PRISM: FROM FEASIBLE GEOMETRY TO ROUTING CAPACITY IN LEARNING TO OPTIMIZE

Anonymous authors
Paper under double-blind review

ABSTRACT

Many learning-to-optimize methods struggle once the constraint set is curved, non-convex or disconnected: they either cannot guarantee hard feasibility there or lose solution quality. Feasibility alone is also not enough: a parametrisation can be feasible everywhere yet unable to represent the solution it is trained toward, and then no training closes the gap. For ray parametrisations, which place the output on a ray from an interior anchor, reachability is set by two budgets, anchors A and feasible segments K per ray: with one anchor and one segment, a second component or the far side of a hole is out of reach; each set admits a region of sufficient (A,K) budgets, and some are provably insufficient. PRISM realises such budgets: one shared network proposes ray candidates from an anchor bank, a parameter-free oracle marks each ray's feasible segments, and with valid anchors and an exact or certified oracle every hard-routed output is feasible. A measurement pass reads a data-supported (A,K) configuration off the feasible training labels and certifies its held-out coverage, without claiming the smallest budget. Under stated regularity assumptions we prove compact-set approximation, single-anchor obstructions, and a coverage-to-gap existence bound. On ten non-convex geometries the measured budgets cover at least 98% of held-out targets while single-anchor controls are starved on every disconnected set; on fifteen parametric problems PRISM meets each row's feasibility criterion and is best on twelve (the three legacy AC optimal power flow rows are coordinate-surrogate, box-only-residual evidence, not full-AC feasibility); on case30, conditional on a certified initial root, each accepted update certifies a continuous, strictly feasible root path with exact cost decrease.

1 INTRODUCTION

A learned solver can return a feasible point while its parametrisation excludes the desired solution. Training can improve the prediction within the reachable set, but cannot remove that structural restriction. Learning-to-optimize (L2O) uses networks to map problem parameters to solutions (Amos, 2025; Kotary et al., 2021; Liu et al., 2024), with applications in dispatch (Huang et al., 2022; Zamzam & Baker, 2020; Zhao et al., 2023) and beamforming (Niu et al., 2021). Existing methods enforce feasibility through solver layers, repair, or constrained parametrisations (Sec. 2). We ask a complementary question for ray-based representations: how many anchors and feasible segments suffice, and how can samples guide that choice?

Under ray coordinates, reachability has two axes (Fig. 1A). From an interior anchor, a ray can leave and re-enter the feasible set: K counts the available feasible segments. Additional anchors restart this count: A counts the available charts. The far side of an annulus can require a second segment; a separate disk requires another anchor at $K{=}1$. We call sets covered by these rays (A, K)-*star-coverable* (Def. 3.2). Two disjoint disks admit both $(2, 1)$ and $(1, 2)$, so a geometry determines a *region* of sufficient budgets, not one canonical pair. Some budgets cannot cover the set: on C nested annuli, no single anchor suffices at any $K \leq C$ (Thm. 4.2), regardless of training.

PRISM realises a chosen budget with one shared network and a parameter-free ray oracle (Fig. 1C). The coordinate-descriptor configuration conditions the trunk on the emitting anchor; fixed banks also admit slot-indexed heads (Sec. 3.2). With valid anchors and an exact or certified oracle, every hard-routed output is feasible; the soft training mixture need not be. A measurement pass uses feasible training labels to select an operating configuration with a held-out coverage certificate (Sec. 3.3),

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

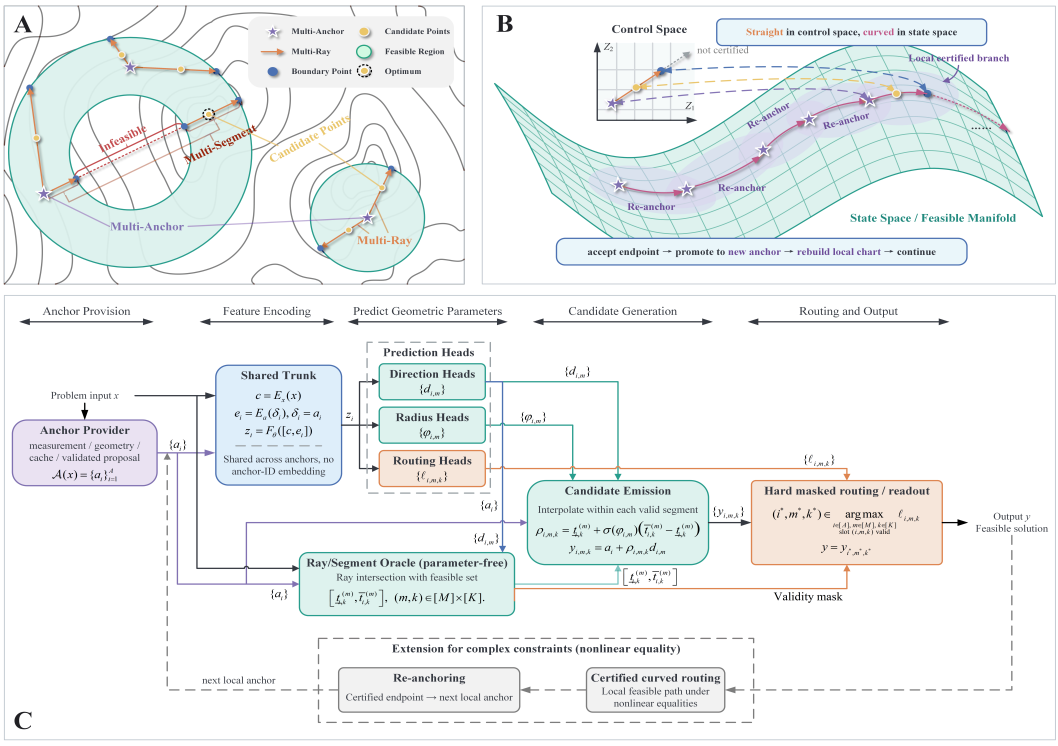


Figure 1: **PRISM: routing capacity (A), shared realisation (C), and a nonlinear-equality extension (B).** Read A, then C, then B. **(A)** Anchors (stars) and feasible ray segments (arrows) define reachable candidates (yellow); the dashed proposal crosses the hole. Re-entering the annulus requires another segment (K); the separate disk needs another anchor at $K{=}1$ (A). **(C)** Coordinate-descriptor realisation: the provider returns $\{a_i\}$, the shared trunk reads x and $\delta_i = a_i$, and the oracle returns each predicted ray's feasible segments. The radial head places a candidate inside a segment; routing selects a valid slot (soft mixture during training, hard arg-max at inference; Sec. 3.2). Slot-indexed heads use the same oracle, emission and routing. **(B)** A control-space segment lifts to a curved state path. A certificate establishes a feasible prefix, and a certified, cost-improving endpoint may become the next local anchor (Sec. 3.4): a conditional extension, not a generic refinement loop.

not the minimum of the budget region. Geometry identifies sufficient budgets, the certificate bounds target coverage, and a capacity probe compares trainability among near-equal-coverage choices. Thus reachability and trainability guide configuration separately (Sec. 5.1). For nonlinear equalities, the case30 specialisation replaces a ray interval with a certified local feasible path (Fig. 1B; Sec. 3.4).

Contributions. **(i) Routing capacity.** We define (A, K)-star-coverability and its sufficient-budget region, extending the convex single-chart $(1, 1)$ ray setting of RAYEN (Tordesillas et al., 2026). The separation theorem identifies single-anchor obstructions on disjoint convex components at $K{=}1$ and on nested annuli under insufficient segment budgets (Thm. 4.2). **(ii) A data-configured method.** One shared realisation supports multiple anchor–segment budgets; feasible labels select anchors and segment counts with held-out coverage guarantees, while the capacity probe distinguishes configurations that cover similar target mass (Sec. 3.3). **(iii) Guarantees and evidence.** Under the stated regularity conditions, approximation holds off switching-set neighbourhoods for compact covers (Thm. 4.1, Cor. A.2). With uniform Lipschitz and diameter bounds and the partial-cover approximation assumptions, measured coverage yields an expected-gap *existence* bound (Thm. 4.3), not a guarantee for the trained predictor. Controlled geometries separate reachability from trainability; on fifteen parametric problems PRISM meets each row's reported feasibility criterion and is best under the reported score on twelve (the three legacy AC-OPF rows are coordinate-surrogate, box-only-residual evidence; Sec. 5.2). On case30, the certified path variant returns feasible, cost-decreasing root paths for the 54 of 64 report loads with certified initial roots, at seconds per trajectory (Sec. 5.3).

2 Related Work

Enforce, repair, or restrict. Existing approaches differ in how they enforce feasibility and in the geometric classes their interfaces admit; our comparison distinguishes these mechanisms from the anchor–segment budget studied here. *Soft enforcement* penalises violation (Kim & Kim, 2024; Park & Van Hentenryck, 2023; Huang et al., 2022). *Post-prediction repair* moves a free prediction until it is feasible — by unrolled correction (Donti et al., 2021), bisection towards a learned interior point (Liang & Chen, 2025), violation minimisation with convergence guarantees (Nguyen & Donti, 2025), or projection through a linearised, locally quadratic or KKT model (Lastrucci & Schweidtmann, 2025; Roy et al., 2026; Iftakher et al., 2026) — and pays an inner feasibility computation at every inference. *Structural parametrisations* fix a geometric interface instead: convex solver layers (Amos & Kolter, 2017; Agrawal et al., 2019) keep an online solve; HardNet (Min et al., 2025) handles input-dependent affine and convex sets and already shows that hard constraints and universal approximation coexist; gauge maps (Tabas & Zhang, 2022) and HomeoProj (Liang et al., 2023) push a canonical domain onto $\mathcal{Y}_x$ and so require a ball-like topology (mechanism table Tab. 5 and the $L_{0.5}$-ball case in App. D.3, App. D.11).

Ray coordinates, and the delta from a single chart. RAYEN (Tordesillas et al., 2026) maps a latent direction to $y = a + \rho d$ with a closed-form maximal step for convex sets from one fixed interior point, the convex $(A, K) = (1, 1)$ case; HOP (homeomorphic polar learning; Deng et al., 2025) carries the same parametrisation to star-shaped sets, still $(1, 1)$, and serves as our $M{=}1$ single-chart baseline. PRISM makes the routing capacity itself the object (Sec. 3). Interior points also serve Bisection Projection, there as a repair *reference*, here as coordinate charts whose number decides what is representable at all; acquiring a per-instance anchor is an online cost not included in the reported timings (App. D.15). "Ray chart" is informal usage; the formal object is Def. 3.2. AC-OPF learners (Fioretto et al., 2019; Huang et al., 2022; Zamzam & Baker, 2020; Kim & Kim, 2024; Park & Van Hentenryck, 2023) penalise the power-flow equality or fix one topology; our core stays application-agnostic (App. B).

3 Method

Problem setting. We consider parametric problems

$$\min_{y\in\mathbb{R}^d} \; f(y;x) \quad \text{s.t.} \quad g_i(y;x) \le 0,\; i = 1,\dots,m, \qquad h_j(y;x) = 0,\; j = 1,\dots,p, \tag{1}$$

with parameter $x \in \mathcal{X}$ (e.g. a load profile) and feasible set $\mathcal{Y}_x = \{y : g(y;x) \le 0,\; h(y;x) = 0\}$. PRISM (Fig. 1C) operates from interior *anchors* supplied by an *anchor provider* $\mathcal{A}(x) = G_{\text{anchor}}(x) = \{a_i(x)\}_{i=1}^{A}$ whose only feasibility contract is validity, $a_i(x) \in \operatorname{relint}(\mathcal{Y}_x)$ (the measurement-pass provider additionally returns $\widehat{K}$ and a coverage certificate); a measurement pass over feasible samples (Sec. 3.3), a cached bank, a geometric rule or a validated proposal (App. D.15) all instantiate it. The geometric class PRISM covers is defined by what the anchors' rays reach.

Definition 3.1 (Segment-aware star; radial accessibility)**.** Interiors and directions are relative to $\operatorname{aff}(\mathcal{Y}_x)$, with direction space V_x and unit sphere $\mathbb{S}_x$; for $Ay = b$ this gives $V_x = \ker A$, and without equality $\mathbb{S}_x = S^{d-1}$; the ambient reading would be vacuous on the equality benchmarks (App. D.2). For $a \in \operatorname{relint}(\mathcal{Y}_x)$ and $d \in \mathbb{S}_x$, let $T_{a,d} := \{t \ge 0 : a + td \in \mathcal{Y}_x\}$, a finite union of closed intervals for piecewise-smooth boundaries. *Counting convention.* Its *segments* are the maximal closed intervals of $T_{a,d}$ — maximal in $\mathcal{Y}_x$, not within one component — numbered $1, 2, \dots$ in increasing t, tangential single points included. Let $s_a(y)$ be the number of the segment containing $\|y - a\|$ on the ray towards y, with $s_a(a) := 1$; the *crossing index* of y is $s_a(y) - 1$. Then $\operatorname{Star}_K(a) := \{y \in \mathcal{Y}_x : s_a(y) \le K\}$, and $\mathcal{Y}_x$ is *radially K-accessible* from a if $\mathcal{Y}_x = \operatorname{Star}_K(a)$.

Definition 3.2 ((A, K)-star-coverability)**.** $\mathcal{Y}_x$ is (A, K)*-star-coverable* if there exist anchors $a_1, \dots, a_A \in \operatorname{relint}(\mathcal{Y}_x)$ with $\mathcal{Y}_x = \bigcup_{k=1}^{A} \operatorname{Star}_K(a_k)$.

Radial 1-accessibility is star-shapedness, $K{=}1$ a classical star cover, and the joint class is strictly larger than either axis. The budgets measured below are $A \le 22$, $K \le 2$; the held-out coverage each one attains is reported per geometry (Sec. 5.1, App. D.6).

The ray-cover budget region. No geometry singles out one budget; what it does single out is the *ray-cover budget region*

$$\mathcal{B}(\mathcal{Y}_x) \;:=\; \{(A,K)\in\mathbb{N}^2 : \mathcal{Y}_x \text{ is } (A,K)\text{-star-coverable}\}, \qquad \mathcal{P}(\mathcal{Y}_x) := \operatorname{Pareto\,min}\mathcal{B}(\mathcal{Y}_x).$$

$\mathcal{B}$ is upward closed, so its content sits on the boundary $\mathcal{P}$: two disjoint disks put $(2,1)$ and $(1,2)$ on it, and C nested annuli around a positive inner radius put $(1,C{+}1)$ on it and exclude every $(1,K)$ with $K \le C$ (Thm. 4.2; App. A.2). This region belongs to the feasible set; the measurement instead returns an *operating configuration* $(\widehat{A},\{\hat{a}_k\},\widehat{K})$ supported by data and a held-out certificate, with no claim that it minimises $\mathcal{B}(\mathcal{Y}_x)$ (Sec. 5.1).

For parametric sets we require a *regular* cover: a fixed direction space $\operatorname{aff}(\mathcal{Y}_x) = b(x){+}V$, continuous anchor maps $a_k(x) \in \operatorname{relint}(\mathcal{Y}_x)$ realising the cover, and indexed intervals of locally constant count whose endpoints stay nondegenerate and jointly continuous in (x,d) off a closed empty-interior exceptional set (Def. A.1, App. A.1); the closed-form oracles of Sec. 3.1 satisfy it wherever the constraint data depends continuously on x.

3.1 Ray-Coordinate Representation

Given an anchor $a(x) \in \operatorname{relint}(\mathcal{Y}_x)$ and a unit direction $d(x;a) \in \mathbb{S}_x$, PRISM parametrises

$$y(x;\theta) \;=\; a(x) \;+\; \rho(x;\theta)\cdot d(x;\theta), \qquad \rho \;=\; \underline{t} \;+\; \sigma(\varphi)\,\big(\overline{t}-\underline{t}\big), \tag{2}$$

where $[\underline{t},\overline{t}]$ is an oracle-certified feasible interval of the ray $\{a + t\,d \,:\, t \ge 0\}$ (Sec. 3.1) and $\sigma(\varphi) \in (0,1)$ is the network's radial coordinate. The direction is the *total* read-out of the raw head output w: $\mathsf{N}_V(w) := P_V w/\|P_V w\|$ if $P_V w \neq 0$, else a fixed $e_V \in V \cap \mathbb{S}$ (P_V the orthogonal projector onto V; if $\dim V = 0$ the set is the point $a(x)$, returned directly) — no extra head, anchor or search. On bounded sets Eq. 2 is used as written for training and evaluation; unbounded sets use the $\tan$ radial map of the next paragraph, which does not change the feasibility argument. Restricted to oracle-certified intervals, the image of the parametrisation lies entirely inside $\mathcal{Y}_x$: the anchor and oracle supply the geometry, the network supplies only coordinates relative to that geometry.

Oracle. The intersection oracle takes the instance geometry (the constraint data of x), an anchor a and a unit direction d, and returns the ordered endpoints of the first K segments of the ray slice $T_{a,d}$ (Def. 3.1). It has no parameters; for every benchmark class the endpoints are closed-form or one-dimensional root finds (App. D.1). An oracle is *exact* when every returned interval is a maximal feasible segment and *certified* when each is proved to lie in $\mathcal{Y}_x$; feasibility needs only the latter.

Two extensions. Unbounded sets use a $\tan$ radial map, so $\varphi \in (0,1)$ covers $[0,\infty)$; linear equalities $Ay = b$ project directions into $\ker A$ with $\mathbf{P}_A = I - A^{+}A$, so $A(a+\rho\widetilde{d}) = b$ for every ρ and the oracle handles the remaining inequalities inside the slice (nonlinear equalities: Sec. 3.4).

3.2 Routing over Anchors and Crossings

A single radial coordinate reaches only the *first* interval of the ray slice, so realising a budget (A,K) means routing over segments and anchors, not directions alone. PRISM routes over A interior anchors $\{a_k\}$ — measured from feasible samples (Sec. 3.3) or supplied per instance (App. B) — each carrying M direction–radius hypotheses from a shared trunk. For each head the oracle returns the ordered endpoints $\underline{t}_k^{(m)} \le \overline{t}_k^{(m)}$ of the first K maximal feasible intervals of the whole slice T_{a,d_m} (Def. 3.1), equal endpoints allowed where the ray is tangent to the boundary and absent segments masked. Every (anchor, head, segment) triple gives one candidate by interpolation inside its segment,

$$y_{m,k} \;=\; a \;+\; \Big(\underline{t}_k^{(m)} + \sigma(\varphi_m)\big(\overline{t}_k^{(m)} - \underline{t}_k^{(m)}\big)\Big)d_m, \qquad (m,k) \in [M]\times[K], \tag{3}$$

so the joint candidate set has size AMK, and under an exact or certified oracle every one of them is feasible by construction. A single joint gate $g \in \Delta^{AMK-1}$, a Gumbel-softmax (Jang et al., 2017) over the AMK logits with absent segments masked, routes among them: $y = \sum g_{m,k}\,y_{m,k}$, soft at training with $\tau = 1$, hard arg-max at inference. The guarantee therefore attaches to the candidates and to every hard-routed output, not to the soft training mixture, which on a non-convex set can transiently leave the set; straight-through routing restores the training-forward guarantee at a small cost in gap (App. D.12). All reported metrics use hard routing.

Anchor-aware shared trunk. One network, shared by every anchor (Fig. 1C), reads a descriptor δ_i of the anchor it emits from:

$$c = E_x(x), \quad e_i = E_a(\delta_i), \quad z_i = F_\theta([c, e_i]), \quad (w_{i,m}, \varphi_{i,m}, \ell_{i,m,k}) = H_\theta(z_i), \tag{4}$$

with $d_{i,m} = \mathsf{N}_V(w_{i,m})$ and a_i the origin of every candidate in Eq. 3. The descriptor is a configuration choice: the anchor's coordinates $\delta_i = a_i$ (default; only then is the bank an input that can change after training) or, for a fixed bank, its slot index, realised as slot-indexed heads $H_{\theta,i}(F_\theta(c))$; for $A{=}1$ with a *fixed* anchor the constant code is absorbed into the predictor's bias (the fixed-anchor specialisation). Static-suite and equality rows take the better reported descriptor mean per problem (retrospective, App. D.14); other rows have a fixed $A{=}1$ anchor.

