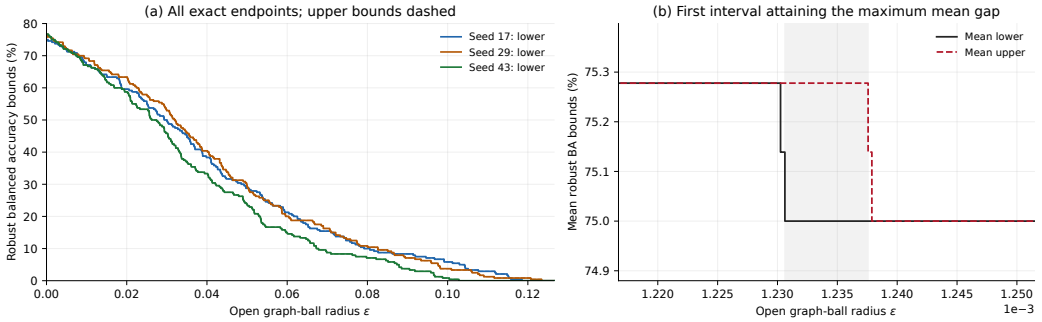

Figure 4: Lower and upper robust-BA bounds constructed from every exact endpoint. Left: separate frozen models; the bounds nearly overlap at this scale. Right: a magnified view of the first interval attaining the largest gap for the equally weighted three-model mean. The mean gap is at most $2/720$, or 0.277778 percentage points, over all radii, while each individual model's maximum is $2/240$, or 0.833333 percentage points. Neither denominator denotes independent experimental subjects. Plotting uses floating coordinates after all interval comparisons and summary statistics have been computed exactly.

D.5 Lower and upper robust-accuracy bounds

For the same 720 intervals in Table 2, we bound the robust balanced accuracy of each frozen model on the full evaluation set. For seed s, let $L_i \le r_i^\star \le U_i$ be an established lower endpoint and the distance to a verified opposite-prediction graph. For an open graph ball of radius $\epsilon > 0$, define

$$\underline{\mathrm{BCA}}_s(\epsilon) = \frac{1}{2}\sum_{c=0}^{1}\frac{1}{n_c}\sum_{i:y_i=c}\mathbf{1}[\hat{y}_i = y_i,\ L_i \ge \epsilon], \tag{13}$$

$$\overline{\mathrm{BCA}}_s(\epsilon) = \frac{1}{2}\sum_{c=0}^{1}\frac{1}{n_c}\sum_{i:y_i=c}\mathbf{1}[\hat{y}_i = y_i,\ U_i \ge \epsilon]. \tag{14}$$

Then $\underline{\mathrm{BCA}}_s(\epsilon) \le \mathrm{BCA}_s^\star(\epsilon) \le \overline{\mathrm{BCA}}_s(\epsilon)$. In particular, an attained opposite graph at $U_i < \epsilon$ rules out robustness of a correct center in the open ball. At $\epsilon = U_i$ that particular graph is on the excluded boundary, so the upper indicator uses $\ge$, not $>$. At zero we report nominal BA separately. Incorrect nominal predictions remain in the denominators and never contribute to either bound. The same inequalities hold after taking the arithmetic mean over the three fixed models.

An exact event sweep, rather than a coarse radius grid. All lower and upper endpoints are parsed as rational numbers. For each correct center its contribution to the envelope width is supported on $(L_i, U_i]$. We evaluate every distinct endpoint and every interval between consecutive endpoints; the curves are constant on intervals $(a, b]$. The independent checker also evaluates the endpoint b and an interior rational point directly from all 720 rows. Figure 4 shows the resulting bounds. Their maximum gap is $2/240$ (approximately .833333 percentage points) for each model and $2/720$ (approximately .277778 percentage points) for the three-model mean. The percentages here are rounded displays of exact rational differences; these deterministic widths measure verification uncertainty.

Table 9 reports selected radii and participant bootstrap intervals. At these radii, the two bounds coincide exactly; agreement only at these selected radii would not by itself establish the maximum-gap claim above.

D.6 A reconstructed EEG graph counterexample

The example is selected by a fixed rule: the first correctly classified case in ascending (seed, global window index) order. This yields seed 17, window 780, participant 66, with nominal class 0 (eyes open) and adversarial prediction 1. Selection does not maximize radius or visual

ϵ	Mean lower (%)	Mean upper (%)	Participant CI (%)
0.00	76.11	76.11	[67.78, 84.58]
0.01	67.78	67.78	[59.17, 76.81]
0.02	60.56	60.56	[51.66, 70.14]
0.05	27.36	27.36	[18.89, 36.39]
0.10	3.47	3.47	[0.28, 7.64]

Table 9: Selected exact envelope values and descriptive 95% percentile bootstrap intervals. Resampling uses 2,000 draws of 20 participants, keeping all 12 windows and the paired predictions of all three seeds for each sampled person. The lower and upper bootstrap intervals coincide at these radii. These pointwise intervals describe participant sampling variation and are separate from the deterministic proof/witness band; they are not a simultaneous confidence band or an independent confirmation.

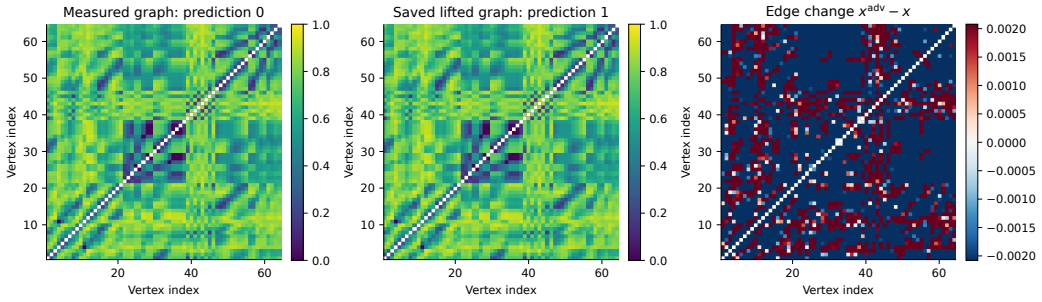

Figure 5: The measured graph and its saved lifted adversary, using the same 0–1 edge scale, followed by their difference with its own symmetric color scale. Vertex indices refer to the stored channel order, not named electrode locations. The nominal and adversarial graphs can look almost identical while the hard Boolean prediction differs. The permitted domain consists of independent edge births: the adversarial matrix is not claimed to be a realizable coherence matrix, EEG recording, or physiological intervention.

contrast. The saved interval is

$$[L, U] = [0.002073077635,\ 0.002078731578] \quad \text{(decimal display only)}.$$

An independent rational Kruskal pass on all 2,016 stored edge weights recovers all 1,953 target H_1 births exactly; an independent evaluator of all 32 strict affine predicates, 32 Boolean nodes, and the inclusive final vote verifies the opposite prediction. In exact arithmetic,

$$\|x^{\mathrm{adv}} - x\|_\infty = \|B(x^{\mathrm{adv}}) - B(x)\|_\infty = \|(1 - B(x^{\mathrm{adv}})) - (1 - B(x))\|_\infty = U.$$

The model export is expressed in ordered birth coordinates B; converting to lifetimes $1-B$ preserves this distance and requires the corresponding affine coefficient change if used as the model input. This is an equality for the saved witness, not a proof that U is the exact minimum adversarial radius. The witness changes 2016 edges and 1953 birth coordinates.

Figure 5 displays the nominal graph, its lifted counterexample and their difference. The witness record stores all rational edges, births, predicate margins, Boolean bits and the vote.

D.7 Computational cost

Verification runs on a shared CPU host with a 32-CPU quota, one numerical library thread per worker and concurrent cases. The measured times describe this execution environment. Shared feature and margin preparation precedes case processing; warm-witness construction, search and replay are accounted for separately below. Integer-scaled and Fraction arithmetic implement the same rational predicates and vote; final replay uses the Fraction evaluator.

The population search allowance is 900 seconds; ten unresolved intervals receive additional 3,600-second attempts, closing the remaining ten after the initial 710. Prior-evidence replay consumes the

new allowance. Final proof/lift replay and controller overhead are additional; the primary watchdog normally allows 120 extra seconds. A returned valid interval can remain unresolved. Allowances, stopping at tolerance, and actual time are distinct; these are not universal end-to-end time caps. Table 10 reports measured elapsed time, including failed and unresolved attempts, by operation. Shared preparation costs a further 24.26 seconds. For the 671 integer-evaluation initial cases, median/P90 elapsed time is 65.4/405.0 seconds; for the 25 Fraction-evaluation cases it is 469.4/598.9 seconds. The 24 cases also used for head controls have median/P90 68.8/196.2 seconds. These disjoint case groups do not support a paired comparison of arithmetic implementations.

Table 10: Computation for the 720 model/window certificates. Total is summed elapsed task time, including failed initial attempts, not CPU time or the wall time of the concurrent experiment. Range is over attempts. Warm witnesses and Boolean bounds are preparation costs, not additional model/window cases. Input recovery overlaps one initial case. Twelve initial-search timings use a worker clock rather than a valid controller clock.

Operation	Attempts	Total (s)	Range (s)
Initial DAG search	720	117403.2	8.0–949.7
Additional refinement	10	3568.9	86.7–825.4
Input recovery	1	87.5	87.5–87.5
Warm witnesses	720	9436.8	9.0–36.0
Boolean-bound preparation	1440	33602.5	12.7–39.2

For twelve cases (four Fraction and eight integer workers), elapsed time is measured by the worker pipeline clock rather than a valid controller clock. These measurements include worker computation but do not establish end-to-end deadline compliance. Invalid clock measurements are excluded; the twelve cases and their checked intervals remain included. Replay-only recovery is charged separately. Two failed exits occur in the 24-case head-control group; one case needs the separately reported input recovery. Boolean preparation has one failed exit. No case is discarded. Replay-only resumptions are additional to the initial-attempt totals: 85 Boolean-preparation resumptions have median/P90 16.3/20.9 seconds, and 101 integer-search resumptions have median/P90 16.0/25.5 seconds. Search allowance and checked interval quality remain separate from these timing measurements.

D.8 What the certificate protects

The perturbed input is an independently variable edge-birth vector in the fixed graph cube. No constraint forces a perturbed coherence matrix to arise from an EEG signal or estimator. Thus neither signal-space robustness, physiological plausibility nor clinical validity follows from these graph balls. A stable erroneous prediction remains an error and contributes zero to BCA. Rational conflicts and witnesses are checked; Boolean and strict-SMT UNSAT rely on the backend (Appendix C.1).

E Controlled validation of the reduction

E.1 Algorithm names and computation accounting

A predicate is a Boolean function; its real-valued affine expression before thresholding is a predicate score. We use “clause” for a DNF conjunction, and “input counterexample” only for a reconstructed, checked input. A solver assignment alone is not such a counterexample. Table 11 distinguishes the compared procedures; Table 12 specifies which costs enter each comparison. EEG uses the separate contracts in Appendix D.7.

Exact denotes a proved decision or optimum of the stated problem, not merely a completed run. Repeated timings and graph copies are not independent datasets. Counts of solved tasks and radius quality remain separate from elapsed time.

Table 11: Fixed algorithm vocabulary. Each comparison uses the same predictor, admissible domain and input metric within its own matrix.

Name	Computation
Independent-fact bound	R_{Bool}, discarding joint predicate dependencies
Ordered SMT	Opposite-answer search over the exact ordered image
Raw-input SMT	Opposite-answer search retaining graph-input variables
Lazy refinement	Exclude unrealizable opposite fact assignments using checked constraints
Projected SMT	Ordered SMT after exact elimination of unread coordinates

Table 12: Timing contracts. Ratios are interpreted within a row, not across rows; a separate audit is not silently included in a search budget.

Matrix	Charged	Separate/excluded
690 graph/query cases	Encoding, search, lift, replay	Imports; second audit
Coordinate projection	Encoding, search, completion, PH replay	Common imports/startup, report writes

E.2 Exact verification of ordered interval cells

We use designed query families to examine when exact reduction changes computational cost or improves an independent-fact bound. Workloads reflect exploratory analysis; configurations and input/model identities are fixed before comparative timing. The resulting comparisons are conditional on these workload families.

A finite-state comparator. Let $c \in \mathcal{C}_\beta$ be the nominal lifetime vector. In this subsection only, the F facts are coordinate tests $[z_{k_j} \bowtie t_j]$, with $\bowtie \in \{>, \geq\}$. General affine predicates use the SMT path, not this interval-cell algorithm. Each truth assignment induces coordinate intervals I_i (possibly open at either endpoint). Repeated tests of the same coordinate are intersected. Write l_i, u_i for the resulting endpoints and $S = (\prod_i I_i) \cap \mathcal{C}_\beta$. Empty intervals or order-infeasible assignments are discarded. Structural padding is the singleton $\{0\}$.

Proposition 8 (Distance to a nonempty ordered interval cell)**.** *For any nonempty S of this form and decreasing c,*

$$\inf_{z \in S} \|z - c\|_\infty = \max\bigl(0, \max_i(l_i - c_i), \max_i(c_i - u_i)\bigr) =: \delta_S.$$

The infimum is attained exactly when S intersects the closed ball of radius δ_S. Minimizing over opposite-answer assignments gives the exact infimal flip distance; the empty minimum is $+\infty$.

Proof. The coordinate distances give the lower bound. For $i \leq j$, nonempty ordered feasibility gives $u_i \geq l_j$. At $\delta = \delta_S$, each of the four comparisons between $u_i, c_i + \delta$ and $l_j, c_j - \delta$ holds: use original feasibility, the coordinate distance bounds, and $c_i \geq c_j$ respectively. Thus the closed intervals $[l_i, u_i] \cap [c_i - \delta, c_i + \delta]$ have an ordered feasible point. Because S is a nonempty convex intersection of open/closed halfspaces, every point of this closed relaxation is approached from S by interpolation with any fixed point in S. Its distance is therefore the infimum even when strict facets exclude the minimizing point. Feasibility with the original strictness flags at δ determines attainment. Finally, a finite union's distance is the minimum of its cell distances. □

The feasibility algorithm propagates suffix maxima of lower endpoints; when endpoints tie, a strict lower endpoint propagates strictness. Each upper endpoint must exceed that propagated bound, or equal it with both endpoints closed. A witness is obtained by increasing strict lower endpoints by half the smallest positive gap between distinct endpoints and then taking suffix maxima. These operations respect every upper endpoint. Enumeration costs $O(2^F(\beta \log \beta + |q|))$ with this implementation, where $|q|$ is formula length. It is exponential in the vocabulary, not a polynomial-time verifier for arbitrary programs. Appendix E.3 evaluates the scaling implementation up to $F = 16$ under explicit time limits.

No SMT solver is used on the enumeration path. Its target vector is lifted by Theorem 1 and replayed with union-find and an independent $\mathbb{F}_2$ boundary computation. External PH replay compares

identical rounded inputs. To report *closed*-ball lower bounds in the coordinate-threshold controls, an attained infimum δ is reported with lower bound $\max(0, \delta - 10^{-4})$; an unattained infimum itself is a safe closed bound. The exact open-ball radius, attainment flag and actual counterexample distance are exported separately.

Raw-input comparator. For every coordinate threshold $L_k > t$, the raw encoding uses the cycle rank of edges with birth below $1 - t$ (Proposition 2), retaining every input edge and every shared threshold dependency. Inclusive predicates retain their inclusive comparisons. It does not query independently perturbed facts.