What each axis buys. Neither axis is a free substitute for the other: at $K{=}1$ no number of rays leaves the anchor's component, and one anchor on C nested rings needs $K \geq C{+}1$ (Thm. 4.2); on disconnected sets anchors are the cheaper and, empirically, the more trainable tool (Sec. 5.1). The M direction heads are a trainability device, not a third geometric axis (App. D.9).

Training. We minimise $\mathbb{E}_x[f(y(x;\theta);x) + \lambda_{\text{aux}}\mathcal{L}_{\text{aux}}(x;\theta)]$, where the auxiliary term pulls each anchor's best head towards the direction from that anchor to the training label and prevents single-head collapse. No violation penalty is needed: feasibility comes from the oracle, not the loss. Hyperparameters are in App. E.

3.3 Anchor Providers and the Measured Configuration

For a static geometry the anchor provider is a *measurement pass* that uses only what training already has: feasible samples and the ray oracle (algorithm and practical rules in App. C). **(1) Components:** the Euclidean minimum spanning tree of the samples, with every edge whose chord leaves the set severed by one oracle call; this finds chord-connected components without a density threshold and errs only toward extra anchors. **(2) Anchors:** within each component a depth-restricted greedy set cover places strictly feasible anchors at deep points, a sample s being reachable from a at budget K iff $\|s - a\|$ lies in one of the first K intervals of the ray $a \to s$ (Def. 3.1). **(3) Certificate:** the fraction of samples in a *held-out* split of the training labels, never used for anchor placement, that is reachable from some anchor; Thm. 4.3 turns it, minus a deviation term, into a confidence lower bound on reachable mass under the instance distribution — not a geometric certification of $\mathcal{Y}_x$. **(4) Segment versus anchor budget:** the pass searches $K \in \{1, 2\}$ and buys larger crossing numbers with anchors (Rings-8: 22 anchors at $K{=}2$); among configurations within 0.05 of the best certificate, a *capacity probe* — a short training run scored on a held-back slice of the training split, never the reporting set — decides. The certificate ranks reachability and the probe trainability; on the spiral the two rankings differ, and the probe's choice trains $38\times$ better (Sec. 5.1). Speculative acquisition and bank augmentation are further provider interfaces with validity, not performance, guarantees (App. D.15).

3.4 Extension to Nonlinear Equality Constraints

Under nonlinear equality constraints a straight control segment can lift to a curved path in the full state space (Fig. 1B). In our case30 specialisation the power-flow equations define a local root branch over the 11 controls, and PRISM replaces the ray interval by a *certified local feasible path* $\Gamma_{a,d}(t)$, $t \in [0, \bar{h}]$, on that branch. **Certified curved routing** keeps a segment feasible: from a certified root a and a control direction d, a certifier proves that the lifted path exists on one continuous root branch and satisfies the equalities and every operational inequality strictly along the whole prefix, not at sampled points. **Re-anchoring** decides where the next segment starts: in MOVE mode a certified endpoint becomes the next local origin, $a_{k+1} = q(t^\star)$, only after an exact cost gate proves improvement; FIXED mode keeps its initial base; a rejected call retains the current root. Neither is a global chart or a generic refinement loop. Together they form the geometric layer of our AC-OPF specialisation; certificate, composition rules and failure semantics are in App. B.6 (Thm. B.6, Prop. B.3), the evaluation in Sec. 5.3.

4 THEORY

Notation. We prove what an (A, K) budget can and cannot reach and what a measured configuration certifies. The statements rest on valid anchors, the oracle and hard routing, not on how the trunk is conditioned. For the fixed-geometry statements, let $Y \subset \mathbb{R}^d$ be compact and (A, K)-star-coverable (Def. 3.2) with anchors $\{a_k\}_{k=1}^A$, direction space $V := \mathrm{aff}(Y) - \mathrm{aff}(Y)$ and sphere $\mathbb{S} := \{v \in V : \|v\| = 1\}$; we write $\underline{t}_s(d) \le \bar{t}_s(d)$ for the endpoints of the s-th interval of the ray slice $T_{a,d}$ (Def. 3.1).

Theorem 4.1 (Universal approximation on a fixed (A, K)-star-coverable set)**.** *Let $\mathcal{X}$ be compact and let $Y \subset \mathbb{R}^d$ be a* fixed *(A, K)-star-coverable set with anchors $a_1, \dots, a_A$. Assume **(R1)** for each k the interval count of $T_{a_k,d}$ is locally constant and the endpoints continuous and nondegenerate outside a closed set $N_k \subset \mathbb{S}$ with empty interior; **(R2)** $y^\star : \mathcal{X} \to Y$ is piecewise continuous over finitely many regions $\{\mathcal{X}_j\}$, with $y^\star(\overline{\mathcal{X}_j}) \subset \mathrm{Star}_K(a_{k_j})$ for some anchor assignment and $\inf_{\mathcal{X}_j} \|y^\star - a_{k_j}\| > 0$; **(R3)** on each $\mathcal{X}_j$ the segment index of $y^\star$ is piecewise constant over a finite subdivision, up to a closed empty-interior exceptional set, and the* target-transversality *set $E_j := \{x \in \mathcal{X}_j : d_j^\star(x) \in N_{k_j}\}$, where $d_j^\star(x) := (y^\star(x) - a_{k_j})/\|y^\star(x) - a_{k_j}\|$, is closed with empty interior. Then for every $\varepsilon, \eta > 0$ there are parameters of a* PRISM *with A anchors, budget K, and M heads per anchor, M being the largest number of refined regions a single anchor is assigned (Step 1 of App. A.1), such that $\|y(x;\theta) - y^\star(x)\| \le \varepsilon$ for all x outside the η-neighbourhood of the switching set $S := \bigcup_j \big(\partial\mathcal{X}_j \cup (\text{segment exceptions}) \cup E_j\big)$, while $y(x;\theta) \in Y$ for* all *x.* (Proof: App. A.1.)

Parametric extension. The same conclusion holds for regular compact parametric covers (Def. A.1; Cor. A.2).

The guarantee is ε-accuracy off the η-collar of the switching set and feasibility at every x, collar included. Both statements assume an *exact or certified* oracle and a *compact* feasible set. Neither covers the tan-mode extension for unbounded sets nor the AC-OPF pipelines — the legacy linearised oracle (Prop. A.4), evaluated per sample only under its legacy box-residual criterion, and the certified local paths of App. B.5, which carry their own conditional certificate; those are method extensions, not covered cases.

Theorem 4.2 (Reachability separation)**.** ***(i)** If $\mathcal{Y}_x = \bigsqcup_{j=1}^C Y_j$ is a disjoint union of $C \ge 2$ compact convex sets and $a \in \mathrm{relint}(Y_{j_0})$, then $\mathrm{Star}_1(a) \subseteq Y_{j_0}$: a single-anchor* PRISM *with segment budget 1 emits only in the anchor's own component, for every parameter setting. **(ii)** Let $\mathcal{Y}_x = \bigcup_{j=1}^C Y_j$ be C nested annuli $Y_j = \{y \in \mathbb{R}^d : R_j \le \|y\| \le R_j + w\}$ with $d \ge 2$, $w > 0$, $R_1 > 0$ and $R_{j+1} > R_j + w$, and let $a \in \mathrm{int}(Y_i)$. Every $y \in Y_j$ then has crossing index $\ge |j - i|$, so a single-anchor* PRISM *with budget K emits nothing in rings with $|j - i| \ge K$. Moreover this anchor's segment count is exactly $\max_{y \in \mathcal{Y}_x} s_a(y) = i + C$, so $\mathrm{Star}_K(a) = \mathcal{Y}_x$ iff $K \ge i + C$; hence no single anchor covers $\mathcal{Y}_x$ at any $K \le C$, and $(1, C{+}1) \in \mathcal{P}(\mathcal{Y}_x)$. In both cases, if the optima put mass μ on the unreachable set U, $R_x := \bigcup_k \mathrm{Star}_K(a_k)$ denotes the reachable set, and the family has separation margin $\gamma := \inf_{x:\, y^\star(x) \in U} \inf_{y \in R_x} \big[f(y;x) - f(y^\star(x);x)\big] > 0$ — the least objective price of being confined to the reachable set when the optimum is not in it — then $\mathbb{E}[\mathrm{gap}] \ge \mu\gamma$ for every θ.* (Proof: App. A.2, with the numerical audit of the lower bound.)

Part (ii) is the impossibility measured in Sec. 5.1, which also marks the geometries the theorem deliberately does *not* cover.

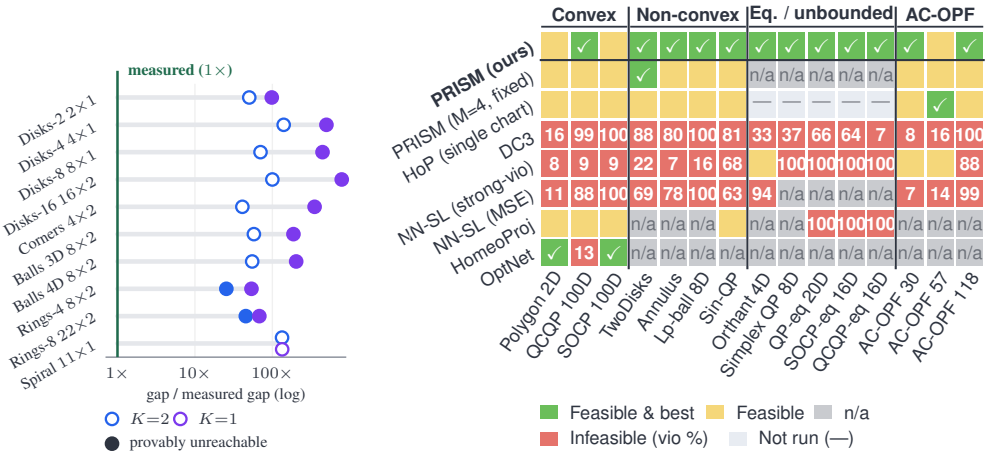
Theorem 4.3 (Measured-configuration guarantees)**.** ***(i)** Fix the anchors before drawing the held-out sample, and write $R_x := \bigcup_k \mathrm{Star}_K^{\mathcal{Y}_x}(a_k(x))$ for the reachable set — the same exact reachability relation as Def. 3.1, which is what the implementation tests. Draw $x_1, \dots, x_m \overset{\text{i.i.d.}}{\sim} \mathcal{D}$ with labels $y_i^\star = y^\star(x_i)$, and let $\hat{c} := \frac{1}{m}\sum_i \mathbf{1}\{y_i^\star \in R_{x_i}\}$. Then w.p. $\ge 1 - \delta$, $\mathbb{P}_{x \sim \mathcal{D}}\big(y^\star(x) \in R_x\big) \ge \hat{c} - \sqrt{\ln(1/\delta)/(2m)}$ (at $m = 400$, $\delta = 0.05$). Drawing the held-out sample from the* instance *distribution is what lets Part (ii) read u off this bound; in L2O the training labels are exactly such a sample. Because Sec. 3.3 compares C configurations on the same sample, we report the union-bound version $\sqrt{\ln(C/\delta)/(2m)}$, which at $C = 2$ costs 0.007 of certified coverage (App. A.3). **(ii)** If $f(\cdot; x)$ is L_f-Lipschitz and $\mathrm{diam}(\mathcal{Y}_x) \le \bar{D}$, both uniformly in x, if $\mathbb{P}_{x \sim \mathcal{D}}(y^\star(x) \in R_x) \ge 1 - u$ — the value Part (i) certifies at confidence $1 - \delta$ — and the* partial-cover *form of Thm. 4.1 (Lem. A.3) holds with accuracy ε off a collar of probability ν, then some θ achieves $\mathbb{E}[f(y(x;\theta)) - f(y^\star(x))] \le L_f\varepsilon + (u + \nu)L_f\bar{D}$.* (Proofs: App. A.3.)

Part (ii) runs through a partial-cover version of Thm. 4.1, since a measured configuration is certified only on $R_x \subseteq \mathcal{Y}_x$ (Lem. A.3): (i) converts the certificate into coverage, (ii) coverage into an expected-gap guarantee that pays $L_f \bar{D}$ on the uncovered mass. A greedy set-cover bound on the number of anchors is in App. A.3. We state no generalisation bound: the hard arg-max router makes $\theta \mapsto y(x;\theta)$ discontinuous where gate logits tie, exactly where a Lipschitz-in-θ hypothesis would be needed.

5 Experiments

Setup. The 25 benchmark rows are 10 static geometries (Sec. 5.1) and 15 parametric problems spanning convex, non-convex, equality/unbounded and AC-OPF classes (App. D.1), reporting the optimality gap $\Delta := \mathbb{E}_x[f(y(x;\theta);x) - f^\star(x)]$ and violation rate $\mathbb{V} := \mathbb{P}_x[\max_i\{g_i, |h_j|\} > 10^{-5}]$ over 5 seeds for the parametric suite; the static suite uses 3 seeds. The twelve non-AC parametric problems are scored by their own objectives; the three AC-OPF rows are *legacy surrogate rows* (coordinate surrogate, box-only residual; App. B.3), with AC-OPF cost and feasibility claims made only in Sec. 5.3. Baselines are HoP (Deng et al., 2025) (single-chart $(1,1)$ ray baseline; our $M{=}1$ re-implementation), DC3 (Donti et al., 2021), NN-SL with a strong violation penalty and with pure label matching, HomeoProj (Liang et al., 2023) and OptNet (Amos & Kolter, 2017).

5.1 Geometry and Measurement: Does the Measured Budget Predict What Is Needed?



(a) **Spending the budget wrong is measurable.** One row per geometry, labelled with the measured configuration; markers are a single-anchor control's gap *relative to* it (reported best-of-descriptor row vs saved single-anchor runs: cross-configuration ratios). *Filled* = budgets that provably cannot reach part of the set (Thm. 4.2); hollow = not excluded. Absolute gaps in Tab. 9.

(b) **Breadth across fifteen parametric problems.** Green with a tick = feasible and lowest gap among feasible methods (the fixed $M{=}4$ row is a PRISM ablation); red cells carry the violation rate. Feasible = mean violation $< 5\%$ over five seeds; the AC-OPF columns are legacy surrogate rows (App. D.1). Same best-of-family rule for PRISM and HoP (App. D.10); equality cells add the retrospective descriptor choice; the SOCP-100D cell is the numerical active-manifold configuration, developed after the ray results and outside the train-tail rule (App. D.13); per-cell statistics in Tab. 7.

Figure 2: **Geometry (a) and breadth (b).** Selection rules and scope are stated in each panel.

The 10 controlled geometries of Fig. 2(a) (per-sample projection objectives, Tab. 9) have known feasible sets, so the returned configuration can be compared with the geometry that produced it: does too little budget show, and does enough budget also train well? **Too little budget is measurable.** The $1{\times}1$ row ($A{\times}K$) is the Thm. 4.2(i) witness: feasible everywhere, starved on every disconnected geometry. For Rings-8 at $K = 2$, the geometric audit reports 22.5% mean single-anchor reachable mass (App. A.2). Separately, the best reported single-anchor control has gap 32.8, versus 0.72 for the measured 22 anchors under the same recipe — a $45\times$ cross-configuration ratio involving both routing budget and descriptor.

The measured configurations distinguish coverage from trainability. On these benchmarks the pass finds every component of the disconnected geometries and, with K searched only in $\{1, 2\}$ (Sec. 3.3), buys any further crossing demand with anchors (Rings-8: 22×2; spiral: 11×1); the held-out coverage of every measured configuration is ≥ 0.98, mostly 1.0, and under one training recipe, with either descriptor, PRISM is best-feasible on $10/10$ at zero violation across all 30 runs per descriptor, while every penalty baseline is infeasible on 7–100% of samples (App. D.6, App. D.14). Coverage does not decide quality: on the spiral the smallest budget that reaches every held-out target (coverage 1.0) trains $38\times$ worse than the probe's choice (11×1, coverage 0.98). The single-anchor $K{=}2$ disk controls are not excluded by Thm. 4.2; their 41–$140\times$ gap ratios (Fig. 2(a)) alone do not establish a reachability obstruction.

Equal coverage, unequal solutions. With the coordinate descriptor the bank is an input, so one trained network can be run on another bank without retraining (Tab. 1). Changing the bank need not preserve the quality of a frozen predictor: on Disks-16 the measured bank and a same-size random bank both have held-out coverage 1.00, but gaps of 0.101 and 3.236; an eight-times larger random bank lowers the gap to 0.132. These are bank replacements at fixed weights, not augmentation, not a single-factor ablation and not a claim that more anchors always help: recovery is incomplete on Disks-2, while the enlarged banks beat the measured coordinate bank on Rings-8 and the spiral. Coverage is an empirical held-out fraction, not a whole-set certificate; finite-precision violations under bank shift are audited in App. D.14.

Table 1: **Fixed network, changed anchor bank.** Frozen coordinate-descriptor checkpoints, trained on the measured bank and evaluated on replacement random interior banks of $1\times$ and $8\times$ the measured size (three seeds; five random banks per condition and seed; seed means of gap and empirical held-out coverage). The measured column is the coordinate configuration, not the best-of-descriptor headline; its coverage is 1.00 except 0.99 on Rings-8 and the spiral. Full sweeps and finite-precision violation audit: App. D.14.

	gap			coverage	
Geometry	meas.	$1\times$	$8\times$	$1\times$	$8\times$
Disks-2	0.069	2.780	0.807	0.95	1.00
Disks-4	0.059	2.323	0.206	0.95	0.99
Disks-8	0.061	2.737	0.123	0.81	0.99
Disks-16	0.101	3.236	0.132	1.00	1.00
Corners	0.056	2.416	0.122	0.83	1.00
Balls 3D	0.140	3.790	0.228	1.00	1.00
Balls 4D	0.077	2.511	0.218	0.79	1.00
Rings-4	0.285	0.789	0.324	0.95	1.00
Rings-8	0.724	1.326	0.386	0.94	1.00
Spiral	0.082	0.506	0.061	0.78	1.00

Table 2: **Solution quality at matched trunk width.** The recorded width-restricted PRISM configurations on the seven non-AC tasks of the width-control study (five-seed mean gaps), against the best competitor meeting the reported violation criterion. Equal trunk width does not imply equal parameter count, compute, selection or a single-factor ablation, and these are not necessarily the headline configurations (Sin-QP headline 0.771). †Numerical active-manifold SOCP configuration; ray results and development history in App. D.13.

	PRISM trunk width		
Problem	512	256	best competitor
Polygon-2D	0.0038	0.0040	OPTNET **0.0000**
QCQP-100D	34.67	**34.59**	HOP 34.97
SOCP-100D†	3.47	3.56	OPTNET **3.15**
TwoDisks	0.503	**0.537**	HOP 0.820
Annulus	0.121	**0.148**	HOP 0.321
$L_{0.5}$-ball	0.3645	0.3746	HOP **0.3719**
Sin-QP	0.814	**0.595**	HOP 1.329

5.2 BREADTH: DOES THE REPRESENTATION SURVIVE HETEROGENEOUS CONSTRAINTS?

Dominance summary. PRISM meets each row's reported feasibility criterion on all 15 problems and is best under the reported score on 12 (Fig. 2(b); scorecard in App. D.4; the three legacy AC-OPF rows are surrogate/box-residual rows, not full-AC feasibility evidence). Two losses are to exact convex solver layers: Polygon-2D, where the layer solves a QP per instance, and SOCP-100D (Tab. 2, next paragraph); the third loss is a near-tie with the single-chart HOP baseline on AC-OPF-57 (0.0064 vs 0.0061). PRISM and HOP use the same train-tail best-of-family rule within the ray family (App. D.10); the equality rows take the retrospective better descriptor per task (App. D.14); the SOCP-100D cell is a later task-specific configuration developed after the ray results and outside that rule (App. D.13).

Quality at matched trunk width. Tab. 2 compares the width-restricted configurations with the best feasible competitor: at the baselines' width the advantage over HOP survives on TwoDisks, the annulus and Sin-QP (AC-OPF-118: App. D.13), shrinks to a near-tie on QCQP-100D and disappears on the star-shaped $L_{0.5}$-ball, where one anchor already sees the whole set. On SOCP-100D the numerical active-manifold configuration gives a gap of 3.47 ± 0.07 with no violation above the 10^{-5} criterion, against 3.15 ± 0.05 for OPTNET; the original ray configuration (16.2) and the implementation are in App. D.13. Direct-ray inference uses one MLP forward plus the closed-form oracle, 32–67 μs per instance on the timed non-convex toys (App. D.7), not a cost of the certified case30 mode.