Exact connectivity constraints and cost. For each threshold subgraph, vertex i has a representative $r_i \in \{0, \dots, i\}$ and a nonnegative integer depth below n. Active edges equate representatives. If $r_i = i$, depth is zero; otherwise i must have an active neighbor with the same representative and smaller depth. Descending chains reach their uniquely named root, so disconnected components cannot be merged. Active edges prevent splitting a component. Conversely, minimum-index representatives and a spanning forest provide a satisfying assignment for every actual subgraph. The number of roots is therefore its exact component count. Raw-edge SMT uses this exact connectivity encoding. With F distinct threshold tests, the implemented explicit cell-membership disjunctions require $O(Fn + F^2 m_E)$ construction work up to formula representation costs. This does not bound SMT search time. Reported method time includes encoding and input replay; nominal search allowances and actual elapsed time are distinguished in Appendix E.3.

E.3 Frozen competing-query and logic-scaling protocols

The verification protocol fixes 690 graph/query cases and 4,140 method/budget arms before comparative evaluation. It keeps the same filtration, metric, cap $1/8$ and tolerance 10^{-4}. Queries read four lifetime coordinates; all remaining cycle-rank coordinates remain part of the ordered domain. Graph family changes therefore need not keep the full cycle rank equal to four.

Competing clauses. Ten seeds each define four generated signed-DNF cases and four constructed proof-switch cases, each realized on three graph families and two sizes. These are 80 program/input identities and 480 graph/query cases, not 480 independently sampled programs. Structural checks reject constants, impossible or redundant clauses and functions equivalent on the feasible domain to a single conjunction. Distinct thresholds and graph realizations do not imply distinct syntactic DNF structures. In this cohort there are 44 syntactic DNF structures, 80 fully specified threshold formulas and 48 feasible truth tables. The 120 logic-scaling formulas are distinct; graph scaling reuses ten formulas across ninety graph realizations.

The constructed stratum uses the consensus-sensitive three-fact formula $(a \wedge c) \vee (b \wedge \neg c)$, with switching margins .005–.014 and positive-evidence margins .1. It is not the four-predicate construction of Proposition 1; Boolean rewriting can remove its fixed-clause obstruction. A closed radius-.04 ball preserves the answer but no single clause. It uses three active facts from eight declared thresholds; we report these counts separately. The generated stratum uses two to four clauses of two to four signed literals, and no outcome-based selection. All generation attempts and feasible truth-table identities are retained. Seeds 100–101 are structural pilots, separate from the ten evaluation seeds.

Logic and graph scaling. The logic matrix varies the declared fact count $F \in \{4, 8, 12, 16\}$ and clause count in $\{2, 4, 8\}$, with ten seeds and a fixed theta graph. Its clauses specify distinct feasible full fact states and use all F facts; their lengths are F, unlike the shorter clauses of the competing cohort. The graph matrix fixes eight facts and four clauses, and varies three families and sizes 16, 32 and 64. Enumeration, ordered SMT and raw-edge SMT receive the same nominal .1- and one-second search allowances. Enumeration remains exponential in F; timed-out partial enumeration returns no positive global lower bound. An exact reference is computed separately only for $F \leq 8$, not supplied to the timed methods.

Execution and uncertainty. Each timed arm is an isolated process with one thread and a 4-GiB address-space limit. Process-start imports are excluded equally; encoding, verification, lift and replay are included in the method time. A second independent audit after return is excluded. Timings are descriptive measurements on a shared host with fixed processor affinity, not exclusive-hardware

benchmarks. The SMT implementation reports some solve/replay costs inclusively; absent component timers are not inferred by subtraction. Larger budgets do not replace smaller budget results. Recorded method time can exceed the nominal allowance; the timing columns report actual elapsed computation rather than asserting hard deadline compliance. All timeouts and failures remain in the accounting. Repeated graph realizations are paired; seed means, rather than individual graph copies, provide the uncertainty units.

Table 13: Enumeration minus ordered-SMT lower radius, paired by seed. Intervals are descriptive 95% Student intervals over ten seed means, not multiplicity-adjusted significance tests. Repeated graph realizations are averaged within each seed. Tiny one-second differences are due to finite tolerance and boundary handling, not unequal exact distances.

Cohort	Budget (s)	Mean difference	95% interval
Generated	.1	.008624	[.005966,.011281]
Generated	1	-.000009	[-.000029,.000012]
Proof-switch	.1	.032760	[.025010,.040511]
Proof-switch	1	-.000076	[-.000076,-.000076]
16 facts, 8 clauses	.1	-.021979	[-.037024,-.006934]
16 facts, 8 clauses	1	.000002	[-.000028,.000033]

Table 13 shows that increasing the budget closes the competing-query gap between enumeration and ordered SMT, while short-budget logic scaling can favor SMT. The degenerate interval for the constructed one-second stratum reflects identical bounds across seeds, not certainty about a broader query population.

Results across regimes. Table 14 retains every graph/query case and both budgets. Aggregation pools only the clause counts at each fact count or the families at each graph size. It exposes the increasing cost of raw graph variables and the short-budget crossover between enumeration and SMT. For the most complex logic stratum, 16 facts and 8 clauses, enumeration and ordered SMT take 317.4 and 122.3 milliseconds at the one-second budget; at .1 seconds their timeout counts are 10/10 and 4/10. Thus the pooled logic rows do not imply a uniform preference for enumeration.

Table 14: Controlled reduction across both budgets and all scaling regimes. Each method cell gives mean lower radius / mean milliseconds / timeout count. Logic rows pool the three clause counts; graph rows pool the three families, with equal counts per stratum. Aggregates use the recorded stratum means; all cases, including timeouts, remain in the denominators.

Cohort	Budget (s)	Cases	Enumeration	Ordered SMT	Raw SMT
Generated	0.1	240	0.0200 / 51.9 / 19	0.0114 / 97.5 / 180	0.0000 / 130.0 / 240
Generated	1	240	0.0215 / 51.1 / 0	0.0215 / 152.7 / 0	0.0169 / 848.9 / 115
Proof-switch	0.1	240	0.0987 / 46.2 / 3	0.0659 / 96.4 / 175	0.0000 / 127.4 / 240
Proof-switch	1	240	0.0999 / 44.7 / 0	0.1000 / 137.0 / 0	0.0926 / 808.7 / 95
Logic: 4 facts	0.1	30	0.0231 / 9.7 / 0	0.0199 / 87.3 / 17	0.0000 / 117.5 / 30
Logic: 4 facts	1	30	0.0231 / 10.7 / 0	0.0232 / 110.7 / 0	0.0232 / 408.3 / 0
Logic: 8 facts	0.1	30	0.0296 / 11.3 / 0	0.0168 / 93.5 / 20	0.0000 / 119.7 / 30
Logic: 8 facts	1	30	0.0296 / 14.8 / 0	0.0297 / 128.4 / 0	0.0297 / 671.3 / 2
Logic: 12 facts	0.1	30	0.0307 / 34.2 / 1	0.0197 / 95.8 / 22	0.0000 / 117.8 / 30
Logic: 12 facts	1	30	0.0309 / 46.2 / 0	0.0310 / 116.5 / 0	0.0309 / 948.5 / 11
Logic: 16 facts	0.1	30	0.0026 / 100.0 / 28	0.0221 / 89.1 / 16	0.0000 / 113.7 / 30
Logic: 16 facts	1	30	0.0350 / 241.4 / 0	0.0351 / 125.5 / 0	0.0277 / 1033.4 / 28
Graph: 16 vertices	0.1	30	0.0326 / 13.4 / 0	0.0181 / 94.7 / 20	0.0000 / 119.0 / 30
Graph: 16 vertices	1	30	0.0326 / 12.0 / 0	0.0327 / 135.5 / 0	0.0326 / 808.7 / 6
Graph: 32 vertices	0.1	30	0.0326 / 17.9 / 0	0.0159 / 100.0 / 23	0.0000 / 141.8 / 30
Graph: 32 vertices	1	30	0.0326 / 16.9 / 0	0.0327 / 202.7 / 0	0.0102 / 1038.8 / 30
Graph: 64 vertices	0.1	30	0.0326 / 23.0 / 0	0.0060 / 109.0 / 30	0.0000 / 185.4 / 30
Graph: 64 vertices	1	30	0.0326 / 27.3 / 0	0.0327 / 232.8 / 0	0.0000 / 1108.9 / 30

E.4 A QUERY-COORDINATE PROJECTION CONTROL

We freeze a constructed matrix before outcomes: cycle ranks $4, 16, 64$, three seed indices, formulas reading either two or all coordinates, and closed decision radii $.01, .025, .05$. Three timing repetitions and full/projected SMT give 324 records. The frozen formula is a disjunction of consecutive

two-literal clauses, alternating positive and negative strict coordinate thresholds at distance .025 from the nominal values. Inputs use a shared-vertex square graph, with a separate exact boundary-rank replay for each SAT input witness. These are supplied formulas, not learned programs or natural-data tasks.

All 324 decisions finish, full and projected answers agree, and all 216 SAT witnesses replay. On the nine sparse-readout input identities, the median relative reduction of median decision time is 55.2%; the corresponding all-coordinate control is 0.5%. Each identity aggregates the same three radii and three repeats. Across the sparse stratum, charged time totals 1.164 seconds for full SMT and .456 seconds for projected SMT. These shared-host, millisecond-scale timings are descriptive; seed indices repeat across cycle ranks. Their interpretation is limited to these fixed workloads and execution conditions.

The geometric conclusion is exact variable elimination with preserved answers and reconstructible witnesses. The observed timing benefit depends on unused coordinates; it does not improve the exact radius.

F Further theoretical results and scope boundaries

F.1 Anchored isometric sections and full H_0 profiles

For the fixed graph let $r = |V| - c(G)$ and $\beta = |E| - r$. Write $T_k = \{0 \le a_1 \le \cdots \le a_k \le 1\}$. Kruskal's pass accepts r and rejects β edges. Their nondecreasing weight vectors are respectively $D(x) \in T_r$ and $B(x) \in T_\beta$. D contains the graph-edge H_0 merge deaths, including zeros; $B = 1-L$ contains the full formal H_1 births. For a disconnected graph, terminal completion adds $c(G)-1$ fixed H_0 deaths at one, which can be appended without changing distances. Essential components are not variable coordinates. Throughout this subsection and the next, $\|\cdot\|_\infty$ is the coordinate metric, rather than diagram matching to the diagonal.

Both D and B are 1-Lipschitz: accepted and rejected counts below a threshold are the graphic-matroid rank and nullity of the threshold edge set. Threshold inclusions for inputs at distance δ shift these counts by at most δ, so every corresponding order statistic moves by at most δ. Concatenating selected coordinates preserves this bound. For a fixed classifier on H, write $R_G(x)$ for its input flip infimum and $R_{T_k}(H(x))$ for its feature flip infimum over T_k.

Let $H = B$ and $k = \beta$ for full H_1, or $H = D$ and $k = r$ for full H_0. Fix a nondecreasing nominal edge order $e_1, \dots, e_m$, resolving ties once, and let $j_1 < \cdots < j_k$ be its rejected positions for B or accepted positions for D. For H_1, this strengthens Theorem 1 by fixing one family of preimages that preserves all pairwise distances.

Theorem 2 (Anchored single-profile lift and exact certificates)**.** *For every $a \in T_k$, let $\delta = \|a - H(x)\|_\infty$ and set*

$$s_x(a)_{e_j} = \max\big(\{0, x_{e_j} - \delta\} \cup \{a_i : j_i \le j\}\big). \tag{15}$$

Then $H(s_x(a)) = a$, $s_x(H(x)) = x$, and

$$\|s_x(a) - s_x(a')\|_\infty = \|a - a'\|_\infty \quad (a, a' \in T_k).$$

Consequently, for every $\rho \ge 0$,

$$H\big(\{y \in [0,1]^E : \|y - x\|_\infty \le \rho\}\big) = T_k \cap \{a : \|a - H(x)\|_\infty \le \rho\}, \qquad R_G(x) = R_{T_k}(H(x)). \tag{16}$$

The final equality holds for any fixed classifier on H, independently of how it was trained.

Proof. The constructed edge values are nondecreasing in j. At a selected position j_i, all preceding targets are at most a_i and $x_{e_{j_i}} - \delta \le a_i$, so the maximum is a_i. For any j, every included target obeys $a_i \le x_{e_{j_i}} + \delta \le x_{e_j} + \delta$. Thus the new vector lies in the cube and within δ of x. The old edge order remains a valid sorted order; Kruskal makes the same combinatorial choices in that order. Alternative refinements of new ties give the same numerical accepted and rejected profiles, as seen from the threshold rank/nullity counts. Hence $H(s_x(a)) = a$.

As functions of a, both δ and each target coordinate are 1-Lipschitz. The maximum in (15) is consequently 1-Lipschitz, while the coordinate stability established above supplies the reverse inequality

after applying H. The anchor identity follows by setting $\delta = 0$. Stability gives one ball inclusion; the lift gives the other. Taking the infimum over opposite-label targets yields equality of flip radii, including cases where that infimum is unattained. □

Operationally, a nearest or nearly nearest opposite-label H_1 or H_0 vector can be turned into a graph at exactly the same distance. The theorem removes geometric conservatism; it does not guarantee that an optimization solver finds that vector within a time budget. The input domain remains the full edge cube.

F.2 Joint H_0/H_1 under order consistency

Consider $h(x) = (D(x), B_{k_1}(x), \ldots, B_{k_s}(x))$, where $1 \le k_1 < \cdots < k_s \le \beta$. The product $A = T_r \times T_s$ is an outer domain: it keeps ordering within each profile but omits constraints between them. Write $R_A(h(x))$ for the flip infimum over A. Stability gives the sound lower bound $R_A(h(x)) \le R_G(x)$. A target consistent with the nominal interleaving has an exact lift. Fix the nominal edge order and its accepted/rejected positions as above.

Theorem 3 (Selected-coordinate lift on a nominal order face)**.** *Let each selected coordinate $h_i(x)$ correspond to nominal edge position p_i. On*

$$F_x = \{z \in [0,1]^{r+s} : z_i \le z_k \text{ whenever } p_i < p_k\},$$

define $\delta = \|z - h(x)\|_\infty$ and

$$y_{e_j} = \max\big(\{0, x_{e_j} - \delta\} \cup \{z_i : p_i \le j\}\big). \tag{17}$$

Then $h(y) = z$ and $\|y - x\|_\infty = \delta$. In particular, every opposite-label $z \in F_x$ supplies a graph witness satisfying

$$R_A(h(x)) \le R_G(x) \le \|z - h(x)\|_\infty . \tag{18}$$

Proof. Order consistency makes the maximum equal z_i at position p_i. The sequence is nondecreasing, and every included target is bounded above by $x_{e_j} + \delta$, exactly as in Theorem 2. Thus the nominal Kruskal order realizes the selected coordinates, with distance at most δ. Stability makes that distance equal to δ. The outer-domain inclusion proves the lower bound. □

The face retains the nominal interleaving of the selected H_0 and H_1 coordinates. Searching only this face can miss a closer counterexample in another order region. Thus the displayed interval is sound, but does not assert global equality of graph and joint-feature flip radii.

F.3 A boundary of unrestricted joint exactness

The order restriction is substantive: even a globally attainable target can cost more to realize than its distance in the concatenated feature metric.

Proposition 9 (Why separate sorted cones do not suffice)**.** *Even a feasible full joint target can require more graph movement than its PH distance. Consider a triangle with edge weights $(0, .8, .9)$ and a pendant bridge of weight .1. Its joint profile is $D = (0, .1, .8)$, $B = (.9)$. The target $D' = (0, .2, .8)$, $B' = (.3)$ has PH distance .6, but its nearest graph preimage is at distance .7.*

Proof. The largest triangle edge is the unique rejected edge, so all target triangle weights must be at most .3. Therefore the .8 accepted weight must belong to the bridge, requiring movement .7. This is attainable by assigning triangle weights $(0, .2, .3)$ in the nominal order and bridge weight .8. The lower and upper costs agree. □

F.4 Co-ordered interval boxes and query scope

The following sufficient condition extends uniform input balls without claiming an image formula for arbitrary unequal edge intervals. The graph, terminal completion and zero-padded lifetime conventions are those of Theorem 1.