Anchor descriptor. Under the same recipe, banks and seeds, coordinates improve Rings-8 from 0.953 (slot-indexed) to 0.724 but worsen the spiral from 0.037 to 0.082; on the moving-anchor equality tasks only SOCP-eq improves consistently ($1.231 \rightarrow 1.124$, $5/5$ seeds). The descriptor changes inductive bias and the bank interface, not accuracy uniformly (App. D.14).

5.3 NONLINEAR EQUALITIES: CERTIFIED LOCAL PATHS ON CASE30 (MOVE/FIXED)

Certified coverage and protocol. 64 sealed report loads, the same for every method; a certified initial root was available for 54 of them — an availability, not a physical feasibility rate, and the ten unavailable loads stay in every denominator. Six certified methods (MOVE/FIXED base $\times$ numeric-gradient direction or one of two frozen trained policies) run on every root with a budget of 8 update calls; all 324 trajectories spent their budget, with 1–8 accepted calls each (66 accepted all 8), so every certified load enters the comparison; trajectories are not independent loads. Every accepted update carries the conditional certificate of App. B.5: continuous same-branch root path, every inequality strict, exact cost decrease.

Economic quality. Against a per-load local reference solve, the mean generation-cost gap on the 54 certified roots is 1.428% for MOVE-GRAD, 1.496% for FIXED-GRAD and 1.550% / 1.666% for the two MOVE-NN seeds (Tab. 3). Against a corrected, bounded 30-epoch full-constraint DC3 adaptation (not a tuned reproduction; no path certificate), on exactly the loads it returned (20 and 26 of 64, against our 54) the same-seed certified MOVE-NN is cheaper on $20/20$ and $26/26$ loads (0.547% vs. 4.986%; 0.996% vs. 5.787%). Re-verifiable continuous feasible paths and real cost improvement on the certified loads cost seconds of certification per trajectory (no speed ratio is claimed); the learned direction gave no stable improvement over the numeric rule (App. B.7).

6 DISCUSSION AND LIMITATIONS

Routing capacity versus model size. With the coordinate descriptor, the same network processes each anchor, so a larger bank need not require more learned parameters (Eq. 4). In the per-bank retraining study (Tab. 20), increasing fixed random banks from $1\times$ to $4\times$ the measured size on Rings-8 leaves the coordinate models at 118k parameters, whereas slot-indexed models grow from 123k to 296k. This does not make inference cost independent of A: at fixed M and K, more anchors require more ray evaluations and provide up to AMK candidate slots. Together with the fixed-weight bank replacements in Tab. 1, this separates geometric configuration from network size, while solution quality and computational cost remain bank-dependent.

Limitations. Where an exact solver layer is affordable, PRISM is competitive but not exact: its convex gaps are those of a learned parametrisation, not of a solver (SOCP-100D: App. D.13). The measurement pass is only as good as its samples; sets needing unbounded crossing numbers fall outside (A, K)-star-coverability at any fixed budget; a state-conditioned refinement loop gave no gain beyond its accept/reject wrapper (App. D.14); generic re-anchoring is absent; and the AC-OPF certificate is conditional on a certified initial root ($54/64$ availability) and costs seconds per trajectory.
Conclusion. PRISM treats the routing capacity a feasible geometry demands as an object to be defined, bounded and measured, and keeps three questions apart: whether a budget of anchors and segments covers the geometry (the (A, K)-cover definition, with the separation theorem marking budgets that cannot), which targets the samples show a chosen budget to reach (the held-out certificate), and how well a covered budget trains (the capacity probe). For nonlinear equalities a certified local path replaces the ray interval on case30, under its own conditions.

AI USE STATEMENT

The research idea, problem formulation, benchmark and experimental design, theoretical claims and the interpretation of results are the authors' own. Large language models (GPT-6 Astra, Claude Fable 5 and 5.1, and Claude Opus 5) were used as assistants in three ways. (i) Code. The code base builds on the authors' existing repository; LLMs assisted, from specifications written by the authors, in extending and revising parts of it, such as new experiment configurations, evaluation scripts and baseline adaptations. The authors reviewed and tested this code. (ii) Theory. The authors supplied the proof ideas; LLMs assisted with intermediate reasoning and with checking individual steps, and every proof in Apps. A and B was re-derived step by step by the authors. (iii) Writing. LLMs assisted with drafting, restructuring and polishing the text, with figure layout, and with literature search and citation checks. LLMs were not used to generate data (benchmark instances are sampled by our scripts from author-specified distributions, App. D.1), to propose hypotheses, or to interpret results. The authors take full responsibility for the content.

REPRODUCIBILITY STATEMENT

Benchmark definitions and violation criteria are specified in App. D.1 and App. D.6; hyperparameters and hardware are listed in App. E, and parametric-suite seeds in App. D.8. The measurement pass, configuration-selection rules and SOCP-100D development history are documented in App. C, App. D.10 and App. D.13, respectively. Apps. B.5–B.7 specify the certified case30 objects, protocol and comparator. The code, the per-seed result files (`experiments/megatable.csv`, `experiments/static_suite_results.csv`) and the certificate records referenced in App. B.5 will be released publicly upon publication.

REFERENCES

Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J. Zico Kolter. Differentiable convex optimization layers. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/9ce3c52fc54362e22053399d3181c638-Paper.pdf.

Brandon Amos. Tutorial on amortized optimization, 2025. URL https://arxiv.org/abs/2202.00665.

Brandon Amos and J. Zico Kolter. OptNet: Differentiable optimization as a layer in neural networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 136–145. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/amos17a.html.

Vasek Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.

George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

Ke Deng, Hanwen Zhang, Jin Lu, and Haijian Sun. Hop: Homeomorphic polar learning for hard constrained optimization, 2025. URL https://arxiv.org/abs/2502.00304.

Priya Donti, David Rolnick, and J Zico Kolter. Dc3: A learning method for optimization with hard constraints. In *International Conference on Learning Representations*, 2021.

Ferdinando Fioretto, Terrence W. K. Mak, and Pascal Van Hentenryck. Predicting ac optimal power flows: Combining deep learning and lagrangian dual methods, 2019. URL https://arxiv.org/abs/1909.10461.

Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991. ISSN 0893-6080. doi: https://doi.org/10.1016/0893-6080(91)

90009-T. URL https://www.sciencedirect.com/science/article/pii/089360809190009T.

Wanjun Huang, Xiang Pan, Minghua Chen, and Steven H. Low. Deepopf-v: Solving ac-opf problems efficiently. *IEEE Transactions on Power Systems*, 37(1):800–803, 2022. doi: 10.1109/TPWRS.2021.3114092.

Ashfaq Iftakher, Rahul Golder, Bimol Nath Roy, and M.M. Faruque Hasan. Physics-informed neural networks with hard nonlinear equality and inequality constraints. *Computers Chemical Engineering*, 204:109418, 2026. ISSN 0098-1354. doi: https://doi.org/10.1016/j.compchemeng.2025.109418. URL https://www.sciencedirect.com/science/article/pii/S0098135425004211.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax, 2017. URL https://arxiv.org/abs/1611.01144.

Minsoo Kim and Hongseok Kim. Self-supervised equality embedded deep lagrange dual for approximate constrained optimization, 2024. URL https://arxiv.org/abs/2306.06674.

James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, and Bryan Wilder. End-to-end constrained optimization learning: A survey, 2021. URL https://arxiv.org/abs/2103.16378.

Giacomo Lastrucci and Artur M. Schweidtmann. ENFORCE: Nonlinear constrained learning with adaptive-depth neural projection, 2025. URL https://openreview.net/forum?id=6XdUI9Cibc.

Enming Liang and Minghua Chen. Efficient bisection projection to ensure neural-network solution feasibility for optimization over general set. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 37071–37099. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/liang25b.html.

Enming Liang, Minghua Chen, and Steven Low. Low complexity homeomorphic projection to ensure neural-network solution feasibility for optimization over (Non-)Convex set. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 20623–20649. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/liang23a.html.

Ya-Feng Liu, Tsung-Hui Chang, Mingyi Hong, Zheyu Wu, Anthony Man-Cho So, Eduard A. Jorswieck, and Wei Yu. A survey of recent advances in optimization methods for wireless communications. *IEEE Journal on Selected Areas in Communications*, 42(11):2992–3031, 2024. doi: 10.1109/JSAC.2024.3443759.

Youngjae Min, Anoopkumar Sonar, and Navid Azizan. Hard-constrained neural networks with universal approximation theorem, 2025. URL https://openreview.net/forum?id=3Imf21Jvwh.

Hoang Nguyen and Priya Donti. Fsnet: Feasibility-seeking neural network for constrained optimization with guarantees. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 39670–39708. Curran Associates, Inc., 2025. doi: 10.52202/085713-1323. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/3874e2be479a9d4e94d4514046c1f934-Paper-Conference.pdf.

Guanchong Niu, Qi Cao, and Man-On Pun. Qos-aware resource allocation for mobile edge networks: User association, precoding and power allocation. *IEEE Transactions on Vehicular Technology*, 70(12):12617–12630, 2021. doi: 10.1109/TVT.2021.3076353.

Seonho Park and Pascal Van Hentenryck. Self-supervised primal-dual learning for constrained optimization. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth*

Symposium on Educational Advances in Artificial Intelligence, AAAI'23/IAAI'23/EAAI'23. AAAI Press, 2023. ISBN 978-1-57735-880-0. doi: 10.1609/aaai.v37i4.25520. URL https://doi.org/10.1609/aaai.v37i4.25520.

Bimol Nath Roy, Rahul Golder, and MM Faruque Hasan. Nlpopt-net: A learning method for nonlinear optimization with feasibility guarantees, 2026. URL https://arxiv.org/abs/2605.00260.

Daniel Tabas and Baosen Zhang. Computationally efficient safe reinforcement learning for power systems. In *2022 American Control Conference (ACC)*, pp. 3303–3310, 2022. doi: 10.23919/ACC53348.2022.9867652.

Jesus Tordesillas, Victor Klemm, Jonathan P. How, and Marco Hutter. Rayen: Imposition of hard convex constraints on neural networks, 2026. URL https://arxiv.org/abs/2307.08336.

Ahmed S. Zamzam and Kyri Baker. Learning optimal solutions for extremely fast ac optimal power flow. In *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 1–6, 2020. doi: 10.1109/SmartGridComm47815.2020.9303008.

Tianyu Zhao, Xiang Pan, Minghua Chen, and Steven Low. Ensuring DNN solution feasibility for optimization problems with linear constraints. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=QVcDQJdFTG.

A PROOFS

A.1 PROOF OF THEOREM 4.1

Step 1 (region refinement). By (R3), refine $\{\mathcal{X}_j\}$ to finitely many disjoint open regions $\{\mathcal{X}'_l\}_{l\le L}$, discarding the closed empty-interior exceptional sets into the switching set S, so that on each $\mathcal{X}'_l$ the anchor index k_l and segment index s_l are constant and $y^\star$ is continuous; the discarded boundaries and exceptional sets are exactly the S of the statement, so taking the $\mathcal{X}'_l$ open costs nothing. Each region receives one candidate slot (anchor k_l, head m_l, segment s_l), so anchor k needs $\#\{l : k_l = k\}$ heads and $M := \max_k \#\{l : k_l = k\}$ suffices; regions that share an anchor and have disjoint shrunk supports may be served by one head, so this is an upper bound. Openness and boundaries are relative to $\mathcal{X}$. Write $\mathcal{X}'^{(\eta)}_l := \{x \in \mathcal{X}'_l : \operatorname{dist}(x, S \cup \partial\mathcal{X}'_l) \ge \eta\}$, a compact set.

Step 2 (continuous polar targets). On $\mathcal{X}'_l$ set $d^\star_l(x) := (y^\star(x) - a_{k_l})/\|y^\star(x) - a_{k_l}\|$ and

$$\varphi^\star_l(x) := \frac{\|y^\star(x) - a_{k_l}\| - \underline{t}(x)}{\overline{t}(x) - \underline{t}(x)} \in [0,1], \qquad \underline{t}(x) := \underline{t}_{s_l}(d^\star_l(x)),\ \overline{t}(x) := \overline{t}_{s_l}(d^\star_l(x)).$$

$d^\star_l$ is continuous on $\mathcal{X}'_l$ (denominator $\ge r_l > 0$ by (R2)); its range over $\mathcal{X}'^{(\eta)}_l$ is a compact $D_l \subset \mathbb{S}\backslash N_{k_l}$ — this is exactly what the target-transversality clause of (R3) buys, since $\mathcal{X}'^{(\eta)}_l$ was built to avoid E_l, and without it a target could ride along a bad direction for a whole region — so by (R1) $\underline{t}, \overline{t}$ are uniformly continuous on a compact neighbourhood of D_l, with $c_l := \min_{\mathcal{X}'^{(\eta)}_l}(\overline{t} - \underline{t}) > 0$ and $\bar{T} := \max_l \max_{\mathcal{X}'^{(\eta)}_l} \overline{t} < \infty$. Hence $\varphi^\star_l$ is continuous on $\mathcal{X}'^{(\eta)}_l$. The degenerate segments that Def. 3.1 counts are directions where $\overline{t} = \underline{t}$; $c_l > 0$ excludes them here, which is what (R1)'s nondegeneracy outside N_k buys.

Step 3 (head approximation). Tietze-extend the polar targets of Step 2 from the compact set $\mathcal{X}'^{(\eta)}_l$ to continuous maps $\tilde{d}_l : \mathcal{X} \to V$ and $\tilde{\varphi}_l : \mathcal{X} \to [0,1]$; they agree with $(d^\star_l, \varphi^\star_l)$ on $\mathcal{X}'^{(\eta)}_l$, but $\tilde{d}_l$ need not have unit length, and may vanish, elsewhere. By universal approximation on compact domains (Cybenko, 1989; Hornik, 1991), for any $\delta \in (0, \frac{1}{2})$ there are parameters whose raw direction output w_l and radial output $\hat{\varphi}_l \in (0,1)$ satisfy $\sup_{\mathcal{X}} \|w_l - \tilde{d}_l\| \le \delta$ and $\sup_{\mathcal{X}} |\hat{\varphi}_l - \tilde{\varphi}_l| \le \delta$ (the sigmoid's open range approximates $[0,1]$-valued targets to any accuracy; exact attainment is not needed by Step 4). The direction is read out by the total map of Sec. 3.1, $\hat{d}_l := \mathsf{N}_V(w_l)$, with P_V the identity when there is no equality constraint and the null-space projector $\mathbf{P}_A$ of Sec. 3.1 when $V = \ker A$. Every estimate below is used only on head l's own support $\mathcal{X}'^{(\eta)}_l$ (a head whose support is empty is simply skipped). There $\tilde{d}_l = d^\star_l \in V$ with $\|d^\star_l\| = 1$, and projection only helps: $\|P_V w_l - d^\star_l\| = \|P_V(w_l - d^\star_l)\| \le \|w_l - d^\star_l\| \le \delta$, so $\|P_V w_l\| \ge 1 - \delta > 0$, the fallback branch of N_V is never taken on $\mathcal{X}'^{(\eta)}_l$, and $\|\hat{d}_l - d^\star_l\| \le 2\delta/(1-\delta) \le 4\delta$. Off $\mathcal{X}'^{(\eta)}_l$ — in the collar, but also inside another head's support, where $P_V w_l$ may well be zero — no bound is claimed for head l: N_V still returns a unit vector, which is all Step 6 needs. On another head's support Step 5 selects its assigned slot; within the collar the selected slot is unrestricted. (That the numerical implementation realises N_V at the branch point is a separate check and not part of this proof.)

Step 4 (error propagation). On $\mathcal{X}'^{(\eta)}_l$, slot l emits $y_l = a_{k_l} + [\underline{t}(\hat{d}_l) + \hat{\varphi}_l(\overline{t}(\hat{d}_l) - \underline{t}(\hat{d}_l))]\hat{d}_l$, while $y^\star = a_{k_l} + [\underline{t}(d^\star_l) + \varphi^\star_l(\overline{t}(d^\star_l) - \underline{t}(d^\star_l))]d^\star_l$. Let ω be a common modulus of continuity of $\underline{t}, \overline{t}$ on the neighbourhoods of $D_1, \dots, D_L$ (the maximum of the finitely many moduli), and take δ small enough that 4δ is below the distance from D_l to N_{k_l}. Telescoping,

$$\|y_l - y^\star\| \le 3\,\omega(4\delta) + \delta\,(\bar{T} + \omega(4\delta)) + 4\bar{T}\delta,$$

using $|\hat{\varphi} - \varphi^\star| \le \delta$, the endpoint moduli, and $\underline{t}(d^\star) + \varphi^\star(\overline{t}(d^\star) - \underline{t}(d^\star)) \le \bar{T}$ for the $\|\hat{d} - d^\star\|$ term. Choose δ so the right side is $\le \varepsilon$, uniformly over the finitely many l.

Step 5 (routing). The joint gate holds one logit per (anchor, head, segment) slot. The target logit function — $+1$ on $\mathcal{X}'^{(\eta)}_l$ at slot (k_l, m_l, s_l), -1 elsewhere — is piecewise constant with all discontinuities inside the η-collar; a continuous network approximates it within $\frac{1}{4}$ uniformly on the closed set $\bigcup_l \mathcal{X}'^{(\eta)}_l$, so on each $\mathcal{X}'^{(\eta)}_l$ the correct slot's logit leads by $\ge \frac{1}{2}$ and the inference arg-max

selects slot l. Validity masking cannot interfere: slot (k_l, s_l) is valid on $\mathcal{X}_l^{(\eta)}$ because its interval is nondegenerate (Step 2).

Step 6 (feasibility everywhere). Every candidate lies in an oracle-certified interval of its own ray, hence in $\mathrm{Star}_K(a_k) \subseteq Y$. Two facts make this an everywhere statement rather than a conditional one: under the total read-out N_V every head has a unit direction at every x (Step 3), and since $a_k \in \mathrm{relint}(Y)$ some ball of radius $r_k > 0$ around a_k in $\mathrm{aff}(Y)$ lies in Y, so the first segment of every head's ray exists and has length $\geq r_k$ in *every* direction; slot $(k, m, 1)$ is therefore always valid, later segments that the oracle does not report are masked, and the mask can never empty the gate. The hard gate then selects one candidate, so $y(x;\theta) \in Y$ for all x, including the collar. (In the parametric case of Cor. A.2 read Y as $\mathcal{Y}_x$ throughout: the oracle reports intervals for the instance's own geometry.) □

Regular parametric covers. The parametric corollary below uses the following definition, stated in full here (the main text gives its summary in Sec. 3).

In L2O the feasible set moves with the instance, so we also name the regularity a *family* needs for the ray coordinates to vary continuously with x. Write $[\underline{t}_{k,s}(x,d), \bar{t}_{k,s}(x,d)]$ for the s-th feasible interval of the ray leaving $a_k(x)$ in direction d.

Definition A.1 (Regular parametric (A, K)-star cover)**.** A family $\{\mathcal{Y}_x\}_{x\in\mathcal{X}}$ of compact sets over a compact parameter set $\mathcal{X}$ is a *regular parametric (A, K)-star cover* if **(i)** the affine hulls share one x-independent direction space, $\mathrm{aff}(\mathcal{Y}_x) = b(x) + V$, so that $\mathbb{S}_V := \{v \in V : \|v\| = 1\}$ is a fixed sphere; **(ii)** there are continuous maps $a_k : \mathcal{X} \to \mathbb{R}^d$ with $a_k(x) \in \mathrm{relint}(\mathcal{Y}_x)$ that *realise* the cover, $\mathcal{Y}_x = \bigcup_{k=1}^{A} \mathrm{Star}_K(a_k(x))$; and **(iii)** for each k the interval count is locally constant, the indexed intervals stay nondegenerate, and the endpoints $\underline{t}_{k,s}, \bar{t}_{k,s}$ are *jointly* continuous in (x, d) outside a closed set $N_k \subset \mathcal{X} \times \mathbb{S}_V$ with empty interior.

Condition (iii) is the parametric form of (R1) of Thm. 4.1; condition (i) keeps the direction sphere fixed rather than x-dependent, covering full-dimensional families and any fixed linear equality. Families whose affine hull *rotates* with x need a continuous projector family and we do not treat them. The closed-form oracles of Sec. 3.1 satisfy it wherever the constraint data depends continuously on x and the crossings stay nondegenerate.