Proposition 10 (Exact image of a co-ordered box)**.** *Index edges in an arbitrary fixed order $1, \dots, m_E$. Let $0 \le l_j \le u_j \le 1$ with both $(l_j)_j$ and $(u_j)_j$ nondecreasing. Let $j_1 < \cdots < j_\beta$ be the cycle-closing positions of a greedy spanning-forest pass in that order. For $A = \prod_j [l_j, u_j]$,*

$$L(A) = \{z \in \mathcal{C}_\beta : 1 - u_{j_i} \le z_i \le 1 - l_{j_i} \text{ for every } i\}.$$

Proof. For an edge-birth vector v, let $b_i(v)$ be the ith smallest cycle-closing birth, including births at one. At threshold t, the number of such births at most t is the cycle rank of the graph of edges with $v_j \le t$. Adding edges never decreases cycle rank. Consequently $v \le w$ coordinatewise implies $b_i(v) \le b_i(w)$. Since both endpoint vectors follow the chosen order, including any tie refinement by that order, $b_i(l) = l_{j_i}$ and $b_i(u) = u_{j_i}$. Every $y \in A$ therefore satisfies the displayed endpoint inequalities.

Conversely, let z belong to the displayed set and put $b_i = 1 - z_i$. Set

$$y_j = \max\big(\{l_j\} \cup \{b_i : j_i \le j\}\big).$$

Each included b_i is at most $u_{j_i} \le u_j$, so $l_j \le y_j \le u_j$. Both the running maximum and l_j are nondecreasing. At j_i, earlier b_k are at most b_i and $l_{j_i} \le b_i$; hence $y_{j_i} = b_i$. The same edge order is valid for y, so the greedy pass has the same cycle-closing positions and $L(y) = z$. New ties do not affect the multiset of cycle-closing births. Disconnected components and isolated vertices only change the cycle rank. Births at one give the required zero coordinates. If $\beta = 0$, both sides contain just the empty vector and the construction reduces to $y = l$. □

Taking $l_j = \max(0, x_j - \epsilon)$ and $u_j = \min(1, x_j + \epsilon)$ in nominal edge order proves the uniform-ball image theorem. Choosing $\epsilon = \|z - L(x)\|_\infty$, and combining the construction with the forward coordinate bound, proves the anchored minimum-distance identity. This establishes a sufficient larger box class; it is not a characterization of all boxes with an exact interval image. The unequal-box counterexample in the main text remains outside its hypotheses.

F.5 The order construction for a finite matroid

The graph proof uses the nullity profile of its graphic matroid. The same construction holds for any fixed finite matroid; this identifies the combinatorial scope of the argument rather than asserting that persistence fiber geometry is new.

Let a finite matroid on a ground set E have rank function rk, and put $\beta = |E| - \mathrm{rk}(E)$. For weights $x \in [0,1]^E$, process elements in nondecreasing weight order, using a fixed tie refinement, and keep an element exactly when it increases the rank of the previously kept set. Denote the weights of the β rejected elements by $b_1(x) \le \cdots \le b_\beta(x)$ and set $L_i^{\mathsf{M}}(x) = 1 - b_i(x)$. This is a nullity-profile statistic; for a general matroid it is not being identified with a persistence diagram.

Proposition 11 (Ordered image of a matroid nullity profile)**.** *For the clipped uniform ℓ_∞ ball $\mathcal{B}(x, \epsilon)$,*

$$L^{\mathsf{M}}(\mathcal{B}(x,\epsilon)) = \{z \in \mathcal{C}_\beta : \|z - L^{\mathsf{M}}(x)\|_\infty \le \epsilon\}.$$

Every target has a preimage at its distance from $L^{\mathsf{M}}(x)$. The co-ordered-box identity of Proposition 10 also holds with the rejected positions of this greedy pass.

Proof. For any subset A, the greedy pass accepts a basis of A when it has processed precisely A, hence rejects $|A| - \mathrm{rk}(A)$ elements. This nullity is nondecreasing under inclusion because adjoining an element increases rank by zero or one. Threshold-set inclusions for two weight vectors at distance at most ϵ therefore show $|b_i(x) - b_i(y)| \le \epsilon$. They also show that tied blocks have a tie-independent multiset of rejected weights.

Fix the nominal element order and its rejected positions j_i. For a proposed ordered target, put $a_i = 1 - z_i$ and assign

$$y_j = \max\big(\{l_j\} \cup \{a_i : j_i \le j\}\big).$$

When l, u follow that order and $a_i \in [l_{j_i}, u_{j_i}]$, the assignment is nondecreasing, belongs to $[l, u]$ and satisfies $y_{j_i} = a_i$. The greedy decisions depend only on the ordered elements and the matroid, so the same pass rejects the same positions. This proves reverse inclusion. Uniform clipping supplies co-ordered endpoints; applying it at $\epsilon = \|z - L^{\mathsf{M}}(x)\|_\infty$ and using the forward bound proves minimum distance. When $\beta = 0$, use the unique empty vector and distance zero. □

The construction requires a sort and at most $|E|$ independence-oracle queries to identify the rejected positions, followed by linear reconstruction. This is not a polynomial-time bound when the oracle itself is expensive. For a graphic matroid, the oracle pass is implemented by a spanning forest. The argument does not extend a graph's filtration to early two-dimensional fillings or variable deaths: these need a different feature map.

Relation to inverse persistence. Order-invariant barcode templates and local barcode coordinates are established in Leygonie et al. (2022); polyhedral fibers and their reconstruction are studied in Leygonie & Tillmann (2022); Leygonie & Henselman-Petrusek (2024). Here the quantitatively specified property is the image of every clipped uniform ball and distance to the fiber of a prescribed nullity profile, anchored at an arbitrary input, including ties. It follows from the explicit order construction above. We neither infer this equality for general barcodes nor claim priority for a general inverse-persistence principle.

F.6 Observation constraints and minimum-cost preimages

The ordered-box identity requires its declared Cartesian input domain. Side information instead gives a feasible set of inputs; it must not be replaced by one fitted reconstruction. This subsection gives an exact finite-cell extension, not another ordered-box identity. It specializes the polyhedral geometry of persistence fibers (Leygonie & Tillmann, 2022; Leygonie & Henselman-Petrusek, 2024) to our output vector and an input-anchored cost.

Keep the fixed graph, terminal-completion filtration and zero-padded lifetime map L of Theorem 1. Let $\mathcal{U} \subseteq [0,1]^{m_E}$ be a nonempty compact rational polytope, $x_0 \in \mathcal{U}$, and W a diagonal matrix with positive rational diagonal entries. Define $\delta_W(x_0, x) = \|W(x - x_0)\|_\infty$. Weights and units are fixed before comparison. A compact extended formulation with real latent variables and rational constraints is allowed; the cost and L depend only on the edge coordinates x. Geometric identities below allow real x_0 and target z. Rational algorithms and bit-complexity statements additionally require rational x_0, z, whose encoding lengths count toward the input size.

$$\text{observations and justified bounds} \longrightarrow \text{admissible inputs } \mathcal{U} \longrightarrow \text{exact pairs } \{(x, L(x)) : x \in \mathcal{U}\} \longrightarrow \inf_{x \in \mathcal{U}:\, q(L(x)) \neq q(L(x_0))} \delta_W(x_0, x).$$

Figure 6: Observation-constrained verification of the same frozen query. Physical containment in $\mathcal{U}$ requires external justification; the geometric result is conditional on this declared set. Its order constraints can reduce the number of required cells, but not improve the true flip-distance infimum relative to an exact solver given the same information.

Figure 6 summarizes the constrained-input verification problem; its feasible relation keeps the observation set intact. For an edge permutation π, let $C_\pi = \{x \in [0,1]^{m_E} : x_{\pi(1)} \leq \cdots \leq x_{\pi(m_E)}\}$. Run the greedy forest pass in this order and let $j_1 < \cdots < j_\beta$ be the cycle-closing positions. Put $A_\pi(x)_i = 1 - x_{\pi(j_i)}$ and $\mathcal{U}_\pi = \mathcal{U} \cap C_\pi$.