Corollary A.2 (Parametric families)**.** *Let $\{\mathcal{Y}_x\}_{x\in\mathcal{X}}$ be a regular parametric (A, K)-star cover of* compact *sets (Def. A.1), which supplies (R1) in the form of Def. A.1(iii). Let $y^\star$ satisfy (R2)–(R3) with the anchor maps $a_{k_j}(\cdot)$ in place of fixed anchors, with $\inf_{\mathcal{X}_j} \|y^\star(x) - a_{k_j}(x)\| > 0$, and with the target-transversality set of (R3) read parametrically as $E_j = \{x : (x, d_j^\star(x)) \in N_{k_j}\}$. The same conclusion holds, with $y(x;\theta) \in \mathcal{Y}_x$ for every x.* (Proof: App. A.1.)

Proof of Corollary A.2. The geometric parts of the argument carry over with the following replacements and checks: $a_k \mapsto a_k(x)$ and $\underline{t}_{k,s}(d) \mapsto \underline{t}_{k,s}(x,d)$, likewise for $\bar{t}_{k,s}$.

Step 2. The polar targets become $d_l^\star(x) = (y^\star(x) - a_{k_l}(x))/\|y^\star(x) - a_{k_l}(x)\|$, continuous because $y^\star$ and a_{k_l} are and the denominator is bounded away from 0 by the corollary's hypothesis, and

$$\varphi_l^\star(x) = \frac{\|y^\star(x) - a_{k_l}(x)\| - \underline{t}_{k_l,s_l}(x, d_l^\star(x))}{\bar{t}_{k_l,s_l}(x, d_l^\star(x)) - \underline{t}_{k_l,s_l}(x, d_l^\star(x))},$$

continuous by Def. A.1(iii) composed with $x \mapsto (x, d_l^\star(x))$. Since $\mathcal{X}_l^{(\eta)}$ is compact and disjoint from the switching set, the graph $\{(x, d_l^\star(x)) : x \in \mathcal{X}_l^{(\eta)}\}$ is a compact subset of $(\mathcal{X} \times \mathbb{S}_V) \setminus N_{k_l}$, so the interval width attains its minimum there and

$$c_l := \min_{x\in\mathcal{X}_l^{(\eta)}} \left[\bar{t}_{k_l,s_l}(x, d_l^\star(x)) - \underline{t}_{k_l,s_l}(x, d_l^\star(x))\right] > 0.$$

No separate uniform-width assumption is needed; compactness supplies it, which is the reason $\mathcal{X}$ is required compact.

Step 4. The modulus of continuity is taken over the same compact subset of $\mathcal{X} \times \mathbb{S}_V$ rather than over $\mathbb{S}$ alone. Nothing else changes: the comparison is between $d_l^\star(x)$ and $\widehat{d}_l(x)$ at the *same* x, so no approximation in the x argument enters.

Step 6. Feasibility is unchanged, since the emitted point lies in an interval the oracle reports for the instance's own geometry.

Step 3. Because V is fixed (Def. A.1(i)), the same projector P_V and the same fallback vector e_V serve every x, so the total read-out N_V and the per-head bounds on $\mathcal{X}_l^{(\eta)}$ are unchanged.

Steps 1 and 5 concern regions of $\mathcal{X}$ and the routing gate and are untouched. □

Remarks. With $A = 1$ this is the radially K-accessible case; with $K = 1$, the sup-norm analogue of star-cover universality for densities. The proof touches the constraint set only through the oracle interface (interval endpoints) — no convexity or smoothness beyond (R1). The theorem gives ε-accuracy off the η-collar of S and feasibility everywhere; no approximation guarantee is claimed inside the collar. Compactness of $\mathcal{X}$ is used twice and is not cosmetic: Step 1 needs $\mathcal{X}_l^{(\eta)}$ compact, and Step 3 invokes universal approximation on compact domains.

A.2 Proof of Theorem 4.2

Part (i). Let $y \in Y_m$, $m \neq j_0$. The segment $[a, y]$ meets the open gap between components: otherwise $Y_{j_0} \cup [a, y] \cup Y_m$ would connect two disjoint compacts. The first interval of $T_{a,\mathrm{dir}(y-a)}$ ends on ∂Y_{j_0}, so $\|y - a\|$ lies beyond it: $y \notin \mathrm{Star}_1(a)$. Since every emission lies in $\mathrm{Star}_1(a)$ (the K=1 case of the architecture invariant), the output stays in Y_{j_0}; integrating the margin γ over the event of mass μ gives the gap bound. □

Part (ii): crossing-index lemma. Along the line $t \mapsto a + tv$ consider the *squared* radius $q(t) := \|a + tv\|^2$, a quadratic with $q''(t) = 2\|v\|^2 > 0$, hence strictly convex; since $R \mapsto \sqrt{R}$ is increasing, $t \mapsto \|a + tv\|$ is unimodal with the same minimiser. So on $[0, \|y - a\|]$ it decreases to the closest approach, then increases. For each radius R strictly between its one-sided extremes it therefore crosses the sphere $\|\cdot\| = R$ exactly once on the outward leg. Each intermediate ring between ring i and ring j thus contributes at least one maximal feasible interval of the ray strictly before the interval containing y, giving crossing index $\geq |j - i|$; a trajectory whose closest approach dips inward first only crosses more bands. So $y \notin \mathrm{Star}_K(a)$ whenever $|j - i| \geq K$, and the gap bound follows as in part (i). The bound is attained by some $y \in Y_j$ — any target the radius reaches monotonically — but not by every one, which is what the segment count below measures. A numerical audit over 800 anchor–target pairs at $C \in \{4, 8\}$ found zero violations of this bound; at $C = 8$, $K = 2$ the reachable feasible mass from a single anchor is 22.5% on average (25% from the best ring).

Part (ii), continued: the segment count of one anchor. We show $\max_{y \in \mathcal{Y}_x} s_a(y) = i + C$ for $a \in \mathrm{int}(Y_i)$.

Upper bound. By unimodality, $t \mapsto \|a + td\|$ is non-increasing and then non-decreasing on $t \geq 0$ (possibly increasing from $t = 0$). The preimage of one band $[R_j, R_j + w]$ therefore has at most one component on each branch, so ring j contributes at most two segments. For $j > i$ it contributes at most one: $\|a\| \leq R_i + w < R_{i+1} \leq R_j$, so the band lies strictly above the starting radius and cannot be met on the non-increasing branch. Summing, $T_{a,d}$ has at most $2i + (C - i) = i + C$ segments, for every d; counting tangential single points does not affect this bound.

Attainment. Take $d = -a/\|a\|$, so $\|a + td\| = |\|a\| - t|$. The inward leg meets rings $i, i - 1, \ldots, 1$ and the outward leg meets rings $1, 2, \ldots, C$, which is $i + C$ segments; $a \in \mathrm{int}(Y_i)$ makes the first one nondegenerate and $R_1 > 0$ keeps the two ring-1 segments from merging across the origin. The target $y = -(R_C + w)\, a/\|a\| \in Y_C$ lies in the last of them, so $s_a(y) = i + C$.

Hence $\mathrm{Star}_K(a) = \mathcal{Y}_x$ iff $K \geq i + C$. Minimising over the anchor's ring gives $(1, K) \in \mathcal{B}(\mathcal{Y}_x)$ iff $K \geq C + 1$; and since $A \leq 1$ forces $A = 1$, no element of $\mathcal{B}(\mathcal{Y}_x)$ dominates $(1, C + 1)$, so $(1, C + 1) \in \mathcal{P}(\mathcal{Y}_x)$. □

Remark (one anchor per ring at K=2 does not suffice). In $\mathbb{R}^2$, take $C = 2$, $Y_1 = \{1 \leq \|y\| \leq \frac{21}{20}\}$ and $Y_2 = \{\frac{11}{10} \leq \|y\| \leq \frac{23}{20}\}$, and let $S = \{\|y\| = \frac{9}{8}\} \subset \mathrm{int}(Y_2)$. For any single $a \in \mathrm{int}(\mathcal{Y}_x)$ and any $y \in S$ with $s_a(y) \leq 2$, the segment $[a, y]$ has minimum radius at least 1: otherwise an inner-ring anchor needs at least three segments to reach y and an outer-ring one at least four. So $\|a + y\| \geq 2$, while $\|a\|, \|y\| \leq \frac{23}{20}$, and the midpoint estimate forces $\cos\theta \geq \frac{271}{529} > \frac{1}{2}$ for the angle θ between a and y. Each anchor therefore covers an arc of S contained in an angular sector of half-width $< \pi/3$,

so *any two* anchors — one per ring or not — leave part of S uncovered, so $(2,2) \notin \mathcal{B}(\mathcal{Y}_x)$ here. One anchor per ring at $K{=}2$ is therefore not sufficient in general; we claim nothing about larger anchor counts.

A.3 Proof of Theorem 4.3

Lemma A.3 (Partial-cover universality)**.** *Write $R_x := \bigcup_{k=1}^{A} \mathrm{Star}_K^{\mathcal{Y}_x}(a_k(x)) \subseteq \mathcal{Y}_x$ for the reachable subset of the instance's own feasible set (fixed geometry is the case $R_x \equiv R$), and let $G := \{x : y^\star(x) \in R_x\}$. Assume the* same oracle regularity *on $\mathcal{Y}_x$ as Thm. 4.1 / Cor. A.2 — R_x is a reachable subset, not a new constraint geometry, and the ray endpoints still come from $\mathcal{Y}_x$ — and assume (R2)–(R3) only on the covered region G, read as: G is a finite union of regions carrying that structure, and the switching set of Thm. 4.1 is enlarged by ∂G. Then for every $\varepsilon, \eta > 0$ there are* PRISM *parameters with $\|y(x;\theta) - y^\star(x)\| \le \varepsilon$ for every $x \in G$ outside the η-collar, while $y(x;\theta) \in \mathcal{Y}_x$ for* all *x — including $x \notin G$, where no accuracy is claimed and the gate simply returns some feasible candidate.*

Proof of Lemma A.3. Steps 1–5 of App. A.1 never use that the anchors cover the whole set: they only need, on each region, an anchor/segment assignment for the target and the regularity of the corresponding intervals, which is the hypothesis here restricted to G. Run them on the regions of G. Step 6 is unchanged — the total read-out N_V gives every head a unit direction at every x, and $a_k(x) \in \mathrm{relint}(\mathcal{Y}_x)$ gives a first segment of positive length for that instance — and yields feasibility everywhere, since every candidate still lies in an oracle-certified interval of $\mathcal{Y}_x$ regardless of whether $x \in G$. □

Part (i). The anchor maps — hence the reachable sets R_x — are fixed relative to the held-out draw, so with $x_i \overset{\text{i.i.d.}}{\sim} \mathcal{D}$ the indicators $\mathbf{1}[y^\star(x_i) \in R_{x_i}]$ are i.i.d. Bernoulli(p) with $p := \mathbb{P}_{x\sim\mathcal{D}}[y^\star(x) \in R_x]$. Hoeffding's inequality applied to their mean $\hat{c}$ gives $\mathbb{P}\big(p < \hat{c} - t\big) \le e^{-2mt^2}$; set $t = \sqrt{\ln(1/\delta)/(2m)}$. If instead the configuration is chosen from a finite family $\{R_{1,x}, \dots, R_{C,x}\}$ fixed before the draw, by comparing their empirical coverages on the same sample, apply the bound to each $R_{c,x}$ with failure probability δ/C and take a union bound: with probability $\ge 1-\delta$, *every* $R_{c,x}$ — hence the selected one — satisfies $p_c \ge \hat{c}_c - \sqrt{\ln(C/\delta)/(2m)}$. Since the guarantee is simultaneous over the family, the rule used to pick among them is irrelevant to its validity. □

What the bound is about. It bounds $p := \mathbb{P}_{x\sim\mathcal{D}}(y^\star(x) \in R_x)$, a property of the target distribution, not membership of $(\widehat{A}, \widehat{K})$ in $\mathcal{B}(\mathcal{Y}_x)$: it says nothing about parts of $\mathcal{Y}_x$ that carry no mass under $\mathcal{D}$. Nor does $\hat{c} = 1$ give $p = 1$; at $m = 400$, $\delta = 0.05$ and $C = 2$ the deviation term is 0.068, so a perfect held-out count gives $p \ge 0.932$ at 95% confidence, under the sampling and exact-membership conditions above.

Greedy anchor count. If some A^* anchors cover all n selection samples and each greedy round picks a candidate covering at least γ times the best marginal coverage, then $T \ge (A^*/\gamma)\ln(1/\varepsilon')$ rounds reach coverage $\ge 1 - \varepsilon'$ (Chvatal, 1979).

Proof of the greedy bound. At any round, the optimal cover's A^* sets jointly cover all u_t currently uncovered samples, so some single candidate covers at least u_t/A^* of them; the γ-approximate pick covers at least $\gamma u_t/A^*$, whence $u_{t+1} \le u_t(1 - \gamma/A^*)$ and $u_T \le n(1-\gamma/A^*)^T \le ne^{-\gamma T/A^*}$ (Chvatal, 1979). Our implementation restricts each round's pool to the deepest uncovered samples plus the uncovered centroid; $\gamma < 1$ quantifies that restriction and is not estimated theoretically.

Part (ii). Split the expectation on the event $\{y^\star(x) \in R_x,\ x \notin B_\eta\}$, whose probability is $\ge 1-(u+\nu)$ by Part (i), where Lem. A.3 and Lipschitzness bound the integrand by $L_f\varepsilon$, and its complement of probability $\le u + \nu$, where the integrand is at most $L_f \bar{D}$ because both points lie in $\mathcal{Y}_x$. □

A.4 Curvature-Safe Pullback

The linearised oracle of the legacy AC-OPF pipeline (App. B, Eq. 8) is exact for affine constraints and approximate otherwise. The following gives a sufficient condition under which the pullback keeps the emitted point inside the true feasible set. It is stated here rather than in the main text because no result in the paper depends on it: the closed-form classes solve the true intersection and never invoke

it, and the AC-OPF pipeline does not evaluate its curvature bound, so feasibility there is verified empirically rather than certified.

Proposition A.4 (Curvature-safe pullback). *Let each $g_i(\cdot\,; x)$ be C^2 along the ray with $d^\top \nabla_y^2 g_i(a + t\,d)\,d \le H$ for all $t \in [0, \widehat{\rho}^\star]$ and some $H \ge 0$, and let $\gamma := \min_i\big(-g_i(a)\big) > 0$ be the anchor margin, and assume $0 < \widehat{\rho}^\star < \infty$ (the ray does leave the linearised set). Write $y = a + \beta\widehat{\rho}^\star d$. Then for any $\beta \in (0, 1)$,*

$$\tfrac{1}{2}\beta^2(\widehat{\rho}^\star)^2 H \le (1-\beta)\gamma \implies g_i(y; x) \le 0\ \forall i, \qquad \tfrac{1}{2}\beta^2(\widehat{\rho}^\star)^2 H < (1-\beta)\gamma \implies g_i(y; x) < 0\ \forall i.$$

For affine constraints ($H = 0$) the strict form holds for every $\beta < 1$, and the linearised crossing is the true one. (Proof: App. A.4.)

Proof. Fix i and write $s_i := \nabla_y g_i(a)\cdot d$, $\rho_\beta := \beta\widehat{\rho}^\star$. Taylor's theorem with the Lagrange remainder along the ray gives, for some $\xi \in (0, \rho_\beta)$,

$$g_i(a + \rho_\beta d) = g_i(a) + \rho_\beta\, s_i + \tfrac{1}{2}\rho_\beta^2\, d^\top \nabla_y^2 g_i(a + \xi d)\, d \;\le\; g_i(a) + \rho_\beta\, s_i + \tfrac{1}{2}\rho_\beta^2 H.$$

Case $s_i > 0$: by definition of the linearised estimate, $\widehat{\rho}^\star \le \widehat{\rho}_i = -g_i(a)/s_i$, hence $\rho_\beta\, s_i \le \beta\,\widehat{\rho}_i\, s_i = -\beta\, g_i(a)$ and

$$g_i(a + \rho_\beta d) \;\le\; (1 - \beta)\, g_i(a) + \tfrac{1}{2}\beta^2(\widehat{\rho}^\star)^2 H \;\le\; -(1-\beta)\gamma + \tfrac{1}{2}\beta^2(\widehat{\rho}^\star)^2 H \;<\; 0,$$

where the last inequality is the assumed condition: under the non-strict hypothesis this yields $g_i(y; x) \le 0$, and under the strict hypothesis $g_i(y; x) < 0$, which is exactly the dichotomy in the statement. *Case $s_i \le 0$:* the linear term is non-positive, so $g_i(a + \rho_\beta d) \le g_i(a) + \tfrac{1}{2}\beta^2(\widehat{\rho}^\star)^2 H \le -\gamma + (1-\beta)\gamma < 0$ for $\beta \in (0, 1)$, using the assumed condition again. For affine g_i, $H = 0$ and both cases give $g_i \le (1 - \beta)g_i(a) < 0$ for every $\beta < 1$. □

B AC-OPF SPECIALISATION

This appendix separates three AC-OPF pipelines (App. B.3): the *legacy* first-order pipeline behind the surrogate rows of Tab. 7 (Apps. B.1–B.2, kept as a historical record and carrying no current feasibility or speed claim), a fast verified numerical service, and the *certified* local-path mode of Sec. 5.3 (Apps. B.5–B.7).

The AC-OPF problem on an n-bus system is

$$\begin{aligned} \min_{P_g, Q_g, V_m, V_a} \quad & \textstyle\sum_i c_{2,i} P_{g,i}^2 + c_{1,i} P_{g,i} + c_{0,i} \\ \text{s.t.} \quad & S_{\text{bus}}(V) - (S_g - S_d(x)) = 0, \quad \text{(power flow, complex equality)} \\ & \underline{P}_g \le P_g \le \overline{P}_g,\ \underline{Q}_g \le Q_g \le \overline{Q}_g, \\ & \underline{V}_m \le V_m \le \overline{V}_m,\ |S_f|, |S_t| \le \overline{S}_{\text{br}}, \end{aligned} \tag{5}$$

with bus complex voltages $V \in \mathbb{C}^n$ (parametrised by magnitude V_m and angle V_a), generator real/reactive power P_g, Q_g, the bus demand $S_d(x)$ as the parametric input x, and branch flow bounds $\overline{S}_{\text{br}}$. The power-flow equation 5 is a non-convex, transcendental *equality* that classical L2O methods cannot enforce directly. We handle it by *partial-output completion*: PRISM produces only y_{part}, the dependent variables are recovered by a Newton power-flow solve, and the inequality Jacobian used by the intersection oracle is propagated through the dependence by the implicit-function theorem.

B.1 LEGACY PIPELINE: PARTIAL-OUTPUT DECOMPOSITION AND IFT JACOBIAN

The decision vector is split into a *partial* component (predicted by PRISM) and a *dependent* component (solved from the power flow):

$$y_{\text{part}} = (P_g^{\text{non-slack}},\ V_m^{\text{gen}}) \in \mathbb{R}^d, \qquad y_{\text{dep}} = (P_g^{\text{slack}},\ Q_g,\ V_m^{\text{load}},\ V_a) \in \mathbb{R}^{D-d}. \tag{6}$$

For IEEE 30/57/118 buses, $d = 11, 13, 107$, respectively. Given y_{part} and demand $S_d(x)$, the power-flow equation 5 is completed for y_{dep} by a Newton iteration on V on the converged local branch

(a converged completion, not a global-uniqueness statement); we use `pypower`'s `runpf` with a cached Y_{bus} matrix for speed.

All inequalities are evaluated on the *full* $y = (y_{\text{part}}, y_{\text{dep}})$. The ray-intersection oracle thus needs $\partial g/\partial y_{\text{part}}$. By the implicit-function theorem applied to $h(y_{\text{part}}, y_{\text{dep}}; x) = 0$:

$$\frac{\partial y_{\text{dep}}}{\partial y_{\text{part}}} = -\Big(\frac{\partial h}{\partial y_{\text{dep}}}\Big)^{-1} \frac{\partial h}{\partial y_{\text{part}}}, \qquad \boxed{\frac{\partial g}{\partial y_{\text{part}}} = \frac{\partial g}{\partial y_{\text{part}}}^{\text{direct}} + \frac{\partial g}{\partial y_{\text{dep}}} \frac{\partial y_{\text{dep}}}{\partial y_{\text{part}}}}. \tag{7}$$

The Jacobian $\partial h/\partial V$ is provided in closed form by `pypower`'s `dSbus_dV`; branch-flow derivatives come from `dSbr_dV`. Assembling equation 7 is a single sparse linear solve of size $D-d$, performed once at the per-sample anchor. This replaces the central-difference Jacobian (which required d extra PFs) with a single sparse solve.