Lemma 1 (Closed-cell lifetime map)**.** *For every $x \in C_\pi$, $L(x) = A_\pi(x)$, including tied births, terminal zero coordinates and $\beta = 0$.*

Proof. Before time one, an edge creates an H_1 class exactly when it closes a cycle in the forest pass; no class dies before terminal completion. Inside a tied birth block, the number of cycle-closing edges is the cycle-rank increment across the block, independent of its internal order. All these edges have the same birth, so the sorted birth multiset is unchanged by tie refinement. Rejected edges born at one supply zero coordinates; terminal completion creates no additional positive intervals. Consequently $A_\pi(x) = A_\tau(x)$ on closed-cell overlaps. This is equality of lifetime vectors, not of named persistence generators. □

Proposition 12 (Constrained image and anchored distance)**.** *Write* $\operatorname{graph}(L|_{\mathcal{U}})$ *for* $\{(x, L(x)) : x \in \mathcal{U}\}$*. Then*

$$\operatorname{graph}(L|_{\mathcal{U}}) = \bigcup_\pi \{(x, z) : x \in \mathcal{U}_\pi,\ z = A_\pi(x)\}.$$

For a supplied target, the restricted fiber distance is

$$D_{x_0,\mathcal{U}}(z) := \inf_{x\in\mathcal{U}:L(x)=z} \delta_W(x_0,x) = \min_{\pi} \min_{\substack{x\in\mathcal{U}_\pi \\ A_\pi(x)=z}} \delta_W(x_0,x),$$

with value $+\infty$ for an empty fiber or empty inner feasible set. Every finite minimum is attained. For rational x_0, z, each inner problem is a rational linear program; every feasible inner problem admits a rational minimizing input.

Proof. Every input has a nondecreasing edge order, and Lemma 1 gives the exact map on every containing cell. The union therefore neither removes inputs nor adds spurious pairs. For fixed z, each feasible set is closed in compact $\mathcal{U}$. The continuous distance attains its minimum there; a finite union preserves attainment. Introduce $t \geq 0$ and constraints $-t \leq w_e(x_e - x_{0,e}) \leq t$ for every edge. Minimizing t with the cell, domain and target equalities is an exact LP formulation, not a relaxation. With rational data a finite optimal LP has a rational optimizer. □

How much order information is sufficient? A pair e, f is *reversing* if $\mathcal{U}$ contains an input with $x_e < x_f$ and another with $x_e > x_f$. Let S contain all endpoints of reversing pairs and let $b_{\mathrm{ord}} = |S|$. The subscript distinguishes this edge count from a Boolean fact vector. Refine permanently equal coordinates by a fixed index order. All other nonreversing comparisons are universal weak orders on $\mathcal{U}$; after identifying permanently equal coordinates, they form a partial order.

Proposition 13 (Bounded order ambiguity)**.** *Linear extensions of these fixed comparisons supply a covering family of at most $b_{\mathrm{ord}}!$ closed cells for $\mathcal{U}$.*

Proof. Sort any feasible input numerically, refining numerical ties by a linear extension of the fixed comparisons. This covers every input. Every pair with an endpoint outside S has a fixed comparison. Thus two total extensions inducing the same relative order on S agree on every pair and coincide. Restriction to S is injective, giving at most $b_{\mathrm{ord}}!$ extensions. Infeasible extensions may be discarded. □

This counts a sufficient selected cover, not every closed cell meeting $\mathcal{U}$: all $m_E!$ closed cells contain an all-equal input. Let T_{LP} bound the rational LP cost, including coefficient bit lengths, nominal point, target and metric. Pairwise order certification and fixed-target lifting cost conservatively

$$O\big(m_E^2 T_{\mathrm{LP}} + b_{\mathrm{ord}}!(m_E^2 + T_{\mathrm{LP}})\big).$$

This is factorial in b_{ord}, not polynomial; partial-order width is not a substitute, since two long chains already have binomially many interleavings. An unresolved comparison must remain ambiguous, giving a conservative $\widehat{b}_{\mathrm{ord}} \geq b_{\mathrm{ord}}$. Enumerating all $m_E!$ permutations does not earn this parameterized bound. Nor is this a complexity bound for arbitrary queries.

Same answer, metric and physical scope. For any frozen Boolean query q, grouping opposite-answer inputs by z gives

$$\inf_{\substack{x\in\mathcal{U} \\ q(L(x))\neq q(L(x_0))}} \delta_W(x_0,x) = \inf_{z:q(z)\neq q(L(x_0))} D_{x_0,\mathcal{U}}(z).$$

An opposite-answer infimum need not be attained. A radius no greater than this infimum certifies the open ball relative to $\mathcal{U}$; $R = 0$ remains the no-positive-certificate sentinel. For a convex affine-query branch containing strict inequalities, establish strict feasibility after intersection with $\mathcal{U}$ and the current order cell, before optimizing its weak closure. Interpolation with a strictly feasible point proves equality of the infima. A boundary optimizer is not automatically an opposite-answer witness. The formulation retains input variables in its LPs; it does not assert that arbitrary side information permits their elimination into an ordered box. Whether measurements justify $\mathcal{U}$, including attacked sensors and storage dynamics, is an external assumption, not a consequence of these propositions or a good fit.

F.7 A SMALLER EXACT REPRESENTATION: SPANNING-TREE REGIONS

The order-cell extension need not retain comparisons that cannot affect the lifetime map. Here we give a coarser exact representation under the same terminal-completion filtration. The minimum-spanning-tree criterion is classical; polyhedral persistence fibers and their reconstruction are established (Leygonie & Tillmann, 2022; Leygonie & Henselman-Petrusek, 2024). The following formulation preserves the minimum input cost while using fewer regions when the cycle rank is small.

Let $G = (V, E)$ be connected, $m_E = |E|$, and $\beta = m_E - |V| + 1$. For a spanning tree $T \subseteq E$ and $e \in E \setminus T$, let $P_T(e)$ be the unique path in T joining the endpoints of e. Let $\tau = (e_1, \dots, e_\beta)$ order $E \setminus T$. Define the closed region

$$C_{T,\tau} = \mathcal{U} \cap \{x_f \le x_e : e \notin T,\ f \in P_T(e)\} \cap \{x_{e_1} \le \cdots \le x_{e_\beta}\}.$$

Let $t(G)$ denote the number of spanning trees. Latent variables in a compact rational extended formulation of $\mathcal{U}$ are retained.

Proposition 14 (Cost-preserving spanning-tree cover)**.** *The regions $C_{T,\tau}$ cover $\mathcal{U}$. On each region, $L_i(x) = 1 - x_{e_i}$, including tied and terminal births. The affine maps agree on overlaps. For rational nominal input and target, the minimum weighted ℓ_∞ cost of a target preimage equals the minimum of the corresponding region LPs. The number of candidate regions is at most*

$$t(G)\beta! \le \binom{m_E}{\beta}\beta! \le m_E^\beta.$$

Proof. The fundamental-cycle criterion says that T is a minimum spanning tree exactly when no complementary edge can replace a heavier edge on its tree path. These are the displayed path inequalities. Every input admits an MST and an ordering of its complementary edges, proving coverage. Kruskal's forest pass rejects precisely the complementary birth multiset. Each rejection births an H_1 class, and all classes die at terminal completion. With tied weights, every MST contains the same number of edges in each weight prefix: the rank of that prefix in the graphic matroid. Thus the accepted and complementary weight multisets do not depend on the chosen MST. This proves the affine map and overlap consistency without identifying named cycles across regions. For a fixed target z, intersecting a region with $z_i = 1 - x_{e_i}$ is a closed subset of compact $\mathcal{U}$. The usual epigraph inequalities for the weighted distance give an LP attaining every finite minimum. The finite cover preserves that minimum; empty fibers have value $+\infty$. A tree is determined by its β-edge complement, proving the combinatorial bound. □

The bound permits polynomial work at fixed β, with additional polynomial per-candidate construction and LP costs; it is an XP bound, not fixed-parameter tractability or a bound for arbitrary Boolean-query optimization. The full-order $b_{\mathrm{ord}}!$ cover can be smaller when observation constraints fix most orders. For K_4, the structural cover has 96 regions instead of 720 full orders. A strict opposite-answer set can still have an unattained distance infimum; Proposition 14 does not replace strict feasibility or attainment checks.

F.8 CONSERVED FLOWS: EXACT IMAGES AND MINIMUM-COST LIFTS

The Cartesian input domain of Theorem 1 is essential. This section gives a separate, restricted result for a coupled domain and weighted additive cost. It does not extend that theorem to arbitrary networks.

Domain and filtration. Let a directed graph be the union of $r \ge 2$ internally vertex-disjoint directed paths $P_1, \dots, P_r$ from a source s to a sink t, each with at least two edges. There are no other vertices or edges. The underlying undirected graph is simple. Fix rational edge capacities $0 < u_e \le 1$, positive rational costs w_e, and a feasible rational nominal nonnegative flow f with throughput T. Admissible flows g conserve flow at every internal vertex and have the same throughput T. There are no additional equality constraints. Consequently $g_e = a_i$ on P_i,

$$0 \le a_i \le U_i := \min_{e \in P_i} u_e, \qquad \sum_{i=1}^{r} a_i = T.$$

Write a_i^0 for the nominal path flow and $W_i = \sum_{e \in P_i} w_e$. The input cost is $c_w(f,g) = \sum_e w_e |f_e - g_e| = \sum_i W_i |a_i^0 - a_i|$. On the underlying undirected graph, vertices are born at zero, edge e at $1-g_e$, and the full simplex is inserted at time one. Before time one, the complex is one-dimensional. Let $L(1-g)$ be its decreasingly ordered global H_1 lifetimes, padded with zeros to length $r-1$. Directed conservation is a constraint on the input, not a definition of persistent homology.

Proposition 15 (Conserved path image and minimum-cost preimage)**.** *For a decreasing nonnegative target $z \in \mathbb{R}^{r-1}$, define*

$$a(z) = \left(T - \sum_{j=1}^{r-1} z_j,\ z_1, \ldots, z_{r-1} \right).$$

Let $U_{(1)} \geq \cdots \geq U_{(r)}$ be the capacities sorted in decreasing order. The exact attainable lifetime domain $\mathcal{Z}_{T,U}$ is specified by

$$a_1(z) \geq z_1 \geq \cdots \geq z_{r-1} \geq 0, \qquad a_j(z) \leq U_{(j)} \quad (j = 1, \ldots, r).$$

For $z \in \mathcal{Z}_{T,U}$, its minimum input cost is

$$D_f^{\mathrm{flow}}(z) := \min_{g: L(1-g)=z} c_w(f,g) = \min_{\pi:\, a_{\pi(i)}(z) \leq U_i \ \forall i} \sum_{i=1}^{r} W_i |a_i^0 - a_{\pi(i)}(z)|, \tag{19}$$

where the left minimum is over admissible flows and π ranges over permutations of $\{1, \ldots, r\}$. Assigning $g_e = a_{\pi(i)}(z)$ for $e \in P_i$ reconstructs a minimizing input. Outside $\mathcal{Z}_{T,U}$, set $D_f^{\mathrm{flow}}(z) = +\infty$.

Proof. At any preterminal threshold, each path is either entirely present or absent. The union of $k \geq 1$ present paths has cycle rank $k-1$; with no present path the rank is zero. Thus the sorted path flows $a_{(1)} \geq \cdots \geq a_{(r)}$ produce lifetimes $(a_{(2)}, \ldots, a_{(r)})$, including padded zeros. Conservation reconstructs the missing largest value as $a_{(1)} = T - \sum_j z_j$. Every preimage therefore assigns exactly the multiset $a(z)$ to the labelled paths.

Such an assignment respects capacities exactly when $a_j(z) \leq U_{(j)}$ for all j. Necessity follows because if the jth largest strength exceeds the jth largest capacity, fewer than j paths can receive the j largest strengths. Sufficiency follows by matching both lists in decreasing order. For any admissible assignment, conservation and capacities hold, its lifetimes equal z, and its edge cost is the sum in (19). Conversely every admissible preimage gives one such assignment. Minimizing over this finite nonempty set proves the equality and attainment. □

Equation (19) is a standard minimum-cost assignment problem (Kuhn, 1955). After reading the graph, the implemented primal-dual assignment routine takes $O(r^3)$ rational arithmetic operations and $O(r^2)$ stored costs; sorting capacities decides target feasibility in $O(r \log r)$. Rational operations have their usual bit costs. These bounds concern one specified target, not arbitrary Boolean-query optimization.

Corollary 3 (Homogeneous path costs and capacities)**.** *If every $W_i = W$ and every $U_i = U$, put $z^0 = L(1-f)$. For every $z \in \mathcal{Z}_{T,U}$,*

$$D_f^{\mathrm{flow}}(z) = W \left(\|z - z^0\|_1 + \left| \sum_{j=1}^{r-1} (z_j - z_j^0) \right| \right). \tag{20}$$

Proof. For $b \geq b'$ and $c \geq c'$, $|b-c| + |b'-c'| \leq |b-c'| + |b'-c|$. Uncrossing assignments therefore makes decreasing-to-decreasing matching optimal under common weights. Common capacities make that assignment feasible. The removed largest path strength changes by $-\sum_j (z_j - z_j^0)$ because both strength vectors sum to T. Separating this coordinate from the remaining sorted absolute differences gives (20). □

The correction follows from sorted absolute-distance geometry and conservation. It gives an explicit input-ball image for this flow domain: for $\rho \geq 0$, the attainable targets at cost at most ρ are precisely

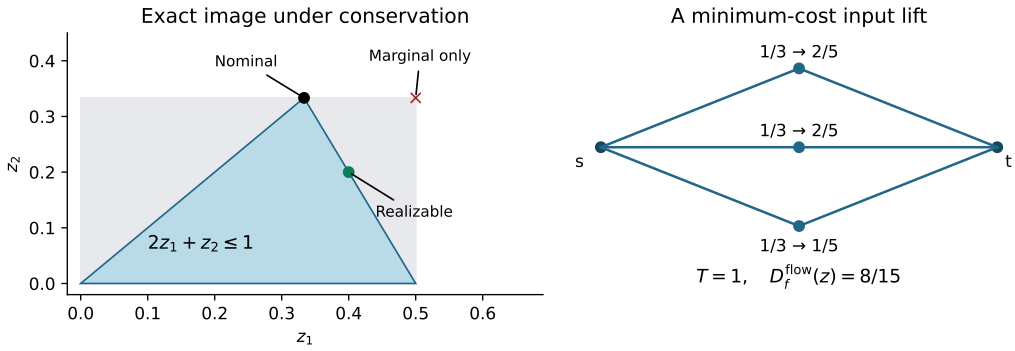

Figure 7: Analytic conserved-flow example, not fitted data. Left: the exact joint lifetime domain (blue) is smaller than the coordinatewise enclosure (gray). Right: an input realizing $(2/5, 1/5)$ at minimum weighted edge cost $8/15$. The common throughput is unchanged.

$\{z \in \mathcal{Z}_{T,U} : D_f^{\text{flow}}(z) \le \rho\}$. For any frozen Boolean query q on this domain, the exact flip-cost infimum is

$$\inf_{z \in \mathcal{Z}_{T,U}:\, q(z) \neq q(z^0)} D_f^{\text{flow}}(z).$$

Target minima are attained; an opposite-answer set defined using strict inequalities can instead have an unattained infimum. Certificate balls use strict input cost below the infimum. A zero infimum supplies no positive certificate.

Why joint constraints and path labels matter. With three paths, $T = 1$ and $U = 1$, the attainable domain is $z_1 \ge z_2 \ge 0$, $2z_1 + z_2 \le 1$. The separate coordinate maxima are $1/2$ and $1/3$, but their pair is unattainable. With two edges per path, unit edge costs and nominal strengths $(1/3, 1/3, 1/3)$, target $(2/5, 1/5)$ costs $8/15$. If the first path is additionally fixed at $1/3$, this target becomes impossible: none of its required strengths $(2/5, 2/5, 1/5)$ equals $1/3$. Figure 7 illustrates this joint domain and the minimum-cost lift. Additional sensor equalities change the feasible set.