Anchor source of the legacy rows. The per-instance interior anchor of the three legacy rows came from an anchor service that runs before the network and outside it: a nearest-neighbour lookup over margin-OPF training solutions, checked for strict interiority under the box-residual criterion and, on failure, repaired by a short chain of numerical fallbacks ending in a margin-OPF solve. This service is not the measurement pass of Sec. 3.3, its cost is not part of any reported timing, and the certified mode of Sec. 5.3 does not use it.

The assembled Jacobian $J \in \mathbb{R}^{m \times d}$ is the same object passed to a linearised polytope intersection oracle of the form of Sec. 3.1:

$$\rho^\star \;=\; \min_{i:\,(J_i \cdot d) > 0} \frac{-g_i(a)}{J_i \cdot d} \cdot \beta, \tag{8}$$

with $\beta = 0.95$ pullback. We are explicit about what this does and does not give. The pipeline behind the three legacy rows runs this first-order oracle only: it does not estimate the curvature bound H of Prop. A.4, so it does not certify that $\frac{1}{2}\beta^2(\rho^\star)^2 H \le (1-\beta)\gamma$ holds, and the $\beta = 0.95$ pullback is a margin chosen empirically rather than one the proposition licenses. What the legacy rows record instead is a *box residual*: the violation metric stored for the AC-OPF columns of Tab. 7 is evaluated in the normalised (P_g, V_m) coordinates after the power-flow completion and is 0 on all evaluated samples of case-30/57/118, but it does not evaluate the full AC constraint set. Those rows therefore support neither a guaranteed nor a fully verified AC feasibility claim; they are legacy breadth evidence. The full-constraint evidence is the certified mode of App. B.5. The structural guarantee of Sec. 3.1 applies to the exact closed-form oracles, which compute the true ray crossings — for affine constraints because the linearisation is exact, and for the quadratic, conic and geometric classes because the intersection is solved in closed form — and therefore do not rely on Prop. A.4 at all.

B.2 Legacy Pipeline: Equality Projection on AC-OPF

The power-flow equation 5 is a *nonlinear* equality constraint; the linear null-space projector of Sec. 3.1 does not directly apply. Instead, we handle the equality by completion + IFT Jacobian as described above: y_{dep} is completed by Newton-PF from y_{part}, so $h(y) = 0$ holds to solver tolerance whenever the completion converges on the selected regular branch, for whatever y_{part} is predicted; this legacy pipeline is evaluated only under its box-residual criterion (App. B.3). The IFT Jacobian $\partial g/\partial y_{\text{part}}$ then linearises the inequalities in the partial-output space, where the standard PRISM polar+intersection machinery applies.

For benchmarks with purely *linear* equality (QP-eq, SOCP-eq, QCQP-eq, simplex QP), we use the simpler null-space projection of Sec. 3.1: $\mathbf{P}_A = I - A^\top (AA^\top)^{-1} A$ projects directions, and the per-sample anchor a is selected on the affine slice $\{y : Ay = b\}$ by min-norm $a = A^\top (AA^\top)^{-1} b$.

B.3 Three Layers of AC-OPF Evidence

The AC-OPF material in this paper comes from three pipelines that must not be read as one. **(L1) Legacy surrogate rows.** The three AC-OPF columns of Tab. 7, Fig. 2(b) and Tab. 6 score the normalised-coordinate surrogate $\frac{1}{2}\|y - y^\star\|^2$, and their recorded violation metric is a *box residual in the normalised* (P_g, V_m) *coordinates*, not an evaluation of the full AC constraint set. They are

kept as breadth evidence only; they do not support a full AC feasibility claim, and no timing of that legacy pipeline is reported. **(L2) Fast verified numerical pipeline (case30).** A numerical service that returns a point only after it passes a full-state check (power-flow residual $\leq 10^{-9}$, all operational inequalities, control error $\leq 10^{-9}$): on the 64 sealed report loads it returned $57/64$ verified points at a mean generation-cost premium of 2.214% over a per-load local reference solve, in 2.204 s of service time for the batch of 64 against 28.795 s for the native reference solves on the same machine and workload. This is a verified numerical service, not a formal geometric guarantee, and its timing is *not* inherited by L3. **(L3) Certified re-anchored paths (case30).** The mode of Sec. 5.3: every accepted update carries a machine-checked certificate of a continuous same-branch root path with all operational inequalities strict and an exact cost decrease. Its evidence is the subject of the rest of this appendix; the implementation contract behind each hypothesis is listed in Tab. 4.

B.4 Scope of the Certified Object

Everything below is *conditional*: it states what an accepted and independently re-verified certificate proves for the fixed rational case30 model. It does not assert that a certificate exists for every load, base root or direction, nor that a re-anchoring or cost gate will succeed. Two objects must not be confused. The *certified object* is the exact root of the rational model inside a stated box, together with a continuous root path along a control segment; the *float export* is a binary64 representative of that root, checked separately by the numerical FINAL contract and never claimed to satisfy the equalities exactly. On the sealed report, 54 of 64 loads had a certified initial root; the remaining 10 are typed `UNAVAILABLE`, which is a coverage statement about this bounded service, not a physical feasibility rate and not a proof of infeasibility. A trajectory marked `COMPLETE_8_STEPS` spent its budget of 8 update calls; the number of *accepted* moves per trajectory ranged from 1 to 8. The learned direction only proposes; validity of the statements below neither depends on nor implies any superiority of the learned policy over the numeric gradient rule. No timing of the earlier fast numerical service is inherited by the certified mode.

B.5 Certified Local Feasible Paths: Objects

(D1) Fixed exact model. The case30 data (admittances, generator and voltage limits, branch limits, quadratic cost) are rational numbers frozen once from the reference floats (`ExactCase30`); the model manifest hash is written into every root record (`case_manifest_sha`). The independent verifier re-encodes the model in its own class (`VCase`). Loads $(P_d, Q_d) \in \mathbb{Q}^{30} \times \mathbb{Q}^{30}$ are bound to every record by hash.

(D2) State, controls, scalings. $q = (e_{1..30}, f_{\text{nonref}}, P_g, Q_g) \in \mathbb{R}^{71}$; $u \in \mathbb{R}^{11}$ = five non-reference P_g set-points and six generator voltage magnitudes; $D_q \in \mathbb{Q}^{71}_{>0}$, $D_u \in \mathbb{Q}^{11}_{>0}$, m_u; $u = m_u + D_u \odot z$.

(D3) Squared system. $F(q; u, P_d, Q_d) \in \mathbb{R}^{71}$ stacks the 60 real/imaginary power-balance rows, the five rows $P_{g,i} - u_i$ and the six rows $e_k^2 + f_k^2 - u_k^2$. F is a polynomial of degree ≤ 2 in (q, u) with rational coefficients; F, F_q are evaluated exactly in ball arithmetic and F_q is affine in q.

(D4) Control segment. For $z_0 \in \mathbb{Q}^{11}$, a frozen (dyadic) direction $d \in \mathbb{Q}^{11}$ and $t \in [0, \bar{h}]$: $u(t) = m_u + D_u \odot (z_0 + t\, d)$. This is a segment in control space; it is not an exponential map or a geodesic.

(D5) Cell. $(t_L, t_R, t_c, w, p, C_0, C_1, R, \text{prec})$ with rational $t_L \leq t_c \leq t_R$ and $w \in \mathbb{Q}$ chosen by outward rational rounding so that $[t_L, t_R] \subseteq [t_c - w, t_c + w]$; dyadic predictor $p(s) = \sum_{k \leq 2} p_k s^k$; dyadic $C(s) = C_0 + C_1 s$; radius vector $R \in \mathbb{Q}^{71}_{>0}$ (the float actually used) with exact scaled radius $r_i = R_i / D_{q,i}$. Write $s = t - t_c$, $\mathcal{B} = \{\xi : |\xi_i| \leq r_i\}$, $q = p(s) + D_q \odot \xi$ and $\mathcal{A}(s, \xi) = I - C(s) F_q(p(s) + D_q \xi; u(t_c + s)) D_q$.

(D6) Certified quantities. Rigorous ball-arithmetic upper bounds $b_i \geq \sup_{|s| \leq w} |(C(s) F(p(s); u(t_c + s)))_i|$ and $B_{ij} \geq \sup_{|s| \leq w,\, \xi \in \mathcal{B}} |\mathcal{A}(s, \xi)_{ij}|$ (the latter as $|I - P_0 - N_0| + |P_1 + N_1| w + \sum_{k \geq 2} |P_k| w^k$ with the exact polynomial matrices P_k and the box terms N_a). `contraction` holds iff

$$b_i + (Br)_i < r_i \;\; \forall i, \qquad \kappa := \max_i (Br)_i / r_i < 1, \tag{9}$$

with exact rational r and strict comparisons of ball upper ends. `feasible` holds iff, over the whole box and $|s| \leq w$, each of the 166 operational inequalities $g_j(q) < 0$ has a negative rigorous upper

bound, e_{ref} a positive lower bound and each requested voltage magnitude a positive lower bound. A cell is *accepted* iff both hold.

B.6 Conditional Statements

Lemma B.1 (one cell: contraction and feasibility)**.** *If a cell satisfies equation 9, then for every $s \in [-w, w]$ there is exactly one $q^*(s) \in p(s) + D_q\mathcal{B}$ with $F(q^*(s); u(t_c + s)) = 0$; $s \mapsto q^*(s)$ is continuous; $C(s)$ and $F_q(q^*(s))$ are invertible. If the cell is also* `feasible`*, $q^*(s)$ satisfies all* 166 *inequalities strictly, $e_{\mathrm{ref}} > 0$ and the requested voltage magnitudes are positive, for every $s \in [-w, w]$.*

Proof. Fix s; $K(\xi) = \xi - C(s)F(p(s) + D_q\xi; u(t_c + s))$. By the mean value theorem on the segment from 0 to $\xi \in \mathcal{B}$, $K(\xi) = -C(s)F(p(s)) + \int_0^1 \mathcal{A}(s, \theta\xi)d\theta\,\xi$, so $|K(\xi)_i| \le b_i + (Br)_i < r_i$ and $K(\mathcal{B}) \subset \mathcal{B}$. For $\xi, \xi' \in \mathcal{B}$, $|K(\xi) - K(\xi')| \le B|\xi - \xi'|$ componentwise, hence in $\|\eta\|_r = \max_i |\eta_i|/r_i$, $\|K(\xi) - K(\xi')\|_r \le \kappa\|\xi - \xi'\|_r$, $\kappa < 1$. Banach gives a unique fixed point $\xi^*(s)$, continuous in s because the family is continuous in s and uniformly contracting. For every q in the box $\|I - C(s)F_q(q)D_q\|_{r\to r} \le \kappa < 1$, so $C(s)F_q(q)D_q$ is invertible (Neumann series) and hence each factor $C(s)$, $F_q(q)$ is invertible; therefore $C(s)F(q^*(s)) = 0 \Rightarrow F(q^*(s)) = 0$ — this step is what makes the preconditioned residual a root certificate rather than a certificate of $CF = 0$. Any other root in the box would be a second fixed point. Feasibility holds because the bounds were taken over the whole box, which contains $q^*(s)$. □

Lemma B.2 (seam: root identity between adjacent cells)**.** *Let two accepted cells be adjacent at $t = t_R^{(1)} = t_L^{(2)}$. If the box of one cell at t contains the box of the other, or a third box centred at a Newton point satisfies equation 9 and contains both, then $q^{*(1)}(t) = q^{*(2)}(t)$ and the concatenated path is continuous.*

Proof. Both roots lie in a box in which the root is unique by Lemma B.1 (with $w = 0$ for the common box). □

Proposition B.3 (certified prefix)**.** *Let ρ_0 be a certified root record (a cell with $w = 0$ at z_0 whose box contains the unique root q_0 of $F(\cdot; u(0))$ with all inequalities strict) and let accepted cells $c_1, \ldots, c_m$ cover $[0, \bar{h}]$ with accepted seams, the first seam-linked to ρ_0. Then there is a continuous $q^* : [0, \bar{h}] \to \mathbb{R}^{71}$ with $q^*(0) = q_0$, $F(q^*(t); u(t)) = 0$, $F_q(q^*(t))$ invertible and all* 166 *inequalities strict for every $t \in [0, \bar{h}]$, unique inside the union of the cell boxes.*

Proof. Lemma B.1 on each cell over its covering domain $[t_c - w, t_c + w] \supseteq [t_L, t_R]$; Lemma B.2 at each seam and at $t = 0$. The outward covering places the nominal interval endpoints inside proved domains. □

Proposition B.4 (re-anchoring)**.** *Let $t_{\mathrm{new}} \in [0, \bar{h}]$ lie (exactly) in the proved domain of a cell c of an accepted prefix and let ρ_1 be a new point root (cell with $w = 0$, own C,* `contraction` *and* `feasible`*) whose box and the box of c at t_{new} are nested. Then ρ_1 represents $q^*(t_{\mathrm{new}})$, its control is exactly $z_1 = z_0 + t_{\mathrm{new}}d \in \mathbb{Q}^{11}$, and it is strictly feasible.*

Proof. Uniqueness of the root in the containing box (Lemma B.1). □

Proposition B.5 (exact cost gate)**.** *Let $[\ell_0, h_0]$ and $[\ell_1, h_1]$ be the exact dyadic cost bounds of ρ_0, ρ_1 (ball arithmetic of the rational quadratic cost over the root boxes, ends read exactly) and $\tau = \max(10^{-6}, 10^{-8}|(\ell_0 + h_0)/2|) \in \mathbb{Q}$. If $h_1 < \ell_0 - \tau$ in exact rational arithmetic, then $\mathrm{cost}(q(\rho_1)) < \mathrm{cost}(q(\rho_0)) - \tau$. The verifier additionally requires the saved bounds to contain its recomputed bounds.*

Theorem B.6 (finite composition: what an accepted trajectory proves)**.** *Consider a trajectory with certified initial root ρ_0 and a budget of* 8 *update calls. In each call a direction d and a requested length are proposed by a numeric gradient rule or by a frozen network; the base root is ρ_0 in FIXED mode and the current root in MOVE mode. The call is* accepted *iff a prefix as in Prop. B.3 was certified from the base root, a re-anchoring as in Prop. B.4 succeeded at one of the parameters $\bar{h}2^{-k}$, $k \in \{0, 1, 2, 3\}$, and the gate of Prop. B.5 holds against the current root; otherwise (no contraction,*

infeasible, seam failure, budget exhausted, Newton failure, unresolved or failed gate) the current root is unchanged. Then, along every trajectory, (i) each accepted root is the unique root of the fixed rational model in its stated box at an exactly known control; (ii) consecutive accepted roots are joined by a continuous, everywhere strictly feasible root path; (iii) the exact costs of accepted roots decrease strictly by at least τ per accepted call; (iv) the final record is bound to the last accepted root. The number of accepted calls is between 0 *and* 8*.*

Proof. Induction over calls with Props. B.3–B.5; rejected calls leave the state unchanged, so the invariant "current root certified, strictly feasible, at an exact control" is preserved. □

Remark B.7 (conditional coverage and failure semantics). No statement above guarantees that a certificate with $\bar{h} > 0$ exists for a given load, base and direction, that a re-anchoring or gate succeeds, or that an initial root is available. Radii come from a finite menu, cells from adaptive bisection under a box budget, and every failure is a typed status (`UNAVAILABLE`, `BUDGET`, `CERTIFIED_NOT_BETTER`, `UNRESOLVED_OVERLAP`, ...) that keeps the previous certified root. Coverage on the sealed report was $54/64$ initial roots; all 6 methods completed their 8-call budget on those 54; accepted calls per trajectory ranged 1–8. The straight-ray results of the main text are not transferred to these implicit paths.

Remark B.8 (evidence layers). Interval inequalities and feasibility bounds were computed by the certifier and re-derived by an independent verifier (separate process, own model encoder, own preconditioner proposals); all 672 delivered chains have a passing final receipt (674 receipt records; two record a historical endpoint-binding failure beside its corrected rebinding). A read-only audit additionally replayed the exact metadata (domains, radii scaling, control updates, cost gates, bindings) but did not rerun the interval computations; the layers are not two independent interval reproductions. Float exports were checked only by the numerical FINAL contract.

B.7 Protocol and Results

Protocol. 64 sealed report loads of case30 (the same loads for every method), a certified initial root obtained from a numerical anchor for 54 of them, and six certified methods run on every root with a budget of 8 update calls each: the base is either fixed at the initial root (FIXED) or moved to the last accepted root (MOVE), and the direction comes either from a numeric cost-gradient rule (GRAD) or from one of two frozen $8{,}972$-parameter policies trained with the same recipe from seeds 580045 and 580145 (NN). All 324 trajectories spent their full call budget; the number of *accepted* calls per trajectory ranges from 1 to 8 (66 trajectories accepted all 8), so "8 calls" is a budget, not an acceptance count. Quality is the generation-cost gap to a per-load local reference solve (not a global optimum), reported as the mean over the 54 certified roots; the ten unavailable loads stay in every denominator.

Table 3: **Certified re-anchored mode on case30**, 64 sealed report loads, 54 certified initial roots. Gap = mean generation-cost gap to the local reference (%); "paired" columns restrict to the loads on which the corrected DC3 adaptation of the given seed returned a FINAL-valid state (n=20 and n=26). Time is the conditional mean per complete 8-call trajectory on two CPU threads; it includes a state Jacobian/IFT solve in every call and mixes records from two runs, so it is a certificate cost, not a deployment latency, and is not comparable to the L2 service time of App. B.3.

Method	returned	gap (%)	median (%)	paired n=20	paired n=26	s/traj.
MOVE-GRAD	$54/64$	1.428	1.341	0.439	0.595	0.95
FIXED-GRAD	$54/64$	1.496	1.419	0.490	0.662	4.13
MOVE-NN, seed 580045	$54/64$	1.550	1.290	0.547	0.676	1.52
MOVE-NN, seed 580145	$54/64$	1.666	1.430	0.949	0.996	1.50
FIXED-NN, seed 580045	$54/64$	1.616	1.432	1.035	1.045	3.91
FIXED-NN, seed 580145	$54/64$	1.558	1.295	0.919	0.954	3.43
DC3 corrected adapt., seed 580045	$20/64$	—	—	4.986	—	—
DC3 corrected adapt., seed 580145	$26/64$	—	—	—	5.787	—

Comparison with a feasibility-enforcing network. The comparator is our own adaptation of DC3 (Donti et al., 2021) to the full case30 constraint set (completion through the power flow, gradient

correction, 30 epochs, two seeds, the official loss weights); a decoding defect in an earlier run was corrected before these numbers were produced. It is a *30-epoch, full-constraint, bounded adaptation*, not a reproduction of the original method at its tuned training budget, and its correction settings changed the pass counts by nothing on the selection split. Its returned states were audited numerically with the same full-state contract as ours, but the point-root certificate objects for those states were not retained (0 saved), so it does *not* carry a path certificate comparable to L3. Comparing only on the loads each DC3 seed actually returned, the corresponding certified MOVE-NN has lower cost on $20/20$ and $26/26$ loads (mean gap 0.547% vs. 4.986% and 0.996% vs. 5.787%), while returning $54/64$ against $20/64$ and $26/64$. An assisted variant that falls back to the existing numerical anchor when DC3 fails returns $54/64$ at 5.25% and $54/64$ at 5.54%; it is a different front end and is listed separately. None of this is a statement about feasibility-enforcing networks in general.

Learning gate (negative, retained). The pre-registered gate compares the improvement of MOVE-NN over FIXED-NN per seed on the 54 common loads against a threshold frozen before the run: $+0.063$ pp against 0.187 pp for seed 580045 (MOVE better on $19/54$ loads) and -0.102 pp against 0.189 pp for seed 580145 (MOVE better on $16/54$). Both fail and the signs differ, so no learning promotion is claimed; the frozen MOVE policy of seed 580045 is moreover the *initial* checkpoint of its selection, so its result is not attributable to training. The numeric gradient rule MOVE-GRAD has the lowest mean gap of the six. The stable result is the geometry — certificate, re-anchoring and the numeric direction — not the learned direction.

The role of re-anchoring in the nonlinear-equality extension. On a static or affine geometry a straight ray from a valid anchor is exact, so the generic method routes over a fixed bank and Sec. 4 applies. Under the power-flow equality a straight control segment lifts to a curved state path whose feasible length is not known in advance, so the AC specialization certifies a local path $\Gamma_{a,d}(t)$, $t \in [0, \bar{h}]$, and, in MOVE mode, takes the certified, cost-gated endpoint as the next local origin, $a_{k+1} = q(t^\star)$; FIXED mode keeps its base and is also a tested certified mode (Tab. 3). Re-anchoring is thus a step of this specialization: not a mathematical necessity, not a component of the generic method, not covered by the straight-ray theorems, and not adopted elsewhere.

Cost scope. The conditional times in Tab. 3 are the cost of certifying, on two CPU threads, with a state IFT/gradient solve inside every call; no unified, repeated end-to-end deployment timing of the certified mode exists, and no speed ratio against a reference solver is claimed for it. The L2 service ratio (0.077 of the native reference service time on the same batch workload) belongs to another pipeline and is not inherited.

Table 4: **Implementation contract** of the certified specialization: each hypothesis of Props. B.3–B.5 and Thm. B.6 is implemented by a named routine of the certifier and re-checked by the independent verifier; the evidence layer says whether the check is an interval inequality, an exact rational replay of stored metadata, a float-only numerical check, or an availability statement without a theorem. Source hashes and receipt paths are in the certificate records released upon publication.

Hypothesis / check	certifier routine (re-checked by the verifier)	evidence layer
fixed exact rational case30 model; manifest hash bound in every root	R58/exact_case.py:ExactCase30 (L35), manifest (L68)	interval / exact
71-row squared system F (60 H + 5 Pg + 6 Vmˆ2), loads bound by sha	R58/exact_case.py:ArbCase.F_polys (L144), Fq_ball (L153)	interval
control segment u(t)=mid+Du(z0+t d), exact rational control update z1=z0+t d	certify_r59.py:Certifier.u_of (L52), recenter z_new (L203ff)	exact metadata
outward covering [tL,tR] within [tc-w,tc+w]; exact in_domain	certify_r59.py:w_cover_of (L30), in_domain (L35)	exact metadata
true q-radius R (float used) and exact scaled r=R/Dq in the inequality	certify_r59.py:Certifier.radii (L68)	interval / exact
contraction b+Br<r, kappa<1 with C(s)=C0+C1 s and exact polynomial matrices	certify_r59.py:Certifier.CellProof.__init__/test (L70-110)	interval
strict feasibility of 166 inequalities, e_ref>0, requested Vm>0 over the whole segment	R58/exact_case.py:ArbCase.ineq_polys (L171), poly_upper/lower	interval
seam: containment or common uniqueness box	certify_r59.py:Certifier.seam (L156), box_at (L151)	interval + exact metadata
recenter: t_new in proved domain; branch identity by nested boxes	certify_r59.py:Certifier.recenter (L203)	interval + exact metadata
exact cost gate upper(new) < lower(old) - tau in Fraction arithmetic; saved bounds	certify_r59.py:Certifier.cost_gate (L220), cost_bounds_exact	exact metadata
composition: accepted ⇒ certified root ∧ connected prefix ∧ recenter in domain	arms_r59.py:certified_step (L98), run_certified (L130)	exact metadata / composition
proposes only (direction, rho from frozen Policy58 or numeric gradient)	arms_r59.py:direction_and_rho (L56)	code path
export checked separately by numerical FINAL contract; certified object is the exact root	arms_r59.py:numeric_export59 (L92)	float numerical
budget & typed unavailability (8 calls; 1000-box reservations)	certify_r59.py:Budget (L38), try_menu (L114)	availability (no theorem)

C THE MEASUREMENT PASS

Inputs: feasible samples $X = \{y_i\}_{i=1}^n$ (training labels), a ray oracle returning sorted boundary crossings from arbitrary interior points, the problem's violation function g, caps $A_{\max}$ and the segment grid $\{1, 2\}$. All measurement calls run at $\beta = 1$ in double precision — the *exact* oracle for the closed-form classes, and the *unpulled linearised* oracle on AC-OPF.

1. **Interiorisation.** Labels produced by projection sit on the boundary (often a float rounding *outside* it). For each sample, propose two candidate landings per probe direction — the exit reading (step $\frac{1}{4}$ of the first crossing) and the fold reading (step past a first crossing within numerical noise, into the following interval) — verify each landing with g, and move the sample to the feasible landing of largest free reach; keep it unmoved if none verifies.
2. **Components.** Build the Euclidean MST of the nudged samples; sever every edge whose first-exit distance from one endpoint toward the other is shorter than the edge (tolerance at sample-spacing scale). Connected components of the surviving graph are the component budget K_0; on thin curved sets this measures chord-connectivity, over-counting only in the direction that costs extra (certificate-audited) anchors.
3. **Greedy anchors (per component, largest first).** Candidate pool = deepest uncovered samples (boundary depth = min oracle exit over 64 directions) plus the uncovered centroid; pick the candidate that newly covers the most uncovered samples, where sample s is covered by anchor a iff $\|s - a\|$ lies in one of the first K intervals of the ray $a \to s$; repeat until component coverage ≥ 0.99 or $A_{\max}$. Verify every anchor strictly feasible with g; drop failures.
4. **Certificate.** Fraction of held-out feasible samples covered by ≥ 1 anchor. Write y_i^{ref} for the label actually tested. Under the sampling and pre-fixed-family conditions of Thm. 4.3(i), the count bounds the probability that y^{ref} is reachable. Reading that as a bound on reachable *optima* — which is what Part (ii) needs — requires that the tested label *is* the optimum, $y^{\mathrm{ref}} = y^\star$, and that membership in the first K intervals is evaluated for that same point; agreement only up to a tolerance does not transfer the event, because reachability is a membership condition that two points closer than the tolerance can fail to share. Solver provenance alone does not establish either; we report the coverage of the evaluated labels and do not claim a certified optimum.
5. **Capacity probe.** Among grid configurations within 0.05 of the best certificate, train each for 10 epochs on a 512-sample subset and score the objective on a train-tail slice; keep the best. The probe never sees the reporting validation split.

On the static suite the pass costs $\sim$1 s per problem (closed-form oracles) and 20–70 s per AC-OPF instance (one power flow plus one analytical Jacobian per source point). Code: `models/auto_anchor.py` in the code release (upon publication).

Reachable-mass audit for Thm. 4.2(ii). Exhaustive scan on nested annuli: 800 anchor–target pairs, zero violations of the crossing-index bound, minimum slack 0 (tight). Single-anchor reachable feasible mass:

rings C	K=1	K=2	K=4
4	0.135	0.367	0.810
8	0.138	0.225	0.468

D EXTENDED EMPIRICAL RESULTS

D.1 BENCHMARK DEFINITIONS AND METRICS

Problems and metrics. We evaluate on a 15-problem suite spanning four categories: **(i) Convex** (3 problems) — Polygon-2D (6-sided polygon QP), QCQP-100D, SOCP-100D; **(ii) Non-convex** (4) — TwoDisks (union of two disjoint disks), Annulus (non-simply-connected), $L_{0.5}$-ball-8D (cusped at axes), Sin-QP ($\min\ \frac{1}{2} y^\top Q y + 30\, p^\top \sin y$ on a 4-vertex polygon: a non-convex objective on a convex feasible set. The sinusoid spans a full period across the polygon, so 84% of optima lie on the boundary but only 45% at a vertex — for the majority the optimum is interior to an edge, and the reference labels are globally optimal, verified by multi-start SLSQP and a 10^5-point brute-force sweep); **(iii)**

Non-compact / Equality (5) — Orthant-4D (unbounded non-negative orthant), Simplex-QP-8D ($\mathbf{1}^\top y = 1, y \geq 0$), QP-eq-20D (linear equality + box: DC3 canonical benchmark), SOCP-eq-16D (equality + SOCP cones + box), QCQP-eq-16D (equality + quadratic inequalities + box); **(iv) AC-OPF, legacy surrogate rows** (3) — IEEE 30, 57, and 118-bus optimal power flow; in these rows the score is a coordinate surrogate and the violation metric a box-only residual in normalised coordinates, so they do not evidence full AC feasibility (App. B.3); the current AC-OPF evidence is App. B.5. For each problem we report optimality gap $\Delta := \mathbb{E}_x[f(y(x;\theta);x) - f^\star(x)]$ relative to a reference solver (SLSQP / `cvxpy` for non-AC problems, `pypower`'s OPF for AC-OPF). On the three AC-OPF rows of Tab. 7, f is the normalised-coordinate surrogate $\frac{1}{2}\|y - y^\star\|^2$ that all methods are trained against, not the generation cost; the certified generation-cost results are in Sec. 5.3; no timing is reported for the legacy pipeline. We also report violation rate $\mathbb{V} := \mathbb{P}_x[\max_i\{g_i(y;x), |h_j(y;x)|\} > 10^{-5}]$ (strict equality and inequality combined). On the twelve non-AC problems the gap is measured under the problem's *own* objective — for Sin-QP that is $\frac{1}{2}y^\top Qy + 30\,p^\top \sin y$, not a linearisation — and against the same reference for every method. A negative gap means a method beat that reference's value. Where the reference is not known to be a certified global optimum, the sign alone does not establish infeasibility, so we read it together with the violation rate. On Sin-QP, for instance, DC3 reaches -17.1 while violating 81% of instances, and NN-SL-strong -8.8 at 68%; the sinusoid rewards precisely the exits that the constraint forbids. All cells are averaged over 5 seeds; we report mean $\pm$ std.

Baselines. We compare PRISM against six baselines: HOP (Deng et al., 2025) (single-chart ray baseline, $M = 1$), DC3 (Donti et al., 2021), NN-SL-strong (supervised with strong violation penalty), NN-SL-MSE (pure label match), HOMEOPROJ (Liang et al., 2023), and OPTNET (Amos & Kolter, 2017). Each cell is averaged over 5 seeds; we report mean $\pm$ std.

D.2 WHY INTERIORS ARE RELATIVE

Under a linear equality $Ay = b$ the interior of $\mathcal{Y}_x$ in the ambient $\mathbb{R}^d$ is empty, so reading int ambiently would make Def. 3.1 and Def. 3.2 vacuous on every equality benchmark we report — not merely inapplicable. An ambient direction would also leave $\mathrm{aff}(\mathcal{Y}_x)$ immediately, so the ray slice would be a single point. Taking both relative to $\mathrm{aff}(\mathcal{Y}_x)$, with directions drawn from $V_x = \ker A$, is the mathematical convention of Sec. 3.1: N_V projects the raw direction onto V_x and normalises it, using the fixed unit vector at zero. This is the read-out used by Thm. 4.1 in the equality case; numerical branch handling is a separate implementation question. Without equality constraints $V_x = \mathbb{R}^d$ and everything reduces to the usual reading.

D.3 HOW FEASIBILITY IS OBTAINED, ACROSS THE CLOSEST WORK

Detail moved from Sec. 2. Post-prediction repair — DC3's fixed correction steps (Donti et al., 2021), Bisection Projection's bisection towards a learned low-eccentricity interior point on general compact sets (Liang & Chen, 2025), FSNET's violation minimisation to convergence under smoothness and Polyak–Łojasiewicz conditions (Nguyen & Donti, 2025), and the linearised, locally quadratic or KKT projections of ENFORCE (Lastrucci & Schweidtmann, 2025), NLPOPT-NET (Roy et al., 2026) and KKT-HardNet (Iftakher et al., 2026) — makes hard feasibility attainable for general nonlinear constraints at the price of an inner feasibility computation at every inference. Gauge maps and HOMEOPROJ require sets homeomorphic to a ball, which excludes our disk-union, annulus, unbounded and equality-constrained problems; the $L_{0.5}$-ball is subtler — homeomorphic but not diffeomorphic to a ball, while an invertible neural network (INN) is a diffeomorphism — and the fit leaves 2.3% of mapped points outside the set (App. D.11).

D.4 FULL PER-PROBLEM RESULTS

Fig. 2(b) in the main text carries the claim; this is the record behind it — the aggregate scorecard (Tab. 6), then every cell, both metrics, five seeds.

Table 5: **How feasibility is obtained, across the closest work.** Columns are read from each paper's own method description; none of them reports performance for methods we did not run. "Online feasibility search" means the forward pass contains an inner loop or solve whose job is to reach feasibility. "Geometry represented" asks what object the method makes explicit about the feasible set; "capacity measured" asks whether the amount of that structure is inferred from data rather than fixed by hand. The single-chart ray baseline HOP is listed too: its distance to PRISM is exactly the last two columns.

Method	Feasibility source	Online feas. search	Capacity measured	Geometry represented
Bisection Proj. (Liang & Chen, 2025)	interior point + bisection	✓	✗	interior-point eccentricity
FSNET (Nguyen & Donti, 2025)	violation minimisation	✓	✗	—
NLPOPT-NET (Roy et al., 2026)	local constrained projection	✓	✗	local quadratic model
ENFORCE (Lastrucci & Schweidtmann, 2025)	adaptive-depth projection	✓	✗	local constraint manifold
HardNet (Min et al., 2025)	enforcement layer	partly	✗	convex feasible set
RAYEN (Tordesillas et al., 2026)	ray parametrisation	✗	✗	first ray interval
HOP (Deng et al., 2025) (single chart)	ray parametrisation	✗	✗	one fixed chart, $(A, K)=(1, 1)$
PRISM (ours)	ray-interval coordinates	✗[†]	✓	finite crossing cover (A, K)

[†]In the emission layer. On parametric nonlinear problems, acquiring a per-instance interior anchor is a separate online step whose cost is not included in the reported timings.

Table 6: **Breadth scorecard.** Feasible = the row's reported violation criterion ($< 5\%$); best-feasible = feasible *and* lowest gap among feasible methods. Denominators differ because methods differ in what they can express: the last column is how many of the fifteen problems each method structurally supports (for HOP: was run on). Three of the fifteen are legacy AC-OPF tasks scored by a coordinate surrogate with a box-only residual metric, not generation cost or full AC feasibility (App. D.1); two of PRISM's twelve best-feasible results are those surrogate tasks.

Method	Feasible	Best-feasible	Applicable
PRISM (ours)	**15/15**	**12/15**	15/15
HOP (single chart)	10/15	1/15	10/15
OPTNET	2/3	2/3	3/15
HOMEOPROJ	4/7	0/7	7/15
NN-SL-strong	3/15	0/15	15/15
DC3	0/15	0/15	15/15

D.5 CAPABILITY MATRIX

Tab. 8 records, for the methods we ran, which benchmark classes each one structurally supports; it is the source of the applicability column of Tab. 6.

D.6 VIOLATION RATES

Per-seed violation rates behind Tab. 7 and Tab. 9 will ship with the code release (upon publication) (`experiments/megatable.csv`, `experiments/static_suite_results.csv`); the seed means are in those tables and in Tab. 11. PRISM is 0% on every cell of both tables (for the three legacy AC-OPF columns under the box-residual metric of App. B.3); underlined entries mark the methods above the 5% threshold.

How the threshold is applied. The 5% feasibility threshold is applied to the *mean* violation rate over the five seeds, which is the usual aggregation but is not the worst case. Two cells are called feasible on the mean while an individual seed exceeds 5%, and both belong to NN-SL-strong: the unbounded orthant (mean 3.5%, worst seed 8.5%) and AC-OPF-57 (mean 2.8%, worst seed 5.5%). We keep the mean-based rule rather than switching to the worst seed, because switching would move both cells into the infeasible column and so would flatter the comparison; readers who prefer the strict reading should treat NN-SL-strong as feasible on $1/15$ rather than $3/15$. The distinction never touches PRISM, whose violation rate is exactly 0 on every seed of every cell — the property is per sample, not per average.

Table 7: **Megatable** on the 15-problem benchmark (5 seeds, mean±std). Each cell: **line 1** = optimality gap val_gap; **line 2** = $\mathrm{vio_rate} \,/\, \mathrm{max_vio}$. **Bold** = lowest val_gap among methods with vio$< 5\%$ ("feasible & best"). Underlined = infeasible (vio$\geq 5\%$). "n/a" = method does not apply (Tab. 8). The three AC-OPF columns are *legacy surrogate rows*: scored by the normalised-coordinate surrogate $\frac{1}{2}\|y - y^\star\|^2$ that all methods train against, with a box-only residual as violation metric; they evidence neither generation cost nor full AC feasibility (App. D.1, App. B.3). PRISM's row is selected among $M \geq 2$ routing configurations on a train-tail slice (App. D.10), so the $M{=}1$ HOP row is a genuine competitor; the three equality cells additionally take the retrospective better of the two reported anchor descriptors (App. D.14), and the SOCP-100D cell uses the numerical active-manifold configuration of App. D.13, developed after the ray results and outside the train-tail rule (the ray configurations gave 16.05–16.17; the $M{=}4$ row shows one). PRISM is bold on $12/15$; the three losses are OPTNET's exact convex layer on Polygon-2D and on SOCP-100D (3.47 vs 3.15), and a near-tie with the single-chart HOP baseline on AC-OPF-57. This is the audit-level record; the claim-level summary is Fig. 2(b) and Tab. 6.

(a) Convex and non-convex feasible sets

Method	Polygon 2D	QCQP 100D	SOCP 100D	TwoDisks	Annulus	Lp-ball 8D	Sin-QP
PRISM (ours)	+0.0040±0.00 0% / 0	**+34.67±0.19** 0% / 0	+3.47±0.07 0% / 3.8e-06	**+0.537±0.01** 0% / 1.2e-06	**+0.121±0.00** 0% / 0	**+0.364±0.00** 0% / 0	**+0.771±0.10** 0% / 0
PRISM (M=4, fixed)	+0.0064±0.00 0% / 0	+34.67±0.19 0% / 0	+16.17±0.12 0% / 0	+0.503±0.04 0% / 1.4e-06	+0.140±0.01 0% / 0	+0.365±0.00 0% / 0	+0.813±0.06 0% / 0
HoP (single chart)	+0.022±0.00 0% / 0	+34.97±0.37 0% / 0	+16.09±0.14 0% / 0	+0.820±0.10 0% / 1.3e-06	+0.321±0.02 0% / 0	+0.372±0.00 0% / 0	+1.33±0.11 0% / 0
DC3	-0.016±0.00 16% / 1.2e-01	-0.562±1.20 99% / 88.46	-203.82±2.63 100% / 1.7e+02	-0.930±0.01 88% / 1.68	-0.348±0.00 80% / 1.04	+0.484±0.06 100% / 17.29	-17.14±0.00 81% / 6.3e-01
NN-SL (strong-vio)	+0.0016±0.00 8% / 1.0e-02	+71.92±3.52 9% / 12.02	+18.68±0.83 9% / 8.4e-01	+2.84±0.03 22% / 9.4e-02	+0.628±0.18 7% / 3.4e-01	+0.415±0.02 16% / 4.6e-01	-8.81±0.12 68% / 3.6e-01
NN-SL (MSE)	-0.0052±0.00 11% / 6.5e-02	+13.56±0.20 88% / 55.71	+2.19±0.08 100% / 6.21	-0.324±0.03 69% / 1.25	-0.139±0.03 78% / 7.4e-01	-0.144±0.02 100% / 10.35	-0.581±0.17 63% / 3.5e-02
HomeoProj	+0.0036±0.00 0% / 0	+153.81±0.46 0% / 0	+33.37±0.04 0% / 0	*n/a*	*n/a*	*n/a*	+7.66±0.27 0% / 0
OptNet	**+0.0000±0.00** 0% / 0	+57.92±1.06 13% / 4.2e-03	**+3.15±0.05** 0% / 8.0e-06	*n/a*	*n/a*	*n/a*	*n/a*

(b) Non-compact and equality-constrained problems

Method	Orthant 4D	Simplex QP 8D	QP-eq 20D	SOCP-eq 16D	QCQP-eq 16D
PRISM (ours)	**+0.014±0.00** 0% / 0	**+0.067±0.00** 0% / 0	**+3.21±0.04** 0% / 0	**+1.12±0.04** 0% / 5.1e-07	**+2.19±0.05** 0% / 0
HoP (single chart)	—	—	—	—	—
DC3	+0.0018±0.00 33% / 0	+0.061±0.01 37% / 0	+4.15±0.18 66% / 1.09	+1.09±0.01 64% / 3.2e-01	+1.27±0.04 7% / 1.3e-01
NN-SL (strong-vio)	+0.040±0.01 3% / 0	+0.478±0.01 100% / 0	+2.44±0.07 100% / 2.3e-01	+3.17±0.03 100% / 5.5e-02	+0.954±0.02 100% / 1.7e-01
NN-SL (MSE)	-0.653±0.02 94% / 0	*n/a*	*n/a*	*n/a*	*n/a*
HomeoProj	*n/a*	*n/a*	+1.98±0.06 100% / 2.3e-01	+1.50±0.05 100% / 5.4e-01	+0.792±0.03 100% / 1.6e-01

(c) Legacy IEEE AC-OPF surrogate rows (case 30 / 57 / 118; coordinate surrogate, box-only residual)

Method	AC-OPF 30	AC-OPF 57	AC-OPF 118
PRISM (ours)	**+0.077±0.00** 0% / 0	+0.0064±0.00 0% / 0	**+0.012±0.00** 0% / 0
PRISM (M=4, fixed)	+0.079±0.00 0% / 0	+0.026±0.01 0% / 0	+0.012±0.00 0% / 0
HoP (single chart)	+0.079±0.00 0% / 0	**+0.0061±0.00** 0% / 0	+0.016±0.00 0% / 0
DC3	+0.076±0.01 8% / 1.8e-02	+0.0052±0.00 16% / 2.7e-03	+0.028±0.00 100% / 4.1e-03
NN-SL (strong-vio)	+0.083±0.00 2% / 4.5e-03	+0.0072±0.00 3% / 5.2e-04	+0.043±0.00 88% / 2.0e-03
NN-SL (MSE)	+0.077±0.00 7% / 2.2e-02	+0.0052±0.00 14% / 3.0e-03	+0.028±0.00 99% / 2.1e-03

Table 8: **Capability matrix, among the methods evaluated in our benchmark.** Of these, PRISM is the only one whose assumptions hold on every benchmark class. The table reports capabilities we exercised, not a survey of the literature. "Hard feasible" denotes *structural* feasibility, which holds on the exact- and certified-oracle classes; the legacy AC-OPF rows record a box-only residual and are not covered by it, while the certified case30 path mode carries its own conditional certificate (App. B.5). "not run": the single-chart baseline was not evaluated on the five equality/unbounded rows ("—" in Tab. 7), so no capability is claimed there; this is the applicability count of Tab. 6.