Keeping the block indices. Connect k parallel-path graphs $H_1, \dots, H_k$ in series: the sink of H_i is the source of H_{i+1}, and there are no other shared vertices. Block H_i has r_i paths. The endpoints have fixed throughput T; all other vertices conserve flow. Retain the same capacities, filtration and additive edge cost. Let f_i be the restriction of f to block H_i, and let $L^{(i)} : [0,1]^{E_i} \to \mathcal{C}_{r_i - 1}$ denote its local ordered lifetime map, where E_i is the edge set of H_i. Define the *block-indexed lifetime vector* $\Lambda(g) = (L^{(i)}((1-g)|_{E_i}))_{i=1}^k \in \prod_i \mathcal{C}_{r_i - 1}$. The block indices come from the fixed candidate graph, not class labels.

Corollary 4 (Exact block-indexed lift)**.** *For a specified block-indexed target $z = (z^{(1)}, \dots, z^{(k)})$,*

$$\min_{g:\Lambda(g)=z} c_w(f, g) = \sum_{i=1}^{k} D_{f_i}^{\text{flow}}(z^{(i)}).$$

Minima are over admissible flows, with $+\infty$ for infeasibility. After the decomposition is validated, a minimum-cost input is reconstructed in $O(|E| + \sum_i r_i^3)$ rational arithmetic operations.

Proof. Every block carries the same fixed throughput. Its path-flow constraints involve no other block's edge variables. Conversely, admissible flows in the blocks concatenate to a globally admissible flow because their terminal injections cancel at shared vertices. The feasible set is a Cartesian product of block flow sets, and the edge cost is additive. Apply Proposition 15 independently and concatenate its minimizers. □

This statement does not make arbitrary Boolean-query verification polynomial. In particular, removing the block indices requires an additional assignment of the target lifetimes to blocks. That assignment can already be hard.

Forgetting block indices. Let $\beta_i = r_i - 1$ and $\beta = \sum_i \beta_i$. Define $\mathsf{U} : \prod_i \mathcal{C}_{\beta_i} \to \mathcal{C}_\beta$ by sorting the multiset union of all block coordinates in decreasing order, retaining multiplicity. Cycle ranks add across the articulation vertices; all bars die at one, so the global lifetime map is $L(1-g) = \mathsf{U}(\Lambda(g))$. The two inverse problems specify different target information.

Proposition 16 (Global lifting by target allocation)**.** *For a global target $z \in \mathcal{C}_\beta$, let $\mathcal{A}(z) = \{\zeta = (\zeta^{(i)})_i : \mathsf{U}(\zeta) = z,\ \zeta^{(i)} \in \mathcal{C}_{\beta_i}\}$. Then*

$$D_f^{\mathrm{global}}(z) := \min_{g : L(1-g) = z} c_w(f, g) = \min_{\zeta \in \mathcal{A}(z)} \sum_i D_{f_i}^{\mathrm{flow}}(\zeta^{(i)}),$$

where flow minima are over admissible inputs and an infeasible cost is $+\infty$.

Proof. Restrict any global preimage to its blocks. The resulting local target belongs to $\mathcal{A}(z)$, and each block cost is at least its minimum lift cost. Conversely, concatenate minimum lifts of any feasible allocation using Corollary 4. Its global target is z and its cost is the sum of local minima. There are finitely many allocations: at most $\beta!/\prod_i \beta_i!$ before identifying repeated coordinates. Taking the minimum proves both inequalities and reconstructs an input whenever the finite cost exists. □

The allocation identity is a finite-union decomposition, not a polynomial bound in β. A capped enumeration returning a feasible input provides an upper bound unless optimality has been established. A frozen global query is still $q(\mathsf{U}(\Lambda(g)))$; replacing it by a newly fitted indexed query would change the predictor.

Finite decision problem. An instance consists of a positive integer k and $z \in \mathbb{Q}^{3k} \cap \mathcal{C}_{3k}$, with integers and rational numerators/positive denominators encoded in binary. The graph G_k is the specified series of k four-path blocks with two edges per path. Capacities are one, throughput is two and the supplied nominal is $f_e = 1/2$. The question is whether an admissible flow has global target z.

Proposition 17 (Erasing block indices can make inversion NP-complete)**.** *For a series of four-path blocks, with two edges per path, unit capacities, throughput two, and nominal flow $f_e = 1/2$, deciding whether a specified global ordered lifetime target has an admissible input is NP-complete. No extra equality beyond flow conservation is required.*

Proof. Use 3-PARTITION (Garey & Johnson, 1979): positive integers $b_1, \ldots, b_{3k}$ and B satisfy $B/4 < b_j < B/2$ and $\sum_j b_j = kB$; the question is whether they partition into k triples of sum B. Construct a series of k four-path blocks and use the nominal $f_e = 1/2$ on every edge. This is always feasible at throughput two. Specify the global lifetime multiset $\{b_j/B : j = 1, \ldots, 3k\}$.

At every preterminal time the blocks meet only at articulation vertices, so their cycle ranks add. Hence the global lifetime multiset is the union of their local triples. In block i, write its four sorted path strengths as $a_{i1} \geq \cdots \geq a_{i4}$. Its local lifetime sum is $2 - a_{i1} \geq 1$, because $a_{i1} \leq 1$. The total target sum is k, so every local triple must sum to one and every a_{i1} equals one. Thus a feasible input partitions the b_j into triples of sum B. Conversely, each such partition gives a feasible block with strengths $(1, b_{j_1}/B, b_{j_2}/B, b_{j_3}/B)$. The graph has $8k$ edges and polynomial encoding length. For any instance of the decision problem, a witness allocates each of the $3k$ target indices exactly once to k triples. Reconstruct the fourth strength of each block as two minus its triple sum, and check that it is at least the triple's maximum and at most one. These rational arithmetic checks have polynomial bit complexity in the input length and exactly characterize a feasible flow, proving membership in NP. The extra restrictions on the b_j are used in the reduction, not in the definition of general instances. □

For unit edge costs, every feasible target generated by the preceding 3-PARTITION reduction (in particular, each target entry is strictly between $1/4$ and $1/2$, and their sum is k) has minimum input cost exactly $2k$: each block contributes $2[(1-1/2)+\sum_{j=2}^{4}(1/2-a_{ij})] = 2$. Thus the hard question is joint realizability, even when the finite minimum cost is known in advance. This is a standard 3-PARTITION encoding, not a claim of a new generic complexity technique. This cost identity does not apply to arbitrary targets on the same graph: its nominal target has cost zero.

G Reproducibility and research assistance

G.1 Models, certificates and experiment records

The evaluated model is specified by its input coordinate convention, standardization folded into rational affine coefficients, selected Boolean operations and wiring, and threshold vote. Prediction and verification use that same export. Each certificate record identifies the model and nominal input, exact rational lower and upper endpoints, stopping status, and the accepted opposite-answer target and reconstructed graph. Replaying a witness checks input admissibility, recomputed PH, distance and the deployed answer; verification trust assumptions are stated in Appendix C.1.

The EEG protocol and statistical units are specified in Appendix D. Controlled workloads, budgets and timing definitions appear in Appendix E. Tables distinguish designed query identities, graph realizations, model/window cases and computational attempts. Source and data identifiers link records to their generation and replay procedures; they do not replace mathematical correctness checks.

G.2 LLM assistance

LLM-based research and coding assistants were used to inspect theoretical proofs, review related literature, implement verification code, inspect saved artifacts, and revise manuscript text and figures. Reported numerical results are bound to saved experiment records; exact arithmetic checks and solver runs, not generated prose, supply their computational evidence. This assistance does not constitute independent formal verification. Responsibility for the scientific claims, citations, code and final submitted text remains with the authors.