	Hard feasible	Convex only?	Disconn. sets	Non-simply connected	Non-compact	Equality constraint	Online speed
DC3	✗	no req.	ok	ok	ok	via comp.	fast
NN-SL	✗	no req.	ok	ok	ok	ok	fast
OPTNET	✓	✓	✗	✗	partial	via KKT	slow
HOMEOPROJ	✓	no req.	✗	✗	✗	✗	medium
HOP (single chart)	✓	no req.	✗	✗	not run	not run	fast
PRISM (ours)	✓	no req.	✓	✓	✓	✓	fast

Table 9: **What the measured budget buys.** Ten fixed geometries, 3 seeds, mean val_gap (lower better). The top row is the configuration the measurement pass *returns*; the two rows below are single-anchor controls at the same training recipe; the PRISM row is the retrospective per-geometry minimum of the two reported descriptor means (both in App. D.14; separate from the train-tail selection of App. D.10). PRISM violates on no cell; the penalty baselines violate on 7–100% of samples and are reported in App. D.6. Ratios to the measured row are plotted in Fig. 2(a).

measured $A \times K$	**Disks-2** 2×1	**Disks-4** 4×1	**Disks-8** 8×1	**Disks-16** 16×2	**Corners** 4×2	**Balls 3D** 8×2	**Balls 4D** 8×2	**Rings-4** 8×2	**Rings-8** 22×2	**Spiral** 11×1
PRISM (measured)	**+0.036**	**+0.037**	**+0.061**	**+0.10**	**+0.038**	**+0.14**	**+0.077**	**+0.26**	**+0.72**	**+0.037**
single-anchor 1×2	+1.83	+5.24	+4.28	+10.14	+1.54	+8.06	+4.24	+6.49	+32.83	+4.92
single-anchor 1×1	+3.57	+18.70	+27.08	+79.80	+13.23	+26.14	+15.66	+13.70	+49.16	+4.94

Tab. 11 gives the per-cell violation rates behind Tabs. 9 and 10 (seed means; source file `experiments/static_suite_results.csv`): PRISM is 0% on all 30 runs, DC3 7–100%, NN-SL-strong 43–100%, NN-SL-MSE 12–37%.

D.7 INFERENCE TIMING ON THE NON-CONVEX TOYS

Tab. 12 reports per-instance inference wall-clock for PRISM (batched forward, $M=4$, CPU) against per-instance SLSQP (`scipy`, warm-started at the anchor; for the disk-union we solve one SLSQP per disk and take the better, which favours the baseline), 200 held-out instances per problem. PRISM's cost is a single MLP forward plus the closed-form intersection oracle, and is essentially geometry-independent; the solver cost grows with the difficulty of the non-convex geometry. No timing is reported for the legacy AC-OPF pipeline; the certified mode's conditional cost is in App. B.7.

D.8 PER-SEED BREAKDOWN

Tab. 7 reports mean ± std over 5 seeds $\{42, 123, 456, 789, 1011\}$; the per-seed values are in the result file `experiments/megatable.csv` of the code release (upon publication). Tab. 7 displays a non-zero standard deviation on seven PRISM cells: QCQP-100D (± 0.19), Sin-QP (± 0.10), SOCP-100D (± 0.07), QCQP-eq-16D (± 0.05), QP-eq-20D (± 0.04), SOCP-eq-16D (± 0.04) and TwoDisks (± 0.01). The remaining cells show ± 0.00 at the table's two-decimal precision, which for the small-gap rows (Polygon-2D, the three AC-OPF rows) does not resolve their relative spread; the per-seed file gives the exact values. PRISM meets the reported violation criterion on every seed of all 15 rows: structurally on the exact-oracle problems; on the three legacy AC-OPF rows, the recorded zero is the box-only residual of App. B.3, not full-AC feasibility.

Table 10: **Static suite, penalty baselines.** Same ten geometries and recipe as Tab. 9. Every entry is infeasible (violation rate $\geq 5\%$); negative gaps are what leaving the feasible set buys.

	Disks-2	Disks-4	Disks-8	Disks-16	Corners	Balls 3D	Balls 4D	Rings-4	Rings-8	Spiral
DC3	-0.86	-0.38	+5.98	+28.84	+5.35	+11.22	+0.91	-0.001	-0.009	-0.031
NN-SL-strong	+3.62	+7.50	+13.95	+34.32	+11.06	+15.36	+8.30	+0.37	+1.16	+1.61
NN-SL-MSE	+0.027	-0.10	-0.12	-0.30	-0.036	-0.28	-0.20	-0.023	-0.069	-0.004

Table 11: **Static suite, violation rates** (% of validation samples with $\max_i g_i > 10^{-5}$; mean over three seeds).

	Disks-2	Disks-4	Disks-8	Disks-16	Corners	Balls 3D	Balls 4D	Rings-4	Rings-8	Spiral
PRISM (measured)	0	0	0	0	0	0	0	0	0	0
DC3	28	24	52	100	55	93	37	7	10	16
NN-SL-strong	98	93	100	100	62	99	99	43	53	54
NN-SL-MSE	12	22	21	24	21	29	26	21	37	13

D.9 PRISM FAMILY ABLATION: M AND HYPERPARAMETERS

The headline "PRISM (ours)" row of Tab. 7 reports, per problem, the variant selected on the train-tail slice of App. D.10 from the PRISM family $\{M = 2, 4, 8\} \times \{\lambda_{\text{aux}} = 10, 100\}$ at a 256 or 512 trunk, all at the same 40-epoch budget as the baselines. The $M{=}1$ configurations are *not* in this pool — they are the HOP baseline, and they get their own row under the same best-of-family rule, so the selection freedom is symmetric. The per-variant val_gap is reported in the second row, "PRISM (M=4, fixed)". Tab. 13 shows the full M-ablation across all problems.

A radial dead zone we removed. Our implementation also offers a *boundary margin* β_b, reparametrising the radial coordinate as $\rho_{\text{scale}} = \text{clamp}\big(\sigma(\varphi)(1 + 2\beta_b) - \beta_b,\, 0,\, 1\big)$ so that φ could saturate exactly onto a segment endpoint. It is a trap. For $\sigma(\varphi) < \beta_b/(1 + 2\beta_b)$ the clamp is active and the gradient with respect to φ is identically zero, so a radial coordinate that drifts into that region can never leave: $\rho \equiv 0$ and the model emits the *anchor* on every input. We measured it — at $\beta_b = 0.05$, 5 of 183 runs collapsed, each returning a bit-identical degenerate val_gap (1.8197 on AC-OPF-57, which is exactly the gap of predicting the anchor everywhere; 3 of 5 seeds there). A controlled sweep showed the margin buys $\sim$13% at $\lambda_{\text{aux}} = 10$ and nothing at $\lambda_{\text{aux}} = 100$, so we set $\beta_b = 0$ everywhere and report only $\beta_b = 0$ runs. The episode is instructive for the paper's own thesis: the collapsed models were observed to remain feasible throughout — they emit the anchor, which is interior — while their optimality was destroyed. Structural feasibility and trainability are independent properties, and only the first is something a parametrisation can guarantee, and then only when the oracle itself is exact or certified.

Observations. The optimal M varies by problem: on AC-OPF case-30 all configurations lie within one standard deviation of the $M{=}1$ column, $M{=}4$ helps on intermediate cases, and $M{=}8$ gives the best multi-head result on QCQP-100D and SOCP-100D. The $\text{v}2$ configuration (smaller λ_{aux}) wins on problems where the optimum lies precisely on the feasible boundary (Polygon, $L_{0.5}$-ball, Annulus, AC-OPF-30/118); the $\text{v}3$ configuration (strong λ_{aux}) wins where the oracle direction is informative (TwoDisks). Sin-QP is the sharpest case for the direction loss: decomposing the excess objective along the gated ray into a *routing* part (wrong direction) and a *saturation* part (right ray, wrong radius), routing accounts for 55–89% of it in every configuration. Raising λ_{aux} from 10 to 100 cuts the median angle between the gated direction and the true one from 8.7° to 1.6° and the gap from 1.887 to 0.814 (Tab. 13). The single robust choice across the suite is $\text{v}3$ with $M{\in}\{4, 8\}$; the best $M{\geq}2$ configuration beats the better $M{=}1$ column on nine of ten problems, the exception being AC-OPF-57, where the $M{=}1$, $\lambda_{\text{aux}}{=}100$ configuration is best (0.0061 vs 0.0066).

Table 12: **Per-instance inference time** (μs, CPU, batch 64) on the non-convex toy problems, 200 instances.

Problem	PRISM	SLSQP	Speedup
TwoDisks	67 μs	3,978 μs	59×
Annulus	60 μs	10,189 μs	170×
$L_{0.5}$-ball-8D	32 μs	5,303 μs	168×

Table 13: **M-ablation.** Mean val_gap over 5 seeds for each PRISM variant. Best per row in **bold**. The `v2` variant uses a 512×512 MLP with $\lambda_{\text{aux}} = 10$; `v3` is identical but with $\lambda_{\text{aux}} = 100$. $M \leq 8$ columns use the baselines' 256×256 trunk. All entries are 40-epoch, $\beta_b = 0$ runs.

Problem	M=1 λ=0	M=1 λ=100	M=2	M=4	M=4 v2	M=4 v3	M=8
Polygon-2D	0.0411	0.0219	0.0077	0.0084	**0.0038**	0.0064	0.0079
QCQP-100D	40.28	34.97	—	—	39.06	34.67	**34.59**
SOCP-100D	24.95	16.09	—	—	20.93	16.17	**16.06**
TwoDisks	1.649	0.820	0.573	0.558	0.912	**0.503**	0.537
Annulus	0.321	0.335	0.160	0.162	**0.121**	0.140	0.157
$L_{0.5}$-ball 8D	0.450	0.372	—	—	**0.3645**	0.3650	0.375
Sin-QP	4.929	1.329	0.685	**0.595**	1.887	0.814	0.771
AC-OPF case-30	0.0789	0.0840	—	—	**0.0766**	0.0792	0.0942
AC-OPF case-57	1.457	**0.0061**	—	—	0.0141	0.0260	0.0066
AC-OPF case-118	11.80	0.0162	—	—	**0.0108**	0.0117	—

D.10 CHOOSING THE REPORTED VARIANT

PRISM has a small configuration family (routing budget $M \in \{2, 4, 8\}$, auxiliary weight $\lambda_{\text{aux}} \in \{10, 100\}$, trunk width $\in \{256, 512\}$), and Tab. 7 reports the best member per problem — ordinary model selection. The question is what data does the choosing.

Our pipeline builds a single 80/20 split (`data/dataloader.py`, fixed generator seed), so there is a training split and a validation split and nothing else. Taking an argmin over nine candidates on the validation split and then quoting that same split's number would let the winner's luck into the reported result: with nine draws, the best of them is optimistic even if every candidate is equally good. The M=1 family has two candidates and is barely affected; the nine-candidate family is.

We therefore carve the missing third split out of the training data. Let π be the pipeline's permutation; the reporting split is $\pi[0.8N :]$ and is never touched here. Within $\pi[: 0.8N]$ we hold back the last 15% as a *selection slice*, train every candidate on the remainder, and rank them on that slice. The winner's ordinary 5-seed `val_gap` — computed the usual way, on the untouched 20% — is what the table prints. Both families are selected this way, so the freedom is symmetric. This is the same discipline the measurement pass already applies to its capacity probe (Sec. 3.3), which scores candidate configurations on a train-tail slice for exactly this reason.

The correction is real but small. Honest selection picks a *different* variant on 7 of the 10 parametric problems, and raises PRISM's reported gap by $+0.031$ on average — concentrated almost entirely in Sin-QP ($0.595 \to 0.771$, $+30\%$), with $\leq +0.08$ everywhere else and exactly 0 on three problems where the same variant wins either way. The M=1 family is unchanged on all ten. **No dominance cell changes**: every problem PRISM wins under validation-set selection it still wins under train-tail selection, and every problem it loses it still loses. We report the train-tail numbers throughout for this M/width/auxiliary-weight choice. The later per-problem choice between the two anchor descriptors (App. D.14) is a separate, retrospective minimum over two reported means and was not run through this selection slice. The SOCP-100D active-manifold configuration (App. D.13) lies outside the ray family and was added after these selections; it was not run through the selection slice either, and its development history is reported there.

D.11 GIVING HOMEOPROJ A FAIR SHOT ON THE CUSPED L_p-BALL

A first attempt to run HOMEOPROJ on the $L_{0.5}$-ball crashed, and a crash is not evidence about geometry. We therefore did the work the cell deserved, and report what happened.

(1) There was no mapping. No pretrained INN for this problem existed — the training directories contained loss records but no saved model. We added L_p-norm support to the generic INN pretrainer and fitted one.

(2) The fit is numerically obstructed, as the geometry predicts. At the default schedule the loss diverges to `NaN` within a few hundred steps. This is the cusp made numerical: $\partial|y|^{1/2}/\partial y$ is unbounded at the axes, so the feasibility penalty has unbounded gradient exactly on the boundary the INN must learn. With a $10\times$ smaller step and gradient clipping the run survives longer, reaches its best fit at 2.3% of mapped points landing *outside* the set, and then diverges anyway.

(3) Two bugs in the baseline's own code were in the way, and we fixed them. Its condition-tensor extraction returned the raw constraint parameter whenever a key named `b` was present; for this problem `b` is a 0-dimensional scalar that collates to shape $[\mathrm{batch}]$, so an entire minibatch was read as one wide condition vector, producing the shape error that had been recorded as "crashes on dimension mismatch". Fixed, and verified not to disturb the vector-valued case (Polygon-2D reproduces to 10^{-4}).

(4) With all of that, it runs — and collapses. Using the best finite checkpoint, HOMEOPROJ trains to completion and emits a *constant*: gap 0.6221 on all three seeds, bit-identical, where predicting the origin on every input scores 0.6213. We therefore do not report 0.622 as a performance number; the honest entry for this cell is that the method is applicable in principle, that its bijection cannot be fitted to a cusped set, and that the resulting model does not train. PRISM reaches 0.3645 here, but the comparison worth making is structural rather than numerical: star-shapedness about an interior anchor is preserved by cusps, diffeomorphism to a ball is not.

D.12 STRAIGHT-THROUGH ROUTING: CLOSING THE SOFT-MIXING GAP

During training the joint gate is a soft Gumbel-softmax, so the training-time forward $y = \sum_j g_j y_j$ is a convex combination of feasible candidates and can leave a non-convex set (Sec. 3.2). Two questions follow: how often does that actually happen, and what does removing it cost?

How often the soft forward is infeasible. Instrumenting the standard (soft) training run and evaluating a train-mode forward on the validation inputs after every epoch, the mixed output violates the constraints on 18% of samples on average on Disks-8 (peak 51% across epochs), 45% on Rings-8 (peak 63%), and 17% on the spiral (peak 88%); three seeds each. The caveat of Sec. 3.2 is real, not hypothetical. The hard-routed forward of the *same* checkpoints is 0% throughout on every closed-form oracle; on the spiral's marched oracle a single mid-training epoch peaked at 0.25% (final models: 0%) — the marching tolerance, not the parametrisation. All reported metrics in this paper use hard routing.

What straight-through costs. Straight-through routing (`gumbel_softmax(hard=True)`: the forward commits to one feasible candidate, the backward uses the soft distribution) makes every training-time forward feasible *whenever the underlying oracle is exact or certified*, with no other change to architecture or schedule. For approximate or marched oracles it removes the soft-mixing failure mode but not oracle numerical error — which is exactly the residual visible above on the spiral. Tab. 14 compares final quality on the eight non-convex benchmarks where mixing can matter, three seeds each. On six of eight, ST is within one standard deviation of soft routing; on TwoDisks and Rings-8 it costs 11% and 43% in gap — still $3\times$ and $24\times$ ahead of the best feasible alternative on those geometries. Practitioners who need the training-forward guarantee (e.g. safe online exploration) can switch it on with one flag; we keep soft routing as the default and the hard guarantee at candidate and deployment level.

Table 14: **Straight-through vs. soft routing** (optimality gap, mean±std over 3 seeds; soft column over 3–6 seeds). Both columns report the hard-routed evaluation, whose violation is strictly 0% on every cell; the difference between them is that straight-through additionally makes every *training-time* forward feasible under an exact or certified oracle, which soft routing does not (App. D.12).

Problem	Soft	ST	Problem	Soft	ST
TwoDisks	$0.432\pm.02$	$0.481\pm.02$	Disks-8	$0.068\pm.02$	$0.091\pm.04$
Annulus	$0.101\pm.02$	$0.126\pm.01$	Disks-16	$0.160\pm.08$	$0.152\pm.06$
QCQP-100D	$34.67\pm.19$	$34.64\pm.19$	Rings-8	$0.953\pm.06$	$1.365\pm.14$
$L_{0.5}$-ball	$0.363\pm.00$	$0.362\pm.00$	Spiral	$0.037\pm.01$	$0.047\pm.00$

D.13 Capacity-Matched Trunk

The main-table PRISM v2/v3 variants use a 512×512 trunk while every baseline uses 256×256 (Tab. 21); the $M\in\{2,4,8\}$ variants and the slot-indexed static-suite runs already use 256×256 (the coordinate static-suite trunks are narrower, 120–152 wide). To isolate what the extra width buys, Tab. 15 restricts PRISM to the baselines' 256×256 trunk on all ten parametric benchmarks, five seeds each. Feasibility is untouched — violation is strictly 0% on every cell — structural on the exact-oracle problems and box-residual-checked per sample on the legacy AC-OPF rows, but in neither case learned. Gaps move by less than one standard deviation on Polygon-2D, QCQP-100D and $L_{0.5}$-ball — on QCQP-100D the 256 configuration is in fact the better one, and on Sin-QP and AC-OPF-57 it is the best configuration we have — by 0.09 on SOCP-100D, and by 7–40% on the rest. Of the seven cells PRISM wins at 512, **five survive at matched capacity**: QCQP-100D, TwoDisks, Annulus, Sin-QP and AC-OPF-118. Two do not, both to a fairly-configured HOP and both by a hair — the $L_{0.5}$-ball (0.3746 vs 0.3719, 0.7%) and AC-OPF-30 (0.0885 vs 0.0789). Both are geometries where a single anchor already sees the whole feasible set, so they are exactly the cells the (A, K) account says the routing should not be winning on merit; at 512 the extra width, not the routing, was carrying them. The cells where routing is structurally necessary — the disconnected, the holed, the multi-basin — are untouched by trunk width, and no feasibility conclusion depends on it at all.

SOCP-100D: ray plateau and the numerical active-manifold configuration. SOCP-100D is the convex cell where every ray configuration stops far from the exact layer: the reported ray row 16.06, the $M=4$ v3 configuration 16.17 ± 0.12 over five seeds, and the selected HOP row 16.09 ± 0.14 (the $\lambda_{\text{aux}}=100$ single-chart member; the $\lambda_{\text{aux}}=0$ member gives 24.95), against 3.15 ± 0.05 for OPTNET and 28.1 for emitting the anchor. Two read-only observations on the stored data describe this plateau without bounding it. At a slack tolerance of 10^{-3}, every stored label has ≥ 19 of the 20 cones near-active (mean 19.8); we do not claim that all 20 are exactly active, nor infer a codimension. Rotating the exact anchor-to-label direction of 400 labels by 1°, 5° and 10° moves the ray exit to 0.88, 0.58 and 0.38 of $\|y^\star\|$ and costs gaps of 2.6, 10.8 and 16.6, and all 20 logged ray runs reach within 5% of their best validation gap by epoch 3–10 of 40. Exploratory prototype variants (deeper or per-instance anchors, re-anchoring hops, objective-selected heads, numeric-gradient hops) stayed between 10.5 and 12 on a contiguous split; these are the variants we tried, not an exhaustive search, and none of these observations is a lower bound.

The configuration reported in Tabs. 7 and 2 treats the near-active cones as a manifold M and routes on it numerically. A proposal anchor a_M is the training-label mean retracted onto M (measured once from the training split and stored with the checkpoint); the predicted direction is projected onto the tangent space of M at a_M and scaled by a predicted length ρ (softplus, unbounded); the point $a_M + \rho d$ is retracted onto M by ten Gauss–Newton minimum-norm steps on the 20 equalities (residual about 10^{-14} in float64). The retracted point is a *proposal*, not the output: the emitted point is the exact-oracle ray from the strictly interior centre a_{int} (the origin) toward the proposal, at the fixed fraction $1-10^{-6}$ of the exit distance (no learned radial logit in this configuration). For a finite, nonzero final ray direction and a strictly interior a_{int}, an exact exit oracle yields a feasible emission by the direct-ray argument. This special implementation has no zero-direction fallback, so no all-input total read-out guarantee is claimed for it. The small Gauss–Newton residual is not a certificate of a continuous feasible path, intermediate iterates need not be feasible, and this is not a second instance of the case30 certificate. Training uses the objective loss only, no auxiliary or violation term, otherwise the megatable recipe (40 epochs, Adam 10^{-3}, batch 64, the pipeline's $80/20$ split, best

epoch on the validation split as for every method). Five seeds give 3.47 ± 0.07 (512×512 trunk) and 3.56 ± 0.03 (256×256); no report sample exceeds the 10^{-5} violation criterion, while 1–2.5% of report samples show a positive float32 residual of at most $3.8\cdot10^{-6}$, so the exact-oracle guarantee and the measured zero rate are different statements. Development history: this configuration was designed after inspecting the ray results; its retraction depth, anchor, trunk widths and the absence of an auxiliary term were chosen on prototype runs whose validation samples overlap the pipeline's reporting split, and the near-active statistic used all 2,000 labels. The reported five-seed numbers are therefore development-exposed rather than a held-out confirmation, and were not run through the train-tail rule of App. D.10, which selects within the ray family only. Cost: one MLP forward, ten batched 20×20 solves and one oracle call, 3.0 ms per instance at batch 1 on two CPU threads (2.0 ms for the $M=4$ ray configuration in the same measurement) and 0.4 ms per instance at batch 400; OPTNET's recorded cost on this problem, 2.3–5.3 ms per sample, comes from a different run, so no speed ratio is claimed. A 60-epoch prototype run reached 3.09 and a 200-epoch cosine-schedule run 2.78; neither is used, because the baselines were trained for 40 epochs. We did not observe a constant near-active set on QCQP-100D (5.8 of 50 constraints on average, varying per instance), so the configuration was not applied there.

Table 15: **Capacity-matched control** (optimality gap, mean±std, 5 seeds). "PRISM 512" = best variant on a 512×512 trunk; "PRISM 256" = best variant restricted to the baselines' 256×256 trunk. Baseline = best gap among competing methods with vio $< 5\%$ (violation rate in parentheses where nonzero); PRISM is 0% on every cell. **Bold** = the matched-capacity PRISM beats the best feasible competitor. †Numerical active-manifold configuration (paragraph below); the ray configuration gives 16.17 / 16.05 at the 512 / 256 widths.

Problem	PRISM 512	PRISM 256	Best feasible competitor
Polygon-2D	0.0038	0.0040	OPTNET **0.0000**
QCQP-100D	34.67	**34.59**	HOP 34.97
SOCP-100D†	3.47	3.56	OPTNET **3.15**
TwoDisks	0.503	**0.537**	HOP 0.820
Annulus	0.121	**0.148**	HOP 0.321
$L_{0.5}$-ball 8D	0.3645	0.3746	HOP **0.3719**
Sin-QP	0.814	**0.595**	HOP 1.329
AC-OPF 30	0.0766	0.0885	HOP **0.0789**
AC-OPF 57	0.0141	0.0064	HOP **0.0061**
AC-OPF 118	0.0108	**0.0150**	HOP 0.0162

D.14 ANCHOR DESCRIPTOR, AND ANCHORS AS A RESOURCE

Two descriptors, one network. Eq. 4 lets the anchor enter the shared trunk either through its coordinates ($\delta_i = a_i$) or, for a fixed bank, through its slot index, realised as slot-indexed heads $H_{\theta,i}(F_\theta(E_x(x)))$. For $A=1$ with a fixed anchor the constant anchor code can be absorbed into the predictor's bias, so the slot-indexed implementation is the corresponding fixed-anchor specialisation (no claim of parameter-by-parameter equivalence at equal width is made); this covers nine of the fifteen problems of Tab. 7. On the static suite and the three equality-constrained tasks they differ, and Tab. 16 reports both under the same recipe, seeds, oracles and banks: the measured banks are re-derived deterministically and match the recorded A, K_0 and certificate on every cell (Rings-4 was recorded under the earlier anchor budget $k_{\max}=8$, which the re-derivation reproduces exactly). Widths of the coordinate trunk are fixed before any result so that the parameter count matches the slot-indexed network (0.94–$1.05\times$); a 256-wide coordinate trunk with 2.7–$4.7\times$ the parameters is reported as a sensitivity check only. Tabs. 9 and 7 report, per problem, the better of the two descriptors (coordinates on Disks-8, Disks-16, Balls 3D, Balls 4D, Rings-8 and SOCP-eq; slot-indexed elsewhere), as their captions state. This descriptor choice is a retrospective per-problem minimum of the two reported descriptor means (same seeds, same reporting split); no separate selection run was made, and it is a different step from the train-tail selection of M, width and auxiliary weight in App. D.10, which was applied to the slot-indexed configurations. Both descriptor means are listed in Tabs. 16 and 17.

Table 16: **Anchor descriptor on the static suite** (three seeds each, seed means of the validation gap; ratio = coordinate / slot-indexed). Coordinates win on the $A \geq 8$ banks except Rings-4 and lose on the $A \leq 4$ geometries and the spiral; the wider trunk does not remove the losses, so they are not a parameter-matching artefact. Both descriptors are best-feasible on $10/10$ at zero violation on every cell.

Geometry	$A \times K$	slot-indexed	coordinate	ratio	width	params ratio	coordinate, 256 wide
Disks-2	2×1	0.036	0.069	1.91	120	1.04	0.062 (4.7× params)
Disks-4	4×1	0.037	0.059	1.57	120	0.98	0.085 (4.4× params)
Disks-8	8×1	0.068	0.061	0.90	128	1.00	0.070 (3.9× params)
Disks-16	16×2	0.160	0.101	0.63	144	0.99	0.130 (3.1× params)
Corners	4×2	0.038	0.056	1.47	120	0.98	0.082 (4.4× params)
Balls 3D	8×2	0.149	0.140	0.94	136	1.05	0.119 (3.6× params)
Balls 4D	8×2	0.108	0.077	0.72	136	1.03	0.065 (3.6× params)
Rings-4	8×2	0.255	0.285	1.12	128	0.97	0.339 (3.8× params)
Rings-8	22×2	0.953	0.724	0.76	152	0.96	0.936 (2.7× params)
Spiral	11×1	0.037	0.082	2.21	128	0.96	0.068 (3.8× params)

Table 17: **Moving-anchor equality tasks, five seeds** (same seeds, data and recipe as the slot-indexed rows of Tab. 7; the anchor is the min-norm point of the per-sample slice, already a linear function of x, so the coordinate descriptor adds inductive bias rather than information).

Task	seeds	slot-indexed gap	coordinate gap	ratio	cells better	max. viol.
QP-eq-20D	5	3.2064	3.3888	1.057	0/5	2.4e-07
SOCP-eq-16D	5	1.2308	1.1238	0.913	5/5	5.1e-07
QCQP-eq-16D	5	2.1884	2.2921	1.047	0/5	4.8e-07

Anchors as a resource. With the coordinate descriptor the bank is an input, so a network trained on the measured bank can be run on any set of interior points. Tab. 18 evaluates the Tab. 16 checkpoints, without retraining, on banks of $\frac{1}{2}$–$8\times$ the measured size drawn uniformly from oracle-verified interior points of a fresh sample set (five draws each), and reports the gap together with the held-out coverage of the drawn bank at the cell's own visibility tolerance (an empirical fraction over 400 held-out targets, the measurement's certificate notion, not a formal path certificate). Reachability is necessary, not sufficient: a random bank of the measured size already reaches 78–100% of held-out targets, yet costs 27–$45\times$ in gap on the disk unions; at 4–$8\times$ the certificate is 0.99–1.00 everywhere and the gap returns to within 1.3–$3.5\times$ of the coordinate network's own measured-bank gap on six of the seven disk unions (Disks-2 excepted), and below it on Rings-8 and the spiral, whose measured banks are therefore not the best operating points on those two geometries. A replaced bank is a new bank, not an augmentation of the measured one, and the recovery at $4\times$ is not uniform across geometries (thinning, anchor-quality and anchor-randomised controls: Tab. 19). Latency is linear in A (batched, one A100): Disks-8 $63\,\mu$s at $A{=}8$ to $488\,\mu$s at $A{=}64$, Rings-8 179 to $1608\,\mu$s from $A{=}22$ to 176. The parameter count does not change with the bank (84k on Disks-8 for every A) for the coordinate-shared configuration only, whereas slot-indexed heads grow with it (84k to 137k) and must be retrained for each bank (Tab. 20); a larger bank still costs oracle calls and latency.

Violation audit. Across the 1065 (checkpoint, condition) records behind Tabs. 18–20, 62 have at least one draw above the 10^{-5} violation criterion of Tab. 9; the largest violation on any exact-oracle geometry is 4.7e-05, and no exact-oracle record exceeds 10^{-4}. They concentrate in low-coverage conditions (48 of the 62 in one-anchor, half-masked, sub-measured or quartile banks; 4 of the 174 records at 4–$8\times$ banks are also flagged). A diagnostic on the worst record (Disks-16, one anchor kept) attributes its excess to float32 rounding of the ray–boundary root when the router saturates $\sigma(\varphi)$ at a segment end and the emitted point sits on the boundary; this was checked on that record, not on every flagged draw. The spiral, whose oracle is a numerical march-and-bisection search, has 6 records with a single sample at up to 2.5e-04. Every measured-bank cell of Tab. 16 is at zero violation, as in Tab. 9.

Table 18: **One checkpoint, many banks.** Coordinate-descriptor networks trained on the measured bank, evaluated without retraining on random interior banks of 1–8× the measured size (five draws; seed means). Certificate = fraction of held-out targets reachable from the drawn bank at the cell's visibility tolerance, the measurement's own notion.

	gap					certificate				
Geometry	measured	1×	2×	4×	8×	measured	1×	2×	4×	8×
Disks-2	0.069	2.780	2.730	1.254	0.807	1.00	0.95	0.98	1.00	1.00
Disks-4	0.059	2.323	1.250	0.245	0.206	1.00	0.95	0.95	0.99	0.99
Disks-8	0.061	2.737	0.805	0.138	0.123	1.00	0.81	0.91	0.99	0.99
Disks-16	0.101	3.236	0.709	0.221	0.132	1.00	1.00	1.00	1.00	1.00
Corners	0.056	2.416	1.107	0.137	0.122	1.00	0.83	0.91	1.00	1.00
Balls 3D	0.140	3.790	1.397	0.514	0.228	1.00	1.00	1.00	1.00	1.00
Balls 4D	0.077	2.511	1.163	0.231	0.218	1.00	0.79	0.92	0.99	1.00
Rings-4	0.285	0.789	0.543	0.404	0.324	1.00	0.95	0.99	1.00	1.00
Rings-8	0.724	1.326	0.793	0.515	0.386	0.99	0.94	0.99	1.00	1.00
Spiral	0.082	0.506	0.178	0.074	0.061	0.99	0.78	0.94	1.00	1.00

Table 19: **Thinning, anchor quality and anchor-randomised training** (gap; seed means over five draws). Left: the measured bank with 25% / 50% of its anchors masked, a single anchor kept, and A_{meas} anchors drawn from the deepest or the shallowest interior quartile. Right: a variant trained with a fresh random bank of A_{meas}–$2A_{\text{meas}}$ anchors every batch (AR), evaluated on the measured bank and on 8× random banks; it generalises better to small random banks but trails the fixed-bank network on the measured bank by 1.1–2.9× and is not the headline.

Geometry	drop 25%	drop 50%	one anchor	deep quartile	shallow quartile	AR, measured	AR, 8×
Disks-2	0.07	2.64	2.64	3.60	2.22	0.153	0.243
Disks-4	1.77	6.65	13.92	5.24	3.72	0.172	0.193
Disks-8	1.73	5.11	26.56	3.21	2.73	0.147	0.185
Disks-16	1.38	4.69	76.29	2.63	4.13	0.234	0.230
Corners	1.86	4.47	9.74	3.96	3.00	0.155	0.169
Balls 3D	2.50	6.59	27.20	5.99	6.24	0.190	0.362
Balls 4D	1.38	3.66	15.47	2.39	3.29	0.088	0.172
Rings-4	0.68	1.65	9.35	0.69	1.02	0.489	0.131
Rings-8	1.13	2.10	44.49	1.42	1.53	0.976	0.329
Spiral	0.28	0.71	5.96	0.39	0.59	0.096	0.043

A refinement loop that was not retained. A state-conditioned variant $F_\theta([c, e_i, E_s(s_t - a_i)])$, warm-started from the matched trunk and evaluated as a frozen-weight accept/reject recurrence ($T = 1, 2, 4$) on QP-eq and SOCP-eq, improved its accepted best-so-far gap by 5–9% at T=2; two controls on the same checkpoint — zeroing the state input, or permuting it across instances — matched or exceeded that gain. The improvement is therefore attributable to the accept/reject wrapper over repeated proposals rather than to state semantics, and neither the state input nor any generic re-anchoring is part of PRISM. Re-anchoring appears only in the AC-OPF specialization (App. B.5).

D.15 ANCHOR PROVIDER INTERFACE

The provider G_{anchor} returns a valid bank and, per anchor, a provenance tag. *Speculative acquisition* composes a fast proposal with the task's strict validator (finiteness and strict interiority, e.g. $g(a) < -\varepsilon$ and $|Aa - b| \le \varepsilon_{\text{eq}}$) and an exact or trusted fallback; every returned anchor either passed the validator or came from the fallback, so validity is preserved whenever the fallback contract holds, while a miss pays for all three stages and no runtime or objective advantage follows. *Anchor augmentation* $\mathcal{A}' = \mathcal{A}_{\text{base}} \cup \mathcal{A}_{\text{spec}}$ retains every base anchor and its candidate slots; by set inclusion the reachable candidate family cannot shrink, invalid speculative anchors are masked before routing, and nothing follows about objective or runtime. Both were implemented as interfaces with deterministic contract

Table 20: **Retraining per bank versus one checkpoint.** Fixed random banks of 1–4× the measured size (one bank draw, three seeds): slot-indexed heads and the coordinate trunk retrained on that bank, against the AR checkpoint evaluated on unseen banks of the same size; the AR column is a five-draw mean over different random banks of that size, not a paired comparison on the same bank. On Rings-8 the shared trunk wins at large banks (0.44 vs 0.85 at 4×); on Disks-8 slot-indexed heads keep a small edge (0.12 vs 0.14).

Geometry	bank	slot, retrained	coord., retrained	coord. AR, no retraining	slot params	coord. params
Disks-8	1× random	3.413	3.867	2.146	84k	84k
Disks-8	2× random	1.568	1.904	0.795	102k	84k
Disks-8	4× random	0.117	0.143	0.202	137k	84k
Rings-8	1× random	1.329	1.114	1.128	123k	118k
Rings-8	2× random	1.070	0.671	0.619	181k	118k
Rings-8	4× random	0.850	0.444	0.414	296k	118k
Spiral	1× random	0.330	0.354	0.471	87k	84k
Spiral	2× random	0.060	0.064	0.154	109k	84k
Spiral	4× random	0.037	0.049	0.060	151k	84k

tests (fast hit, invalid proposal, empty or non-finite proposal, mask, provenance, order, base-slot retention); on the frozen Disks-8 bank a valid extra anchor left every legacy candidate slot unchanged and kept the hard output feasible, and on QP-eq a nearest-training-anchor proposal was rejected by the equality validator on every instance, so the exact min-norm fallback served all of them.

E Hyperparameters and Implementation

Tab. 21 lists the hyperparameter configuration for every method on every problem. We use a uniform configuration across problems for all PRISM variants (only M and the architectural flag varies).

Table 21: **Hyperparameters.** "LR" = learning rate (constant); "BS" = batch size; "MLP" = hidden sizes of the policy MLP (in addition to a final linear projection to the layout-specific output dimension).

Method	LR	BS	MLP	Epochs	Aux λ	Bound. margin
PRISM (default v3)	10^{-3}	64	512×512	40	100	0
PRISM (v2)	10^{-3}	64	512×512	40	10	0
PRISM ($M\in\{2,4,8\}$)	10^{-3}	64	256×256	40	100	0
PRISM (static, coordinates)	10^{-3}	64	$w\times w$, $120\le w\le 152$	60	10	0
PRISM (static, slot-indexed)	10^{-3}	64	256×256	60	10	0
HOP (M=1, best of $\lambda_{\text{aux}}\in\{0,100\}$)	10^{-3}	64	256×256	40	0/100	0
DC3	10^{-3}	64	256×256	40	—	$\lambda_{\text{vio}}=1$
NN-SL-strong	10^{-3}	64	256×256	40	—	$\lambda_{\text{vio}}=100$
NN-SL-MSE	10^{-3}	64	256×256	40	—	MSE-to-label only
HOMEOPROJ	10^{-3}	64	256×256	40	—	INN pre-trained
OPTNET (cvxpylayers QP)	10^{-3}	64	256×256	40	—	exact QP layer

SOCP-100D active-manifold configuration. Trunk 512×512 (256×256 for the width control), objective loss only (no auxiliary or violation term), Adam 10^{-3}, batch 64, 40 epochs, ten Gauss–Newton retraction steps, proposal anchor measured from the training split, emission by the exact ray from the interior centre at fraction $1-10^{-6}$; developed after the ray results and outside the train-tail selection (App. D.13).

Compute. The anchor service of the three legacy AC-OPF rows (App. B) used a single NVIDIA RTX A6000 (48 GB); all other training is CPU on a 16-core Intel i9 server; the static-suite descriptor comparison, the bank evaluations and their latencies (App. D.14) ran on one A100 80GB per job. No timing of the legacy AC-OPF pipeline is reported. The certified mode of App. B.5 was run on two CPU threads with a fixed budget of 8 update calls per trajectory and per-call interval-box reservations; its policies are the frozen networks named in App. B.7, and no hyperparameter was changed after the report split was sealed.